VOLUME 10, NUMBER 15
NOVEMBER 18, 1996

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

Chromatic's Mpact 2 Boosts 3D

Mpact/3000 Becomes First Media Processor to Ship in Volume

by Yong Yao

Even as manufacturing partners LG Semicon and Toshiba are beginning volume shipments of the first media processor, Mpact/3000, Chromatic has disclosed its forthcoming Mpact 2 design. Announced at last month's Microprocessor Forum, Mpact 2 (code-named M2) is expected to double the overall performance of Mpact/3000 (code-named M1). The gain is much greater on 3D graphics, as new floating-point capabilities position Mpact 2 to meet the 3D-performance requirement for mainstream PCs in 1998.

The performance gains stem mainly from a doubling of internal clock speed, an additional Rambus channel, the addition of floating-point capability, and a new 3D rendering engine. Other enhancements, such as doubling the on-chip cache and supporting 66-MHz PCI, also contribute to the improved performance.

One of the biggest competitive advantages that Chromatic has over other media-processor vendors is the richness of its software offerings. The recent release of the Mpact software, Mediaware Release 1.0, contains 2D/3D graphics, MPEG-1 audio/video decoding, FM/wavetable synthesis, V.34 modem, and full-duplex speakerphone. We expect future revisions of Mediaware to enhance these functions and also add functions such as DVD playback, video editing, and videoconferencing.

Chromatic expects Mpact 2 to be in PCs for the 1997 holiday season. We believe the technical success of Mpact/3000 makes this goal achievable. The company intends to position Mpact 2 at the high end while Mpact/3000 covers the low end. Whether the PC market will accept that positioning remains to be seen, because Mpact/3000 may not have enough performance by the time Mpact 2 is in production.

Mpact 2 Doubles Mpact/3000 Performance

Doubling performance every 12 months is Chromatic's plan for keeping ahead of Intel's processor advances. To meet this goal for Mpact 2, Chromatic will increase the internal clock speed from 62.5 MHz in Mpact/3000 to 125 MHz or higher, replace 4K of cache with 8K, and add a second Rambus channel to reach 1,200 Mbytes/s of total memory bandwidth. The chip supports 300-MHz RDRAMs with an effective transfer rate of 600 MHz. Figure 1 shows the internal block diagram of Mpact 2, which is similar to that of the original Mpact/3000 (see MPR 10/23/95, p. 23).

The faster clock speed is achieved by using a 0.35-micron CMOS process, optimizing critical paths, and decoupling the internal clock from the Rambus interface clock. Unlike its predecessor, the Mpact 2 processor core can run asynchronously from the RDRAM, which provides the freedom to choose the RDRAM and internal clock speeds independently. One drawback to the asynchronous design is additional latency, due to clock synchronization, when accessing RDRAM.

Continued on page 6

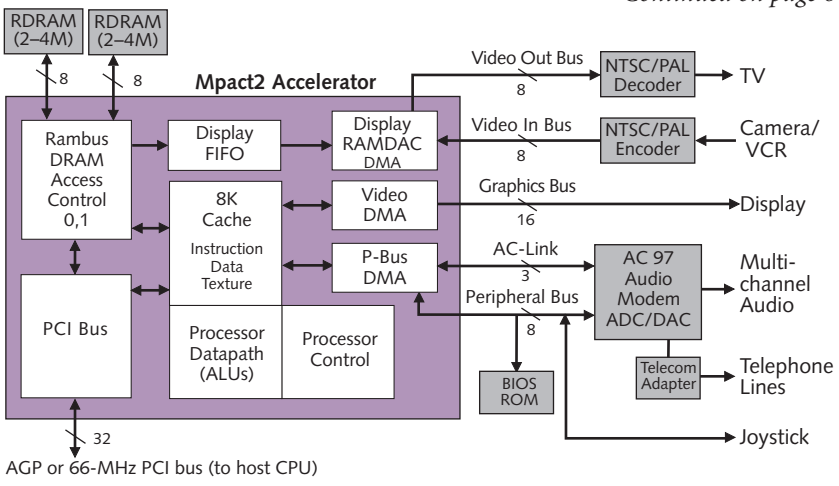

Figure 1. Mpact 2 doubles performance over Mpact 1 by increasing the internal clock speed, doubling the amount of on-chip cache, adding a second Rambus channel, and doubling the speed of the PCI bus. A multimedia subsystem requires few extra chips.

AT A GLANCE

MICROPROCESSOR REPORT

Publisher and Editorial Director
Michael Slater
E-mail: mslater@mdr.zd.com

Editor in Chief
Linley Gwennap
E-mail: linley@mdr.zd.com

Senior Editor
Jim Turley
E-mail: jturley@mdr.zd.com

Senior Analyst
Yong Yao
E-mail: yyao@mdr.zd.com

Senior Analyst
Peter N. Glaskowsky
E-mail: png@mdr.zd.com

Editorial Assistant: Kathy Acuff

Editorial Board
Dennis Allison, Rich Belgard, Brian Case, Jeff Deutsch, Dave Epstein, Don Gaubatz, John Novitsky, Bernard Peuto, Nick Tredennick, John F. Wakerly

Editorial Office
298 S. Sunnyvale Avenue
Sunnyvale, CA 94086-6245
Phone: 408.328.3900 **Fax:** 408.737.2242

Microprocessor Report is published every three weeks, 17 issues per year. Rates are: N. America: $495 per year, $895 for two years. Europe: £375 per year, £645 for two years. Elsewhere: $595 per year, $1,095 for two years. Additional copies in the same envelope: $175 per year in North America, $225 elsewhere. Back issues are available.

Published by
MICRODESIGN RESOURCES

President: Peter Christy
E-mail: pchristy@mdr.zd.com

Business Office
874 Gravenstein Hwy. So., Suite 14
Sebastopol, CA 95472
Phone: 707.824.4004 **Fax:** 707.823.0504
Subscriptions: 707.824.4001
E-mail: cs@mdr.zd.com

World Wide Web: www.chipanalyst.com

Winner, Computer Press Award, 1993, 1994

Printed on recycled paper with soy ink.

THE EDITOR'S VIEW

Clouds Darken Over RISC

Some Will Perish After Introduction of Intel/HP Architecture

RISC processors, unless they're home grown, are becoming less and less interesting to system vendors. NEC's rejection of MIPS chips for its systems (see MPR 10/28/96, p. 5) is another sign of the times; the Japanese company is nominally switching to HP's PA-RISC, but the real destination is the future Intel/HP architecture, IA-64. As IA-64 becomes a reality, we expect other vendors to follow NEC's lead. Ultimately, this trend could result in one or more RISC processor vendors giving up on their architectures.

Nearly every major computer company today is either building systems around x86 chips or using its own in-house RISC processor (or both). Apple, of course, gets its processors from IBM and Motorola, but the Mac vendor's involvement in PowerPC is so intimate that it hardly counts as an exception to this rule. Fujitsu is a staunch SPARC system vendor but consumes many processors from its Ross and Hal subsidiaries. Otherwise, the largest computer vendors making significant use of an externally produced RISC processor are Tandem and Pyramid, midsize players at best.

This situation leaves Digital, HP, Silicon Graphics, and Sun with so-called MORPs: my own RISC processor. The volume of chips such as the 21164, PA-8000, or R10000 is truly miniscule, a few tens of thousands per year. Under the old business model, high-end processors like the R4000 would eventually migrate into low-cost systems, pumping up their lifetime volume. In today's competitive market, recycled high-end chips are no longer adequate for the low end; instead, vendors are designing price/performance products like the R5000 and the PA-7300LC, preventing high-end chips from ever reaching volume price points.

If high volumes are never attained, conventional economics says the cost of developing a beast like the R10000's follow-on must be spread across the miniscule volume of parts ultimately produced, resulting in a huge surcharge per chip. The RISC vendors claim the revenue from their systems is more than enough to cover this cost. Each of the MORP vendors collects several billion dollars per year in RISC system revenue. The cost of maintaining a couple of hundred processor designers is perhaps 1% of that revenue.

These vendors are willing to bear this cost because they gain a competitive advantage: better performance, particularly on floating-point applications, than mass-market processors can provide. Furthermore, the costs of moving their user base to a new architecture are daunting, even if that architecture would eliminate the need for internal CPU development. Thus, no vendor has followed HP's lead in terminating its RISC architecture.

That may change in a few years, however. Today's out-of-order RISC processors are overburdened with the very complexity that RISC was intended to eliminate. Too little of the chip is doing real work; the rest is devoted to keeping two sets of books, in case the software ever conducts an audit to find out what the CPU has really been up to. While it is certainly possible to work with a balky architecture (witness the success of x86), the performance growth rate of RISC architectures will be difficult to maintain.

HP, with Intel's help, aims to solve this problem by moving to IA-64, a clean new architecture that we believe will push much of this complexity back into the compiler, boosting performance by devoting more of the chip to instruction execution. Because of this leap forward, we expect Merced, the first IA-64 processor, to outperform all traditional RISC processors. HP will solve the migration problem by offering compatibility with PA-RISC, probably through some sort of translation/emulation scheme.

Other RISC vendors will then be faced with some difficult choices. Staying with current instruction sets will probably put them at a performance disadvantage to IA-64. To close this gap, they could develop a new instruction set along the lines of IA-64, but such a massive development effort, as well as moving customers to that instruction set, will be more expensive than maintaining the existing product line.

Once the thought of a new instruction set is on the table, another option is to simply adopt IA-64. This move eliminates the cost of in-house development while ensuring performance competitiveness. If buying processors from Intel is unthinkable, vendors could instead band together to develop a competitive instruction set, sharing the development costs and building more volume for the new design.

SGI can easily differentiate its products on the basis of 3D performance, regardless of the underlying CPU; the company is a good candidate to switch to IA-64 or Project 2K, a new instruction set under development by the PowerPC vendors (see MPR 8/26/96, p. 12). It is less clear how Sun or Digital would differentiate their products from those of other IA-64 vendors. Digital seems likely to stick with Alpha until the bitter end, which, given the company's recent return to unprofitability, may be sooner than later. Sun may stay with SPARC or join an anti-Intel alliance. Efforts to rival IA-64, however, are already two years behind; so far, this delay is probably not fatal, but the clock is ticking for these vendors. M

Intel's Katmai, Willamette Surface

Intel's river map rolls onward. Rumors have recently surfaced regarding two new processors planned for 1998. Sources indicate Intel is working on a processor code-named Katmai that will be a minor improvement on Deschutes. (Katmai, pronounced KAT-mye, is named after a river in Alaska; Intel has apparently moved beyond Oregon and California in its naming conventions.)

The key to Katmai appears to be the addition of new MMX instructions, forming a set referred to as MMX 2. No word yet on what the new instructions are, but they may address 3D geometry performance, which MMX does not. One possible addition would be dual floating-point operations similar to those in MIPS V (see page 24). Katmai is said to be planned for a 1H98 debut, about six months after Deschutes first appears.

We believe Intel is moving aggressively to improve the multimedia performance of its processors, particularly for 3D graphics, which is weak today. These moves will both drive another PC upgrade cycle and protect against a significant incursion from media processors (see MPR 10/7/96, p. 3). MMX, as originally announced, is just the first step, and subsequent Intel processors will offer continued improvement.

For the second half of 1998, Intel is preparing a much more significant upgrade, a chip code-named Willamette ("Wil-LAM-ette") after the Oregon river. This device, which we previously called the P68 (see MPR 3/5/96, p. 3), is expected to be based on the P6 core but with significant performance enhancements, perhaps including a larger reorder buffer and an additional integer unit. Willamette will presumably carry forward the MMX 2 enhancements from Katmai. As a result of these changes, we expect Willamette to deliver 30–50% better integer performance per clock cycle than Deschutes and perhaps a bigger improvement on 3D graphics and some other multimedia applications.

Intel is also said to be working on a 100-MHz version of the P6 bus, which is today limited to 66 MHz. The current speed is adequate for systems with four 200-MHz Pentium Pro processors using 512K caches and is likely to support four 300-MHz Deschutes processors if they are each coupled with around 2M of cache. A four-way Willamette server, however, will overwhelm the 66-MHz bus. Even a one- or two-CPU desktop system may need the faster bus for Willamette, since these systems will probably stick with 512K of cache. We expect Intel will support the faster bus speed as an option for Willamette, and possibly for Katmai as well.

Willamette's main role is to fill in the x86 line underneath Merced, the first IA-64 processor. We expect Willamette, along with a subsequent shrink version, will be Intel's mainstream processor throughout 1999 and 2000, giving Merced plenty of time to gain software and system support while easing into the market. —*L.G.*

New MicroUnity Chip for Cable Modems

At last month's Microprocessor Forum, MicroUnity's Craig Hansen revealed the first fruit of the company's strategic repositioning (see MPR 8/5/96, p. 5). To replace its high-power BiCMOS chip designed for generalized signal processing, the company has developed a low-power CMOS device tuned for cable modems. The new device retains the instruction set and much of the microarchitecture of the BiCMOS chip but makes several changes to reduce power consumption.

The first change, of course, is from BiCMOS to CMOS. The original processor ran at 1 GHz; because bipolar circuits consume power whether they are switching or not, it makes sense to switch them as fast and as much as possible. This factor led to the idea of a multithreaded design that switches from one instruction stream to the next on every cycle, overlapping five threads at once.

In a CMOS process, this architecture burns far too much power: a prototype multithreaded CMOS processor was measured at 60 W despite clocking at only 200 MHz. The new chip is a traditional single-threaded design running at 113 MHz. Although the peak performance of this design is obviously greatly reduced, the throughput of a single thread is similar to that of the BiCMOS design, which ran each thread at 200 MHz but encountered greater memory latencies relative to the cycle time.

The new media processor, built in 0.35-micron four-layer-metal CMOS, runs at 3.3 V and consumes just 4 W (maximum) at 113 MHz. Although it has five million transistors, about three million are in the 64K of on-chip cache, and the die size is just 100 mm^2. It retains enough performance to handle various modulation algorithms (such as QAM, DMT, and CDMA), 802.2 bridging, and higher-level services for a cable modem.

An entire cable modem could be built from just the MicroUnity media processor and its associated media codec, 512K of DRAM, 128K of flash, a tuner for the cable input, and an Ethernet connection to the PC. The media codec provides a PCI interface, allowing the use of commodity Ethernet chips. An optional telephone modem can be added through the PCI interface as well.

By leveraging its original media processor design, the company has quickly created a new design, although prototypes have yet to be built. MicroUnity is now seeking a semiconductor partner to fabricate and market the device. With low power consumption and a small die size, the new design is much better suited to the market needs of cable modems and set-top boxes than the original BiCMOS monster. Whether the new device can save MicroUnity from extinction depends on how quickly the market for cable modems increases from its current near-zero state and whether the company has any other irons in the fire. —*L.G.*

Newton First Design Win for StrongArm

Digital's StrongArm has made its first public appearance in the form of Apple's newest Newton. As we predicted (see MPR 2/12/96, p. 1), the SA-110 forms the heart of the new Apple MessagePad 2000, which Apple claims runs "up to 10 times faster than any previous model" of Newton; in fact, the 162-MHz StrongArm chip is rated at 185 MIPS, compared with 18 MIPS for the ARM610 used in previous models. Other than the new CPU, the 2000 is similar to earlier Newtons. Apple expects to ship the new model in 1Q97.

The nominal 160-MHz StrongArm was chosen because its power consumption, at 1.65/3.3 V, most closely matched that of the earlier Newton's 5-V ARM610. The handheld unit is based on a commercial processor, like previous Newtons, and a set of four ASICs built by Cirrus. The 2000 also contains 8M of ROM, 4M of flash memory, and 1M of DRAM. Power is provided by four AA batteries; Apple advertises three to six weeks of useful battery life.

The new unit comes with the usual assortment of productivity applications plus a Web browser and software (with cables) for communicating with desktop Macintoshes and PCs—a notorious shortcoming of previous Newtons. Apple also claims vastly improved handwriting recognition, thanks mostly to the faster processor.

In addition to appearing in the 2000, StrongArm will be at the heart of Apple's licensed Newton reference platforms. Inexplicably, the portable eMate 300, which Apple rolled out on the same day, uses an ARM710 processor.

Digital can now trumpet its first design win for StrongArm, one that has the potential to outstrip total Alpha sales in a very short time. The company has publicly demonstrated network computers and other devices based on StrongArm, but no other customers are willing to identify themselves at this time.

Apple's decision to announce the Newton 2000 months before it is available is widely seen as a pre-emptive strike against a spate of Windows CE announcements this week. The new Newton is expected to be more expensive—perhaps by a factor of two—than the WinCE units, so Apple will again have to convince buyers that Newton's features are worth the price premium. —*J.T.*

Oak Describes Single-Chip DVD Decoder

At last month's Microprocessor Forum, Oak Technology *(www.oaktech.com)* previewed a possible future product for DVD players and similar consumer devices: a single chip combining a new 32-bit RISC processor core with multiple blocks of fixed-function logic.

The new architecture takes a different tack than other programmable multimedia processors like Chromatic's Mpact 2 (see cover story). Instead of using a high-performance processor core to execute multimedia functions directly, the design performs the bulk of these functions in hardwired logic, relying on a fairly simple CPU core to manage the flow of data among the various logic blocks.

Hardwired functions include an MPEG-2 transport stream parser, MPEG-2 video and audio decoders, a Dolby Digital AC-3 audio decoder, and a complete digital-video back end for letterbox and pan-and-scan support.

The programmable RISC engine described by Oak's Wen Hsu consists of a simple 32-bit scalar CPU core with 64 32-bit registers, supporting 32K of instruction space and 1K of data. It operates at 67.5 MHz with a three-stage pipeline. Because of the short pipeline, Hsu claims there is no penalty for any type of branch operation, load, or store.

The chip would connect directly to 2M of SDRAM, an audio DAC, an analog video encoder, an 8/16-bit host CPU, and the DVD drive's controller chip. No other significant components would be needed to make a complete DVD player. This system would cost less than a player based on a set of discrete MPEG-2 audio and video decoder chips.

Oak says that such a chip could be fabricated in 0.35-micron CMOS, operating from a 3.3-V supply with 5-V-tolerant I/O, and packaged in a 160-pin PQFP. No specific product plans were disclosed, however.

