

Chromatic Raises the Multimedia Bar

Mpact Chip First to Combine MPEG with Modem, Audio, and Graphics



by Dave Epstein

With its debut product, startup company Chromatic Research has redefined the multimedia accelerator, introducing the M1 Mpact Media Engine at the recent Microprocessor Forum. This powerful x86 PCI device is the first to combine MPEG-1 encode and decode, MPEG-2 decode, and high-speed modem emulation with audio, video, and graphics acceleration. Chromatic architect Stephen Purcell claims that Mpact will process more than two billion operations per second (BOPS).

Purcell founded the company (originally called Xenon Microsystems) two years ago with Mike Farmwald and David Holt. Wes Patterson, formerly COO at Xilinx, now heads the company, which has raised a total of \$25 million to date from its fab partners and venture-capital financing.

The Mpact chip delivers the feature set of a collection of add-in cards or chips that would cost hundreds of dollars. But according to Purcell, a complete subsystem that includes Mpact, a 16-Mbit RDRAM, and support chips will cost just \$150 and can be placed on the motherboard. The chip expands what Nvidia started (*see 090904.PDF*) but adds MPEG and modem emulation, comprehending essentially all of today's mainstream multimedia functions, as Table 1 shows.

Chromatic expects that the chip will be in systems by mid-96. Using a unique business model, the company has partnered with Toshiba and LG Semicon (formerly Lucky Goldstar) to produce and sell the Mpact chip, hoping to gain quick acceptance in the market.

Chipless Semiconductor Company

Sounds like an oxymoron? In this age of fabless semiconductor companies, chipless is simply the next step. Chromatic designs the chips and even produces the mask data, but it does not sell them. This model is similar to that of MIPS Technologies, and like MIPS, Chromatic takes a modest royalty on each Mpact chip. Unlike MIPS, Chromatic's real business is to make money on the software. The hardware is simply the razor, and the startup the blade company, selling functions and upgrades to the software platform. In this way, Chromatic hopes to avoid the fate of MIPS and remain independent.

Chromatic intends to develop all Mpact software itself and does not plan to release the Mpact instruction set. One advantage of this model is that the chip need not be burdened by the robust protection and debugging

features found in general-purpose CPUs, simplifying the hardware design. The company also doesn't need to immediately create powerful development tools for outside use. The onus rests on Chromatic, however, to supply all software necessary to support PC application standards as they arise and change. The company is looking to partner with ISVs and OEMs to help with this effort.

Both Toshiba and LG Semicon are formidable players with massive manufacturing capabilities. These companies, with help from Chromatic, must convince system makers to incorporate the Mpact chip on their motherboard, since there are no current plans for an Mpact add-in card. Given the size of these manufacturers, even large OEMs should be comfortable with their ability to meet supply requirements. Purcell indicated that some large PC companies have already signed up but would not reveal their names. One hint: the Mpact press release contains a testimonial from Ted Waitt, CEO of Gateway 2000.

Motherboard-Resident PCI Design

Today's high-end PCs have multiple controllers hanging from various buses, such as ISA, VESA, PCI, and serial/parallel ports. Chromatic instead delivers a single high-powered processor that combines advanced high-speed channels for graphics, video, audio, and telephony. By combining all these capabilities into a single PCI-based chip, Chromatic hopes to offer high-end features while significantly reducing system hardware cost.

The Mpact system architecture is an innovative mixture of the host CPU and its memory structure, the Mpact chip (including what amounts to an on-chip cache), the Mpact main memory, and various channel hardware to support the video, audio, and peripheral

Video	Supports JPEG, MPEG-1/2 decode at 30 fps, MPEG-1 encode with accelerated motion estimation, Microsoft Windows 95 MCI API
2D Graphics	Supports GDI, DCI, and DirectDraw with all bitBLTs, ternary ROPs, and hardware cursor
3D Graphics	Supports 3D-DDI and Direct3D with z-buffering, double-buffered rendering, and textured primitives
Audio	Supports MIDI, Wavetable, Waveguide physical modeling, 3D sound, Dolby AC-3, and DirectSound
Fax/Modem	Up to 28.8 kbps; supports DSVD
Telephony	Supports full-duplex speakerphone, Windows 95 TSPI and DirectPlay, and AT+V
Videophone	Supports H.320 over ISDN and H.324 for POTS

Table 1. Chromatic's Mpact processor can perform an impressive range of graphics, video, and audio functions.

devices. The software is split among an x86 host-based resource manager, an Mpack-based real-time kernel, and various drivers for the individual devices.

