Crystal SLIMD Speeds PCI Audio

Innovative Signal-Processing Architecture Offers High Performance, Low Cost

by Peter N. Glaskowsky

Crystal Semiconductor's first entry into the expanding PCI audio market combines an innovative 300-MIPS DSP core with 43K of on-chip SRAM and a powerful DMA engine, eliminating the need for additional local memory while providing ample horsepower for the most demanding audio algorithms, including Dolby AC-3 decoding.

The new CS4610 audio accelerator is built around an original DSP core that shares some features with media processors like Chromatic's Mpact while maintaining a small die size (under 50 mm^2) and a lower price point than most of its competition (only $22 in 10,000-unit quantities).

Crystal has dubbed the core a SLIMD architecture: somewhat long instruction, multiple data. Each instruction word is 40 bits long and dispatches three independent operations to one ALU and two MAC units. Indeed, the CS4610 is a 500-MIPS machine by some standards, since it can perform one 16-bit ALU operation, two 20 × 16-bit multiplies, and two 40-bit adds in each cycle at 100 MHz.

Core Processor Optimized for Stereo Audio

Unlike a conventional DSP core, Crystal's new design is optimized for audio synthesis. The core combines a 16-bit integer ALU with two parallel multiply-accumulate (MAC) units that can simultaneously perform the same function on both samples in a stereo data stream. All three units are controlled by a single 40-bit instruction word.

The SLIMD core, shown in Figure 1, connects a 16 × 16-bit register file (also configurable as 8 × 32), a 16-bit ALU, dual MACs with shift and saturation logic, and controllers for the on-chip memory. The core contains 250,000 transistors, roughly one-third of the CS4610's die.

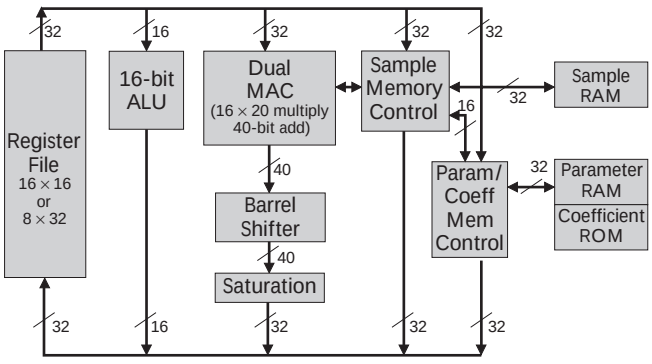

Figure 1. The CS4610's SLIMD core includes registers, an ALU, dual MACs, and local memory. The design is more complex than shown here; there are actually more than 10 separate data buses.

The register file consists of eight register pairs, facilitating the parallel MAC operations. It has four write ports and six read ports, supporting simultaneous operation of the ALU and both MACs.

As Figure 2 shows, four types of local memory are provided, accounting for another third of the die. Firmware modules are downloaded from the host and stored in a 3.5K × 40-bit local program RAM. The host driver selects the appropriate modules to download according to the functions required by the active application.

An additional 12K of parameter RAM stores downloaded data items associated with the current audio algorithms, while 24K of coefficient ROM holds precalculated coefficient data for sample-rate conversion operations, a necessary part of any AC'97 controller. The AC'97 codec interface, AC-Link *(see* **1009MSB.PDF***)*, runs at a fixed 48-kHz sample rate, so audio data from the applications running on the host must be converted from their native rate, which varies from 8 to 48 kHz, to the 48-kHz rate required by AC-Link.

Finally, the CS4610 includes 14K of RAM to hold audio sample data, typically short segments of wavetable data transferred from host DRAM for use in audio synthesis.

Host transfers are managed by a sophisticated DMA engine that can handle up to 96 simultaneous data streams with full scatter/gather support. This high degree of complexity allows Crystal to dispense with local wavetable memory and works even with limited on-chip storage. In classic ISA-based wavetable synthesizers, entire wavetables were