Media processors from Samsung and Philips can also be used in DVD players, but the higher performance and floating-point support in these parts is likely to make them too expensive for this application. Like Oak, Fujitsu has targeted the DVD market for its MMA chip (see page 11), which was announced at the Forum. MMA's more flexible design may make it easier to adapt to the evolving DVD standard, but Oak's simpler design may be less expensive. Given Oak's unwillingness to discuss an actual product, it appears Fujitsu will reach the market first. —*P.N.G.*

TriTech Debuts Highly Integrated 3D Accelerator

The TR25201 3D graphics controller, part of TriTech's new Pyramid3D product family, is the first 3D chip for the personal-computer market to include integrated hardware support for geometry, setup, and rendering in a single device. This combination lets the device offload from the host processor all parts of the 3D pipeline except scene definition.

The result is high-performance and high-quality rendering. TriTech rates the TR25201 at 1M triangles/s (for 25-pixel shaded, textured, fogged, and Z-buffered triangles). The part also provides hardware support for radiosity-based lighting, trilinear MIP-mapped and bump-mapped textures, Phong shading, and complex shading effects such as environment mapping. A PCI interface and integrated 200-MHz RAMDAC are also included, supporting 1,600 × 1,200 resolution.

Microsoft's Direct3D API does not yet allow application developers to take advantage of all the features of Pyramid3D. TriTech is working directly with several ISVs to develop targeted software and with Microsoft to extend Direct3D. Even within the current constraints of Direct3D, the TR25201 is said to offer faster rendering than the other 3D chips in its price range.

Continued on page 21

Chromatic Mpact 2

Continued from page 1

Simply doubling clock speed would not double the overall performance if there were not enough data to feed the engine. Mpact 2 satisfies this need by doubling the external memory bandwidth as well as the amount of on-chip cache. By simply adding a second Rambus channel, Chromatic doubles the memory bandwidth while adding only eight data pins to the package.

As Figure 2 shows, the new cache is divided among a 2K instruction cache, a 2K texture cache, and a 4K data cache. The data cache has six read ports and six write ports. In the first-generation Mpact design, there is no dedicated cache for textures, and the single 4K cache holds both data and instructions with only four read ports and four write ports.

In addition to PCI and other standard hardware interfaces supported by Mpact 1, the new chip adds support for Intel's AC '97. The AC '97 interface (see MPR 7/8/96, p. 4) allows Mpact 2 to use any off-the-shelf AC '97–compliant audio codec, reducing the cost of the audio subsystem.

The PCI interface has been modified to operate at 66 MHz, twice the frequency of the standard PCI bus. This higher speed enables Mpact 2 to store textures in main memory and increases bandwidth to the rest of the system. Few chip sets exist today to drive a 66-MHz PCI bus, so Mpact 2 will operate the bus at 33 MHz if necessary.

By the time Mpact 2 appears, however, chip sets that support the AGP interface (see MPR 6/17/96, p. 11) should be appearing. These chip sets will connect to Mpact 2 using 66-MHz PCI. For this reason, Chromatic describes its chip as AGP-compliant. The initial Mpact 2, however, does not support any of the advanced features of AGP, such as dual clocking and split transactions.

Mpact 2 integrates a 220-MHz RAMDAC on chip. The high-speed DAC supports displays up to 1,600 × 1,200 × 18 with a refresh rate of up to 85 Hz.

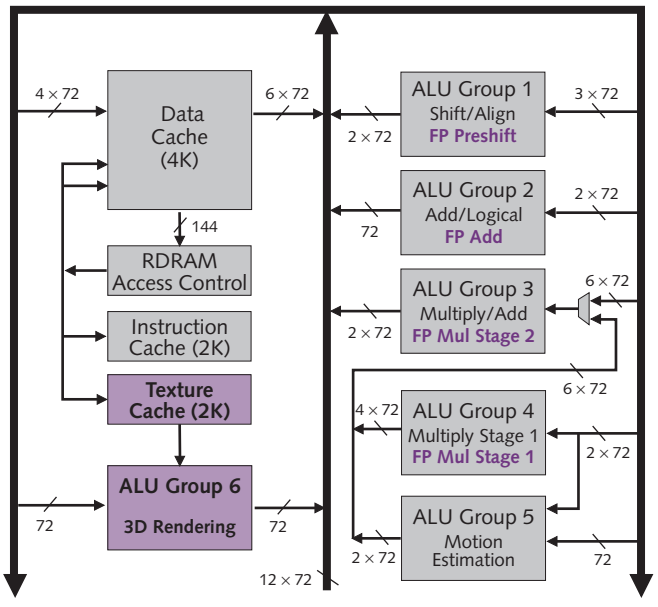

Figure 2. Mpact 2's internal data paths are all 72 bits wide, with a 792-bit crossbar carrying 11 results back to all six ALU groups and to the on-chip caches. Changes from the previous design are highlighted in purple. ALU groups 1–4, for example, are modified to support FP operations.

Mpact 2 Adds Floating-Point Capability

Because Mpact/3000 lacks floating-point capability, the chip's biggest weakness is 3D graphics. Without strong 3D performance, a PC media processor cannot succeed in the long run. At the Forum, architect Stephen Purcell said that improving 3D graphics performance was a major design goal for Mpact 2.

The key to 3D geometry processing is floating-point performance. Instead of adding a separate floating-point unit, Chromatic modified the four general-purpose ALU groups to perform FP operations, as Figure 2 shows. This change has made a minimal impact on die size. All FP operations are performed in single-precision mode, since 3D graphics for a PC rarely use double-precision. To take advantage of the 72-bit data paths, each ALU in Mpact 2 performs two single-precision operations in parallel. This method is similar to the paired-single format in the MIPS V instruction set (see page 24).

To perform a floating-point add or subtract, the operands are first sent to ALU Group 1, which performs a preshift operation to correctly align the operands. This operation uses the existing shifter in that ALU. Using the crossbar, the operands then move to ALU Group 2, which performs the actual add operation. Thus, FP addition is fully pipelined with a latency of two cycles. Similarly, for floating-point multiplication, the operands are first sent to ALU Group 4, then the multiplication is completed in ALU Group 3 on the following cycle. Again, FP multiplication is fully pipelined with a latency of two cycles.

3 stages	Input polygon
3 stages	Generate spans from polygon
3 stages	Generate pixels from spans
10 stages	Apply perspective (divide step) to pixels
3 stages	Generate texture address and present to texture cache
2 stages	Access texture cache and format texel
2 stages	Apply filters
2 stages	Blend/modulate pixel/texel; diffusion, specular lighting
2 stages	Apply fog effect
5 stages	Cluster final pixels

Figure 3. The 3D-rendering pipeline in the Mpact 2 consists of 35 stages, broken down as shown. Such a long pipeline is not a problem for 3D rendering because there are no pipeline hazards.

Each cycle, a pair of single-precision FP adds can be launched along with a pair of single-precision FP multiplies. At 125 MHz, the peak performance of the Mpact 2 design is thus 500 MFLOPS.

Chromatic added about 30% more instructions to the original Mpact instruction set, many specifically for the new floating-point capability and the 3D graphics unit. Unlike on the integer side, there is no direct support for FP multiply-accumulate; this operation is synthesized from a multiply and an add with a total latency of four cycles. There is also no FP divide instruction; this operation can be synthesized by a sequence of multiplies and adds.

New 3D Unit Speeds Rendering

Even the strong performance of the general-purpose Mpact engine is inadequate for high-speed 3D rendering. Chromatic's design philosophy is to use the programmable engine for general processing and for algorithms that may change, but specific hardware may be required to accelerate fixed algorithms. This philosophy led to the addition of a hard-wired motion estimator (ALU Group 5 in Figure 2) in the original Mpact and, in Mpact 2, a new 3D rendering unit.

The 3D unit, ALU Group 6, contains a 35-stage pipeline that runs concurrently with instruction execution. Since 3D graphics operations don't contain any pipeline hazards such as data dependencies or mispredicted branches, the long pipeline provides high throughput without reducing performance. Figure 3 shows this 3D rendering pipeline.

Other hardware improvements also aid Mpact 2 in 3D graphics. The addition of the 66-MHz PCI bus, the dedicated texture cache, and the wider path to external memory all improve 3D rendering performance.

One of the advantages that Mpact 2 has is its flexibility in allocating 3D tasks. The media processor can handle geometry, setup, and rasterization, or let the host processor take over geometry and/or setup. If triangle setup is performed in Mpact, it will be done in firmware using floating-point operations. Chromatic claims Mpact 2 can achieve one million triangles per second with 50-pixel triangles, Gouraud shading, 18-bit Z depth, perspective-correct textures, bilinear filtering, alpha blending, and fog. If delivered, this performance would put Mpact 2 right in the mainstream for PC 3D in 1998.

Although Mpact 2 is not a part of Microsoft's Talisman reference platform (see MPR 8/26/96, p. 5), some of Talisman's features have been incorporated into the Mpact 2 hardware. According to Chromatic, features included are chunking, image and texture compression, and affine transformations. Like Talisman, Mpact can also choose to render particular objects below the frame rate but composite them at the full frame rate. These features may help Mpact 2 support Talisman-enabled 3D games and perhaps become an officially supported Talisman platform in the future.

Price & Availability

Mpact 2, packaged in a 304-pin SBGA, is slated to sample in 1Q97 with volume production in 3Q97. Pricing has not been determined. Contact Chromatic (Sunnyvale, Calif.) at 408.752.9100 or access the Web at *www.mpact.com*; contact LG Semicon (San Jose, Calif.) at 408.432.5024; contact Toshiba (San Jose, Calif.) at 408.526.2612.

Chromatic Releases Mediaware 1.0

Mediaware is the fuel for the Mpact hardware engine. Chromatic recently shipped version 1.0 of Mediaware for Mpact/3000. The release includes an Mpact real-time kernel, Mpact resource manager, 2D/3D graphics, audio, video, fax/modem, and telephony (see sidebar, page 8, for details). Mpact 2 is software-compatible with its predecessor; therefore, the existing Mediaware will run on Mpact 2.

Mediaware 1.0 is designed to run on any Pentium processor. Chromatic is currently revising its software to take advantage of the MMX extensions incorporated in Intel's P55C Pentium and other forthcoming processors. The new code, called MMX Mediaware, is planned to ship in 1H97.

As Figure 4 shows, there are seven Mediaware modules, corresponding to the seven main multimedia functions. These modules, plus the Mpact real-time kernel, reside on the media processor. The rest of Mediaware software—Mpact DirectX drivers, MMX Mediaware, and the Mpact resource manager—run on the host processor.

When Chromatic originally announced its Mpact chip last year, the company expected to include MPEG-1 encoding and full DVD (MPEG-2 video and AC-3 audio) decoding in its initial Mediaware, but these functions have not yet appeared. Chromatic's software efforts were hampered by changing APIs at Microsoft as well as the relatively small size of the company compared with the task at hand.

For future Mediaware releases, we expect Chromatic to add DVD audio and video decoding, 3D positional audio,

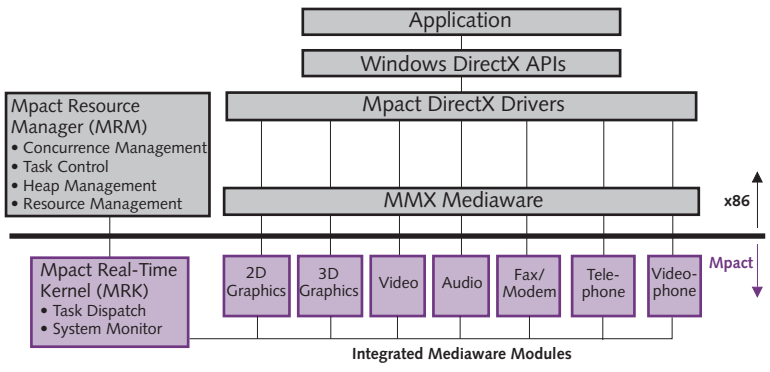

Figure 4. The Mpact Mediaware provides seven multimedia functions. It interfaces to applications using Microsoft's DirectX APIs.

Complete Contents of Chromatic Mediaware Release 1.0

The most challenging problem in supplying a media processor for PCs is not the hardware but the software. While the Mpact processor is a complex piece of work, the associated Mediaware performs a huge array of functions. To deliver all of these functions, Chromatic has roughly twice as many software engineers as hardware engineers.

Without functional software, a media processor is worthless. Of course, none of the emerging media processors has an established software base, so all of this code must be created before the first systems can ship. By bringing its Mpact processor and Mediaware to market first, Chromatic has established a significant lead over other media-processor contenders. The list of software functions below demonstrates the barriers to entry in this market and sets a standard for competitors to match.

Video
Full support for the OpenMPEG command set; real-time MPEG-1 video decode at 30 frames per second, 18-bit color; system layer timestamp-based video/audio synchronization; fully accelerated color-space conversion, bilinear interpolation and filtering; full screen or video in a window; crisp synchronization with graphics output; fully accelerated video CD playback; patented hardware motion-estimation and image processing; PAL/NTSC video input and output.

Audio
DOS and Windows 95 audio through WAVE, MIDI, DirectSound API; MPEG-1 (layers 1 and 2) audio decode; wavetable synthesis, 32 simultaneous voices, complex envelopes and filters; precise audio control through support for key layering, velocity switching, variable keyboard scaling for alternate tunings, sharpness of attack, and speed of decay; general MIDI-standard compatibility with 128 original instruments and 60 drums, 7 additional Roland GS and Yamaha XG drum kits; up to 8 simultaneous play and record channels with different formats and frequencies resampled up to 44 kHz, with 36-bit internal precision; psychoacoustic minimization of quantization noise; Sound Blaster compatibility; Mpact Audio Process Manager for minimizing synchronization overhead between audio tasks; HMI and Miles driver support; industry-standard joystick with MIDI; industry-standard MPU-401 MIDI port.

2D Graphics
Full DOS and Windows 95 GUI acceleration through DirectDraw APIs; VGA register-level compatibility and Super VGA graphics modes; VESA BIOS Extensions 2.0 and UniVBE support; noninterlaced screen resolution up to 1,280 × 1,024 × 18-bit color at 60 Hz; full BitBLT acceleration engine, including transparent BLTs and device bitmaps, with ternary operations; GUI acceleration of two-point line draws and trapezoidal and polygon fills; accelerated YUV conversion, scaling, filtering, clipping, and hardware cursor; full text acceleration engine; full hardware and BIOS support for VESA Display Power Management.

3D Graphics
Full Windows 95 real-time 3D graphics acceleration through Direct3D API; acceleration of 3D-rendering pipeline rasterization; Z-buffering and double-buffered rendering; flat and Gouraud shading; accelerated 3D spans and 3D geometric primitives; diffuse and specular highlighting; true color (24-bit RGB) with decal and color-modulation texture blending; 2×2 ordered dithering; bilinear and trilinear MIP-mapped texture filtering; perspective-corrected texture mapping; support for texel formats including 555 and 565 RGB, 888 RGB and 8888 RGBα, 4-bit- and 8-bit-palletized; full transparency, including alpha maps for textures; full alpha blending support; depth cuing with atmospheric effects; subpixel accurate rendering; buffer/texture memory management through DirectDraw HAL.

Fax/Modem
Full Windows 95 support for Microsoft Unimodem V, TAPI and VCOMM APIs; data modulation up to V.34 bis; V.32 bis support; V.22 bis support, Bell 212A and Bell 103; V.42 bis data compression and V.42/MNP 2-4 error correction; V.8 compliance; fax modulation up to V.17 in answer and originate modes; full fax class 1, 2, and 2.0 command-set implementation; data automoding; DTMF generation, dial tone and busy detection, ring detection, and auto answer; V.14 async-to-sync conversion and RS-232 interface support; full range of baud rate and parity support on DTE interface with memory-to-memory interface.

Telephony
Support for Windows 95 TAPI and VCOMM APIs; adaptive full-duplex speakerphone with acoustic echo cancellation; answering machine and voicemail functionality; outgoing and incoming message support; concurrent DTMF detection; call-progress detection; IMA ADPCM voice compression/decompression; Caller ID support.

Mpact Real-Time Kernel and Mpact Resource Manager
Real-time operating system with oversubscription management; dynamic linking, loading, and profiling of all Mpact threads; RDRAM memory allocation and compaction; host and Mpact communication functions; Mpact dynamic resource allocation and degradation.

video editing, and videoconferencing over POTS, ISDN, and the Internet. With the increased performance provided by Mpact 2, Chromatic can also develop more compelling functions such as cable modem and ADSL-related support.

Mpact Delivers Strong Performance

Figure 5 shows some Mpact performance data provided by Chromatic, all compared to the performance of a P55C-200 system with basic graphics and audio subsystems. Some of the numbers are based on simulations and others on real measurements. Mpact's biggest gain over the P55C alone is in its video-motion estimation, since the current MMX does not include any instructions to aid motion estimation. Therefore, we expect Mpact-based products to perform well for applications such as video authoring, videophone, and videoconferencing, where motion estimation is required.

The figure shows a 2× improvement in overall multimedia performance (a composite of the seven main functions) between Mpact/3000 and Mpact 2. The enhancements in Mpact 2 have a much bigger impact on 3D graphics, delivering an 8× improvement in this area. Note the comparison in Figure 5 is against a P55C system with no 3D hardware support; we expect a P55C with a midrange 3D accelerator (not just a "free-D" chip) to match the performance of the Mpact 2 system. The price of Mpact 2 must be close to that of the 3D chip for the Chromatic part to be competitive.

Mpact "R" Offers Intermediate Solution

Table 1 lists media processors that Chromatic and its partners plan to offer in 1997. There are actually three versions of the M1 planned. The first is called Mpact/3000, reflecting its 3,000 MOPS of peak throughput. The second version, Mpact R/3000, is identical except for the inclusion of an on-chip RAMDAC (the same one carried forward into Mpact 2). The third offers a 20% speedup by clocking the RDRAM at 300 MHz and the CPU core at 75 MHz. This version is called Mpact R/3600, since it produces 3,600 MOPS.

Chromatic is already testing first silicon of the Mpact R design and expects LG and Toshiba to provide samples later this quarter, with production versions in 1Q97. The company had originally hoped to see widespread use of its Mpact/3000 in PCs during this Christmas season, but it has missed this window of opportunity. At this point, we expect many PC makers to wait until the integrated RAMDAC is available, reducing system cost and board footprint compared with the initial implementation. The Mpact R should quickly replace the original Mpact chip in the product line, with the two speed grades offering a price/performance tradeoff. When Mpact 2 reaches the market next fall, it will become the new high-end part.

Mpact Competes Effectively

Today, more than a dozen vendors are involved in designing and/or manufacturing media processors.