Not surprisingly, Rambus cofounder Farmwald, now Chromatic's chief technologist, chose a Rambus memory interface to get the biggest bang for the area and (eventually, if you believe Rambus) buck. A single 16-Mbit RDRAM provides 2M of memory, the standard system size. The Rambus transfers one byte every 2 ns, or one 72-bit word every 16 ns, giving rise to Mpack's 62.5-MHz core clock rate. Two RDRAMs are needed for MPEG-1 encoding or MPEG-2 decoding.

Chromatic calls Mpack a VLIW (very long instruction word) and SIMD (single instruction, multiple data) vector processor. While it is all of these things, it is not a general-purpose CPU but rather a sophisticated, special-purpose multimedia controller. The internal architecture, shown in Figure 1, is more akin to a DSP than a microprocessor. It includes an SRAM (which Chromatic calls a cache) that houses both instructions and data, an instruction unit, five function units called ALU groups, the Rambus controller, and the I/O port controllers.

There is no virtual memory-management hardware, since all programs run in real memory out of the RDRAM. Even the SRAM is controlled completely by software, with little hardware support. Programs set up data and instruction areas of the SRAM and can overlay themselves, taking care not to oversubscribe the SRAM, for there is no hardware protection.

The 4K SRAM is organized into 512 words of 72 bits each. Software partitions this memory into instruction and data areas, with the instruction area configurable to 256, 512, or 1K bytes. The instruction side is a direct-mapped cache with a whopping 128-byte line size. The data side is self-managed, looking more like scratch-pad memory with no real line size. Data is moved in 72-bit words; all "bytes" are 9 bits long for extra precision and

to take advantage of the extra bit in the RDRAM for data instead of parity.

Prefetching is done only by explicit instruction direction, and it is done in address order. Chromatic's programs are written carefully to avoid any prefetching delays. The SRAM has four general-purpose read ports and four write ports, all of which can be accessed simultaneously in a single cycle. One write port is typically dedicated to the RDRAM interface and one to a DMA channel; the other two are general-purpose.

Dual Instruction Execution

Mpack has a straightforward instruction model with a few advanced features. The very long instruction word is an eight-byte instruction pair. The instructions may be three, four, or five bytes in length, with pairing done by hand or using Chromatic's compiler. Resource conflicts may cause the instructions to be executed sequentially, but even so, packing two instructions together helps with code density.

An instruction typically consists of one byte of opcode, two source bytes, and one destination byte. These 9-bit source and destination addresses can access any 72-bit word in the on-chip SRAM. Other instruction forms include a three-byte format for two-operand instructions and a five-byte format that allows four-operand calculations such as multiply-add.

An instruction count register creates a repeated vector that can improve code density and speed inner loops. The vector operations are coded only in eight-byte instruction pairs and have a maximum iteration count of 127. Vector loads transfer data from the RDRAM to the SRAM at 500 Mbytes/s. A vector operation can move as many as 256 bytes.

Branches are simple two-byte immediates, allowing a maximum code size of 1M in the Mpack RDRAM. Two forms of each conditional-branch instruction—branch likely and branch unlikely—allow

static prediction set by the programmer or compiler. Conditionals are only on the sign of a result; there are no other condition codes.

Hundreds of Adders

Figure 2 shows the five function units that Purcell calls ALU groups. Each group consists of essentially an eight-byte (72-bit) arithmetic unit that can operate on one, two, four, or eight bytes at a time. Thus, they can be configured to do eight 9-bit ALU operations in one cycle, giving very good performance at the lower precision required by many multimedia

