Multimedia Chips Complicate Choices

Accelerator Chips, Software Will Come Together in 1996–97

by Jim Turley

Over the past several issues, we have reviewed a number of new media accelerators. This article brings many of these chips together and compares their features, benefits, and futures.

Over the past six months, several companies, new and old, have announced chips that combine audio, video, and graphics acceleration, creating a new category of devices: media processors. Some are destined to be popular in PCs before the end of this year, offering an impressive array of multimedia capabilities at a cost lower than that of a set of discrete audio, video, and graphics chips. Driven by growing demand for multimedia PCs, media processors will eventually displace discrete chips in these systems.

The move to media processors is enabled by Microsoft APIs, particularly the new DirectX interfaces. As Windows 95 applications move to DirectX, it will become much easier to substitute new audio and video hardware while retaining compatibility with existing software. But older applications will remain on users' disks for the next few years, so media processors must offer compatibility with current register-level hardware standards, primarily super VGA and Sound Blaster, to run these applications.

This new class of media processors is coming from industry stalwarts IBM and Philips, as well as from startups Nvidia and Chromatic. Each has taken a different approach to delivering exciting multimedia performance on a PC, and each has its own benefits and drawbacks.

Microsoft Delivers Enabling Software

A typical consumer PC shipping today contains a Sound Blaster–compatible chip for audio, a graphics accelerator chip, possibly an MPEG-1 decoder chip, and a high-speed modem card. Applications written for Windows 3.1 generally use APIs such as GDI to access graphics functions, DCI for video, and the WAV interface for audio. Most game programs, however, avoid these APIs because of their poor performance. Another problem is that these older APIs don't support newer functions such as 3D graphics. Instead, most games run under DOS, avoiding the Windows APIs and accessing the audio and video chips directly.

Programs that do not use the Windows APIs must access the hardware through specific registers and hardware protocols. For audio, most chip vendors have adopted the register-level interface of the original Creative Labs Sound Blaster card, which was not designed to support more advanced audio functions. Graphics chips have a variety of register-level interfaces, so programs must be written to support specific chips (typically, several of the most common). As new graphics chips are developed, existing programs must be modified to support the new devices.

To help solve this problem, Microsoft has created a new set of APIs called DirectX. These APIs are considered part of the 32-bit Win32 (Windows 95) interface but didn't begin shipping until late last year, two months after the release of Windows 95 itself. The new operating system still supports 16-bit applications using APIs such as GDI, although the rarely used DCI is no longer supported.

Today, DirectX is a bundle of four APIs, with two more slated for a future release (*see* **100103.PDF**). These APIs add little performance overhead. Most software developers, including game vendors, are migrating to DirectX as they port their titles to Windows 95.

One problem is the late arrival of Direct3D, caused by Microsoft's false start with 3D DDI, the company's purchase of Rendermorphics, and its rejection of Intel's proposed 3DR interface. Game vendors that had already begun Windows 95 ports of software incorporating 3D effects have been forced to write code that accesses 3D chips directly, since Direct3D only recently became available. Vendors such as Creative Labs and Nvidia have encouraged this tactic in an attempt to establish their proprietary 3D interfaces as industry standards. Thus, the first generation of 3D games bypasses Direct3D, but the ability to support multiple hardware platforms is likely to attract software vendors to the Microsoft API for future versions.

Nvidia Has Early Start; Philips Has Market Muscle

The four media processors announced to date share many characteristics but differ in key ways, as Table 1 shows. Note that there is a span of nearly two years between the availability of Nvidia's NV1 and IBM's Mfast, with the others falling somewhere in between. Thus, Nvidia and Chromatic are likely to have significantly improved their devices by the time Mfast debuts.

Nvidia was the first vendor to unveil a multimedia accelerator (*see* **090904.PDF**); its NV1 processor is now in production. The chip combines a powerful audio DSP with 2D and 3D graphics acceleration. Nvidia has allied itself with SGS-Thomson in an unusual relationship: SGS builds all the parts and also has rights to sell the lower-performance DRAM version (the STG2000), while Nvidia sells the high-end VRAM part itself. By ceding a large portion of the volume to SGS, Nvidia ensures that its foundry will place a priority on building the parts, yet the startup has cleverly kept the higher-profit slice of the market, in addition to the royalties it receives on the parts SGS sells.