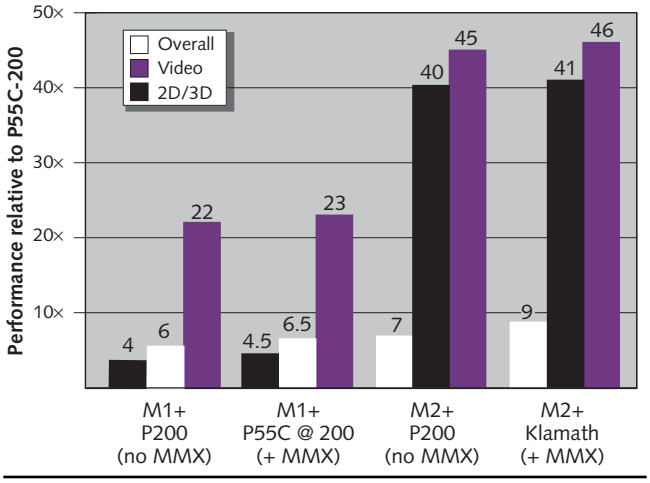

Figure 5. The M1 delivers a performance boost over a basic P55C system on most multimedia functions, as shown by the "overall" bar. The M2 further improves performance, particularly on video encoding and on 3D graphics. In all cases shown, the Mpact runs at 62.5 MHz and the Mpact 2 at 125 MHz. (Source: Chromatic)

The critical questions are how these devices will compete among themselves and how they can effectively compete against software-only and hardwired approaches.

The advantages of a media processor over a hardwired multimedia IC are easy-to-achieve functional integration, flexibility for evolving standards and improving algorithms, and the ability to deal well with real-time applications such as videoconferencing and surfing the Internet. The real-time advantage is due to the lack of real-time support in today's Windows 95 and NT. The programmability of media processors makes it easy to incorporate some kind of real-time kernel that can coexist with Windows. With dynamic resource allocation, a media processor can devote all its horsepower to the current application. This feature offers a potentially lower cost than the hardwired approach. In the hardwired system, if the current application doesn't use a particular function block, transistors dedicated to that function are wasted. The main disadvantages of media processors are programming overhead and design complexity.

A more interesting topic is how the Mpact architecture competes with other media-processor architectures. Like Mpact, Samsung's MSP (see MPR 8/26/96, p. 1) handles seven common multimedia functions. MSP's advantages

	Mpact/ 3000	Mpact R/ 3000	Mpact R/ 3600	Mpact 2/ 6000
MOPS (peak)	3,000	3,000	3,600	~6,000
Internal RAMDAC	No	Yes	Yes	Yes
RDRAM Bandwidth	500 MB/s	500 MB/s	600 MB/s	1,300 MB/s
3D Rendering H/W	No	No	No	Yes
Cache Size	4K	4K	4K	8K
AGP, 66-MHz PCI	No	No	No	Yes
Samples	Done	4Q96	4Q96	1Q97
Production	Now	1Q97	1Q97	3Q97

Table 1. By 2H97, Chromatic and its manufacturing partners plan to bring four Mpact media processors into the market. (Source: Chromatic)

include an open software model, better potential 3D graphics, and a larger initial market. Also in MSP's favor is that Microsoft has chosen MSP as one of the key components for its Talisman reference design. Microsoft's real-time kernel will likely work with MSP first, and it will take extra time to port it to other architectures, such as Mpact, if Microsoft or Chromatic decides to do so.

The Chromatic chip, on the other hand, is available now, costs less than MSP, and takes advantage of the computing power of the host CPU. Unlike Samsung's media processor, which replaces the host CPU when running multimedia tasks, Mpact assists the CPU. Mpact uses its resource manager for dividing multimedia workloads between the CPU and Mpact in real time. Under this cooperative computing environment, Mpact should be able to compete effectively in the PC market against any of the other media processors that have been announced so far.

Mpact's biggest advantage over other media processors is that it is real and in production; the others are simply prototypes. Trimedia says its TM-1 (see MPR 11/13/95, p. 22) will enter production in 1Q97, but TM-1's software will not be completed until the end of 1997, a full year behind schedule. In the PC business, time to market is everything. If Chromatic and its partners execute properly, we don't believe that other media processors for PCs will present any significant threat to Mpact in 1997.

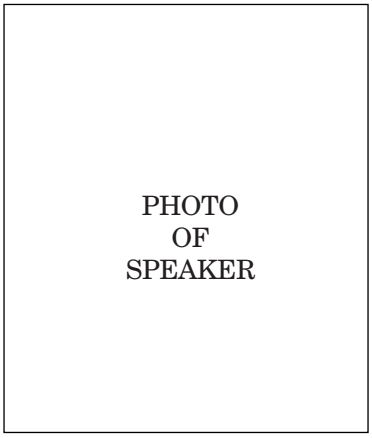

MICHAEL MUSTACCHI

At the Forum, Chromatic founder Stephen Purcell announces the second-generation Mpact 2.

Open ISA Has Advantages

The company's unusual business model includes a closed software environment: Chromatic is the only vendor allowed to develop software for Mpact. This restriction creates competitive disadvantages, especially since most other vendors claim their media-processor architectures are open. But Chromatic and its partners have nothing to lose by saying their architecture is also open, and we believe they will do so in the future.

The potential upside for opening the Mpact instruction-set architecture (ISA) is huge. For example, unlike Samsung and Philips' TriMedia, Chromatic focuses mainly on PC applications today. Mpact could be used for Macs and non-PC applications, but only if software is written for those applications, and Chromatic is not interested in doing so right now. The startup company is already stretched developing the required Mediaware for PCs.

Opening the software environment creates possibilities for third-party ISVs to bring Mpact to non-PC markets. This would allow Chromatic to use its limited software resources for PC applications while expanding the potential base of Mpact processor sales.

Enabling Mpact for non-PC applications is an important way to secure the future of Mpact. Non-PC devices will greatly expand the potential market size and make Mpact an interesting CPU for network computers, set-top boxes, videoconferencing devices, videophones, car navigators, and so on. There are many advantages to staying away from the "Wintel" architecture, where Intel and Microsoft dominate the entire platform. The majority of the Mpact development for non-PC applications will be software-related, but to encourage this, Chromatic must open its ISA.

Future Devices Are Promising

Mpact/3000 has proved Chromatic has a working architecture and not just a paper tiger. It also provides a means for kicking off Mpact software development. We believe that its lack of integration and 3D performance are limiting its acceptance, however. The forthcoming Mpact R/3600 and Mpact 2 parts solve these problems and should help propel Mpact into mainstream PCs.

Now that the Mpact 2 design has taped out, we expect Chromatic and its manufacturing partners to develop Mpact 2 derivatives. An obvious one is a notebook version that adds support for an LCD display and advanced power management. A notebook Mpact 2 would be a more compelling product than its desktop counterpart. First, the Mpact design saves space, power, and cost compared with a flock of fixed-function sound, graphics, video, and modem chips. Second, due to thermal restrictions, notebooks tend to use less powerful host CPUs. Therefore, the host needs more help to achieve strong multimedia performance.

Chromatic has already started the design of its third-generation Mpact, codenamed M3. We expect the M3 will be fully Talisman-compliant, since Mpact must support the PC entertainment applications for which Talisman is designed. In addition, transforming from its own real-time kernel to the Talisman real-time kernel can be beneficial in providing a standard software interface for applications.

By late 1998 or 1999, the IEEE 1394 (FireWire) interface will be important in PCs. Therefore, integrating a 1394 interface into the M3, or at least one of its derivatives, would be useful. Besides adding 1394, some M3 derivatives may include new features such as a video codec or frame-buffer memory on the chip.

To succeed in PCs, Mpact must offer a significant performance gain as a complement to MMX. Mpact is good at real-time events, isochronous applications, and video encoding and decoding. Together with an MMX host processor, Mpact offers significant value over the MMX-only or fixed-function approaches for future multimedia-ready PCs. This value should gain Mpact design wins on both PC motherboards and add-in cards. M

Fujitsu Aims Media Processor at DVD

MMA Combines Long-Instruction-Word Core, Integrated Peripherals

by Peter N. Glaskowsky

Hoping to claim a share of the growing consumer multimedia market, Fujitsu's Shunsuke Kamijo described the company's new multimedia assist (MMA) processor at last month's Microprocessor Forum. The new architecture features a two-pipeline long-instruction-word (LIW) core capable of executing up to six 16-bit integer operations simultaneously for a peak rate of 1.08 GOPS at 180 MHz. This is a higher clock rate than most other announced media processors, the result of a design that favors simplicity over sophistication.

The first chip in the family adds two 8K SRAMs; graphics, DMA, and SDRAM controllers; and several integrated peripherals to the MMA core. Fujitsu did not announce pricing or availability; we expect to see this product in limited sampling by the end of the year, with production volumes available in the first half of 1997.

Unlike other recently announced media processors, MMA is intended for consumer electronics, such as intelligent televisions and DVD players. MMA can perform DVD decoding, modem functions, and videoconferencing in software (but not all at once). MMA provides no floating-point support, 3D acceleration features, or PCI interface, making it unattractive for the PC market.

MMA Core Includes Five Execution Units

The heart of MMA is its LIW core with two pipelines and five execution units, as Figure 1 shows. Each 64-bit instruction word contains a pair of 32-bit instructions. One instruction is always dispatched to the primary pipeline and the other to the secondary pipeline. This simple LIW is much less sophisticated than the VLIW architecture of Philips' TriMedia (see MPR 11/13/95 p. 22). In TriMedia, each instruction word contains up to five operation slots, and there are 27 execution units. Each unit is fully pipelined, allowing the TriMedia core to sustain five operations per clock in long code segments.

In MMA, one pipeline contains a 32-bit ALU and the load/store unit. The other contains a second ALU along with multiply-accumulate and divide-shift units.When the primary ALU is used in combination with the multiply-accumulate unit in the secondary pipeline, the MMA core can achieve a peak execution rate of three 32-bit or six 16-bit operations per cycle.

MMA's instruction-pairing rules are relatively simple. All five execution units are fully pipelined, yielding single-cycle throughput, so one instruction can be dispatched to each pipeline in each clock cycle. Division and modulo operations are an exception: these take 33 to 36 cycles to complete and stall both pipelines while they execute. Taken branches require three cycles; one branch-delay slot is provided, which may be used for any single-cycle operation.

Fujitsu also provides 27 instructions for multimedia operations such as MPEG decoding and surround-sound processing. These include variations of ADD, SUB, and MAC instructions with signed and unsigned saturating arithmetic as well as 16-bit SIMD versions of most instructions.

Although the load/store unit supports byte-size data types, the ALUs support only 16- and 32-bit data. Other media processors support byte operations, allowing twice the peak performance when dealing with byte-oriented data types like RGB pixels or 8-bit sound.

MMA provides an unusual saturation mode. A source register may specify the bit at which saturation occurs. This mode allows 32-bit operations to perform 24-bit saturating arithmetic, for example. This should prove useful, since few multimedia data types are exactly 16 or 32 bits in length. For example, MPEG decoding uses 9-bit values, and AC-3 audio requires 24-bit precision. Saturation to an arbitrary (non-power-of-two) value, however, is not supported.

Register File Matched to Pipeline Requirements

The MMA core includes a 32 × 32-bit multiported register file. It supports five simultaneous reads and three writes, enabling sustained single-cycle throughput on the inner loops of common multimedia functions.

This register set is smaller than in most competing media processors. For example, TriMedia has 128 32-bit registers. These larger register sets are especially useful for 3D acceleration. MMA's register set was designed for algorithms like the inverse discrete cosine transform (IDCT) in MPEG decoding. These algorithms typically have working sets that will not fit in a register file; instead, registers are used to store control values and coefficients, and data storage depends on fast access to memory.

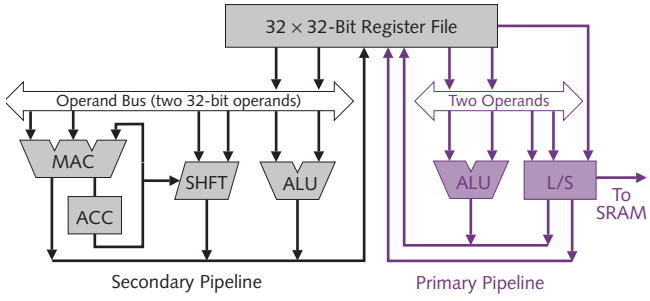

Figure 1. The MMA core includes two pipelines that operate in parallel. The primary pipeline supports ALU and load/store operations, while the secondary pipeline can perform multiply-accumulate, shift, and ALU operations.

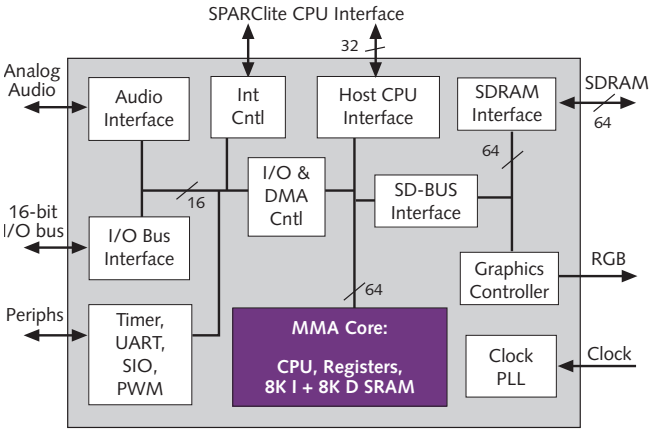

Figure 2. To support consumer multimedia applications like intelligent television, MMA includes many functional blocks in addition to the programmable MMA core.

SRAMs, Not Cache, Meet Core Bandwidth Needs

Rather than relying on caches to reduce average memory access penalties, MMA provides an 8K SRAM instruction store plus an 8K SRAM data store. Each SRAM is organized as 1K 64-bit words, and both operate at the full 180-MHz rate of the core. The instruction SRAM can provide a 64-bit dual instruction word to the core on each clock, while the data SRAM can transfer one, two, or four bytes to or from the register file in each cycle.

Unlike caches, the SRAMs must be managed by software. Transfers between the SRAMs and the local SDRAM can be performed only by the on-chip DMA controller, which is interlocked with the core. During a DMA transfer, the core is stalled. The DMA controller supports normal block transfers at the maximum rate of the DRAM interface, as well as rectangle transfers, a way to realign raster-oriented video data. DMA transfers are controlled to byte boundaries, a necessary feature given the unpredictable block lengths of digital audio and video data.

In effect, MMA trades the complexity of a cache controller for added software complexity. This tradeoff would be unacceptable in a general-purpose processor, but MMA is designed to execute simpler multimedia code that can be optimized for this architecture. A similar scheme is used by Chromatic's Mpact, although Mpact's software "caches" are multiported and offer much higher net throughput.

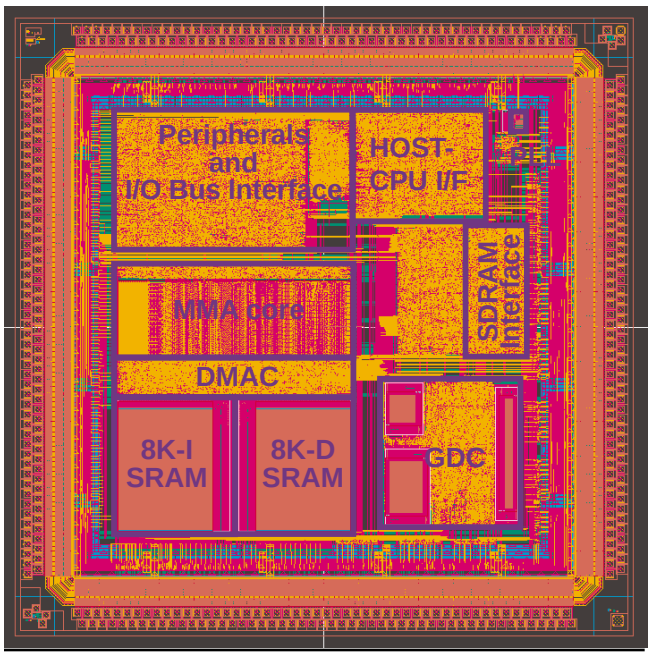

Figure 3. This is a die plot of MMA, which is fabricated in a 0.35-micron three-layer-metal CMOS process. The die size is 77 mm^2, with 1.3 million transistors. The MMA core is only 4.3 mm^2 in size.

SDRAM Interface Sustains 1-Gbyte/s Bandwidth

Figure 2 shows the internal organization of the first MMA product. The MMA core includes the LIW engine, SRAMs, and DMA controller. The integrated SDRAM controller manages up to 32M of SDRAM operating at up to the 180-MHz pipeline rate. While current SDRAMs do not support this high speed, future parts will. At 180 MHz, the peak data rate is 1.4 Gbytes/s, but sustained rates will be lower due to bus turnaround delays and page misses.

In more realistic implementations, the SDRAM interface will run at half the core speed, or 90 MHz, yielding 720 Mbytes/s peak and about 500 Mbytes/s sustained throughput. The SDRAM interface supports four-word burst transfers and is fully pipelined, so subsequent reads to the same DRAM page do not cause wait states.

Integrated Peripherals Adapt MMA to TV

To support intelligent televisions, MMA includes a graphics display controller (GDC) module. In addition to interlaced and noninterlaced NTSC and VGA resolutions, the GDC also supports a "wide-VGA" mode of 860 × 480 pixels that can be used to enhance quality on high-end televisions. Such televisions offer greater horizontal than vertical resolution, so typical "square pixel" display modes like 640 × 480 do not achieve the best possible visual quality. MMA's wide-VGA support allows processor-generated content like the graphical user interface to be displayed directly in the higher resolution, while some digital video content, such as wide-screen-mode DVD playback, can be scaled to fill the wider effective screen size.

The GDC manages three display windows, each with a separate frame buffer in the local SDRAM. The windows can be positioned and overlapped arbitrarily. Frame buffers can contain pixels in the YCrCb color space for MPEG decoding, RGB for user-interface displays, or a 15-color-plus-transparency mode for processor-generated captions and simple graphics. As the screen is drawn, these pixel types are all converted to standard digital RGB using a color-space conversion engine in the GDC.

The GDC shares access to the SDRAM with the DMA controller and host processor. SDRAM refresh activity can

be synchronized to the display controller, taking place during the horizontal refresh interval. This eliminates the need for deep FIFOs in the display refresh path, since the display controller can always depend on uninterrupted access to the SDRAM during the display period of each scan line.

Other peripherals on the die include a timer, two serial I/O controllers, and a pulse-width-modulation module for motor control, plus interfaces for an audio chip and other off-chip peripheral devices. MMA does not include a RAMDAC, however, which would have been fairly easy to add given the relatively low resolution display modes it supports.