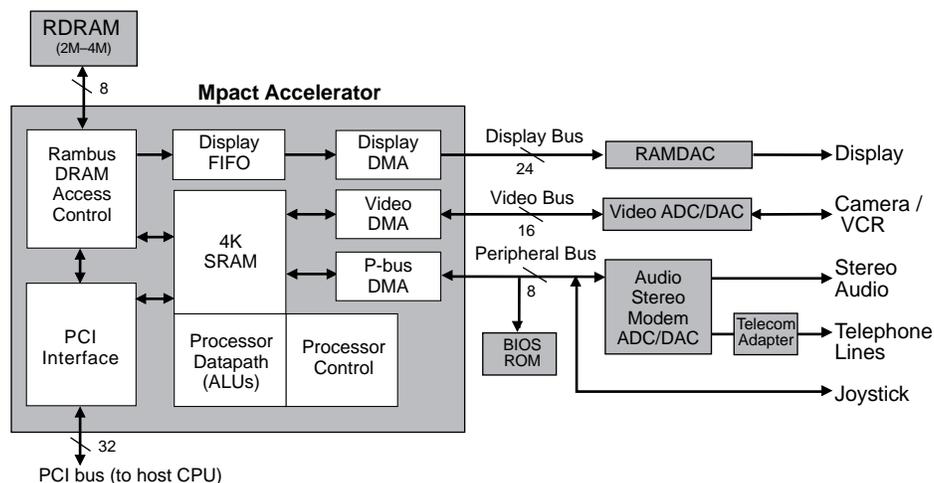


Figure 1. An Mpack subsystem connects directly to the PCI bus and contains the Mpack chip, one or two RDRAM memory chips, and various audio and video support chips.

algorithms. The 9-bit bytes lend themselves nicely to 16-bit audio applications, providing two bits of extra precision for intermediate results. This structure avoids the need to go to 24 bit data, which would use more space and time.

Generally, only one ALU group is active per instruction (two for an instruction pair). Multiplies and special “inner loop” instructions can activate more than one group. To achieve the rated 2.0 BOPS, all four standard ALU groups must be processing eight bytes per cycle at the 62.5-MHz clock speed.

Each group has a specialization. Group 1 is a shift/align unit, while Group 2 is a standard ALU. Multiplication uses Groups 3 and 4: Group 4 produces partial products using Wallace trees, then Group 3 completes the multiplication and adds a third operand. The multipliers can be configured to produce eight $9 \times 9 \rightarrow 18$ -bit multiplies, four $18 \times 18 \rightarrow 36$ -bit multiplies, or two $24 \times 24 \rightarrow 36$ -bit multiplies. Alternatively, Group 3 can be used as a dual three-input adder.

Finally, Group 5 is a specialized motion-estimation unit with some 400 ALUs, most consisting of a few bits. Purcell would not disclose additional details on Group 5 but indicated that it can achieve 20 BOPS, ten times the performance of the other function units combined.

All these units are connected by a crossbar bus that can place any result into any input for the next cycle. This requires a massive 792-bit unidirectional bus with a single source (11 results of 72 bits each) and 19 taps.

The pipeline has just four stages: fetch, decode, execute, and writeback. The leisurely clock rate allows SRAM data to be read and used in the same cycle (no load-use penalty). Splitting multiplication between two ALU groups allows it to be fully pipelined with one-cycle throughput and two-cycle latency. Correctly predicted branches have no penalties, whereas an incorrect prediction costs two cycles. The loop instructions repeat the execute and writeback stages up to 127 times before moving on.

Small Today, Smaller Tomorrow

Chromatic and its semiconductor partners designed the custom chip using an increasingly common cell-based approach. Data paths, along with the SRAM and the Rambus interface (designed by Rambus), were custom designed, while the rest of the chip is standard cell, as Figure 3 shows. A large part of the data path swizzles the 72-bit result buses to their respective destinations, leaving a quite small set of highly customized circuits. The Rambus interface produces the core 62.5-MHz clock, and the other interfaces accept their own asynchronous clocks (33 MHz for the PCI bus, 135 MHz for the display, etc). FIFOs are used for frequency matching.

