TI Introduces Four-Processor DSP Chip

320C80's Integral Crossbar, Parallel Operation Give Two Billion Ops

by Curtis P. Feigel

Texas Instruments has announced a powerful new digital signal processor that is capable, the company claims, of delivering two billion operations per second. The 320C80, also called the MVP (multimedia video processor), becomes the flagship of TI's DSP line.

As Figure 1 shows, the 320C80 contains four of what TI calls ADSPs (advanced DSPs). These fast integer processors each allow a high degree of parallel execution and are intended to manipulate and process pixel and bit-field data. The processors share 50K of on-chip SRAM by way of an integral crossbar switch. A dedicated transfer controller moves data into, out of, and within the chip. Two frame controllers support timing data transfers to video I/O devices. The device also contains a RISC processor, called the MP (master processor), that orchestrates the various processors and contains a floating-point unit.

One of the MVP's strengths is not built into the chip but comes along with it. The new 320C80 is not binary-compatible with previous processors in TI's TMS320Cxx family of DSPs. But these processors, the first of which shipped in 1982, are the choice of more than one hundred third-party vendors—a group that is comfortable with the company and has evolved a large set of tools. In the world of product development, this inertia can count for much. TI has ported the basic tool set so that developers will see the same familiar user interface when they port code to the new chip.

Design Driven by Performance

In the early planning stages of the MVP, TI evaluated SIMD (a single instruction thread operating on multiple data sets) and MIMD (multiple instruction threads operating on multiple data sets) architectures for various audio, video, and graphics algorithms. Rather than build special hardware that directly executes, say, P×64 video compression code, the designers chose a flexible (if not exactly general-purpose) processor. The company estimates that video teleconferencing represents only 5–15% of the MVP's total market.

This decision draws on a lesson learned in the early days of DSP development: the main application then foreseen for DSPs was speech compression and recognition, but now, according to TI, only about 1% of DSPs are used for that purpose. Advancing technology has enabled applications that were not obvious at the time. The company estimates that 10–20% of systems that use the MVP will have more than one of these processors.

Anticipating that the MVP would be built in its 0.6-micron process, TI allowed a budget of four million transistors. Designers then concentrated on fitting as much processing power into the chip as possible. The resulting device has enough performance to be used as a complete P×64 codec, performing H.261 video encoding and decoding at 30 fps (frames per second) as well as G.728 audio compression and decompression, all on one chip. A single chip can perform MPEG-1 encoding on CIF-resolution images in real-time, although multiple chips deliver better image quality. The company estimates that, to perform real-time MPEG-2 encoding at full CCIR 601 resolution, designs will need at least four chips. *(see **080204.PDF**)*