MMA is not designed for general-purpose tasks like control and communications. Instead, it will typically be used as a coprocessor for multimedia tasks, coupled with a general-purpose host CPU like Fujitsu's SparcLite.

The first MMA implementation has a 32-bit SparcLite interface built in. In this configuration, MMA acts as a unified memory architecture (UMA) DRAM controller for SparcLite, storing the frame buffer plus code and data for both processors in the local SDRAM. A typical intelligent TV controller would consist of the MMA, a SparcLite processor, ROM, RAM, a RAMDAC, and a few analog interface components.

Software Development Tools

Fujitsu has made a special effort to enable third-party software development, but it has not announced specific third-party relationships. Fujitsu commissioned a full suite of development tools from Green Hills (*www.ghs.com*), extending the existing set of SPARC tools to include an assembler and simulator for MMA as well as an MMA-aware version of the Green Hills Multi development environment. Multi gives an MMA programmer a unified environment with separate windows for SPARC and MMA operations.

PHOTO OF SPEAKER

MICHAEL MUSTACCHI

Shunsuke Kamijo describes Fujitsu's first media processor at the Microprocessor Forum.

Fujitsu has also developed a version of Wind River's VxWorks (*www.wrs.com*) that runs on the SparcLite/MMA target system, supporting remote-control debugging operations through VxServ. VxWorks runs only on the SparcLite processor; at this time, Fujitsu has no real-time OS kernel for MMA itself.

Fujitsu is developing its own set of essential software libraries for MMA. At the Forum, Fujitsu's Kamijo showed its development schedule for six library functions: MPEG decode and encode, JPEG decode and encode, JBIG (a lossless still-image compression scheme), and the V.34 modem algorithm. All of these functions are projected to be available by 1Q97.

Benchmarks Demonstrate MMA Performance

Kamijo showed limited benchmarks based on MPEG-1 performance. The 180-MHz MMA can decode video-only MPEG-1 bitstreams at 118.4 frames per second (fps), suggesting that standard 30-fps MPEG-1 requires only 25.3% of the device. When standard 48-kHz audio is included, the decode rate dropped to 94.1 fps, increasing utilization for 30-fps MPEG-1 to about 32%.

Although Fujitsu has not released benchmarks for DVD applications, the company says MMA will be able to perform MPEG-2 video plus Dolby AC-3 audio decoding at 30 fps, meeting the basic requirement for DVD support. This is a critical capability; without DVD capability, MMA could be relegated to the much smaller market for video karaoke players and other MPEG-1 products.

Fujitsu did not discuss performance on other multimedia tasks like V.34 modem operation, but it is unlikely that MMA will be able to support a V.34 modem connection while simultaneously decoding DVD content.

For More Information

Pricing and availability for MMA have not been announced. Contact Fujitsu (San Jose, Calif.) at 408.922.9574 or on the Web at *www.fujitsumicro.com.*

Goals Determine Results

The first MMA is shown in Figure 3. The die size is 77 mm^2, fabricated in a 0.35-micron three-layer-metal process, operating at 3.3 V and packaged in a 352-ball BGA. Based on the MDR Cost Model, the estimated manufacturing cost of this part is $30. The MMA core, a full-custom design, occupies only 4.3 mm^2 of the die. Even at this small size, it offers more than enough processing power for the target applications. Power consumption for the device is very low, at only 600 mW (typical). This compares very well with Trimedia's TM-1 at about 4 W, offering an advantage in consumer applications.

MMA is a well-balanced part that should compete effectively against non-programmable devices in DVD players while also working well in products that require more intelligence. Fujitsu's primary competition will come from dedicated DVD decoders from C-Cube and Oak (see page 5). Hardwired devices may be smaller and less expensive than MMA, making them a better choice for cost-sensitive DVD players, but MMA is more flexible, making it more suitable for an intelligent TV.

In addition, MMA is likely to be much less expensive than PC-oriented media processors like TM-1 due to its smaller SRAMs and lack of floating-point support and 3D acceleration, making it a good fit in the embedded multimedia applications for which it was developed. M

TurboSparc Offers Low-End Upgrade

New Fujitsu Chip Plugs into MicroSparc-2 Systems for Performance Boost

by Linley Gwennap

Eyeing an installed base of underpowered SparcStation 5 workstations approaching half a million units, Fujitsu has developed a new CPU as a field upgrade for those systems. The TurboSparc processor is also appropriate for new low-cost workstations. Following the MicroSparc tradition, TurboSparc is a highly integrated 32-bit processor with on-chip cache, memory, and SBus interfaces, as Figure 1 shows. The performance improvement is modest, however, moving users from the equivalent of a 486 up to a midrange Pentium. TurboSparc is shipping now, both as a standalone device and in an upgrade kit.

Fujitsu estimates the 170-MHz TurboSparc will deliver 3.5 SPECint95 and 3.0 SPECfp95 (base), although current systems need more compiler tuning to achieve these marks. If achieved, these scores would represent roughly twice the integer performance of a 110-MHz MicroSparc-2 (MS-2) and about 50% better floating-point performance. The gain over slower versions of MS-2 will be even greater. This boost makes TurboSparc attractive as a field upgrade. This performance, however, is well below that of a Pentium Pro, PowerPC 604e, or even a high-end Pentium, all of which sell for about the same price as the $499 TurboSparc.

New CPU Core Boosts Clock Speed

Although Fujitsu builds and sells MicroSparc-2, that chip design was developed and is owned by Sun. For its upgrade chip, Fujitsu chose to develop its own CPU core, deploying a small team in San Jose (Calif.). For cost and time-to-market reasons, the TurboSparc team chose a simple scalar CPU design based on the 32-bit SPARC v8 architecture. In this regard, TurboSparc is similar to MS-2 but quite different from the superscalar 64-bit UltraSparc.

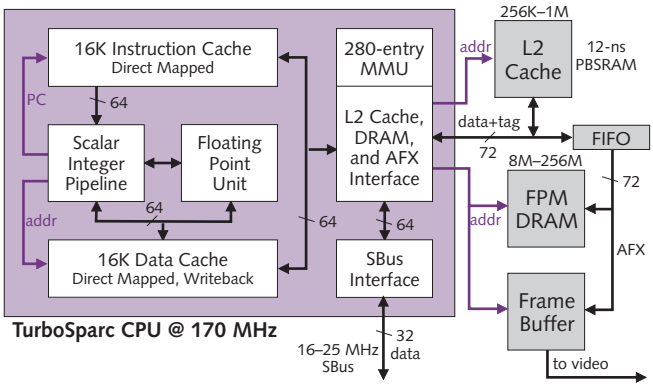

Figure 1. TurboSparc integrates nearly all the memory and system interfaces needed for a complete low-cost workstation.

Because of the similar throughput of the MS-2 and TurboSparc cores, the higher clock speed of the latter part provides a large part of its performance gain. TurboSparc clocks 55% faster than its predecessor. Some of this speed advantage is due to a gate shrink from 0.4-micron CMOS to 0.35-micron.

Most of the speed gain, however, comes from a new pipeline. Whereas MicroSparc-2 uses the classic five-stage RISC pipeline (see MPR 11/15/93, p. 1), TurboSparc extends it to six stages for integer instructions and eight for FP instructions. As Figure 2 shows, a new "resolve" stage, after the data-cache access, checks for any faults from the cache access before proceeding to the writeback stage. This new stage provides more time for the cache access to complete, avoiding a critical timing path. The elimination of branch folding, a feature found in MS-2, also eases the timing and helps achieve better clock speeds.

The other new stages handle FP instructions, forming an integrated integer/FP pipeline that simplifies the control logic. Most FP operations, including ALU ops and multiplication, have a four-cycle latency and thus are complete by the FR stage. FP divide and square root both process two bits per cycle and can take from 8 to 50 cycles to complete; the average is 21 cycles for single-precision operands and 35 cycles for double-precision operands. The FP unit also handles integer multiply and divide operations; multiplication has a seven-cycle latency, while division is the same as a single-precision FP division.

One unusual feature of the new CPU is its ability to handle branches without penalty and without prediction. The instruction cache provides two instructions per cycle, while the CPU consumes only one. On a branch, there is enough bandwidth to fetch from both the taken and non-taken paths until the branch condition is resolved, eliminating any branch penalties, as Figure 3 shows. This method avoids the complexities of branch prediction or branch

Fetch	(F) Fetch two instructions from I-cache
Decode	(D) Decode one instruction and read operands
Execute	(E) Execute integer operation
Memory	(M) Read tags and data from D-cache
Resolve	(R) Check tags, abort data if L1 cache miss
Write	(W) Write result to integer register file
FP Resolve	(FR) Complete FP ALU or multiply (4 cycle latency)
FP Write	(FW) Write result to FP register file

Figure 2. TurboSparc's pipeline adds one stage to MicroSparc-2's for integer instructions and two more stages for FP instructions.

folding and is similar to the technique used by QED in the RM7000 processor (see MPR 10/28/96, p. 36).

As Figure 1 shows, TurboSparc contains 16K each of instruction and data cache. The instruction cache is the same size as in MS-2, but the data-cache size is doubled relative to that chip and is now a write-back cache rather than write-through. These caches are fairly small compared with those of recent microprocessors and are direct mapped, further reducing their hit rate. The MMU is SPARC v8 compliant. It includes a 256-entry TLB for data translations, four more entries for large data pages, a four-entry instruction TLB, and a 16-entry I/O TLB. The total number of TLB entries is four times more than in MS-2, eliminating another performance bottleneck for some applications.

DRAM Interface Now Supports L2 Cache

The second major performance enhancement from MS-2 is the addition of an external level-two cache. MicroSparc-2 is completely limited to the paltry 24K of cache that is available on the chip, relying on direct access to external DRAM for all other memory references. This shortfall creates a significant performance degradation on SPEC95 as well as on many workstation applications.

TurboSparc supports 256K–1M of direct-mapped L2 cache. This cache has the same 32-byte line size as the on-chip caches and uses a write-through policy. It runs at one-half the core CPU speed, requiring 12-ns pipelined burst SRAMs for the 170-MHz processor. An access that hits in the L2 cache stalls the CPU pipeline for 12 cycles. With a 72-bit interface, it takes four cycles to return a full cache line. A one-third-speed cache is also supported, but this choice will reduce performance.

Using ×36 SRAMs, the cache tags are stored side by side with the data. Each cycle, the 72-bit interface returns 64 bits of data, 2 parity bits, and 6 tag bits. After the first two accesses, the complete 12-bit tag can be assembled and checked to see if the access has hit in the cache. This design eliminates the need for separate external cache tags (or internal tag storage). The tags cannot be checked until two cycles after the data is received, but the new R stage allows enough time to abort the writeback if the tag check fails.

DRAM accesses are started in parallel with L2-cache accesses and are aborted if the L2 cache hits. This strategy reduces the duration of a pipeline stall on an L2 cache miss by overlapping the DRAM access. The DRAM interface is configurable for page-mode DRAM of various speeds but does not handle more advanced memories such as EDO or SDRAM. These memory types are not supported in the older SparcStation systems, so Fujitsu did not bother to add them, keeping the design as simple as possible to speed its completion. With 60-ns DRAM, the CPU pipeline stalls for 24 cycles on an access to main memory.

The memory controller supports up to eight banks of 32M each, or 256M maximum. This limitation is similar to that of MicroSparc-2.

Price & Availability

The 170-MHz TurboSparc chip is available now at a list price of $499 in quantities of 1,000. The 160-MHz TurboSparc upgrade kit costs $1,500 in quantities of one. To get more information on TurboSparc, contact Fujitsu (San Jose, Calif.) at 800.866.8608 or access the Web at *www.fujitsumicro.com/sparcupgrade/sparcmicro.html*.

Sun AFX Graphics Supplements SBus Interface

Recently, Sun has added a new graphics interface called AFX. These graphics cards reside on the main memory bus instead of the pedestrian SBus, significantly improving bandwidth. AFX requires adding only a few extra control signals to the existing memory bus, which Fujitsu has done in TurboSparc. This change allows an end user to plug an AFX graphics card into a system that has been upgraded with TurboSparc.

Like MS-2, the new processor supports SBus directly on the chip. The SBus operates at 16–25 MHz, typically one-eighth of the CPU clock speed. Up to six bus masters can be connected to the 32-bit SBus. TurboSparc is fully compatible with the Macio and Slavio chips that supply basic I/O functions in the SparcStation 5 and similar systems.

The integrated memory and bus interfaces make multiprocessor configurations impossible. This fact simplified some aspects of the TurboSparc design. The cache does not support multiprocessor coherency, for example, and the CPU core does not execute certain SPARC v8 instructions for MP synchronization.

Manufacturing Cost Shrinks

TurboSparc is built in Fujitsu's 0.35-micron four-layer-metal CS-60ALE, keeping the die size down to 132 mm^2, relatively svelte for a processor with so much integrated system logic. MicroSparc-2, by comparison, weighs in at 233 mm^2 using the 0.4-micron CS-55 process (see MPR 7/10/95, p. 16).

Although the gate shrink is minor, a bigger gain is seen in the metal layers: the CS-55 metal layers are from a 0.5-micron process, whereas CS-60ALE is a complete 0.35-micron process. Thus, the TurboSparc die size is about what we would expect if the MS-2 die was shrunk to the same 0.35-micron process.

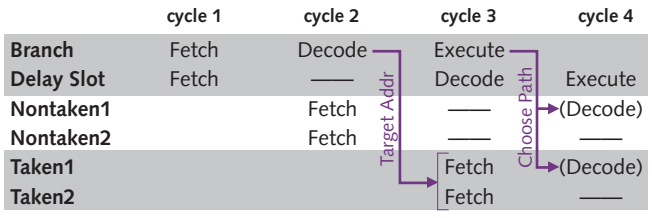

Figure 3. On a branch, TurboSparc fetches from both the nontaken and taken paths. By cycle 4, the branch is resolved, and either instruction can be executed without penalty. (— indicates stall)

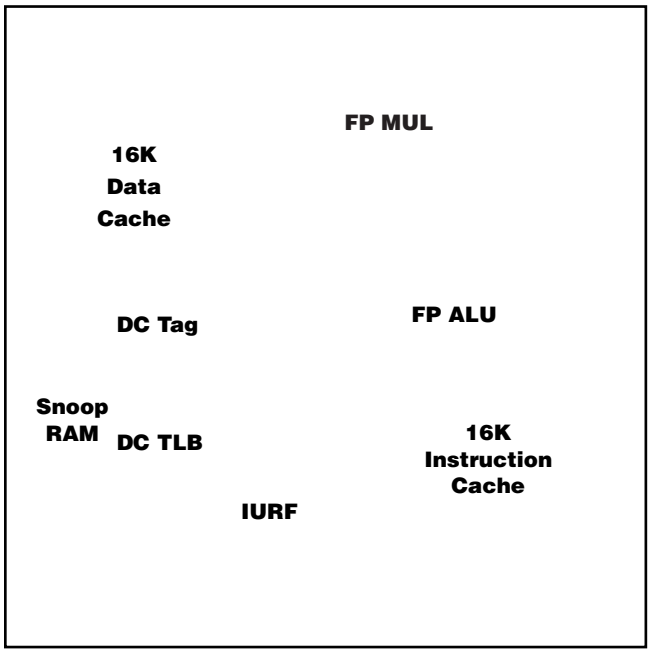

Figure 4. Fujitsu's TurboSparc combines a scalar SPARC CPU with a complete set of system interfaces. Sporting 3 million transistors, the die measures 11.5 × 11.5 mm in a 0.35-micron four-layer-metal CMOS process.

The transistor count of TurboSparc is slightly higher: 3.0 million, compared with 2.3 million for MS-2. Most of the increase is due to the extra 8K of cache, with the remainder in the L2 cache and AFX interfaces. The CPU core has about the same number of transistors as in MS-2. Because the physical layout of MS-2 is rather loose, Fujitsu was able to pack more transistors into the same relative die area. Figure 4 shows the TurboSparc die.

TurboSparc is packaged in a 416-contact plastic BGA. The plastic BGA saves cost compared with the old-style ceramic PGA used for MS-2, despite the extra 95 leads required by the new interfaces. Combining the savings from the plastic package and the smaller die, the MDR Cost Model estimates the cost of building TurboSparc at about $50, a third less than the cost of MS-2. The PowerPC 603e and Pentium chips deliver similar performance at a build cost of $30–$40, but these chips cannot match the integrated system logic of TurboSparc.

Like MS-2, the new chip runs at 3.3 V. The maximum power dissipation is 9 W at 170 MHz, matching the maximum power of MS-2 despite the significantly higher clock speed. Fujitsu paid more attention to moderating power in the new design, adding some gated clocking to keep the chip within the same thermal envelope as its predecessor.

Module Upgrades MicroSparc-2 Systems

Designing a TurboSparc upgrade for MicroSparc-2 systems was no easy task. To gain the necessary performance boost, TurboSparc adds a secondary cache, but this cache is obviously not present on a MicroSparc-2 motherboard. Thus, Fujitsu has designed a module containing a TurboSparc processor and 256K of L2 cache implemented with two 32K×36 SRAMs. The module is a small PC board with a pin-grid array on the bottom that plugs into an MS-2 socket.

The PC board has an odd "L" shape to fit into the existing SparcStation 5 design, carefully avoiding all obstacles. Fujitsu believes this board will also fit into most other MS-2 workstations. The TurboSparc processor has its own fan mounted on the heat sink to ensure adequate cooling. An onboard voltage regulator delivers the extra current needed by the SRAMs and buffers.

For yield reasons, the modules use a 160-MHz TurboSparc; the company is saving the 170-MHz parts for customers buying standalone chips. Fujitsu is marketing the TurboSparc module directly to end users in an upgrade kit that contains the module, documentation, and an extraction tool to remove the MS-2 chip. The kit, which retails for $1,500, also contains a new PROM with the appropriate boot code for the new processor.

From Woeful to Weak

In addition to the upgrade kits, new systems from several small SPARC system vendors are using TurboSparc. Sun, however, has been conspicuously absent among vendors adopting the new chip. This oversight is surprising: with MicroSparc-2 delivering just 1.4 SPECint95 and 1.9 SPECfp95 (base), the performance of Sun's low-end workstation line is quite woeful by current standards.

TurboSparc offers a significant boost but still barely matches the SPEC performance of a good 120-MHz Pentium box on both integer and floating-point code. A high-end Pentium PC will run rings around a TurboSparc workstation on many technical applications while costing less than half as much. Fujitsu argues that SPEC exaggerates the performance difference because Intel's SPEC results rely on far more compiler tuning than Fujitsu's estimates.

