TI Joins With IBM and Intermetrics in Bid to Establish Multimedia DSP Standard

By John Snell, Consultant
Timbre Engineering, San Geronimo, CA

Texas Instruments has joined with IBM and Intermetrics in an attempt to set a new standard, called Mwave™, for incorporating DSPs into PCs. The Mwave alliance includes a new DSP chip from TI, which integrates a simple DSP core with many serial interfaces for multimedia peripherals, software development tools from Intermetrics, and a DSP manager and real-time executive being developed by IBM and TI.

The DSP manager is being integrated into PC operating systems (OS/2, Windows, or DOS) by providing a standard programming interface for MCI-compliant multimedia drivers. (MCI, media control interface, is part of Microsoft's multimedia extensions for Windows.) IBM plans to incorporate the Mwave subsystem in all its future multimedia platforms, and it has stated that future systems will include the DSP on the motherboard. TI will provide a complete system solution, including IBM's DSP manager and real-time executive, to computer and add-in board manufacturers.

DSPs have an opportunity in multimedia computers because such computers require real-time performance beyond the reach of common CPUs and operating systems. A 24-bit DSP (Motorola's DSP56001) was included on the motherboard in the NeXT machine and in Silicon Graphics' Indigo for high-fidelity audio processing and music synthesis, as well as other multimedia functions. Multiple DSP96002, TMS320C40, or i860XP processors as well as video ASICs have been used on add-in cards. Analog Devices, AT&T, Motorola, and TI have all been attempting to find a market for their DSP chips in PCs. Too few technical details have been revealed to adequately compare the Mwave approach to competing solutions, but it is clear that TI has scored a major coup in striking a deal with IBM to support its approach.

The Mwave subsystem is capable of providing the following functions on a PC: Group 3 fax modem, telephone answering, DTMF tone generation, phone call progress detection, full-duplex 9600-baud V.32 data modem, voice compression, 10-to-20-voice musical sampler, and CDROM-XA decoding. The alliance is developing other Mwave applications including speech recognition, text-to-speech, audio time scaling, 10-band stereo equalizer, echo, reverb, and surround-sound concert hall effects, AM, FM, and wavetable music synthesis, as well as sound and image compression.

Mwave DSP Chip

With a simple processor core and a set of tightly-coupled application-related peripherals on chip, the Mwave DSP resembles an embedded controller. TI and IBM have announced the alliance but not any specific products, and the companies declined to provide any substantive details. Sources close to the project have, however, provided a general overview. The DSP core was described as a 16-bit, 17-MIPS Harvard architecture processor, not related to prior TMS320 DSPs. The multiplier accepts two 16-bit operands to produce a 32-bit product in a single instruction cycle, but the chip apparently has a 32-bit ALU with an extended-precision accumulator.

Although it is less expensive, linear computation with 16-bit arithmetic can result in poor low-frequency resolution (tuning) of oscillators and filters, as well as audio fidelity problems with IIR filter structures (such as equalizers). The 32-bit ALU alleviates this problem to some degree, but double-precision multiplication with software will be slow, limiting performance in some applications. The need for tuning accuracy and clean processing in the audio industry has driven the widespread adoption of Motorola's DSP56001, with its 24-bit multiplier and 56-bit accumulators.

The Mwave DSP cost was also reduced by not including RAM on-chip. Two cycles will be required for each double-precision transfer between off-chip RAM and the Mwave's 32-bit ALU. In other DSPs that provide three internal fast RAMs or cache for simultaneous access of an instruction and two operands, cycles are often wasted moving data, parameters, and instructions between off-chip and on-chip memory. However, with no internal RAM, three wide buses to three external fast RAMs would be needed for no-wait-state transfers, resulting in higher chip and system cost.

Instead of relying on local or on-chip memory, the Mwave DSP uses the host's main memory. This solution saves the cost of local DRAM and supporting components, at the expense of delays in accessing the PC's memory across the host system bus. In demonstrating the ability to cope with these delays, IBM has used the DSP to mix 32 channels of digital audio (presumably a short piece) stored in the PC's main memory, in real time at a 44-kHz sample rate.