Chromatic's Mpact Media Engine (*see* **091404.PDF**) is not expected to ship until 2Q96, about nine months behind the Nvidia chip. Mpact is a more powerful part, however, adding video functions like MPEG-1 and -2 decoding, MPEG-1 encoding, and videophone capabilities to the basic audio, 2D, and 3D graphics delivered by Nvidia. The Chromatic part can also act as a 28.8-kbps fax/modem. Chromatic is a design-only firm, turning over manufacturing and marketing duties to Toshiba and LG Semicon. The company is also negotiating with a third, U.S.-based, vendor. Chromatic plans to write all drivers and software itself; the Mpact programming model will not be disclosed.

Philips has displayed its devotion to multimedia acceleration, devoting years to developing its media processor and forming a separate Trimedia division to build and market such products. The first device in this line, the TM-1 *(see* **091506.PDF***)*, is slated to debut in 4Q96, six months behind the Chromatic chip and too late for the 1996 holiday sales season. Philips says its device will perform a range of functions similar to Chromatic's, including audio, 3D graphics, MPEG decoding, and 28.8-kbps modem. One exception is that Philips has not yet discussed MPEG-1 encoding.

Leveraging its presence in the consumer electronics market, Philips designed the TM-1 to serve as the core processor in a set-top box as well. The company has the capability to manufacture and market the chip as a standalone device, in PC add-in cards, and as part of set-top boxes and other consumer devices, thereby guaranteeing a slew of internal design wins. Pricing will be a key issue when competing against Chromatic. Another differentiator is that Philips plans to make its programming model public and provide a C compiler, which lets other vendors write custom drivers and other software for the TM-1.

IBM's media processor targets a performance level that far surpasses that of the other devices: the Mfast chip *(see* **091601.PDF***)* can perform a host of operations, including the monumental task of MPEG-2 encoding. The chip is still under development, however, and will not appear in systems before 2H97, giving the competition ample time to rev their engines.

Audio/Graphics Engines Differ

Nvidia's NV1 includes two separate processors: a powerful audio DSP and a 2D/3D graphics accelerator. Unlike the other media processors, the chip uses a separate engine to handle graphics. This engine provides normal 2D graphics acceleration; Nvidia declines to provide benchmark numbers but claims that its chip matches the performance of high-end graphics accelerators in 8-bit (256-color) modes and surpasses them on 16-bit (16 million-color) modes. There are no provisions for MPEG-1 decoding in the chips, although a software decoder can be used.

Audio functions are handled by the 50-MHz DSP, which can achieve 350 million operations per second (MOPS). Such performance enables the chip to generate 32 channels of CD-quality sound concurrently, even if the sounds were created with different sample rates. It supports simultaneous input and output at sample rates up to 48 kHz. These capabilities require a high-end sound card today and would consume most of the cycles of a 100-MHz Pentium performing this function entirely in software. For legacy software, the chip also performs Sound Blaster emulation.

Unlike Nvidia, Chromatic has taken the approach of building a single powerful multimedia engine that performs both audio and video processing. The main part of the

	Nvidia NV1	Chromatic Mpact	Philips TM-1	IBM Mfast	Pentium-100
Vendor(s)	Nvidia, SGS-Thomson	Toshiba, LG Semicon	Philips	IBM	Intel
Wavetable audio	yes	yes	yes	yes (Mwave)	yes
Sound Blaster compatibility	yes	yes	no	yes (Mwave)	no
28.8-kbps fax/modem	no	yes	yes	yes (Mwave)	no
2D/3D graphics	yes	yes	yes	yes	yes
MPEG-1 decode	no	yes	yes	yes	yes
MPEG-1 encode	no	yes	yes	yes	no
MPEG-2 decode	no	yes	yes	yes	no
MPEG-2 encode	no	no	no	yes	no
H.261 videoconferencing	no	yes	yes	yes	no
Windows drivers	yes	yes	yes	yes	yes
PCI interface	yes	yes	yes	yes	n/a
Memory type	DRAM/VRAM	RDRAM	SDRAM	SDRAM	SDRAM
Memory bus width	32 bits/64 bits	8 bits	32 bits	128 bits	64 bits
Memory bandwidth	100–200 Mbytes/s	500 Mbytes/s	400 Mbytes/s	800 Mbytes/s	~100 Mbytes/s
Core clock speed	50 MHz	62.5 MHz	100 MHz	50 MHz	100 MHz
Core performance (8-bit ops)	350 MOPS*	2 BOPS†	3.8 BOPS‡	20 BOPS	200 MOPS
Relative manufacturing cost	low	low	medium	high	zero
Volume production	now	2Q96	4Q96	2Q97	now

Table 1. A comparison of the four major new media accelerators reveals major differences in the performance of their internal audio/graphics engines and in memory bandwidth. The Nvidia chip is the only one shipping, while IBM's Mfast is still more than a year away. *Audio DSP only; no figures for graphics engine. †Does not include motion-estimation engine. ‡Does not include VLD and image coprocessors.