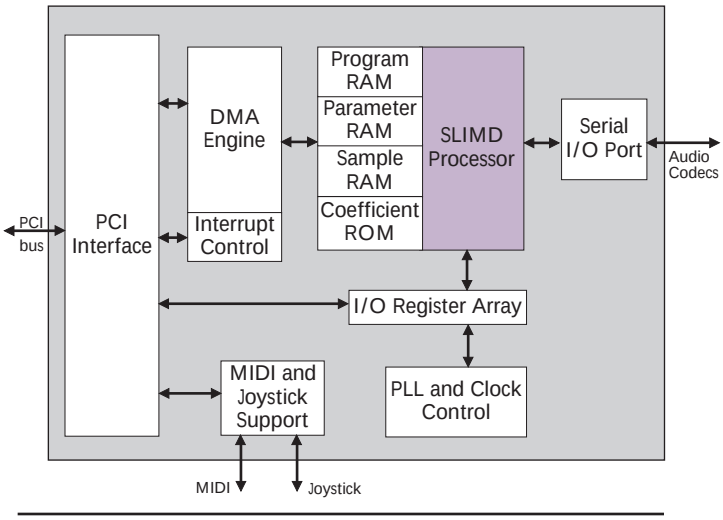

Figure 2. Crystal's CS4610 is built around a 300-MIPS SLIMD processor and includes 43K of on-chip SRAM and a DMA engine plus interfaces to PCI, MIDI, and joystick peripherals as well as an external audio codec.

Price & Availability

The Crystal CS4610 is currently available in sample quantities, with full production scheduled for 3Q97. The device comes in a 100-pin metal quad flat pack (MQFP) and is priced at $22 in 10,000-unit quantities.

For more information, contact Crystal Semiconductor (Austin, Texas) at 512.442.7555, or visit the company's Web site at *www.crystal.com*.

stored in local ROM or RAM, adding significant cost to the audio subsystem. Crystal's DMA engine, combined with a just-in-time DMA-request controller, allows the small on-chip RAM to substitute for more expensive local memory.

ISA Out of Gas; Audio Vendors Board PCI Bus

Crystal's CS4610 joins several PCI audio chips announced last fall, including VLSI's SongBird *(see* **1015MSB.PDF***)*, S3's Sonic Vibes, and Oak's TelAudia3D family. As Table 1 shows, these parts illustrate the wide range of design choices faced by today's audio-chip designers. VLSI's VL82C829 and Oak's OTI-611 both include DSPs to support 3D audio, but neither is said to support AC-3 decoding.

PCI audio usually precludes Sound Blaster compatibility, but most vendors have now developed schemes to support this legacy audio standard. In motherboard applications, S3's chip includes internal Sound Blaster emulation hardware and can be connected to ISA DMA and IRQ lines. Crystal's solution uses a CS423xB ISA sound chip to support Sound Blaster audio as well as the audio codec functions required by the CS4610. Only VLSI offers a PCI-only solution that preserves full Sound Blaster compatibility, but SongBird's Sound Blaster emulation requires core-logic support for the serial-IRQ and distributed-DMA standards, which is not provided by Intel chip sets.

	Crystal CS4610	S3 Sonic Vibes	VLSI Songbird	Oak OTI-611
MIPS	300	n/a	25*	25*
Local memory	No	No	Optional	Optional
Sound Blaster-compatible	External hardware	Optional	Optional	No
AC'97 codec interface	Yes	No	Yes	Yes
3D audio	HRTF	SRS only	HRTF	HRTF
Dolby AC-3	Yes	No	External HW	No
Downloadable firmware	Yes	n/a	No	Yes
Soft modem	Future*	No	No	Host-based
Price	$22	$30	$25	$25*
Package	100-MQFP	160-PQFP	208-MQFP 208-TQFP	160-PQFP

Table 1. Crystal's CS4610 offers much higher performance and supports more audio algorithms than competing sound chips. n/a = not applicable. (Source: vendors, except *MDR estimates)

Oak's OTI-611 is unique in supporting an external modem codec, and its price includes a license for Oak's V.34 host-based soft-modem code. Soft-modem implementations *(see* **101602.PDF***)* are a natural extension of these new PCI audio chips, especially when used with AC'97 codecs that include both audio and modem interfaces.

Crystal has no announced plans to support soft-modem algorithms on the CS4610, but the part has more than enough performance to perform these tasks internally rather than on the host processor, as Oak does. Crystal has said that the CS4610 can be programmed to make use of the modem channels on an AC'97 audio codec. With soft-modem support, the CS4610 would be even more valuable to motherboard and add-in card vendors, and we believe Crystal will take advantage of this opportunity as soon as the soft-modem code can be developed internally or licensed from an outside vendor such as AltoCom (formerly part of General Magic), PCTel, or Motorola.