Since multimedia applications involve real-time streams of data, efficient handling of I/O is critical. Interrupts from peripherals can introduce indeterminant

timing, a problem for real-time applications. The alternative of polling peripherals wastes precious real-time cycles. Traditional DSPs consume additional cycles in transferring data between peripherals and memory. To solve such problems, sources indicate that the Mwave DSP includes an on-chip DMA to free the core DSP of virtually all direct interaction with I/O.

The peripheral interfaces are serial ports like those on other TI DSP chips. The next version of the Mwave DSP may integrate complete interfaces to multimedia peripherals on the chip. Instead of just serial ports, sources expect to see 8- to 12-bit A/D and D/A converters for voice and telephone as well as an S/PDIF (Sony/Philips Digital Interface Format) interface for connecting CD players and DAT (Digital Audio Tape) recorders.

Real-Time Executive and DSP Manager

The Mwave DSP Manager and its associated device drivers run within the PC's OS. Virtual hardware tasks (e.g., fax, musical sampler, JPEG) execute under control of the real-time executive, which runs on the Mwave DSP. The Mwave DSP Manager dynamically links and loads DSP virtual hardware tasks and provides a communication conduit between applications and virtual hardware. Management of the DSP, peripherals and intra-task communication buffers is performed by the DSP manager. It ensures sufficient processing cycles are available before linking, loading, and running real-time virtual hardware tasks. The DSP Manager also keeps track of the DSP subsystem status and reports errors to the application.

The Mwave operating system is a multitasking real-time executive. It provides efficient, synchronized data transfers between tasks, while isolating task software developers from the complications of multitasking. A variable frame rate and dynamic deadline scheduler permit applications with different sample rates to be efficiently mapped into the multitasking real-time environment. The open Mwave environment should allow application developers to tap DSP multimedia tasks with minimal concern about other concurrently running DSP tasks. For example, the designer of a synthesizer task should not have to think about a concurrently running fax-modem task.

The Mwave OS does not include a file system, but instead relies on the PC's operating system to move data to and from the disk. The Micro Channel permits reasonably fast disk drive interfaces, and OS/2 provides stream handlers with fast threads, allowing for bursts of data. With large buffers (e.g. 64 Kwords), stereo audio can be moved between the disk drive and main memory in real time. Performance over the ISA bus under Windows will be slower, but an EISA-bus system should provide decent performance.

The Mwave alliance expects the API (application programming interface) to support floating-point and multiprocessing extensions in future versions of the Mwave DSP. To ensure that today's device drivers will work with future generations of the Mwave DSP, TI is maintaining strict control of the DSP Manager. Virtual hardware tasks will probably change with new chip architectures, but the intent is that higher-level application device drivers will remain the same.

Intermetrics is preparing a Software Development Toolkit, which it claims "will enable developers to easily migrate compute intensive tasks to the Mwave processor. A comprehensive set of programming modules will provide a powerful development environment, including sophisticated debugging capabilities, libraries, and standard tools." Larry Bertram, Intermetrics Director of Multimedia Products, further indicated that their software toolkit will allow programmers to integrate the standard Mwave virtual hardware tasks with host (PC) applications. A C compiler, debugger, and assembler/linker will be available from Intermetrics.

Conclusions

Multimedia is considered pivotal in bringing computing to the mass consumer market. Yet we should request donations for glasses and hearing aids for the half-blind, tin-eared marketeers who set current multimedia PC standards. Were they naive about media requirements or just being cheap? Video is typically jerky and noisy, audio quality is poor, and few of the synthesized music instruments are convincing. Despite added functionality, most consumers will delay this relatively high expense until multimedia computers provide the media fidelity of other current consumer equipment. Will the Mwave standard provide sufficient quality, speed, generality, and ease of development at a reasonable cost? Without more information than has been made available, it is impossible to answer this question, but Mwave clearly has some strengths and some limitations.

The Mwave subsystem may be viewed as a first step in multimedia computing, supporting applications that are most affordable today (compression, speech, fax and data modem applications as well as limited music synthesis and audio processing applications). Future Mwave DSPs will evolve to provide higher precision, floating-point, and multiprocessing features, enabling general-purpose music processing and high-resolution graphics.