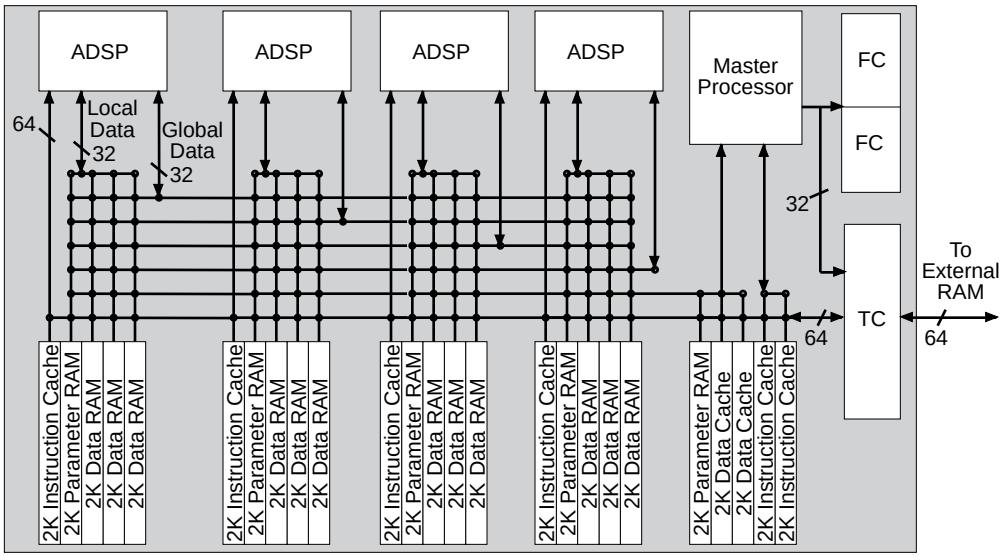

Figure 1. The 320C80's four DSPs and its RISC master processor all share the chip's on-board RAM via a crossbar switch. The transfer controller (TC) acts as an intelligent DMA controller, servicing access requests to external memory space. The Frame Controllers (FC) handle raster-format data.

Since its preview about a year ago *(see **070203.PDF**)*, the MVP has garnered design wins in products under devel-

opment by Xerox, Sony, Matrox, and others. Applications range from color and 3D image processing to document processing and fingerprint recognition. Customers reportedly have already implemented H.261 video decoding and MPEG-1 encoding and decoding.

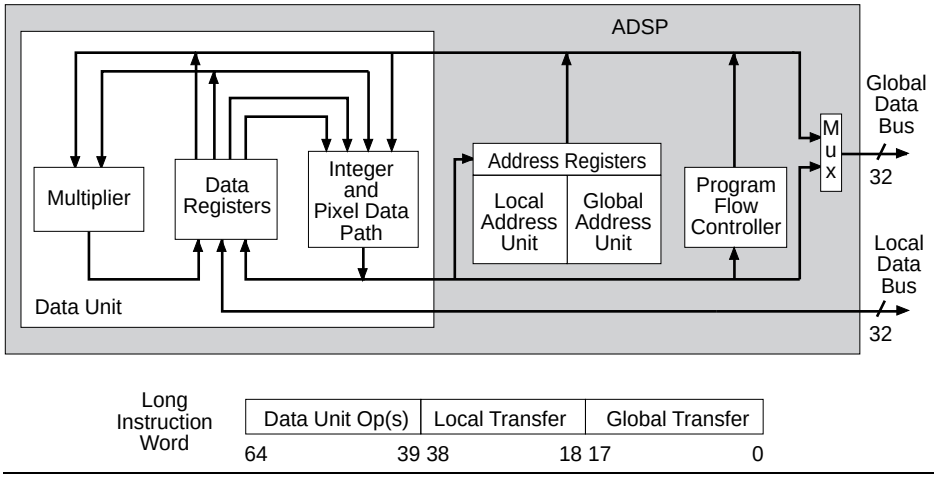

Figure 2. The chip's four ADSPs (advanced digital signal processors) are each divided into three major function blocks. These blocks all work in parallel and contain units that themselves operate in parallel. All functions are controlled by 64-bit LIWs (long instruction words).

Crossbar Offers Flexibility

The MVP's crossbar allows all its processors to operate in parallel, sharing memory while minimizing conflicts. The memory is segmented into 25 blocks of 2K each. The crossbar makes 16 of these blocks accessible from all of the processors. In addition, each ADSP has its own block of local RAM for instruction cache; the MP uses two blocks for its data cache and two blocks for its instruction cache.

The crossbar selects RAM blocks based on the addresses each processor generates. Thus, the crossbar can switch as often as every cycle—access to on-chip RAM can complete in the same cycle. A round-robin priority scheme resolves access conflicts.

As Figure 1 shows, the ADSPs each have a global data bus that can access any of the shared memory blocks and a local data bus that can access only three nearby memory blocks or the processor's parameter RAM memory block. These RAM blocks must be loaded from or stored to external memory by a specific request to the TC (transfer controller). Each processor's parameter RAM is used as a buffer for service requests to the TC—only the TC accesses external memory.