In any case, the audience for TurboSparc remains diehard SPARC advocates who need compatibility with a large installed base of SPARC hardware and software. There are plenty of Sun-only shops around to which TurboSparc is appealing, both as a field upgrade and in new systems. Even some of these diehards, however, are eying the low cost of x86-based systems. Users not tied to SPARC will see little attraction in TurboSparc systems.

Sun's longer-term solution for this price point is UltraSparc-2i (see MPR 10/7/96, p. 1). This processor is slated to exceed the performance of Intel's P6 chips while including a set of integrated system logic similar to TurboSparc's. US-2i cannot provide a field upgrade for MicroSparc-2 systems, however, and is not due to appear in systems until 4Q97, a year from now. Until then, TurboSparc will have to power low-cost SPARC systems from Sun and others. But the new chip may not have enough power to get these systems safely through the wake created by the Pentium and P6 workstations now entering the market. M

ARM Tunes Piccolo for DSP Performance

Full DSP Module to Add Signal-Processing Capacity to ARM7 Chips in 2H97

by Jim Turley

Addressing a growing chorus of users seeking better digital-signal-processing performance, ARM has composed its own signal processor. Named Piccolo (Italian for small; apparently no suitable human appendages were left), the new DSP core will become yet another optional piece of the modular ARM portfolio. Working in concert with an ARM7 core, Piccolo should boost ARM's fortunes with makers of wireless devices, PDAs, modems, and disk drives.

ARM architect Dave Jaggar described Piccolo at last month's Microprocessor Forum. Never a vendor to adhere blindly to convention, ARM developed an unusual new DSP core while preserving the traditional ARM merits of low transistor count, small die size, and modest power consumption. The Piccolo core should be available in 10–12 months, debuting in application-specific devices from an unnamed vendor, possibly in cellular phones or pagers.

Piccolo does not replace the ARM core but rather works in parallel with it, expanding a chip's DSP capabilities. The concept is not new and is reminiscent of Hitachi's SH-DSP (see MPR 12/4/95, p. 10) and Motorola's 68356 (see MPR 6/20/94, p. 9). TI has also combined its TMS320C54x DSP with an ARM7 core.

Register File Backed with Reorder Buffer

Piccolo is an autonomous digital-signal processor sharing some resources with an ARM7-based chip. The DSP executes its own instruction set (which is incompatible with ARM code), uses its own registers, and follows its own control path. The ARM and Piccolo cores communicate solely through a pair of input and output buffers. The ARM core orchestrates overall chip control as well as accessing operands in memory, while Piccolo concentrates on signal-processing loops.

Piccolo has its own register set, independent of ARM's, as Figure 1 shows. The register set is nominally orthogonal, with sixteen 32-bit registers variously referred to either as d0–d15 or as a0–a3, x0–x3, y0–y3, and z0–z3. The upper and lower halves of all sixteen registers can alternatively be addressed as "half registers" for 16-bit operations. The first four registers, a0–a3, double as 48-bit accumulators for some instructions, notably multiply-accumulate.

Piccolo's register file is logically accessible from ARM code, although access is physically mediated by the input and output buffers. The registers are addressed with LDP (load to Piccolo) and STP (store from Piccolo) instructions, which are simple variations of existing ARM coprocessor instructions. These transfer up to 16 bytes of data between Piccolo registers and memory without passing through the ARM register file. Piccolo also has three special registers that hold ID information, status, and the DSP program counter.

The input buffer acts something like a data cache for the Piccolo registers. It is organized as eight 32-bit entries but can double as sixteen 16-bit entries as well.

Input, Output Buffers Decouple Two Cores

The input buffer is not simply a FIFO from ARM to Piccolo. Instead, it is managed more as a reorder buffer, or ROB. Each entry in the ROB is tagged with its destination register as the ARM core loads it. When a Piccolo register is freed, that register is automatically refilled with an entry from the ROB tagged for that register. If more than one ROB entry is tagged for the same register, Piccolo selects the oldest entry. (The oldest entry is identified by its position in the ROB.) If the same register becomes free again, Piccolo selects the next-oldest entry from the ROB destined for that register, and so on.

The ROB allows ARM code to fetch DSP data or coefficients from memory in whatever order is most convenient while allowing the DSP code to consume the items in what-

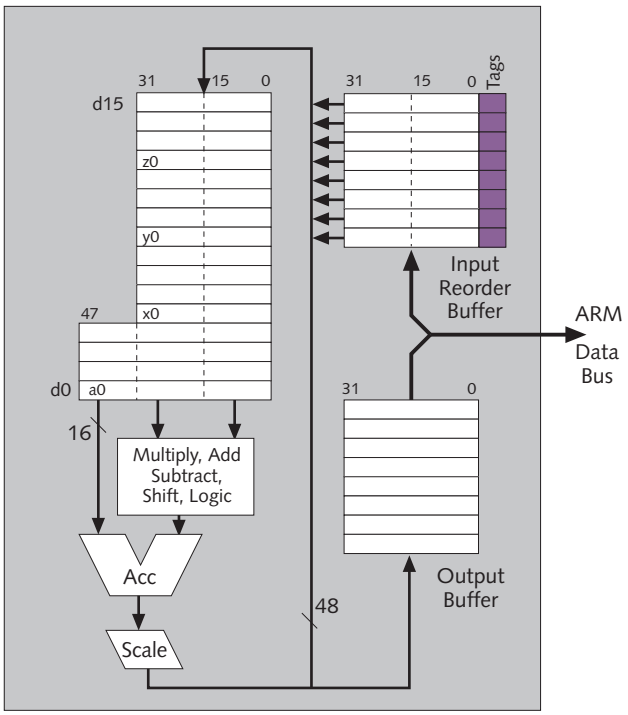

Figure 1. Piccolo includes its own register file and loads and stores its data from two buffers attached to the ARM data bus. The ARM core loads and empties the buffers with coprocessor instructions.

ever order they are required. For example, ARM code can load a block of four data points from memory with a single LDP instruction and write them to the ROB, then load a set of four coefficients from a different area of memory. Piccolo can then pull the data and coefficients from the ROB in pairs.

The output buffer is much simpler than the input ROB. It is a fairly straightforward 8 × 32-bit FIFO that queues results until ARM can retrieve them. The only special feature of Piccolo's output buffer is that it concatenates two consecutive 16-bit entries into a single 32-bit word. Thus, a single 16-bit result will not "fill" the first entry in the FIFO; a second 16-bit result is needed before ARM can remove the entry from the output buffer.

Input, Output Buffers Mediate Communications

If Piccolo tries to read from an empty input buffer, it will stall until ARM loads it. Conversely, if the input buffer is full when ARM tries to load it, ARM will stall until Piccolo removes at least one entry. The same is true of the output buffer: either Piccolo or ARM will stall if the buffer is full or empty, respectively. Jaggar optimistically described this condition as an opportunity to save power.

There is no signaling between ARM and Piccolo; the two cores are interlocked only through these two buffers. For instance, Piccolo cannot interrupt or otherwise indicate to ARM that the output buffer is full. It is the programmer's responsibility to synchronize integer and DSP code so that ARM services Piccolo's data-input and -output needs.

Piccolo is unaware of interrupts, exceptions, or faults that may occur while the ARM chip is running. Hardware interrupts, for example, do not stop Piccolo processing. If, during interrupt processing, the ARM CPU neglects to empty Piccolo's output buffer, Piccolo simply stalls until the output buffer is emptied again. Likewise, Piccolo stalls if its input buffer runs dry.

Piccolo fetches its own DSP code from on-chip or off-chip memory, which is cached in a private instruction cache. This cache is separate from any cache the ARM core might also have. Piccolo's cache is fully associative and holds at least 64 instructions, arranged in four or more 64-byte lines. Although this cache is small in size, ARM claims most common DSP inner loops fit easily into this space.

Programmers wishing to optimize performance can directly manipulate Piccolo's cache. By writing an address to a Piccolo coprocessor register, ARM code can force a Piccolo cache miss at that address, loading a line of 16 DSP instructions. Using this technique, control code can preload DSP code into some or all of Piccolo's cache. To launch DSP code, ARM writes to Piccolo's program counter.

Reordered Refill Reuses Registers

With Piccolo's limited resources and the data-hungry nature of most DSP algorithms, it is important to keep Piccolo's register file as busy as possible. After the contents of a register are used for the last time, Piccolo automatically reloads that register with new data from the ROB. Determining when a value is used for the last time is the programmer's responsibility.

Piccolo's instruction format includes a "refill" bit for each source operand. Programmers must flag the last use of

Mnemonic	Description	Mnemonic	Description	Mnemonic	Description
Arithmetic		**Arithmetic**		**Multiplication**	
ADD	Add	ADDADD	Parallel add, add	MUL	Multiply
ADC	Add with carry	ADDSUB	Parallel add, subtract	SMUL	Multiply, saturating
ADDA	Add and accumulate	SUBSUB	Parallel subtract, subtract	MULA	Multiply-accumulate
CAS	Add conditional	SUBADD	Parallel subtract, add	MULS	Multiply-subtract
SADD	Add, saturating	CMNCMN	Parallel compare (add/add)	SMULA	Multiply-accumulate, saturating
SUB	Subtract	CMNCMP	Parallel compare (add/sub)	SMULS	Multiply-subtract, saturating
SBC	Subtract with carry (borrow)	CMPCMN	Parallel compare (sub/add)	SMLDA	Multiply-accumulate, doubling
SUBA	Subtract and accumulate	CMPCMP	Parallel compare (sub/sub)	SMLDS	Multiply-accumulate, dbl, saturate
SSUB	Subtract, saturating	**Logical**		**Special**	
RSB	Reverse subtract	AND	Logical AND	EMPTY	Mark register for refill
SRSB	Reverse subtract, saturating	ORR	Logical OR	ZERO	Clear selected registers
RSC	Reverse subtract with carry	BIC	Bit clear (Logical AND NOT)	OUTPUT	Force registers to output FIFO
CAS	Conditional add/subtract	EOR	Logical exclusive-OR	**Loop Constructs**	
CASC	Conditional add/subtract w/carry	TST	Logical AND, no register write	RMOV	Set register-mapping parameters
CMP	Subtract w/o register write	TEQ	Logical EOR, no register write	REPEAT	Initiate loop
CMN	Add w/o register write	ASL	Shift left, arithmetic	NEXT	Terminate loop
SABS	Absolute value, saturating	ASR	Shift right, arithmetic	**Branch & Miscellaneous**	
MIN	Find minimum absolute value	LSR	Shift right, logical	Bcc	Branch conditional
MAX	Find maximum absolute value	ROR	Rotate right	SELcc	Select (conditional move)
MINMIN	Parallel minimum, minimum	MOV	Copy 16-bit immediate	SELTTcc	Parallel select, true/true
MAXMAX	Parallel maximum, maximum	CLB	Count leading bits	SELTFcc	Parallel select, true/false

Table 1. The Piccolo DSP core adds an entire new instruction set to ARM processors. A number of parallel instructions can operate on two 16-bit values simultaneously.

each register so a new value can be obtained from the ROB. The last use is specified in the assembler syntax by affixing a caret (^) to the register name. After the last use, hardware moves the oldest pending entry in the ROB to that register in the next cycle. A one-cycle load-use penalty is exacted if the immediately subsequent DSP instruction depends on the new value.

Instruction Set Parallels ARM

Piccolo instructions are encoded differently from ARM instructions; because they are executed in different pipelines, they don't have to be compatible. Like ARM, Piccolo instructions offer in-line operand scaling and optional flag updating. Unlike ARM, conditional execution is not supported; Piccolo uses conditional-branch instructions instead. Taken branches face a three-cycle penalty.

The current core implementation has four pipeline stages (fetch, decode and register read, execute, and write-back). At 40 MHz, a Piccolo-equipped ARM processor can execute 40 MOPS. Counting multiply-accumulate as two operations, or using parallel "split" instructions, Piccolo reaches 80 MOPS.

The complete Piccolo instruction set is listed in Table 1. The instruction set includes the usual add, subtract, and multiply operations (no divide), with both saturating and nonsaturating versions. Piccolo also has a number of "split" functions that work on two 16-bit operands at once, similar to MMX and other new media extensions. Just as with ARM, integer division is a stepped, iterative process that takes 3–19 cycles for 32÷16-bit division.

Multiply-accumulate, add-accumulate, and subtract-accumulate are also supported and use one of the 48-bit accumulators, a0–a3. The large accumulators avoid any possibility of overflow on 16 × 16-bit operations. With a 48-bit accumulator, even repeated multiply-accumulates would not overflow in less than 65,535 iterations, an ample number for any practical application.

The CLB (count leading bits) instruction reports the number of bits by which the source must be shifted left until its two most significant bits differ. That count can then be used for normalizing other numbers.

Every instruction has, conceptually, two possible destinations. Depending on the assembly syntax, Piccolo can store results to the register file, the output buffer, both, or neither. (This last combination, while seemingly pointless, is useful for creating nondestructive comparison or test instructions.)

Writing the results to the register file is the usual case when output values will be reused as input in a subsequent calculation. When the output value will not be reused, it can be written directly to the output buffer for later removal by ARM. Storing results to both the register file and the output buffer allows ARM code to retrieve intermediate results. To specify that the result should be stored to the output buffer, the assembly syntax adds a caret (^) to the destination.

Price & Availability

The Piccolo DSP core has been licensed to a number of semiconductor companies. The first Piccolo-equipped processor is expected to ship in 2H97. For information, contact ARM (Cambridge, U.K.) at 44.1223.400.400, fax 44.1223.400.410 or in the U.S. at 408.399.5190, fax 408.399.8854 or visit the Web at *www.armltd.com.*

Register Remapping Rotates Resources

Like most DSPs, Piccolo has a loop-construct primitive. The REPEAT instruction takes two arguments: the size of the loop and the intended number of iterations, which can be a register value. REPEAT instructions can be nested four deep.

The looping construct also supports a feature whereby Piccolo's registers are remapped after each pass through the loop. The purpose of this unusual feature is to allow DSP loops to operate on a series of coefficients without explicitly addressing different registers. Using the loop-remapping option, up to eight registers shift one, two, or four apparent positions in the register file, modulo the number of registers being remapped. After the final pass through the loop, the registers are back in their original positions.

The code fragment in Figure 2 illustrates how both the register remapping and the ROB refill are used on an FIR filter. The code multiplies a series of data points (d[n]) with an eight-element vector (c[0]–c[7]). Because each data point is multiplied with each of the vector elements, a loop with modulo-8 remapping will automatically access each element in turn. Also, once a data point has been fully accumulated, its register is no longer needed and can be refilled with new data from the ROB. The net result is that Piccolo needs to

```
REPEAT #N/4, x++ n4, y++ n4

  MULA    A0, X0.L^, Y0.L,  A0
  MULA    A1, X0.H,  Y0.L,  A1
  MULA    A2, X1.L,  Y0.L,  A2
  MULA    A3, X1.H,  Y0.L^, A3
```

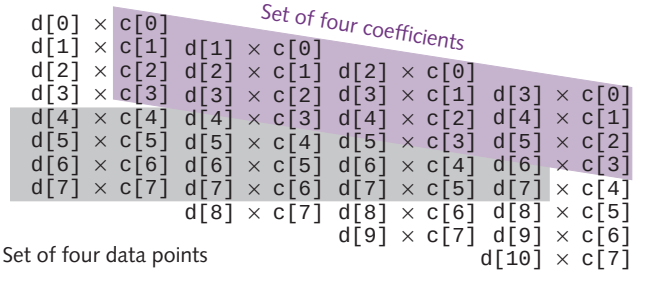

Figure 2. In this example of an FIR filter, a sequence of data points is multiplied with each of eight coefficients in turn. For every fourth pass through the loop, one data point and one coefficient can be replaced with new values. Piccolo's register rotation and register reloading handle both these tasks without explicit code, while ARM handles memory accesses.

load only one new data value on each pass through the loop, which will probably be waiting in the ROB.

Piccolo Debut Set for Mid-1997

During his presentation, Jaggar confirmed that Piccolo development is complete and that some number of both current and unannounced ARM licensees have signed on. Those companies have not been identified, but at least one is said to be working on a Piccolo-based chip for release in mid-1997. The first implementation will be with an ARM7TDMI (that is, an ARM7 core with the Thumb, hardware multiplier, and emulator/debugger modules).

The company claims Piccolo requires somewhat less than 1.5 mm^2 of silicon in a 0.35-micron three-layer-metal process. A basic ARM7 core needs about 2.2 mm^2 in the same process, making Piccolo about two-thirds as large as ARM7.

In a 0.6-micron process, parts are expected to reach 40 MHz, about the same speed as ARM7 in the same process. In 0.35-micron technology, which more aggressive ARM vendors like VLSI are now delivering, Piccolo should hit 66–80 MHz. Jaggar suggested an updated Piccolo with a longer pipeline will reach 120 MHz. This version might be mated to an ARM8 or StrongArm core.

DSP-Like Performance

Piccolo has not been fabricated, so no definite performance figures are available. ARM has simulated Piccolo running at a variety of clock speeds and compared its performance with that of more conventional, albeit aging, parts from AT&T, Motorola, and TI. A 66-MHz Piccolo edged out the 1627, the 56002, and the 320C52 running at 40–70 MHz.

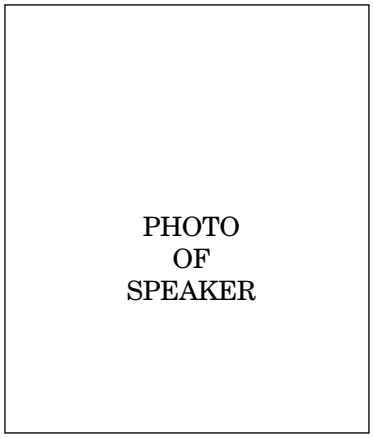

MICHAEL MUSTACCHI

Dave Jaggar of ARM conducts the first disclosure of the Piccolo DSP at the Microprocessor Forum.

Granted, the performance of Piccolo was simulated and the three competitors chosen don't represent the state of the art, but ARM's version seems to hold its own. On the other hand, TI, Motorola, AT&T, and other vendors are all shipping parts now, with 100-MHz and faster DSPs in the works. Before Piccolo ships, a performance gap may open up.

ARM's results indicate that customers with a moderate requirement for DSP speed need not give up performance to use Piccolo. Many embedded DSP applications also require a conventional integer processor for control flow, user interface, protocol handling, and other features. ARM's CPU/DSP combination may be more convenient and more economical than separate parts. It may be more accurate to characterize a Piccolo chip as a DSP with CPU capabilities rather than a CPU with DSP capabilities.