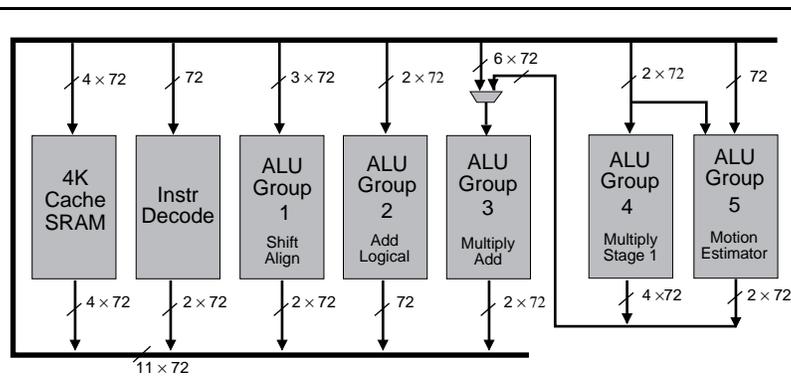


Figure 2. Mpac's internal data paths are all 72 bits wide, with a 792-bit crossbar carrying 11 results back to all five of the function units (ALU groups) and the on-chip SRAM. The instruction decoder takes its input from one of the SRAM ports.

Chromatic designed and laid out the circuits for each manufacturer separately, tuning the layout to the design rules of each. This technique is far superior to the least-common-denominator design-rule method used for the R4000, for example. Chromatic's method allows straightforward compaction and die shrinks, as justified by volume. These shrinks may require multiple rows of bond pads or possibly area bonding, as the first layout is nearly pad limited.

The 100-mm² processor is implemented in 0.5-micron three-layer-metal CMOS and consumes 1.5 million transistors. This first design is expected to be in production in 2Q96. The design is currently going through a shrink to Toshiba's 0.35-micron three-layer-metal process to yield a chip less than 65 mm² with appropriate I/O-pad technology. That tapeout is supposed to occur before year-end, and Chromatic expects full production of the shrink version to begin in early 3Q96, though this

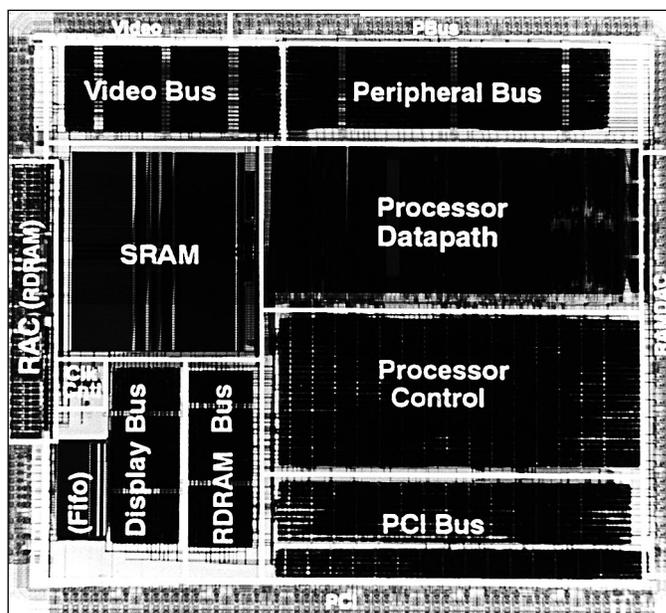


Figure 3. The Mpac die contains 1.5 million transistors and consumes 100 mm² in a 0.5-micron three-layer-metal CMOS process.

Price & Availability

Chromatic expects to sample the M1 Mpack Media Engine in 4Q95, with volume production in 2Q96. The company did not announce a price for the chip but says that the cost of a subsystem including Mpack, a 16-Mbit RDRAM, and support chips should not exceed \$150. For more information, contact Pete Foley of Chromatic Research at 415.254.5826 or check the Web at www.mpack.com; contact Amir Naghavi of Toshiba America at 408.526.2612; or contact Arun Kamat of LG Semicon America at 408.432.5024.

sounds aggressive. The MDR Cost Model places the manufacturing cost of the 0.5-micron version at \$30, shrinking to \$25 for the smaller part.

The 62.5-MHz Mpack operates at 3.3 V but tolerates 5-V inputs. Chromatic has not released power numbers. The package is a 240-pin heat-slugged PQFP.

Versatile Software Environment

The hardware technology is the enabler, but the software must be comprehensive and seamless and can make or break the solution. Chromatic's software architecture is clean and lends itself nicely to upgrades as standards evolve. This "soft" approach sets Mpack apart from fixed-function solutions. It does, however, require that the company deliver all the pieces of that interface; every new PC standard, such as wave-table or Sound Blaster audio, needs to be supported by Chromatic software. The good news is that standard APIs are taking hold in the PC industry.