Eschewing these ambitious goals, VLSI and Oak rely on relatively simple 16-bit DSP cores running at moderate speeds, providing just enough performance to support audio algorithms such as the head-related transfer function (HRTF) that creates the illusion of 3D sound. These parts are not fast enough for Dolby AC-3 decoding, waveguide synthesis, or other demanding applications, requiring the designer either to provide additional hardware such as a separate AC-3 decoder chip or to use the host processor to implement these functions. In contrast, the faster DSP core on the Crystal part provides more than enough performance to offload all audio processing from the host CPU.

Oak and Crystal both support downloadable microcode for their on-chip DSPs, while VLSI does not; SongBird has its code fixed in ROM. Downloadable code permits simple software-only field upgrades.

Crystal Chip Is Best Buy

The parts also differ in their use of PCI. The key advantage of PCI in this application is its low-latency, high-bandwidth access to main memory. All PCI audio chips support bus-master operation, allowing them to access digital audio data in main memory. This is the preferred method for use with Microsoft's DirectSound API, eliminating the need for the host CPU to send data to the audio chip one sample at a time. According to S3, when the audio chip handles this data-movement function, host CPU usage can drop from as high as 38% for an ISA-based solution to less than 1%.

S3 and Crystal believe strongly enough in this method to eliminate all local wavetable storage, reducing the pin count of their parts. Oak and VLSI support optional wavetable ROMs, reducing bus activity when performing the wavetable-based synthesis commonly used for games and MIDI playback.

Crystal's implementation gives it an edge over the other PCI-based sound chips. None support as many simultaneous data streams, limiting them to fewer "voices" and therefore less realistic sound. Oak's TelAudia 3D, for example, supports

only 24 voices and requires a local ROM for wavetable synthesis. With its 96-voice polyphony, the CS4610 exceeds the capabilities of even many high-end audio synthesizers.

Without the need for off-chip memory, the CS4610 provides the most compact solution of available PCI audio chips. The standard package is a 100-pin metal quad flat pack (MQFP). Competing products range from S3's and Oak's 160-pin PQFP to the 208-pin MQFP and TQFP package options on VLSI's VL82C829.

The CS4610's $22 price compares well to the $30 asking price of S3's SonicVibes, which lacks a DSP and does not support full 3D audio, although a simpler SRS surround-sound mode is available. VLSI's SongBird is only $25 with complete 3D-audio support, but to support DVD playback, it must be used with an external AC-3 decoder such as Zoran's $20 ZR38600 or with host-based AC-3 decoding software.

Crystal to Boost Performance

Crystal's new chip offers a variety of configuration options for system designers. For example, designers may choose to provide full Sound Blaster compatibility by combining the CS4610 with Crystal's existing CS423xB ISA-bus audio codec. Alternately, any standard AC'97 codec will provide full Windows compatibility plus DOS game support in a DOS window while eliminating any connections to the ISA bus—a desirable feature in new PCs because of inevitable ISA resource conflicts. With two external Crystal CS4333 stereo DACs, the new part will support full 5.1-channel output for Dolby AC-3 or MPEG-2 audio decoding.

Crystal has shown an ambitious roadmap that includes a 125-MHz (375-MIPS) speed-bump part in 1Q98, followed by an even more complex 150-MHz, 900-MIPS device in 3Q98. That much performance would be more than ample to permit simultaneous execution of every interesting audio algorithm—waveguide synthesis with 3D spatialization plus AC-3 decoding and a soft modem, for example.

Media processors like Chromatic's Mpact, however, are even more powerful, able to handle all the same audio algorithms as well as 2D and 3D graphics, digital video, and more. Some OEMs will appreciate the all-in-one nature of media processors, while others will prefer the flexibility of separate audio and graphics subsystems. The 3D graphics market is evolving very rapidly, with a new generation about every six months, and media processors are unlikely to maintain the same pace, due to their greater complexity. Using an audio chip like Crystal's will allow an OEM to upgrade to newer 3D chips as they become available.

Crystal will also face competition from less expensive PCI audio solutions and, we believe, from AC'97 controllers integrated into future Intel core-logic chip sets. Oak's recently announced OTI-610 provides basic 3D audio for only $15, and pure host-based solutions may need only a $5 AC'97 codec. The greater power and flexibility of the CS4610 should give it a unique advantage in home PCs, however, even as sound chips get cheaper and host processors get faster. MPR