Two Cores Need to Cooperate

Although DSP performance looks good, ARM performance suffers terribly when the DSP is active. According to ARM's own figures, from 33% (during GSM loops) to 87% (during FFTs) of ARM's bandwidth is used to feed Piccolo. In either case, there's little time remaining for a task switch or for meaningful processing.

In an ARM/Piccolo duet, the two cores essentially take turns running. The CPU can run at full bandwidth when the DSP is idle, or vice versa. This characteristic is a good match with a large subset of application requirements. A cellular telephone, for example, might depend on the ARM for the user interface and for number recall while dialing, then switch to heavy DSP usage while the call is in progress. Disk drives provide another applicable example: interface protocols (SCSI, IDE, etc.) can be handled by ARM code while data recovery is handled by the DSP.

Piccolo is less well suited for a handheld organizer or PDA; while Piccolo will make a fine modem or wireless controller, the remaining ARM bandwidth might frustrate users. Consumers will generally want to continue using their system while the modem or wireless link is active, and Piccolo would sap most of the ARM's strength in such circumstances.

Different Bottlenecks

Piccolo bears similarities to both the 68356 and the SH-DSP. Like the Motorola part, Piccolo-equipped chips will execute two separate instruction streams on two different cores at the same time. Unlike the 68356, however, the CPU and DSP cores share an address and data bus, necessitating a lot of cooperation between the two.

The SH-DSP binds the two cores more tightly, as does ARM's approach. In fact, in the Hitachi design, the CPU and DSP execute from a single instruction stream and are not independent at all. Each 32-bit instruction is split, with half going to the DSP and half to the CPU for address generation and flow control.

Both the 68356 and the SH-DSP have separate X and Y memories typical of DSPs but not found in Piccolo. ARM chose instead to be clever with its register file, supplying operands from DSP registers, not local memories. Even though Piccolo can access a pair of registers per cycle, it can reload only one from the ROB. A data-intensive algorithm can starve the register file. In the end, Piccolo is cursed with lower operand bandwidth than other DSPs, and this may prove to be its ultimate weakness.

Like other companies, ARM has responded to a growing need to merge signal processing with arithmetic or control-flow processing. As Hitachi and Motorola have found, the combination can be a popular one, as the nature of embedded designs changes and the demand for wireless appliances increases. With its announcement of Piccolo, ARM can end 1996 on a high note. M

Most Significant Bits

Continued from page 5

Another unique feature of the TR25201 is support for multichip configurations. In theory, up to 10 TriTech devices can be interconnected to share the 3D rendering process. TriTech claims each chip requires only about 10% of a Pentium-100 processor for scene management, eliminating a common bottleneck in most other 3D accelerators.

Two independent 32-bit SGRAM buses support local memory. TriTech believes the dual buses provide better overall performance than a single 64-bit bus. Up to 32M of SGRAM can be attached, enabling high-resolution frame and Z buffers as well as local texture storage.

TriTech *(www.tritech-sg.com)* also offers the TR25202, which lacks the geometry processor and is less expensive than the TR25201. Both parts are packaged in 304-ball BGA packages and are expected to be in production in 2Q97. Pricing is estimated at $70 for the TR25201 and $50 for the TR25202, both in 10,000-unit quantities. —*P.N.G.*

S3 Upgrades Popular 3D Accelerator

After shipping three million units of its original Virge 3D chip, representing 70% of the market for PC 3D adapters, S3 has announced two new parts that offer three times the basic rendering performance and significantly improved display quality at about the same price point. Like their predecessor, these chips offer low-end 3D acceleration at essentially no premium over a fast 2D graphics chip. The new chips lag well behind the performance of more expensive 3D accelerators, however, many of which include setup engines.

The Virge/DX and Virge/GX graphics chips are identical except for their frame-buffer DRAM controllers: the DX supports EDO DRAM, while the GX also supports SDRAM and SGRAM. Both provide perspective correction and trilinear MIP-mapped texturing without a performance penalty. Turning on these quality enhancements on the older Virge resulted in a dramatic drop in frame rate.

The new parts achieve their rated 3× performance improvement using a significantly improved rendering pipeline that also provides better support for large triangles, reducing the host-CPU overhead for setup processing. The S3 chips do not, however, completely offload the 3D setup process, preventing them from matching the performance of midrange and high-end 3D accelerators.

Both the DX and the GX support up to only 4M of local memory, limiting the amount of texture storage to 1.25M in a typical 800 × 600 double-buffered configuration with 16-bit color and Z-buffer support. Given the dramatic reductions in memory prices over the past year, 8M configurations are becoming more popular, although perhaps not at the price points S3 focuses on.

The new parts also include S3's 2D and video accelerator engines, and they are compatible with the software drivers and 208-pin PQFP packages of the earlier Virge chip, as well as S3's Trio family of 2D accelerators. Pricing for the Virge/DX and Virge/GX is $28 and $29, respectively, in 10,000-unit quantities. Both chips are now sampling, with volume production scheduled for 4Q96.

We expect the "free-D" approach embodied in S3's Virge, ATI's 3D Rage, and other low-cost graphics chips will become the standard solution for 3D in low-end systems, while integrated setup and geometry engines will provide key differentiating features for lower-volume midrange and high-end 3D products. —*P.N.G.*

VLSI SongBird Leads Audio Migration to PCI

The VL82C829 SongBird 3D audio accelerator from VLSI is the first PCI-based accelerator to support Microsoft's DirectSound 3D API and legacy SoundBlaster audio, and one of the first audio controllers that connects to an external analog codec using Intel's AC '97 interface (see MPR 7/8/96, p. 4).

SongBird includes a fixed-point DSP to implement a full range of sound effects, including the head-related transfer function (HRTF) algorithm used to create true 3D positional audio using only two speakers. The DSP gives SongBird a significant performance advantage over emerging software-based audio implementations. According to the company, a software-only HRTF implementation requires more than half of a Pentium-166 processor.

The chip is packaged in a 208-pin MQFP or TQFP. It includes interfaces for PCI, optional SRAM for reverb effects, a joystick, an external Dolby AC-3/MPEG audio decoder, and the external codec. Sampling is scheduled to begin in 1Q97 with production in 2Q97. Pricing is $25 each in 10,000-unit quantities. —*P.N.G.*

Motorola 68376 Flies as TouCan

Motorola's new 68376 is the first in a promised series of 68K-based microprocessors to include a CANbus controller. The part is sampling now, with production set for 1Q97. Called TouCan, the line of controllers extends the reach of Motorola's popular 68300 family further into high-end automotive applications. Volvo, for one, has committed to use the part in future models.

The CAN (controller-area network) specification has gained popularity in automotive and industrial applications because of its deterministic response, resistance to EMI, and simple wiring topology. Currently, most vendors implement CAN interfaces with separate chips; the '376 is the first part from Motorola to integrate CAN onto the microprocessor.

The '376 includes the usual set of 68300-family peripherals, including the timing-processor unit (TPU), queued serial module (QSM), configurable timer module (CTM), and 10-bit A/D converter. The part also has a relative abundance of memory: 8K of mask ROM, 4K of general-purpose SRAM, and nearly 4K more SRAM for the TPU. The complex TPU and the large ROM make the '376 attractive for automotive manufacturers. In a 160-lead QFP, the chip lists for $27.16 in 10,000-unit quantities. —*J.T.*

QED's PowerPC 603q Heads for Low Cost

Design Firm Wraps Its First PowerPC Chip; Future of the Part Still Uncertain

by Jim Turley

The inveterate MIPS jockeys at Quantum Effect Design have backed a second horse and developed their first PowerPC microprocessor. At last month's Microprocessor Forum, QED engineer Bryan Chin trotted out the 603q, a midrange PowerPC that delivers performance similar to the 603e's but at half the power and two-thirds the die area.

The 603q was aimed at low-end desktop and high-end embedded designs. Was, that is, until QED's unnamed customer prepared to pull the plug on the project scant weeks after it taped out. Because QED has no PowerPC license of its own, it cannot sell the chip except to a licensed fab. The company is now left with a completed PowerPC design and potentially no customer to sell it to.

PHOTO OF SPEAKER

MICHAEL MUSTACCHI

QED engineer Bryan Chin discusses the PowerPC 603q at the Microprocessor Forum.

The Vanishing PowerPC Client

QED recently shifted its business strategy from a design-for-hire service to that of a fabless microprocessor vendor (see MPR 9/16/96, p. 5). The company's stated goal is to develop high-end embedded CPUs based on MIPS and "other mainstream instruction sets," implying the door is open to more PowerPC designs and perhaps even other architectures. But building and selling PowerPC chips requires a PowerPC license, a nontrivial expense that a small company like QED can ill afford at this stage.

	602	603q	603e	603e
Clock Frequency	80 MHz	160 MHz	120 MHz	240 MHz
Instr Cache Size	4K	16K	16K	16K
Data Cache Size	4K	8K	16K	16K
Power (typ)	1.45 W	1.6 W	3.5 W	4.5 W
Dispatch (sustained)	1 instr	1 instr	2 instr	2 instr
Dispatch (peak)	2 instr	1 instr	3 instr	3 instr
FP Repeat Rate, DP	n/a	4 cycles	2 cycles	2 cycles
FP Repeat Rate, SP	1 cycle	4 cycles	1 cycle	1 cycle
External Bus I/F	Mux'd	Mux'd	Demux'd	Demux'd
IC Process	0.5µ 4M	0.5µ 3M	0.5µ 4M	0.35µ 4M
Die Size	50 mm^2	69 mm^2	98 mm^2	79 mm^2
List Price (1,000s)	$37	n/a	$121	$408
Availability	Now	n/a	Now	Now

Table 1. The 603q does not have the superscalar features of other midrange PowerPCs, but its small die size and modest power requirements make it competitive in high-end embedded systems. (n/a = not applicable) (Source: vendors)

The alternative is to design PowerPC parts solely for licensees, namely Motorola, IBM, Exponential, or Apple. The restrictive terms of Exponential's unusual license agreement might not allow it to build or sell QED's design, and Apple is unlikely to start manufacturing chips any time soon. That leaves IBM and Motorola as QED's sole potential customers, one of which appears likely to back out.

Although QED is mum, other sources suggest the client was Motorola, and ultimately Apple. The rapid pace of PowerPC speed improvements in the past several months coupled with Apple's bad experiences trying to sell low-end Macs in 1996 and the fact that both Motorola and IBM already have embedded PowerPC cores may have all conspired against the 603q.

QED's situation is not hopeless. The company owns all the intellectual property embodied in the 603q, which can be leveraged for a new, similar design. It also now has its own PowerPC verification suite, test vectors, and software.

No-Frills Design Cuts Cost, Speed

The 603q was intended for cost-effective desktop applications, moving into high-end embedded systems over time. The market targets meant that QED could afford to emphasize the integer unit over floating-point performance. The goal was not to beat existing PowerPC chips in a horse race but to provide similar performance at a more competitive price. To reach those goals, QED fell back on its strengths, developing a fast but simple pipeline and surrounding it with lots of cache.

The 603q implements a classic five-stage RISC pipeline. The part is not superscalar, nor does it implement branch prediction or independent load/store, branch, or ALU datapaths. QED's minimalist approach allowed it to squeeze the 603q onto a die just 38% larger than a 602, even though its caches are three times bigger. Table 1 outlines the differences between the 603q and its nearest competitors, the 602 and 603e.

The PowerPC integer core proved to be a true design challenge for a team weaned on MIPS microprocessors. The PowerPC instruction set is much richer than MIPS and "at the far, far end of RISC," according to QED's Chin. String operations, unaligned accesses, and multiple load/store instructions caused the biggest headaches.

Even though unaligned accesses and string operations are not part of the basic PowerPC specification, the 601

supported them, and they are used by many Macintosh applications. To achieve total Mac compatibility, the 603q had to handle these extraneous functions transparently.

According to QED's Chin, the design team was able to trade on its MIPS experience in the FP portion of the 603q, but the integer pipeline had to be developed from scratch. All IEEE 754–compatible FPUs are similar to some extent, so QED was able to reuse designs from the R4600 family. The concept of the merged integer and FP pipelines was taken from the R4650, another QED design.

The goal was to save design time and die area, but Chin said that, in retrospect, very little of either was saved, while a significant amount of FP performance was sacrificed. A more fully featured FPU, like that in the 603, might have boosted FP performance proportionally more than the extra die area it would have required.

Cache Uses Unique 6T SRAM Cell

One trick up QED's sleeve is a new SRAM-cell design. The company developed its own six-transistor SRAM cell, which it has patented. The new cell organization uses about 7–8% less die area and half the power of a conventional SRAM cell, according to the company. This same SRAM design is used for the RM7000's cache (see MPR 10/28/96, p. 36).

Given QED's approach to microprocessor design, it's logical the company would work on shrinking the bulkiest portion of its chips, which is the cache. As the chart in Figure 1 shows, the cache is the single largest consumer of real estate on the 603q. Since the part wasn't intended to be a performance leader, its design goals could best be met by reducing die size as much as possible. The part measures just 69 mm^2 overall.

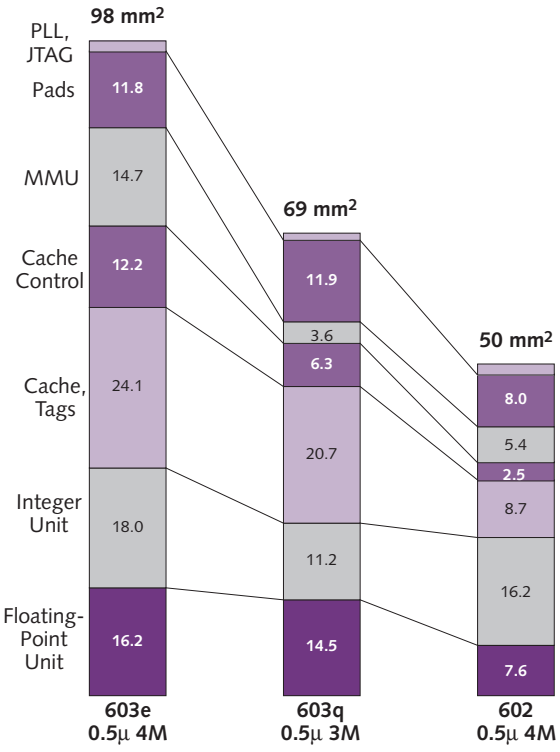

Figure 1. Comparing relative die allocation for the 602, 603q, and 603e highlights QED's ability to shrink the PowerPC integer unit by eliminating superscalar execution and branch handling. (Source: QED)

Price & Availability

Availability of the 603q is uncertain. QED may deliver derivative PowerPC devices at some point in the future. For more information, contact QED (Santa Clara, Calif.) at 408.565.0315; or visit the Web at *www.qedinc.com.*

Performance Matches Slower 603e Chips'

The 603q's bus interface is an amalgam of the 602's and 603's, with the former's multiplexed address/data bus and the latter's control-signal protocol and timing. In its CQFP-160 package, the 603q is not pin-compatible with either of those chips, although control logic or ASICs from a 603-based design could probably be reused.

QED's tests indicate the 603q achieves about 70% of the 603e's performance on the SPEC suite on a per-clock basis. The 160-MHz 603q should therefore deliver about the same integer performance as a 120-MHz 603e. FP performance will be somewhat lower due to the 603q's simplified FPU and single-issue design.

Motorola and IBM recently announced 603e parts running at 240 MHz (see MPR 10/28/96, p. 4). These chips should deliver about twice the integer performance of the 603q, according to QED's figures, and do even better on FP. At about $400 a pop, these new PowerPCs aren't cheap; the 120-MHz 603e is much more in line with most embedded designers' budgets. Unfortunately, no pricing information is available for the 603q, nor is it ever likely to be.

What's Ahead for QED?

Certainly the company is disappointed that its first PowerPC processor may never see the light of day. The project lasted nearly two years and expanded QED's design staff significantly. Its introduction was also timed to coincide with the company's coming-out party, marking its entry into the merchant CPU business.

On the plus side, QED is willing and able to undertake a new PowerPC design armed with its new verification techniques, test vectors, and design tools. Its list of potential clients is obviously very small (i.e., two), but its record of nimble, cost-effective designs could help either IBM or Motorola win a particular customer or market segment.

In the meantime, QED will busy itself finishing its high-end RM7000 design and other MIPS processors. As its past history has shown, the company can survive quite comfortably with just one CPU architecture. M

Digital, MIPS Add Multimedia Extensions

Digital Focuses on Video, MIPS on 3D Graphics; Vendors Debate Differences

by Linley Gwennap

Support for multimedia data types has become nearly pervasive, as Alpha and MIPS have joined the throng of instruction-set architectures with multimedia extensions. At last month's Microprocessor Forum, Digital announced its motion-video instructions (MVI), which will first appear in the 21164PC and 21264 processors next year. Also at the conference, MIPS Technologies rolled out two sets of multimedia extensions. The first, MIPS V, supports parallel floating-point operations and will mainly benefit 3D graphics. A separate set of instructions called MDMX (MIPS digital media extensions) provides broader support for parallel integer operations.

The MIPS V extensions allow two single-precision operands to be stored in a double-precision floating-point register using the new paired-single (PS) format. Several new instructions can then operate on this data in parallel, effectively doubling performance in this mode. Since many 3D graphics applications, as well as some scientific software, use single-precision FP, most of these applications could see a boost from MIPS V. The company would not discuss what processors will implement MIPS V, but we expect the follow-on to the R10000, code-named H1, will do so.

MDMX, also known as Mad Max, is an optional set of instructions similar to Intel's MMX (see MPR 3/5/96, p. 1) in that they define a set of media registers that is mapped onto the FP registers, new data types that store 8- and 16-bit data in parallel in the media registers, and instructions that operate on this data in parallel. MDMX's unique twist is its 192-bit accumulator that allows integer multiplication and accumulation to occur without any overflows or loss of precision.

Digital's additions are more spare, in keeping with the minimalist nature of the Alpha instruction set. Digital engineer Pete Bannon argued that current Alpha processors are fast enough to handle relatively simple tasks like audio mixing and video decoding without any instruction-set extensions, so why add new instructions that could slow the decoders or the execution units? To meet a 2-ns cycle time, simplicity is a requirement. The new instructions are designed to speed video encoding, a much harder problem, without slowing the CPU on other tasks.

These two vendors join HP, Intel, and Sun in adding multimedia extensions to their instruction sets. Large performance gains on multimedia applications, coupled with the relatively small incremental hardware cost, have led to this widespread adoption. Of the major desktop processors, only PowerPC now lacks such extensions, an incredible oversight given Apple's focus on multimedia.