Chromatic's multilevel software architecture is pictured in Figure 4. It can be decomposed into three levels: the application, the driver, and the virtual devices. Additionally, there is a resource manager (not to be confused with Microsoft's RMI) running in the x86 processor and the Mpack real-time kernel (MRK) running in the Mpack chip itself. In this diagram, all but the application are

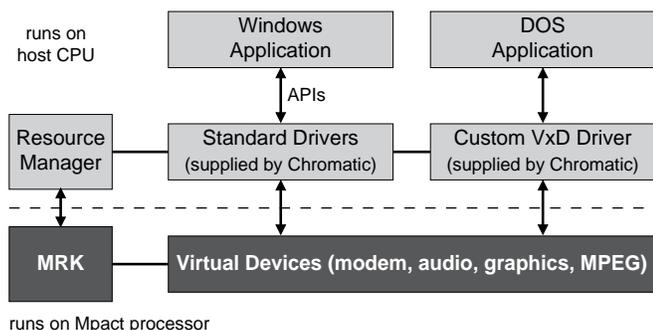


Figure 4. Mpack requires drivers and a resource manager that run on the host CPU while Mpack itself uses a small real-time kernel (MRK) to execute the required tasks.

written, supplied, and supported by Chromatic.

The device drivers (on the x86 side) are the compatibility software, or front end, to the entire Mpack solution. These drivers must strictly adhere to the API, communicating to the virtual devices through MRK and, periodically, to the resource manager. Windows applications generally go through just a few standardized driver interfaces inside Windows, such as GDI for graphics, TAPI.DLL for modems, or MMSYSTEM.DLL for multimedia. These Windows 3.1 and Windows 95 APIs will be fully supported by Chromatic with Mpack's first release, the company claims. DOS applications and less standardized Windows applications may require custom drivers written by Chromatic.

The resource manager manages the RDRAM memory and other resources, maintains priorities, and sets up block and direct-communication transfers. This allocation changes as tasks are created and completed.

The resource manager talks to MRK, a multitasking kernel that juggles interrupts and requests from both the drivers and the hardware nodes operating the physical devices. This kernel allocates the Mpack internal SRAM and performs task switching and task synchronization. MRK makes all devices appear independent to the drivers and maintains real-time response by giving priority to the nearest deadline event. It is also quite clever in its handling of the instruction side of the SRAM cache, overlaying device tasks dynamically.

Finally, the virtual-device tasks that control the actual hardware are quite simple in most cases. The one exception is the audio section, which is complex enough to merit another mini-kernel that handles several processes at once. This section has its own queuing, priority, and task-switching capabilities. Audio deserves this special attention due to the sensitivity of human hearing, which readily detects dropped bits that create clicks or discontinuities in the sound stream.

The software works together to balance the load

Function	Mpack	2M RDRAM	Pentium 100
Kernel (MRK)	5%	1%	1%
14.4-kbps modem	15%	10%	1%
28.8-kbps modem	30%	10%	16%
H.263 video	10%	14%	5%
G.723 audio	5%	1%	1%
MPEG-1 decode, 30 fps w/audio	35%	18%	24%
MPEG-1 encode, 30 fps w/audio	100%*	(4M)	n/a
MPEG-2 decode, 30 fps w/audio	100%*	(4M)	n/a
Full-duplex speakerphone	15%	2%	4%
SVGA 800 × 600 × 18 bpp	25%	60%	12%
SVGA 1024 × 768 × 18 bpp	54%	87%	25%

Table 2. The effect of various combinations of functions can be calculated by summing the utilizations shown here for the Mpack processor, RDRAM memory, and the host CPU. "n/a" indicates not available. (Source: Chromatic except *MDR estimates)

between the host CPU and the Mpack chip. Normally, as much processing as possible is shifted to Mpack, freeing the host CPU. If an extensive combination of applications—for example, a high-resolution graphics stream, a 28.8-kbps modem, and a demanding audio task—are launched at once, the Mpack engine could become overwhelmed. In this case, the resource manager shifts complex graphics operations—such as text acceleration, font caching, and solid fills—to the GDI driver on the host CPU, maintaining the performance of the most critical operations.