62-MHz Mpact engine can carry out two billion operations per second (BOPS), nearly six times the performance of Nvidia's audio processor. (Nvidia does not rate the speed of its graphics engine in operations per second.) Mpact includes a separate motion-estimation unit, used for video compression, that Chromatic claims can generate an additional 20 BOPS, but the company has not provided details on this unit to qualify its claim; we suspect these operations are on smaller data types (a few bits each) and are not comparable to the BOPS claimed by other vendors.

Internally, the Philips TM-1 is much like Mpact, with a single main processing engine capable of up to 38 operations per cycle at 100 MHz, all on 8-bit quantities. The TM-1 adds two on-chip coprocessors for video preprocessing. One handles variable-length decoding (used in video decompression), and the other performs image scaling and color space (YUV to RGB) conversion. The 3.8-BOPS performance of the main engine does not include the performance of these coprocessors.

IBM's Mfast has by far the most outlandish architecture of the bunch. Mfast achieves an incredible 20 BOPS at 50 MHz by incorporating 20 separate 32-bit processors on a single chip. Sixteen are arranged in a 4 × 4 "mesh," which is controlled by the remaining four processors. IBM used several clever design techniques to compact each processor and reduce the amount of wiring required to connect them.

Although the high-performance engine clearly has enough horsepower for audio functions as well as graphics and video, IBM expects that initial implementations will rely on today's Mwave DSP chip for audio and communications (fax/modem), leaving the Mfast chip to focus on graphics and video. This has the benefit of leveraging the Sound Blaster compatibility built into Mwave, but the two-chip design will cost more, not only for the Mwave chip but for its local memory as well.

CPU Overhead a Factor in Processing Power

One drawback to the single-engine design used by most of the accelerators is that the processing load must be balanced between various operations. For example, just providing 2D graphics acceleration for a 1024 × 768 × 16-bit display consumes up to 60% of Mpact's processing capabilities, according to Chromatic; tacking on either MPEG-1 decode or a 28.8-kbps modem operation essentially taps out the chip. Although Mpact can execute all three functions simultaneously, GUI performance degrades significantly in these circumstances. Another issue is that the Chromatic chip requires some significant preprocessing on the host CPU. Performing MPEG-1 decoding along with the high-resolution display, for example, consumes half the processing power of a 100-MHz Pentium, enough to cause noticeable slowing of the system.

Performing MPEG-1 decompression consumes 22% of the TM-1's processor cycles, compared with about 30% for Mpact. Mpact places a significant (25%) burden on a P-100 host CPU for a complex task such as MPEG-1 decoding, while Philips says that TM-1's coprocessors help minimize the impact on the host CPU.

Software Strategies Diverge

For 3D graphics, Nvidia has again diverged from the mainstream by originally supplying a curve-based library for its chip. Traditional 3D graphics construct surfaces from flat polygons, causing curved surfaces to appear distorted or faceted. Nvidia's NURBS (nonuniform rational B-splines) algorithm represents objects as a set of curved polygons, allowing many objects to be displayed with greater realism despite using a smaller number of vertices.

Unfortunately, standard 3D software (including Microsoft's Direct3D API) does not use the NURBS model; for these programs, the NV1 substitutes flat polygons, but these programs then cannot take full advantage of the NV1's performance. Sega is porting some of its game programs to PCs equipped with an NV1 *(see* **0911MSB.PDF***)*, but even that vendor is hesitant to commit to the NURBS model, leaving the utility of this feature in doubt.

Nvidia is now downplaying its NURBS capability, instead supplying DirectX drivers and emphasizing the NV1's ability to accelerate more conventional polygons. The company expects that as programmers become more familiar with NURBS, software will eventually emerge that takes advantage of the NV1's unique features.