MIPS V Boosts FP Performance

Most of these extensions aim to improve performance when handling 8- and 16-bit integers. These data types are common in audio, video, and graphics applications, yet traditional ALUs can operate on only one integer at a time. With most processors now implementing 64-bit data paths, up to seven-eighths of this data path is wasted when operating on small integers. Packing four or eight small integers into a single 64-bit register and operating on them in parallel using a SIMD (single instruction, multiple data) approach greatly increases throughput.

MIPS realized a similar opportunity exists on the floating-point side, yet no other vendor has moved to seize it. Modern microprocessors are optimized to handle double-precision floating-point data, meaning that a 64-bit data path exists in the FPU. Yet single-precision floating-point data is common in many applications, including 3D graphics and signal processing; these applications waste half of the FP data path. The same SIMD approach can double throughput on single-precision data.

MIPS V does just that. The new PS format pairs two single-precision values in each FP register. Table 1 lists the

At the Microprocessor Forum, several instruction-set architects gathered to discuss extensions for multimedia processing, including (l to r) Pete Bannon of Digital, Earl Killian of MIPS Technologies, Ruby Lee of HP, Uri Weiser of Intel, and Marc Tremblay of Sun Microelectronics. At right is moderator Linley Gwennap of MicroDesign Resources.

Modified Instructions	
ADD, SUB, MUL, ABS, MOV, NEG	Basic computational operations
MADD, MSUB, NMADD, NMSUB	Multiply add/subtract
C.cond	Parallel compare
MOVF, MOVT	Conditional move
New Instructions	
LUXC1, SUXC1	Load 8 bytes without alignment
ALNV	Realign data
PLL, PLU, PUL, PUU	Rearrange PS data
CVT.PS.S	Convert to/from PS format

Table 1. MIPS V modifies existing FP instructions to use the new paired-single (PS) data type and adds a few new instructions for packing, rearranging, and unpacking PS data.

existing FP instructions that accept PS operands in MIPS V; these include the basic arithmetic operations. Note that the product of two single-precision operands is still a single-precision value, with the exponent adjusted accordingly. Thus, multiplying two paired-single operands results in a paired-single product.

The table also lists a few new instructions in MIPS V. The LUXC1 and SUXC1 instructions (don't try to pronounce them!) load and store 64 bits at a time regardless of the alignment of the address; that is, the lowest three bits of the address are simply ignored. These FP instructions help load pairs of single-precision operands when the operands are part of a vector that is not aligned to a 64-bit boundary. The ALNV instruction can then properly align the data.

The Pxx instructions are useful for copying the upper or lower half of a PS value to the upper or lower half of another PS value. Finally, a new convert (CVT) instruction creates a PS value from two single-precision values. This instruction takes two FP registers as source operands.

These instructions will add a small amount of circuitry, particularly when compared with the size of a high-performance double-precision FPU, yet they will provide a big performance boost on some frequently used algorithms, such as fast Fouriér transform (FFT) and matrix multiplication.

MDMX Goes Beyond Intel's MMX

The basic features of MIPS' MDMX are similar to those of Intel's MMX. MDMX creates a new set of 32 media registers, each 64 bits wide. To reduce storage requirements, they are mapped onto the floating-point registers. Intel uses the same strategy for its MMX registers, but since x86 has only eight FP registers, there are also only eight MMX registers. MDMX adds a set of eight single-bit condition flags, which map onto the MIPS FP condition flags.

The new registers support two data formats: oct byte (OB) and quad half (QH). For the jargon-impaired, the former refers to eight 8-bit values packed into a single 64-bit register, while the latter consists of four 16-bit values. Like MMX, MDMX can boost performance on many multimedia algorithms by 2–4×.

Table 2 lists the standard MIPS instructions that are modified to operate on the media registers using either the

Modified Instructions	
ADD, SUB, MUL, MIN, MAX, MSGN	Saturating arithmetic
AND, XOR, OR, NOR, SLL, SRL, SRA	Logicals and shifts
ALNI, ALNV	Align vectors
C.EQ, C.LT, C.LE	Compare bytes
New Instructions	
SHFL.op	Shuffle bytes
PICKF, PICKT	Combine vectors
ADDL, SUBL, MULL, MULSL	Store result in ACC
ADDA, SUBA, MULA, MULS	Operate on ACC
RZU, RNAU, RNEU, RZS, RNAS, RNES	Round ACC
RAC, WAC	Read/write ACC

Table 2. The MIPS MDMX extensions modify several instructions to use the new vector data types and add several new instructions for accessing the 192-bit accumulator and for arranging data.

OB or QH format specifier. These include arithmetic, logical, shift, and min/max instructions. The new variations always operate in saturating mode, where overflows and underflows are clamped at the maximum and minimum values, respectively. This mode is useful when operating on pixel or amplitude data. These features are also found in MMX.

MIPS calls these operations "vector-to-vector" arithmetic. MDMX also supports a vector-to-scalar mode. In this mode, a single byte (or halfword) selected from any part of a media register is combined with each byte of the second source register, as Figure 1b shows. A third mode allows the scalar value to be specified as a 5-bit immediate value.

The vector-to-scalar modes come in handy when multiplying a vector by a constant value, which occurs in inverse discrete cosine transformation (IDCT) and many signal-processing algorithms. These modes are also useful for motion estimation and other multimedia algorithms.

New Instructions Arrange Bytes

The shuffle (SHFL) instruction, similar to MMX's PUNPACK, performs eight different byte rearrangements. It can convert data between OB and QH formats with signed and unsigned options and can interleave bytes from two registers. Unlike the PERMUTE instruction in HP's MAX2 (see sidebar, page 27), SHFL does not perform arbitrary reorganization.

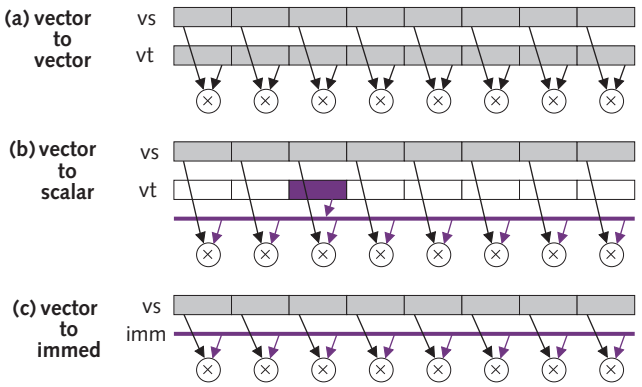

Figure 1. MDMX instructions operate in (a) vector-to-vector mode, (b) vector-to-scalar mode, and (c) vector-to-immediate mode.

For More Information

For more information on these multimedia instruction-set extensions, check the following Web pages. For MIPS V and MDMX, *www.mips.com/ISAV/index.html*. For MAX2, *www.hp.com/wsg/strategies/strategy.html*. For MMX, *www.intel.com/pc-supp/multimed/mmx/index.htm*. For VIS, *www.sun.com/sparc/vis*. (No Web information is yet available for Digital's MVI.)

Other MDMX instructions go beyond the capabilities of MMX. The align (ALNx) instructions can extract eight bytes out of a sequence of bytes in two different source registers. These instructions can realign bytes from an arbitrarily aligned data stream.

Three compare (C.xx) instructions perform a parallel comparison and set the condition flags. There are eight condition flags defined, so each corresponds to a single byte in OB format (only four are used in QH format).

The PICK instructions can then be used to combine the individual values in two registers depending on the contents of the condition bits. If a condition bit is set, the corresponding byte is copied from the first source register; otherwise it is taken from the second source register. The C.xx instruction is similar to the parallel compare (PCMPxx) instruction in MMX, but MMX requires a sequence of three instructions to do the same task as a single PICK.

MDMX Adds Wide Accumulator

The wide accumulator is a completely different approach to multiplication than used in other instruction sets. The fundamental problem: when two integers are multiplied, the product may have twice as many significant bits as the operands. For example, it takes 16 bits to hold the product of two arbitrary 8-bit values. The problem gets worse for a multiply-accumulate operation, common in many signal-processing algorithms. Accumulating, say, one hundred 16-bit products could require 23 bits to avoid any possible overflow.

In MMX, 16-bit data is "promoted" to 32 bits by the multiply-accumulate (PMADDWD) instruction. This restricts the available parallelism once the data is promoted and still doesn't provide enough bits to avoid overflows. Also, there is no corresponding MMX instruction for 8-bit multiply-add.

MDMX solves all of these problems with its 192-bit accumulator. This special register can be partitioned into eight 24-bit values or four 48-bit values, corresponding to the OB and QH formats. In the former case, each 24-bit value can accumulate the product of 256 multiplies of 8×8 bits each. In QH mode, the accumulator can handle the sum of 65,536 multiplies of 16×16 bits each. In both cases, there is no loss of precision and no possibility of overflow.

Table 2 lists several new instructions that use this accumulator. Vector data can be added, subtracted, and multiplied, with the result placed in the accumulator with no loss of precision. While this is most useful for multiplication, vector data can be added and subtracted in nonsaturating mode, with potential overflows and underflows captured in the accumulator for further processing.

Vector data can also be added directly to the accumulator register, or multiplied and accumulated. There are several instructions that shift and round the final result. Finally, the RAC and WAC instructions move data between the accumulator and the media registers, where it can be copied to memory. These instructions, which access the accumulator in three 64-bit chunks, are used for saving state.

Single Accumulator Can Be a Bottleneck

At the Forum, MIPS architect Earl Killian said the accumulator obviated the need for a parallel 32-bit data type, which is supported in MMX and Sun's VIS. This data type is mainly used to provide adequate precision when operating on 16-bit values; the MDMX accumulator provides the same function. A few algorithms, such as Dolby Digital AC-3 audio, require data with more than 16 bits of precision (AC-3 uses 20 bits); MDMX comes up short for these algorithms, although a MIPS V processor can double throughput if the algorithm is converted to single-precision FP data.

While the accumulator provides the advantages noted above, only one set of calculations can use the accumulator at a time. Other architectures can unroll loops and accumulate into several general-purpose registers at once. Thus, the single accumulator can be a bottleneck. Another problem with the accumulator is that the operating system must be modified to save additional state, which could be a problem in an embedded system using an off-the-shelf real-time OS.

The MDMX extensions are more extensive than the MIPS V instructions and will require somewhat more hardware to implement. If the floating-point registers and data path are leveraged, however, the impact is still fairly small. The 192-bit accumulator must be added, but the other changes require only a few extra buses and multiplexers along with minor tweaks to the arithmetic and shift units.

Although MDMX is optional, MIPS expects it to be used in a variety of processors for computer systems and embedded products. A few specialized embedded processors may not implement MDMX for cost reasons.

Digital Boosts Motion Video

The high native performance of Digital's 21164 processor allows it to perform full DVD decoding (MPEG-2 video and AC-3 audio), as well as a variety of lesser tasks, in software without any special instructions. The only significant remaining hurdle, according to Digital's Bannon, is MPEG-2 video encoding, which overwhelms even a 21164. Analysis of MPEG-2 encoding software showed that a single task, motion estimation, consumes more than 70% of the CPU cycles. Improving performance on this one task would solve the encoding problem for Digital.

New Instructions	
PERR Ra, Rb, Rc	The sum of the absolute difference of each of the 8 bytes in Ra and Rb is written to Rc
UNPACK Rb, Rc PACK Rb, Rc	Unpack bytes into halfwords (4) or words (2) Pack words (2) or halfwords (4) into bytes
MIN Ra, Rb, Rc MAX Ra, Rb, Rc	For each of the 8 bytes in Ra and Rb, write the smaller (MIN) or larger (MAX) to Rc

Table 3. Digital's MVI extensions add only a few instructions that are designed primarily to boost video-encoding performance.

Instead of using the floating-point registers (or a new set of registers mapped to the FP side), Digital chose to add a few integer instructions. This simpler approach, also taken by HP, allows existing instructions such as AND to operate on parallel data. Unlike Pentium, Alpha already has 64-bit integer registers, the same width as the FP registers.

As others have done, the MVI extensions define new formats to pack 8-, 16-, and 32-bit data into the integer registers. The number of special instructions for operating on these data types is, however, quite small, as Table 3 shows.

The key instruction is PERR, which calculates the sum of the absolute differences of eight pairs of bytes, a function also implemented in Sun's VIS. This function is at the heart of motion estimation, which involves comparing blocks of pixels (typically 8×8) to find which pair has the smallest difference. A single PERR instruction replaces nine normal Alpha instructions in the motion-estimation inner loop, providing a significant speedup.

No Parallel Arithmetic Instructions

The new PACK and UNPACK instructions convert packed bytes into 16- or 32-bit formats and vice versa. This conversion allows operations on 8-bit data. Digital did not, however, provide any parallel arithmetic instructions except for the new MIN and MAX instructions, which are helpful in video encoding. In a few situations, standard instructions can be used to operate on parallel data, but only if the algorithm avoids all overflows or underflows. This exclusion prevents most algorithms from using parallel arithmetic.

The MVI instructions will help the forthcoming 21264 perform MPEG-2 encoding in real time, a feat beyond the capabilities of any current microprocessor. MVI's impact on other multimedia tasks will be small. Video encoding to less stringent standards, such as MPEG-1 or H.320 video conferencing, will consume a smaller percentage of the enhanced 21164PC than the 21164. With no multiply-accumulate instruction or parallel arithmetic, tasks such as audio processing and communications will see little gain, if any, but these tasks already require little effort from the Alpha chips.

MVI also neglects two performance-critical multimedia tasks: 3D graphics and image processing. Both of these tasks can consume all available performance even on a powerful processor. MVI lacks the dual-FP instructions found in MIPS V, which help 3D. An image-processing application like Photoshop would benefit from parallel arithmetic during pixel manipulation, but MVI has none. This shortfall will reduce the gap between Digital and Intel processors on optimized image-processing applications.

Other Instruction Sets Evolve

At the Forum, HP's Ruby Lee discussed extensions to the PA-RISC architecture that are now shipping as part of the PA-8000 processor. These extensions, called MAX2, are a superset of the original set of extensions (see MPR 1/24/94, p. 16) implemented in the PA-7100LC and PA-7200. These 32-bit processors can operate on two 16-bit halfwords at once; the 64-bit PA-8000 can operate on four halfwords in parallel.

MAX2 also adds two new instructions. MIX, similar to Intel's PUNPACK, interleaves halfwords from two registers, speeding matrix transposition. PERMUTE goes further by allowing any rearrangement, with or without repetition, of the four halfwords in a register. These instructions are useful in the IDCT operation (part of video decoding) and in cryptography, for example.

The VIS instruction set (see MPR 12/5/94, p. 16) has been on the market for about a year; according to Sun's Marc Tremblay, the company is working on improving it in future processors. One feature Tremblay is strongly considering is parallel FP operations, similar to those in MIPS V. Look for these and other new features in future UltraSparc processors.

Intel's Uri Weiser, who led the development of MMX, admitted he is also looking at future upgrades. Although Weiser would not comment on any specifics, the company is widely rumored to be planning extensions called MMX 2, which could begin shipping in processors as early as 1H98 (see page 4).

Vendors Excel in Different Areas

Because multimedia is a diverse concept, comparing these instruction sets is not a simple matter. The best approach is to look at different multimedia applications and assess each instruction set on that application. Each vendor has targeted some applications and neglected others.

For basic manipulation of audio and pixel data, the new extensions (except for MVI) are fairly similar, providing packed 8- and 16-bit data types along with instructions that operate on them with saturating and nonsaturating arithmetic. As Table 4 shows, Sun neglects saturating arithmetic, while HP doesn't support 8-bit arithmetic. Intel offers only two-operand instructions, whereas the RISC architectures support a separate destination specifier. Image processing (e.g,. Photoshop) is the key application in this category.

Rendering 3D images is also image processing and takes advantage of these same basic extensions. Most vendors agree, however, that 3D rendering is best done by a hardware accelerator and not on the CPU.

	MIPS V/ MDMX	Intel MMX	Sun VIS	HP MAX2	Alpha MVI
No. of Registers	32	8	32	31	31
Register Type	MM/FP†	MM/FP†	MM/FP†	Integer	Integer
Parallel Arithmetic	8 × 8 bits 4 × 16 bits	8 × 8 bits 4 × 16 bits	4 × 16 bits 2 × 32 bits	4 × 16 bits	Min/max only
Unsaturating?	No	Yes	Yes	Yes	n/a
Saturating?	Yes	Yes	No	Yes	n/a
Three Operands?	Yes	No (2)	Yes	Yes	Yes
Parallel Multiplies	4 or 8	4	4	None	None
Multiply/Add	8 × 8 → 24 16 × 16 → 48	16 × 16 → 32	8 × 16 → 16	Shift-and-add	None
Vector-to-Scalar?	Yes	No	No	No	No
Parallel Shifts?	Yes	Yes	No	Yes	No
Parallel Average?	No	No	No	Yes	No
Parallel Compare?	Yes	Yes	Yes	No	No
Pack/Unpack?	Yes	Yes	Yes	Yes	No
Interleave?	Yes	Yes	Yes	Yes	No
Permute?	No	No	No	Yes	No
Pixel Error?*	No	No	Yes	No	Yes
Block Load/Store?	No	No	Yes	No	No
Parallel FP?	Yes	No	No	No	No

Table 4. The combination of MIPS V and MDMX offers more features than other multimedia instruction sets, while Alpha's MVI offers fewer. *motion-estimation instruction. †multimedia registers mapped on floating-point register set. (Source: vendors)

but hardwired acceleration has proved useful in other cases. Intel has not addressed this problem.

The CPU's main role in 3D graphics is geometry processing. This processing is typically done using single-precision (sometimes double-precision) floating-point math and thus is unaffected by most multimedia extensions. Only MIPS V, with its paired-single operations, provides a large performance boost for 3D geometry. Sun's specialized 3D array instructions assist some 3D applications.

Many of these algorithms require fast access to large data sets. The block load and store instructions in VIS speed these accesses and are unique among these processors. Of course, access times remain largely a function of the processor's memory bandwidth.

Most audio applications require some type of signal processing, such as FFTs and other matrix arithmetic. The key instructions for these algorithms are multiply and multiply-accumulate. MDMX offers the most flexibility here, particularly with its unique 192-bit accumulator, while MMX is close behind. HP's parallel shift-and-add instruction can multiply by only small constant values. In addition to audio, these instructions are good for modem emulation.

Intel's CISC heritage shows up in its lack of MMX registers: only 8 compared with 32 for the other instruction sets. While more registers can improve performance on a range of applications, the extra registers are particularly helpful for matrix arithmetic. With 32 registers, an entire 4 × 4 transformation matrix can be held in half the register set, leaving the remainder free for working data. Since the same transformation is often used for many operations in a row, this ability greatly reduces cache accesses compared with the limited MMX register file.