Each ADSP can access its instruction cache, its local data bus, and its global data bus in parallel. Should the processor attempt to access nonlocal RAM over the local bus, the access is diverted to the global bus. Accesses to external memory space generate a request to the TC to perform the operation. The MP can perform one instruction and one data access each cycle. In total, the crossbar can support 15 concurrent accesses per cycle.

ADSPs Use Long Instruction Word

The ADSPs process pixel and bit-field data and perform digital signal processing. Figure 2 shows that each contains three separate operation units that support multiple parallel operations in a single cycle, all controlled by a 64-bit LIW (long instruction word). The PFC (program flow controller) increments the program counter, controls execution of instructions, and handles interrupts. It incorporates three zero-overhead loop controllers. Two independent address units can each perform a load or store operation every cycle. One unit has global access to the shared RAM blocks, and the other can access only the local blocks.

The data unit's integer multiplier can perform either one 16×16 or two 8×8 multiplications per cycle, in parallel with one add. Its three-input ALU can perform multiple parallel operations and can generate up to four 8-bit results per cycle. The ALU's data path incorporates a barrel shifter, an expander that can replicate one, two, or four bits to fill a 32-bit word, and bit-detection logic to identify when the leftmost or rightmost bit changes or is a 1 (useful in entropy-encoding algorithms such as Huffman compression).

As Figure 3 shows, the MP is a 32-bit RISC processor with a floating-point unit. Despite TI's involvement with SPARC, the MP does not use that architecture. It accesses its 4K instruction cache and 4K data cache (both four-way set-associative) via the crossbar. Per the RISC credo, its integer pipeline performs all accesses to cache via load and store instructions; integer and logic operations act only on registers. The FPU has separate pipelines for its multiply and add units, although they operate on the same set of registers as the integer unit. Because integer instructions complete every cycle, while floating-point instructions may require up to 33 cycles, the registers are scoreboarded

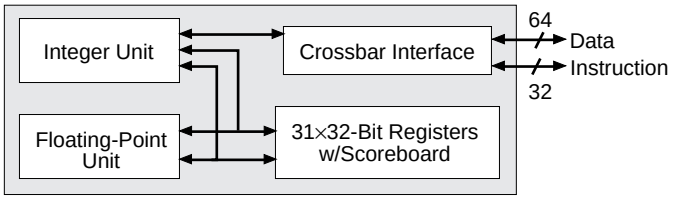

Figure 3. The master processor is not based on any other RISC design. Its architecture is optimized to run code generated from high-level language compilers such as Gnu C.

to ensure operations complete in the proper sequence.

Vector instructions allow the FPU to initiate a single-precision multiply and a single- or double-precision add into the unit's seven-stage pipeline. Results of multiply operations are piped directly into the add unit, which speeds convolution, dot-product, and FFT (fast Fourier transform) calculations. The MP reaches its peak speed of 100 MFLOPS when it initiates a single-precision floating-point multiply and a double-precision add in parallel with a 64-bit load or store.

Transfer Controller Handles External RAM

The TC handles DMA functions autonomously to off-load memory manipulations from the processors. It prioritizes and performs all processor accesses to external memory space, including cache misses. The processors may request data transfers by specifying one-, two-, or three-dimensional source and destination arrays (which TI calls packets).

The TC's memory controller directly interfaces to SRAM, VRAM, and DRAM. A bit in the chip's configuration register determines whether external memory accesses are little-endian or big-endian. Dynamic bus sizing allows the TC to handle data-transfer widths of one, two, four, or eight bytes, configurable on a page-by-page basis. Its maximum data-transfer rate is 400 Mbytes/s. To allow a host processor or another MVP access to the external RAM, the chip's memory control lines can be forced into a high-impedance state.

The MVP also contains two independent frame controllers intended to automatically manage external video capture and display devices. These can directly interrupt the TC to accommodate both interlaced and progressive-scan frames, and they can be programmed to generate PAL and NTSC formats.