Chromatic intends to have a full set of Windows 95 drivers available when the chip is released, and the chip provides Sound Blaster and VGA emulation for legacy software, allowing it to replace existing audio and graphics chips. Thus, if all goes as planned, Mpact should provide a fully compatible multimedia solution. Several significant PC vendors, including Gateway 2000, are rumored to have already signed up to use the Chromatic design.

A major drawback of Philips' TM-1 is its lack of support for legacy software; Philips has no plans to support Sound Blaster or other register-level standards, which could significantly hinder its adoption.

PCI in Common; Memory Interfaces Differ

As Table 1 shows, four different memory types are represented among the four accelerators: conventional DRAM, synchronous DRAM (SDRAM), video RAM (VRAM), and Rambus DRAM (RDRAM). Philips and IBM agree SDRAM is the way to go but differ in how much, while Nvidia's licensing agreement with SGS-Thomson has spawned two versions of the NV1 with different memory interfaces: standard DRAM (STG2000) and VRAM (NV1).

Mpact's local memory relies on RDRAM *(see* **091302.PDF***)* rather than conventional DRAM or VRAM. Several memory makers (including, not coincidentally, Toshiba and LG Semicon) build and sell RDRAM which currently costs a bit more than standard DRAM but less than VRAM. Mpact requires only a single 16-Mbit RDRAM for most functions; MPEG-1 encoding or MPEG-2 decoding requires a second RDRAM chip. The 500-Mbyte/s Rambus interface provides higher bandwidth than either the Nvidia or Philips approaches.

The TM-1 relies on SDRAMs to deliver 400 Mbytes/s to the processor, 20% less bandwidth than Chromatic gets from its RDRAM. The SDRAM requires a 32-bit interface, slightly increasing system cost compared with the 8-bit Rambus.

The first Mfast chip will use a wide, 128-bit interface to SDRAM, providing nearly twice the bandwidth of any of the other devices. The downside is that the minimum configuration requires 4M of 256K×16 SDRAM, increasing system cost compared with the 2M of memory used by competing designs. A proposed version of Mfast with a smaller 2 × 2 mesh might use a narrower 64-bit interface, allowing a 2M configuration but reducing bandwidth to that of the TM-1.

In all four cases, the media accelerators connect directly to PCI and provide connections for an external DAC and audio codec to connect to a video monitor and speakers. No glue logic is required, simplifying system design and reducing the real-estate requirements.

Prices Not All Set Yet

SGS-Thomson has priced its DRAM version of the NV1, the STG2000, at $50, while Nvidia's VRAM version goes for $60, both in 10,000-unit volumes. The prices include Windows 95 drivers and a custom DAC that supports a digital joystick as well as the video display. No firm pricing is yet available on Mpact, but Chromatic expects it to be about the same as that of the NV1, with a full Mpact subsystem costing about $150, including 2M of RDRAM, a RAMDAC, and other logic.

Philips also has not revealed pricing for the TM-1 but claims it will "eventually" be below $50. The die size is not available, but the company's decision to include both floating-point hardware and 48K of cache on the chip will probably force a more expensive die than either Nvidia or Chromatic.

IBM claims the size of the Mfast die is "reasonable," but it seems likely that Mfast will be significantly more expensive to build than any of the other media processors. The version with a 2 × 2 mesh would probably make the chip's cost comparable to that of competitive products but would also bring performance down to a range more likely to be matched by Chromatic or Trimedia.

Because Mfast is so far in the future, IBM has no comment on pricing. It appears likely that the initial Mfast, particularly with the Mwave chip and extra SDRAM included, will cost significantly more than competitive devices but will probably outperform them as well. Assuming IBM can deliver on its performance and schedule promises, the question is whether applications will exist that make users want to pay a premium for this added performance.

Intel's NSP Posits CPU-based Approach

Another approach to accelerating multimedia is to process audio and video data on the host CPU, a technique Intel originally called native signal processing (NSP), but now refers to simply as native services. Ideally, this approach would eliminate the expense of a media processor entirely by taking advantage of a computing resource that is already in the machine—an x86 (or other host) processor. There are, however, several problems with a CPU-only approach.