High Performance, Low Cost

The bottom line is Chromatic's ability to deliver performance beyond that which is available today at a given price point. Purcell pointed out that there is no way that native signal processing, with any processor now available, can come close to the 2.0 BOPS generated by Chromatic's chip. Only systems with expensive add-in cards can approach this performance level.

As Table 2 shows, a reasonable Mpack system consumes only 25–50% of the host x86 processor with all devices running simultaneously. For instance, an 800 × 600 × 18-bit display running MPEG-1 audio and video decode and a simultaneous 14.4-kbps modem uses 47% of a 100-MHz Pentium. This performance reduction is certainly noticeable by a user running other applications, but it greatly outperforms most other solutions available today. By mid-1996, 167- or 180-MHz Pentiums will be available, significantly reducing the performance degradation on the host CPU.

The size of the RDRAM does not become a bottleneck with almost any application involving an 800 × 600 × 18 display, but a 2M system begins to slow things with a 1,024 × 768 × 18 display and any other multimedia applications running. For these higher-definition displays, or systems with MPEG encode or MPEG-2 decode capabilities, a second RDRAM is recommended.

Prices Must Be Driven Down

Chromatic is the first company to offer a truly high-performance multimedia solution with very few chips on the motherboard. Nvidia's NV1 is the only competitor today, but it has no semblance of modems or MPEG-1 decode and doesn't even breathe the words encode or MPEG-2. These capabilities set Chromatic apart, at least for the time being. On the other hand, Nvidia's solution is significantly cheaper and here now. Nvidia's and other combinations of accelerator chips still don't do as much as the Chromatic solution, but they may do

enough at a competitive cost. The coming holiday season will see a large volume of these multimedia solutions, as Compaq and Packard Bell have already committed to MPEG for their home PC lines. Compaq, for instance, is using S3 hardware MPEG and graphics acceleration.

If all goes well, Mpack should begin appearing in PCs around mid-96. These systems should offer performance superior to that of other designs while minimizing cost and physical size. The Mpack chip will support the full range of multimedia functions that will be standard fare by the end of 1996.

One problem is that mainstream users today simply do not need all these functions; applications requiring videophone capability or MPEG-2, for instance, are nearly nonexistent. These applications aren't there because the solutions are new and too expensive, a chicken-and-egg problem. Until applications demand this performance, Chromatic's challenge will be to

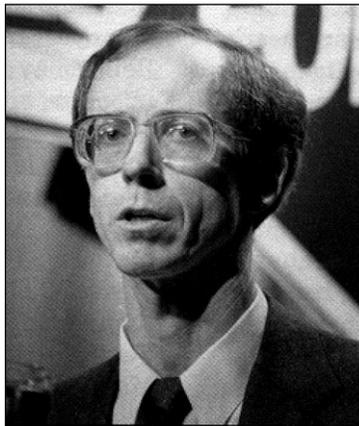
deliver its device at a price with little, if any, premium over mainstream products. The company is taking a bold step and is on the leading edge of video technology. Thus, it is well positioned to take advantage of the adoption of this technology by the application providers and the growth that should follow.

Other vendors, such as IBM and Philips (see **081603.PDF**), are working to develop competitive multimedia accelerators but appear to be behind Chromatic's schedule. Processors with multimedia enhancements, including Intel's P55C, may provide enough power for mainstream applications without any accelerator at all, especially if Microsoft can make Windows deliver real-time response.

The other challenge Chromatic must face is that of software compatibility and support. Companies many times its size have a significant continuing software effort to provide just graphics acceleration, for instance, while Chromatic and its partners must support a broad range of multimedia functions, even those that are newly being defined. This software effort cannot be underestimated.

Chromatic's announcement has surely captured the imagination of the true multimedia PC supplier by offering an integrated solution ahead of the competition. The home user—the primary consumer of multimedia PCs—is the beneficiary; PCs priced for the mainstream will see a dramatic improvement in comprehensive multimedia performance. Add-on accelerator companies had better take note: the bar has been raised. ♦

Dave Epstein was previously VP of engineering at NexGen. He is now an independent consultant and can be reached at 415.493.8332.



Chromatic cofounder Stephen Purcell explains the capabilities of the Mpack multimedia chip.

MICHAEL MISTRACCHI