Integration Solves *n*-Squared Problem

The crossbar switch is one of the most flexible schemes for connecting multiple processors. Its drawback is that n processors require n^2 connections for each data, address, and control line. Because a 64-bit processor might have one hundred lines or more, the number of connections quickly gets out of hand as the number of processors increases. This is why designers of massively parallel machines explore more exotic interconnection topologies such as rings, tori, and hypercubes. But IC fabrication has advanced to the point that TI can solve the n^2 problem by integrating the processors, RAM, and crossbar on one piece of silicon.

Figure 4 illustrates that TI also mitigates the n^2 problem by providing fewer connections than a full crossbar. With expected advances in IC technology, in a few years it should be possible to double the number of processors on the device. Using the same connection pattern would increase the number of switch elements by only about 190%, not the 400% required for a fully connected crossbar. In the current implementation, the crossbar takes about as much space as one of the ADSPs, but much of this area is due to routing limitations using two metal layers. A projected eight-processor version could use an additional metal layer and actually reduce the crossbar's area. TI does hint that an eight-processor version is on its roadmap, but that it will first market devices with fewer than four processors.

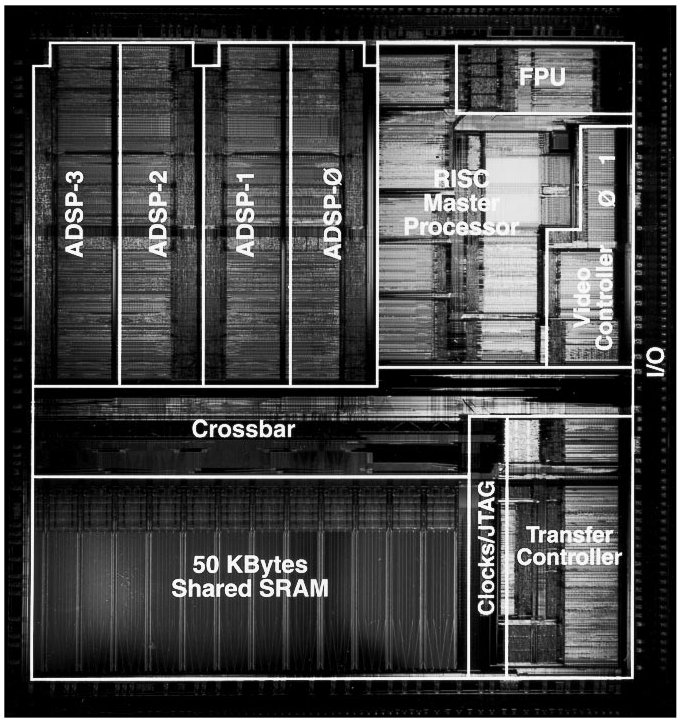

Figure 4. This die photo of the 320C80 shows the partitioning of the device's 4 million transistors. Built in TI's 0.6-micron process, the 18.2 ×18.8-mm device is packaged in a 305-pin ceramic PGA. At 50 MHz and 3.3 V, it consumes a maximum of 7.5 W.

Software May Be Stumbling Block

The real challenge with this chip is to write programs that can take full advantage of all its parallelism. Writing code to run efficiently on a LIW machine is hard enough; coordinating four such processors and debugging the resulting code could be a real headache.

An optimizing compiler may be able to use the hardware with moderate efficiency, but high utilization requires a different programming paradigm—the Linda extensions to C would be a good start. Until that happens, customers are faced with the prospect of hand-coding critical parts of their programs to reach the full potential of this processor.

TI has taken steps toward simplifying the programmer's job by providing a real-time multitasking executive that runs on the MP to manage the chip's parallel-processing tasks. The kernel contains a dispatch routine that schedules multiple tasks to share the MVP's resources. Tasks running on the MP can

pass messages to each other and to the external host. This helps the MVP present a simpler, uniprocessor-like interface to the programmer.