A source of irritation between Intel and Microsoft is the poor real-time capability of Windows, including Windows 95. Real-time response is critical for executing tasks such as audio processing and videoconferencing on the host CPU, and Intel maintains that Windows 95 is not up to the task. This rift led Intel, in conjunction with Spectron, to develop IA-Spox, a real-time kernel running in conjunction with Windows. Microsoft is nonplussed about sharing control of the PC with a second operating system and plans to improve the real-time features of Windows, but any such improvements are not expected before the end of 1996. Chastened by its fruitless experiences evangelizing NSP, Intel will wait out the year until Windows support for real-time functions materializes.

Another issue is that there are significant differences between the way multimedia data and other data are processed. A general-purpose CPU like the current Pentium is simply not efficient in handling digitized audio samples or pixel data. Of course, such a CPU can still slog its way through and deliver some reasonable functions. A 100-MHz Pentium can perform simple audio functions using less than 10% of its total processing power, according to Intel *(see* **090603.PDF***)*. The rather low quality ProShare video (160 × 120 at 10 frames per second) consumes 22% of that processor for simultaneous encoding and decoding. Full MPEG-1 decoding (352 × 240 pixels) requires virtually all of the CPU but reaches 25–30 frames per second. Obviously, these figures will improve on more powerful processors such as the 166-MHz Pentium chips or Pentium Pro.

Multimedia CPUs in the Works

A larger improvement will come when Intel introduces its P55C Pentium later this year. This processor will be the first

For More Information

Contact Chromatic Research (Mountain View, Calif.) at 415.584.5800; fax 415.584.5849; or via the Web at *www.mpact.com.*

Contact IBM Microelectronics (Research Triangle Park, N.C.) at 919.543.4706; fax 919.254.6963.

Contact Nvidia (Sunnyvale, Calif.) at 408.720.6100; fax 408.720.6111.

Contact Philips Trimedia (Sunnyvale, Calif.) at 408.991.3838; fax 408.991.3300.

Contact SGS-Thomson (Carrollton, Texas) at 214.466.7644; fax 214.466.6572.

CPU to implement MMX (multimedia extensions). We expect the enhancements to boost performance on key functions by 2–4×. MMX will also be incorporated into Pentium Pro by 1Q97.

With their recent cross-licensing agreement *(see* **1002ED.PDF***)*, AMD gains access to Intel's MMX technology, creating a powerful, if somewhat uneasy, alliance between the x86 vendors. The first MMX-capable processor from AMD will be its NexGen-based K6, due in 4Q96.

Cyrix announced plans at last year's Microprocessor Forum to incorporate similar extensions of its own *(see* **091403.PDF***)*. However, in light of the recent exchange between Intel and AMD, Cyrix may also adopt MMX to avoid being left out in the cold.

Even with improved microprocessors, however, common yet compute-intensive functions such as MPEG-1 decoding and handling modem protocols will consume a significant portion of the host CPU. Users may notice that their system slows down when they send a fax, a situation they may not tolerate. Furthermore, even with Intel's NSP design, a typical system still requires a 2D or 3D graphics accelerator, a DSP for the modem port, and a Sound Blaster chip for backward compatibility. Thus, the cost savings are not large, whereas the performance penalty is visible. For these reasons, we believe that in the near term NSP will be restricted to low-end PCs, while even midrange systems will include dedicated multimedia hardware.

Ultimately, however, a MMX-enhanced Pentium Pro running at 300 MHz could offer compelling multimedia performance in an NSP environment, delivering 8× or more signal-processing power than today's Pentium-100 systems. Such a processor, expected to reach mainstream PCs by 1998, will perform most of today's common multimedia tasks with a barely perceptible performance loss. Thus, to survive, media processors must offer both a level of multimedia performance well beyond that of the CPUs of their day and hope applications exist which demand that level of performance. We project that in 2–3 years, NSP will be accepted in the mainstream, leaving media processors for high-end systems. In the meantime, media processors have a strong mainstream opportunity.

Multimedia Adoption and Development Ramp Up

As the first media processor to reach the market, the NV1 has defined a new product category. The device itself, however, resembles little more than a high-end audio engine and a 2D/3D graphics engine glued together. While this combination creates significant system-level cost savings, it does not offer the full level of integration seen in the other designs. A follow-on device will use a process shrink to reduce cost, improve performance, and integrate the DAC. A third-generation device, due in 1997, will probably involve a complete redesign, perhaps more along the lines of more-advanced media processors.