Decompressing audio and video, such as MPEG-2 decoding, requires different operations, such as IDCT. While these operations make heavy use of multiply-accumulate, they also benefit from instructions that perform byte rearrangement for matrix transformations. All of the vendors but Digital have added such instructions.

Compressing video relies almost entirely on motion estimation. Both Digital and Sun have supplied SAD (sum of absolute differences) instructions to speed this function. Both MIPS and HP claim the motion-estimation problem can be solved in a variety of ways without resorting to a special hardwired instruction. Until these vendors can demonstrate video compression, evaluating these claims is difficult,

Time to Market Is Important

Instruction-set design is important, but other factors will affect the ultimate success of these new extensions. Software must be changed to take advantage of the new instructions, and for the most part, this change involves rewriting the code by hand to operate in a SIMD fashion; the compiler cannot revise the algorithms. Without new software, the finest instruction sets are useless.

Of the current vendors, HP was the first to deliver multimedia extensions, starting with the PA-7100LC in 1Q94. Today, only a few applications take advantage of these extensions, primarily for MPEG decoding and video conferencing. Sun has shipped the most multimedia-enabled processors to date and is slowly building a software base. Intel has been helping software vendors recode to MMX for more than a year, but the first MMX systems will not ship until 1Q97. The Digital and MIPS extensions are further in the future. Digital's software task is easier: the company will probably release a new video-encoding library and call it a day.

Of course, these extensions will be no more popular than the underlying processors, making MMX the clear winner here. MIPS V will solidify Silicon Graphics' position as the pre-eminent 3D platform. Combining high intrinsic performance and the MVI extensions, Alpha is positioned to become the best video-authoring platform. Digital's lead in image-processing performance, however, may shrink. The biggest loser of all could be PowerPC, which apparently still has its collective head in the sand while other vendors, particularly Intel, are poised to take share in markets where multimedia is important. MP

RECENT IC ANNOUNCEMENTS

PART NUMBER	VENDOR	DESCRIPTION	PRICE/QUANTITY	AVAILABILITY
DSP				
56303	Motorola 714.557.6884	Digital signal processor hits 100 MHz in 0.5-micron three-layer-metal process; device has 24-bit precision and 3.3-V supply voltage.	$32/10,000	Samples—Now Prod.—1Q97
56305	Motorola 714.557.6884	Latest member of 56300 family of 24-bit DSPs is optimized for GSM base stations with filter, Viterbi, and cyclic code coprocessors.	$75/1	Samples—Now Prod.—1Q97
DPL3520A	ITT Intermetall 408.526.2082	Processor for Dolby Pro Logic Surround supports standard channel separations as well as pseudo and panorama modes; in PLCC-68.	$13/10,000	Prod.—Now
MICROPROCESSOR				
µPD780805	NEC 800.366.9782	Microcontroller has stepper-motor interface for automotive instrument clusters; with 40K of mask ROM or 60K of flash memory.	$7/25,000	Samples—Now Prod.—1Q97
INTERFACE				
PCI1131	TI 800.477.8924	CardBus controller meets PC 97 guidelines and supports two 32-bit slots, dual voltage, and Zoom Video; in TQFP-208.	$14.60/150,000	Prod.—Now
78Q2120	TDK 800.624.8999	Ethernet transceiver supports 10/100-Mbps data rates over 10Base-T, 10Base-TX Category 5 wiring; 3.3-V supply; TQFP-80 package.	$25/500	Prod.—Now
TSB12LV31	TI 800.477.8924	Link-layer chip for IEEE-1394 has 8/16-bit bus interface, supports 100-Mbps data rate, and operates at 3.3 V; in 100-lead PQFP.	$9.72/1,000	Prod.—Now
MEMORY				
M5M5V64R16	Mitsubishi 408.730.5900	High-speed SRAM has 10-ns access time, center-mounted V_{CC} and ground pins to reduce noise susceptibility; with 64K×16 capacity.	$17/1,000	Samples—Now Prod.—1Q97
M5M5256D	Mitsubishi 408.730.5900	Low-power SRAM runs on 3.3 V, has 32K×8 capacity, 8-mA active supply current, 0.05-µA standby current, and common I/O pins; in TSOP package.	$2.50/1,000	Prod.—Now
M5M5208 M5M5512	Mitsubishi 408.730.5900	Low-power SRAMs have 2-Mbit ('08) and 512-Kbit ('12) capacity and access times from 45 to 120 ns. Standby current is 0.3 µA.	$9.75/1,000	Samples—Now Prod.—2Q97
M5M5V2132	Mitsubishi 408.730.5900	Synchronous pipeline-burst SRAM for L2 caches operates at 133 MHz; part has 2-Mbit capacity, organized as 64K×32; in TQFP-100 package.	$19/1	Samples—Now Prod.—1Q97
M24128 M24256	SGS-Thomson 617.259.0300	Serial E^2PROMs have large, 128-Kbit ('128) or 256-Kbit ('256) capacity; both run from 1.8-V supply and come in DIP, SO, and SOL packages.	$5/1,000	Samples—Now Prod.—1Q97
71V508	IDT 800.345.7015	Synchronous SRAM has 1-Mbit capacity and zero bus-turnaround capability to eliminate dead cycles between reads and writes.	$12/10,000	Prod.—Now
HY57V16XY10	Hyundai 408.232.8000	Synchronous DRAM (SDRAM) has 16-Mbit capacity, 66- to 100-MHz clock rate, and clock-to-data-out time of 8–10 ns; in 44/50-lead TSOP.	$12.50/1,000	Prod.—Now
MISCELLANEOUS				
STLC7549	SGS-Thomson 617.259.0300	Analog front-end chip for PCs complies with AC '97 specification, with stereo codec, modem codec, switching, and gain; in TQFP-64 package.	$12/10,000	Samples—Now Prod.—1Q97
bq2091	Benchmarq 972.437.9195	Battery-capacity monitor IC handles NiCd, NiMH, and Li-ion batteries using SMbus (smart battery data) control protocol; in 16-lead SOIC.	$5.68/10,000	Prod.—Now
bq2058	Benchmarq 972.437.9195	Lithium-ion battery supervisor controls two external FETs in the charge/discharge path for overvoltage, undervoltage, and overcurrent protection.	$1.91/10,000	Samples—Now Prod.—1Q97

PATENT WATCH

by Rich Belgard, Contributing Editor

The following U.S. patents related to microprocessors were issued recently. Please send comments or questions via e-mail to belgard@umunhum.stanford.edu.

5,535,346

Data processor with future file with parallel update and method of operation
Issued: July 9, 1996
Inventor: Thomas L. Thomas, Jr.
Assignee: Motorola
Filed: July 5, 1994
Claims: 20
A data processor with a future file and a method of operating it. The processor has multiple architecturally visible registers, an execution unit, a rename buffer, and a future file. The rename buffer stores a result of the instruction and copies the result to the architecturally visible registers in response to instruction completion. The future file has multiple entries, each storing the most recent value of the architecturally visible registers. The most recent values are determined with respect to the series of instructions. Each of the entries provides the most recent value to the execution unit as an operand. Each of the entries also stores a value of the architecturally visible registers when an exception occurs.

5,530,941

System and method for prefetching data from a main computer memory into a cache memory
Issued: June 25, 1996
Inventors: Pirmin L. Weisser, et al
Assignee: NCR
Filed: August 6, 1990
Claims: 13
On a bus with multiple bus masters, a broad method of prefetching information predicted to be used from a main memory into a cache memory such that none of the system access requests from the bus masters will be delayed.

5,530,890

High-performance, low-cost microprocessor
Issued: June 25, 1996
Inventors: Charles H. Moore, et al
Assignee: Nanotronics (now Patriot Scientific)
Filed: June 7, 1995
Claims: 10
A stack-based microprocessor containing a central processing unit, a separate DMA processing unit, and an on-chip DRAM controller. The highly integrated microprocessor is designed to reduce system costs by reducing external chip counts.

5,530,825

Data processor with branch target-address cache and method of operation
Issued: June 25, 1996
Inventors: Bryan P. Black, et al
Assignee: Motorola
Filed: April 15, 1994
Claims: 4
A processor with a branch target-cache. The branch target-cache is accessed by the fetch address of the current instruction. The fetch address is a four-instruction address. The cache contains a tag identifying one of the four instructions at the fetch address to which the entry corresponds. The target address becomes the predicted address if the tag matches the offset into the fetch address of the current instruction.

5,530,817

Very large instruction word type computer for performing a data transfer between register files through a signal line path
Issued: June 25, 1996
Inventor: Yoshio Masubuchi
Assignee: Toshiba
Filed: February 22, 1993
Claims: 11
A very large instruction word (VLIW) computer architecture in which the instruction word is divided into operation-field groups. Each field is executed by an independent processor. The computer contains independent register files for each of the respective processors. Data transfer between register files is directed by a command included as an operation for at least one of the processors.

5,530,804

Superscalar processor with plural pipelined execution units, each unit selectively having both normal and debug modes
Issued: June 25, 1996
Inventors: Gregory C. Edgington, et al
Assignee: Motorola
Filed: May 16, 1994
Claims: 42
A superscalar pipelined processor in which each of the superscalar execution units is capable of executing in a debug address space in debug mode and a normal address space in normal mode. The debug mode executes instructions to debug the CPU at substantially the same speed as the normal mode.

OTHER ISSUED PATENTS

5,534,796 *Self-clocking pipeline register*
5,530,891 *System-management interrupt mechanism within a symmetrical multiprocessing system*
5,530,824 *Address translation circuit* M

CHART WATCH: MOBILE PROCESSORS

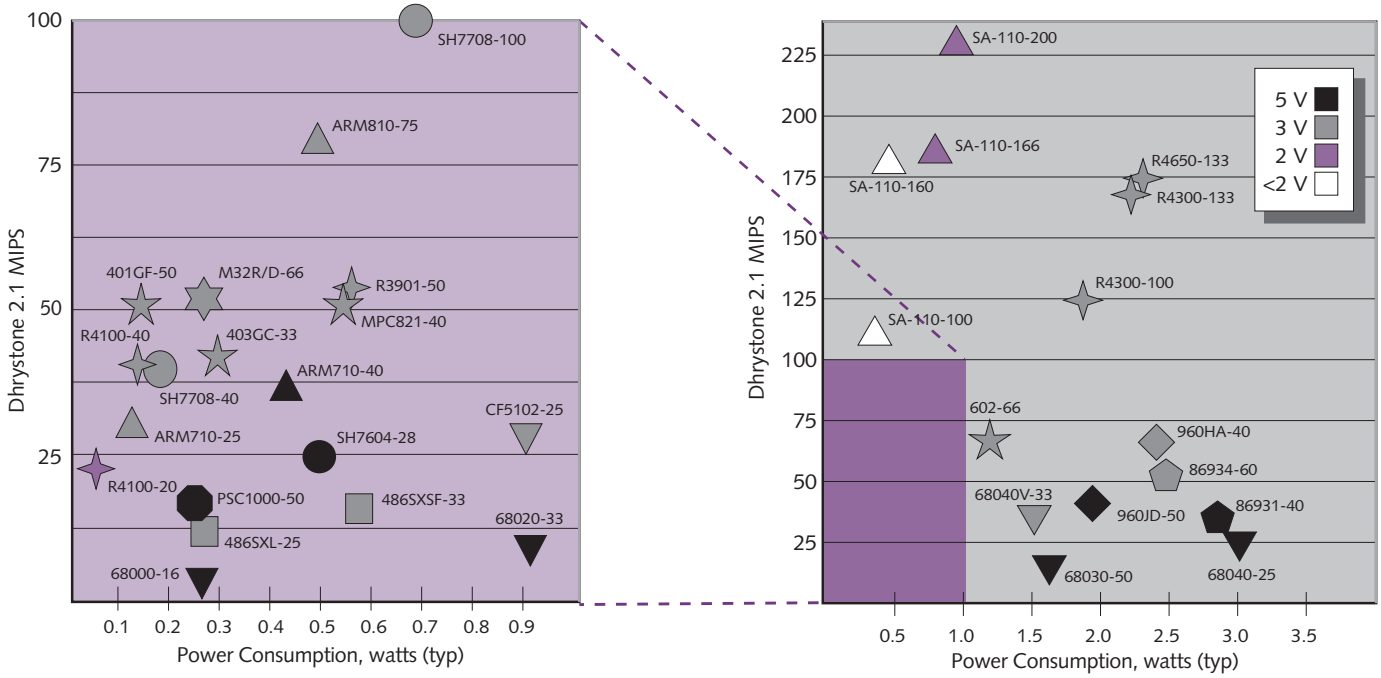

This Chart Watch covers low-power processors for portable and battery-powered systems. The table and the chart in the upper right show the performance/power ratio for a number of embedded CPUs and notebook processors; the chart above is an inset for the lowest power of these processors.

The chart on the right compares x86 and PowerPC processors for notebooks, including relative performance (diamonds) and typical power consumption (bars).

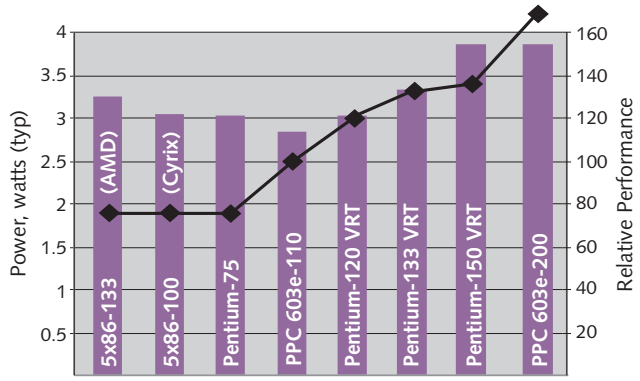

	SA-110	ARM710	SH7604	PPC 401GF	R4100	960SA	CF5102	486SXSF	29040
Vendor	Digital	VLSI	Hitachi	IBM	NEC	Intel	Motorola	Intel	AMD
Clock rate	200 MHz	40 MHz	20 MHz	50 MHz	40 MHz	20 MHz	25 MHz	33 MHz	50 MHz
I/D cache	16K/16K	8K	4K	2K/1K	2K/1K	512/0K	2K/1K	8K	8K/4K
FPU?	No	No	No	No	No	No	No	No	No
MMU?	Yes	Yes	No	No	Yes	No	No	Yes	Yes
Bus width	32 bits	32 bits	32 bits	32 bits	32 bits	16 bits	32 bits	32 bits	32 bits
Bus frequency	66 MHz	40 MHz	20 MHz	50 MHz	20 MHz	20 MHz	25 MHz	33 MHz	25 MHz
MIPS	230 MIPS	36 MIPS	20 MIPS	52 MIPS	40 MIPS*	9 MIPS	27 MIPS	16 MIPS*	67 MIPS
Voltage§	2.0/3.3 V	5 V	3.3 V	3.3 V	3.3 V	5 V	3.3 V	2.7/3.3 V	3.3 V
Power (typ)	900 mW	424 mW	200 mW	140 mW	120 mW	1,100 mW	900 mW	515 mW	1,650 mW
MIPS/watt	239	85	100	371	333	8	30	31	40
MIPS/mm²	4.30	1.04	0.24	2.36	1.60	0.17	n/a	n/a	0.56
Transistors	2,100,000	570,295	450,000	300,000*	450,000	346,000	n/a	n/a	1,200,000
IC process	0.35µ 3M	0.6µ 2M	0.8µ 2M	0.5µ 3M	0.5µ 3M	1.0µ 2M	0.6µ 3M	0.8µ 2M	0.7µ 3M
Die size	50 mm²	34 mm²	82 mm²	22 mm²	25 mm²	51 mm²	n/a	n/a	119 mm²
Est mfg cost	$18*	$9*	$7*	$4*	$8*	$4*	$9*	$15*	$20*
Availability	Now	Now	Now	Now	Now	Now	Now	Now	Now
Price (10K)	$49	$28	$27	$13	$25‡	$13	$25*	$72†	$86

† list price in 1,000's ‡list price in 100,000's §core/bus voltage n/a: information not available (Source: vendors except *MDR estimates)

Berkeley Benchmarks for DSP

The DSP aficionados at Berkeley Design Technology (BDT) have developed a collection of 11 DSP benchmarks and used them to rate the DSP capabilities and performance of a number of standard 32-bit microprocessors. The results will be available in ***DSP on General-Purpose Processors***, expected to run to 500 pages and cover six different popular CPUs: the PowerPC 604/604e, ARM7, SH-DSP, R4650, and Pentium with and without MMX. The results show that some general-purpose CPUs can outperform even dedicated DSP chips on common tasks.

Copies of the report will be available in about six weeks. Single copies will sell for $2,500, with substantial discounts for multiple copies. To order, contact MicroDesign Resources at 800.527.0288 or 707.824.4001, or visit the MDR Web site at *www.chipanalyst.com*.

PDA Expo and Forum Comes to San Francisco

The first annual **Handheld and PDA Expo** will be held at the Dunfey Hotel in San Mateo (near San Francisco) December 4–6. Hosted by PDA Inc., the show focuses on electronic organizers, smart phones, handheld computers, and mobile data technology. Seminars cover cost justification, electronic commerce, and Internet/intranet communications.

Single-day registration starts at $50; a full conference package, including the exposition, forum, and seminars for all three days, goes for $795. For more information or to register, call PDA (San Francisco) at 415.252.8008; fax 415.252.8055; or set your browser to *www.wmsltd.com*.

The Web Has Standards?

Evidently, according to technical publishers O'Reilly & Associates. The second edition of ***World Wide Web Journal*** brings together in print the complete collection of defining Web documents, including Tim Berners-Lee's specifications, the CERN RFCs, and the depository of the World Wide Web Consortium. Specs include HTML, URLs, HTTP, PNG, PICS, and PEP.

The 356-page journal sells for $25. For more information or to order, call O'Reilly (Sebastopol, Calif.) at 800.998.9938 or 707.829.0515, or visit *www.ora.com*.

You've Got Questions? They've Got Answers

The Desktop Computer Encyclopedia contains more than 8,500 definitions, 800 illustrations, and 1,024 pages. The book is a follow-on to Alan Freedman's popular *Computer Glossary* that describes hard disks, SCSI, pixels, RAM, full population, and the three C's of computing.

This ponderous tome is available in both CD-ROM and dead-tree versions for $40, or $60 for one of each. To order, or for more information, call Amacom (New York) at 800.262.9699 or 518.891.1500, or check out *www.amanet.org*.