In a typical DSP application, an MP-resident task acquires ownership of an ADSP. It then must build and initialize the data structures that cause the ADSP to operate. Once an ADSP is set up, processing can be as simple a matter as commanding the TC to begin feeding it data. The kernel provides help, but it doesn't change the fundamental problem: finding opportunities to employ the hardware's parallelism falls on the programmer's shoulders.

At the announcement, the company demonstrated a high-level language debugger displaying the states of the MVP's processors in separate windows. Third-party developers are building simulation and hardware-emulation tools to be ready later this year. These include C++-based functional-device and system simulators running under X-Windows. The optimizing C compilers for both the MP and the ADSPs share a common linker, and they load user code for all the processors as a single executable object file. Although none are available yet, TI also promises a library of video, graphics, and audio primitives.

Stiff Competition in Compression Arena

TI quotes some impressive performance figures for the 50-MHz 320C80, including that it can perform:

- 130K Dhrystone MIPS using the MP only
- JPEG encoding of a 352×240 image (4:2:2 format) in 8 ms
- JPEG decoding of the same image in 6 ms
- DCTs on 8×8 kernels at the rate of 800,000/s
- Generation of shaded polygons of 50 pixels each at the rate of 700,000/s
- P×64 encoding and decoding of CIF frames at 30 fps (frames per second)

Even with this impressive performance, the MVP faces some tough competition in the video arena from IIT's VCP chip *(see **080204.PDF**)*. It's difficult to directly compare these two dramatically different devices; the VCP has numerous hardware function blocks designed specifically to handle video compression and decompression algorithms, while the MVP's multiple DSPs are more general-purpose in nature—an advantage, TI argues, due to ever-improving algorithms.

Each manufacturer claims its device has prodigious power: TI rates the MVP at 2,000 MOPS (million operations per second), while IIT claims 4,000 MOPS for its VCP. Both devices are in the $300–$400 price range. The MVP does not have the built-in video preprocessing and postprocessing capabilities of IIT's chip, which give the latter an advantage in embedded and end-user video applications even though it requires an external DSP to process audio.

Price & Availability

TI projects that the 320C80's price, in 10,000-unit quantities, will fall in the $300–$400 range. Samples are available now, with production due in late 1994. For further information, contact TI at 800.477.8924 or your local TI sales office.

Numerous manufacturers are turning out devices with algorithm-specific hardware that will cost an order of magnitude less than the MVP. HP has extended the instruction set of its PA-RISC processors *(see **080103.PDF**)* and MIPS Technologies is not far behind. Keep in mind that TI has formed an alliance with video-compression leader C-Cube, which indicates that the MVP is not intended to compete for the high-volume, price-sensitive designs.

The MPR Cost Model *(see **071004.PDF**)* estimates that, once the process is mature, each MVP will cost TI about $260. A version of this chip with only one or two processors could be less expensive—one-half to one-third the cost, according to TI.

Design Wins in Exotic Applications

Xerox intends to use the MVP in a document-processing system involving capture, recognition, and indexing of paper forms. Printrak (Anaheim, Calif.) is using the device to implement fingerprint recognition in a sophisticated security system. Both companies were attracted by the MVP's mix of high-performance processing and general-purpose flexibility. Printrak's use of a single MVP allows the company to eliminate 28 circuit boards from its current system.

The MVP, as a more general-purpose coprocessor, could speed color and high-resolution image processing in the prepress industry. The transfer controller's ability to handle multidimensional blocks of data would be welcomed in 3D and virtual-reality applications. The MVP's built-in crossbar makes the chip a natural for network bridges that must convert among multiple protocols—the TC's ability to handle arrays seems an obvious way to manage packetized data.

Although the initial version of the MVP may not find its way into videophones, set-top boxes, or any other consumer multimedia applications, its price is in line with its high performance. Future chips with fewer processors and lower cost may one day hit those targets. The 320C80 comes with something more critical to a high-end processor: a set of tools and a support system, putting designers ahead in the time-to-market wars. Establishing a new architecture is always a challenge, but the MVP has the kind of performance that makes the effort worthwhile. ♦