Nvidia's design is available now and has fewer software issues, so it is a less risky solution. Its overall savings are smaller, particularly now that a number of vendors have recently introduced low-cost 2D/3D chips. The NURBS 3D model is innovative and interesting, but a long time may pass before a significant amount of software uses it, rendering its advantages moot until then.

The Nvidia chip lacks the video capabilities of the more powerful devices, and it has no modem capability at all. Thus, it falls short of being a fully integrated multimedia solution like the others. Nvidia points out that a Pentium PC is capable of MPEG-1 decoding in software, and that a low-cost MPEG-1 decoder chip can be added, if desired, to avoid bogging down the main CPU. Other video functions, such as MPEG-2 decoding and videoconferencing, are infrequently used today. Thus, the NV1 is capable of meeting the needs of most current users and, not incidentally, is the only media processor available today. Clearly, however, Nvidia has much work to do if it wishes to compete with the media processors rolling out in 1996 and beyond.

Chromatic's Mpact appears to be the best choice for PC vendors, offering a broad set of features that eliminate a number of expensive chips from the system. The company promises that samples will be available this quarter, so system designs using Mpact can begin right away.

The Chromatic chip handles extended video functions such as MPEG-2 and videoconferencing. If these features are needed, Mpact is obviously the superior solution, but until there is demand for these features, they offer little differentiation. A more significant differentiator is support for a 28.8-kbps (V.34) modem. This fast modem alone costs around $50 today, although this price is likely to fall by the time Mpact ships. Combined with the audio and graphics support in the Chromatic design, the modem capability makes Mpact very attractive compared with a traditional discrete solution or the Nvidia design. One drawback is the significant impact on host CPU performance, but as CPU power continues to increase, this will become less of a burden.

Philips' TM-1 delivers essentially the same feature set as Chromatic's Mpact, other than its more open software model, which is of little value unless Chromatic fails to deliver its software. A major problem for Philips, however, is its lack of Sound Blaster compatibility. Throughout 1996–97, there will still be a significant amount of legacy software in users' hands, and this software will not run on the TM-1. Ultimately, this will become less of an issue, leaving Philips to compete with Chromatic on price and capability. Whether Philips' floating-point unit provides a significant advantage on 3D graphics performance remains to be seen.

The Philips chip is also six months behind Mpact and is likely to be more expensive. But if Chromatic stumbles, particularly in software support, the TM-1 could step in. The company's internal volume could sustain TM-1 until then.

IBM's Mfast is shooting for a higher level of performance, with support for MPEG-2 encoding and functions that haven't even been defined yet, like MPEG-4 decoding and high-resolution videoconferencing. As with Chromatic, the issue is when these functions will be demanded by end users; if there is little demand, the IBM part will be restricted to niche markets. The proposed lower-cost Mfast could be better suited to the mainstream, but this product is even further out than the initial Mfast, pushing it into 1998.

Future Looks Bright, Dazzling, and Noisy

Except for Nvidia, all these vendors face the common challenge of delivering the necessary software drivers to perform the promised functions. Without these drivers, the advanced capabilities of these processors lie fallow. In addition, they must demonstrate that they can ship in volume products that meet the performance and feature set goals outlined above. Since most are using innovative but unproven designs, they must meet many challenges before reaching these goals.

On the other hand, all the designs except Nvidia's are reprogrammable, whereas the NV1 has a relatively fixed function set. The programmable accelerators will have the advantage of technical adaptability, giving them more agility in the market as audio, video, and graphics standards continue to change and evolve.

It is interesting to note that, although most of these chips are essentially DSPs at heart, none is being offered by a traditional DSP vendor. Makers of standard DSPs have tried to make their way into the PC market for many years, touting DSP's advantages in audio and modem applications. It evidently took a new group of companies, with a different vision and a slightly different technical approach, to turn PC-based DSP into reality.

Part of the delay in moving DSPs into the PC has been due to the lack of standardized software support. The classic chicken-and-egg problem stymied the adoption of any "nonstandard" PC hardware. Now that DirectX promises to give programmers a unified set of APIs, multimedia hardware vendors are free to innovate, as long as they provide API-compliant drivers and all the APIs are shipping.

The concept of the media processor is a good one and, in the grand tradition of PC integration, seems likely to triumph over the current jumble of audio, graphics, and video chips. As other vendors jump into this market, prices will fall and the number of useful options will increase. In the meantime, adopting one of these early media processors will allow system vendors to differentiate their products from their competitors' through lower cost and a larger multimedia feature set.