One of Mwave's primary competitors will be AT&T's DSP3210 processor with its VCOS kernel, function li-

Continued on page 17

FPUs must now share the same die with the CPU, so there is pressure to keep the transistor count from increasing greatly. The net result is longer latencies, and, frequently, lower bandwidth. Often, due to decreased basic cycle times, actual FPU latencies (measured in ns) have held roughly constant, even though latencies in clock cycles have increased. You can see these effects when comparing the newly announced processors to their ancestors. (At first glance, SuperSPARC appears to be the exception, but this is illusory. Since it can execute two *dependent* integer operations in one cycle, at 40 MHz it can be considered as an 80-MHz non-superscalar integer unit with a six-cycle latency and two-cycle bandwidth FPU—a modest boost over the previous four-cycle latency of the SPARCstation 2 at 40-MHz.

The limited transistor budget—and the recognition of Whetstone's misplaced emphasis—probably caused Motorola to drop some of the transcendental and trigonometric instructions from the native instruction set of the 68040. (They are emulated via software traps to maintain software compatibility.) The extra transistors were used to speed up the basic floating-point performance, which is quite respectable. Intel couldn't do that in the i486 due to the primitive state of operating systems for the PC, resulting in a slow on-chip FPU with a lot of microcode.

For a recent example of how a limited transistor budget affects the FPU, compare SuperSPARC and Alpha. Both are implemented in 0.8-micron triple-metal processes with die sizes near the limit of the reticle. (SuperSPARC is BiCMOS, and Alpha has a 0.5-micron effective gate length, but both processes provide about the same number of potential transistor sites.) The SuperSPARC designers clearly were concerned about using inexpensive secondary cache technology, so they devoted large amounts of chip area to the 36-Kbyte primary cache, leaving enough space for only a 75-ns latency FPU (three cycles at 40 MHz). The Alpha designers rely on a high-speed secondary cache with a smaller primary cache. This allows space for a much faster 33-ns latency FPU (5 cycles at 150 MHz).

The Crystal Ball

The battle of the benchmarks will continue. The RISC workstation vendors are moving to single-chip superscalar and superpipelined CPUs with their limited FPUs, while Intel is talking about a P5 (586) with much better FPU performance than the 486.

As device geometries decrease, more and more transistors will allow FPU latencies to continue to decrease relative to integer latencies. Larger transistor budgets will permit latencies to decrease faster than process-driven gate speeds. Soon the law of decreasing returns will set in; Alpha is getting close to the limit for add and multiply, even though its divide performance could be considerably improved. Within the next few years, every desktop microprocessor will have an FPU with a 3 to 5 cycle latency for add and multiply, with radically shorter divide and square root times to be expected thereafter.

In the distant future (5+ years), when the silicon and compiler technologies have advanced to the point where high-performance machines issue three or four instructions every cycle and transistor budgets are dramatically increased, you will see single-chip processors with very fast, dual FPUs. (This trend is visible today in multichip, application-specific processors).

Since their introduction, microprocessor floating-point units have seen tremendous evolution in architecture and implementation. The painfully slow separate coprocessor chips of the early days have been replaced by dedicated, high-performance, floating-point data paths integrated on to the CPU die. The push for high integration, with the CPU, FPU, and cache all on the same chip, has limited the growth in FPU transistor budgets, but future generations will see another spurt of FPU performance increases. ♦

Mwave

Continue from page 11

brary, and development tools. IBM could match the generality of AT&T's solution while beating AT&T's price and performance by releasing the DSP Manager and real-time executive on TI's TMS320C31. Instead, IBM has chosen a less-general, lower-precision DSP core to minimize cost and allow integration of serial ports for peripherals. However, thanks to the Mwave DSP's efficient DMA, many signal processing cycles will be freed from the task of handling real-time I/O.

Although real-time video processing is a key ingredient of multimedia computing, it can be expensive. It is unclear how the Mwave standard will handle any form of real-time video in the future or how it will work with the Intel DVI chips IBM has been using in its multimedia computers. Due to high costs, products based on processors (e.g. iWARP) with the kind of power needed for full-featured multimedia computing cannot be successfully marketed to the masses at present. However, Mwave's dramatic improvement over currently dismal multimedia standards for the PC, along with the potential to support future higher-precision DSPs, may result in Mwave forming a long-lasting foundation. ♦