MICROPROCESSOR REPORT

TigerSHARC Swallows DRAM

Analog Devices' TS201S Includes 3MB of eDRAM

By Markus Levy {7/14/03-02}

Embedded DRAM was the focus of the Analog Devices announcement of its TigerSHARC TS201S at **Embedded Processor Forum 2003**: the TS201S contains 3MB of eDRAM. The TS201S is based on a lineage of DSPs that began with the introduction of the TigerSHARC

architecture, ADI's third generation of DSPs. Key architectural features include its ability to perform both floating-point and 1-, 8-, 16-, and 32-bit fixed-point DSP operations, glueless multiprocessing support, and high-speed link ports.

Like its closest competitor, Texas Instruments' TMC320C64x, TigerSHARC uses a very long instruction word (VLIW) load/store architecture. TigerSHARC executes as many as four instructions per cycle with its interlocking ten-stage pipeline and dual computation blocks. Each block contains a multiplier, an ALU, and a 64-bit shifter and can perform one 32- × 32-bit or four 16- × 16-bit multiply-accumulates (MAC) per cycle.

The TS201S has four dedicated data buses associated with the core's functional blocks (J-ALU, K-ALU, program sequencer, and SoC interface). The TS201S's buses—namely, the JBUS, KBUS, SBUS, and SoC bus—are 128 bits wide and support the core's Harvard architecture (Figure 1). The first three buses (J, K, and S) run at full chip clock rate: at 600MHz, each bus transfers 9.6GB of data per second, more than enough to handle the bandwidth requirements of many applications. Extreme bandwidth demands can be demonstrated using a FIR filter performing sustained multiply-accumulates. The TS201S can perform eight MACs per cycle. Each MAC needs two 16-bit

values (32 bits total), eight times 32 accounts for the 256 bits provided by the J and K buses.

The SoC bus provides a bridge between the chip's memory subsystem, core, and peripherals. It also runs at half the core speed, with the aid of buffers and FIFOs to handle the
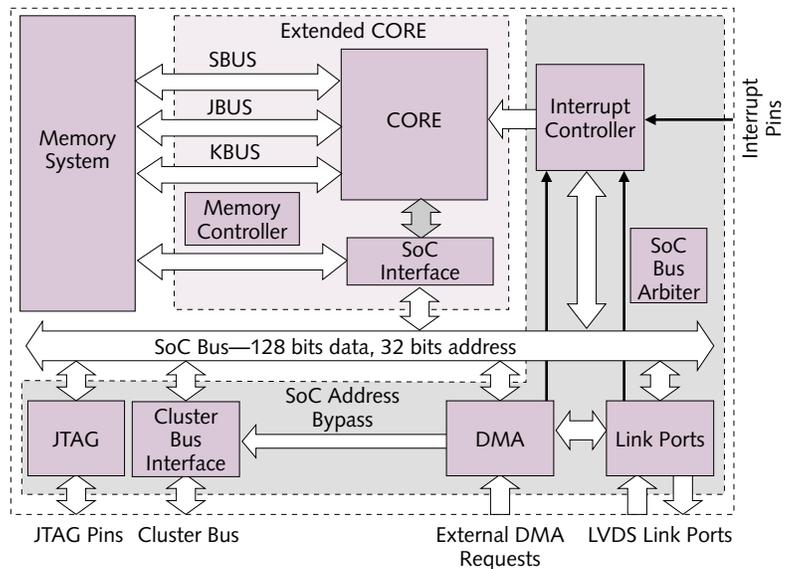

Figure 1. Each of the TS201S's four on-chip buses supports up to 9.6GB of data bandwidth per second. This separate bus structure allows the core to simultaneously access three memory blocks while the chip's DMA is filling the remaining three memory blocks.

Article Reprint

mismatch between on-chip and external speeds. This bus will allow Analog Devices to evolve the product line by attaching various peripherals.

Digesting the TS201's Memory Bandwidth

The TS201S processor contains 24Mb of eDRAM memory, certainly the largest memory array in a DSP that we have seen. This 3MB array is divided into six simultaneously accessible blocks of 4Mb each (128K words × 32 bits) that connect to the four internal buses through a crossbar interface connection. This block-wise configuration enables the processor to perform up to four memory transfers in the same cycle.

Contrast the TS201S's bus arrangement with the previous-generation Tiger-SHARC (the TS101S), which has three buses associated with its memory blocks, as opposed to the core's functional blocks. This architectural difference leads to greater throughput on the TS201S processor, as it allows more-flexible data movement between the functional blocks. Many DSPs support a memory model that allows access to the program and to x- and y-memories simultaneously. The TS201S goes one step further in that it also allows an I/O access to memory in the same cycle. This I/O access permits the chip's DMA to transfer data to and from this fourth memory block while the core continues to run.

Memory access conflicts occur when two functional blocks attempt to access the same internal memory block in the same cycle. When a conflict occurs, extra penalty cycles are incurred for the conflicting access. The extent of the

Kevin Leary, product line director for wireless infrastructure at Analog Devices, discusses details of ADI's TS201S at EPF2003.

penalty partly depends on the way the conflicting accesses stack up on the priority chain. The TS201 organizes priorities in the following order (from highest to lowest):

1. High-priority external transactions: S-bus
2. J-ALU accesses: J-bus
3. K-ALU accesses: K-bus
4. Low-priority external transactions: S-bus
5. Instruction fetches: I-bus (lowest)

When a transaction is stalled, it is held off according to priority. New transactions may bypass old transactions from different sources, if their priority is higher. The extent of the penalty also depends on the length of the transaction from higher-priority transactions, but this is under control of the programmer.

In the worst-case scenario, where all functional blocks are attempting to access a single memory block simultaneously, the penalty is no worse than would be experienced with a traditional memory model using a single bus. On the other hand, the penalty cycles arising from conflicting access can be avoided if the linker tool positions the instruction and data in different memory blocks.

Cache Rescues eDRAM

When eDRAM is compared with external DRAM, the benefits of eDRAM are obvious. Those benefits include reduced latency, dramatically increased bandwidth to main memory, reduced power consumption, reduced space, and the multiple, independent address/data streams.

In general, eDRAM derives reduced latency from its closeness to the CPU core and by the use of small-size memory blocks (0.5MB) but will still incur a penalty of three to six cycles per access. Furthermore, the TS201S's eDRAM operates at half the core speed, which restricts the performance of the chip's high-speed buses. Although adding significant design complexity and die area to the chip, each of the TS201S's six memory blocks is supported by 16K caches and 1K prefetch buffers that potentially eliminate the latency of the chip's eDRAM (Figure 2). ADI claims that, except for a few pathological cases, performance of the eDRAM memory system is only 3% slower than that of one using a standard level-1 cache.

The prefetch buffer does an autoprefetch that anticipates forward or backward reads. Any sequential accesses come from the prefetch buffer, and not from the cache, thereby avoiding cache miss overhead. The cache then takes care of the nonsequential data accesses, such as those seen during the processing of an FFT, where the data is in bit-reversed order.

The TS201 uses a least-recently-used (LRU) cache-replacement policy, with a last-replaced-page (LRP) optimization for embedded DRAM. Using user-accessible
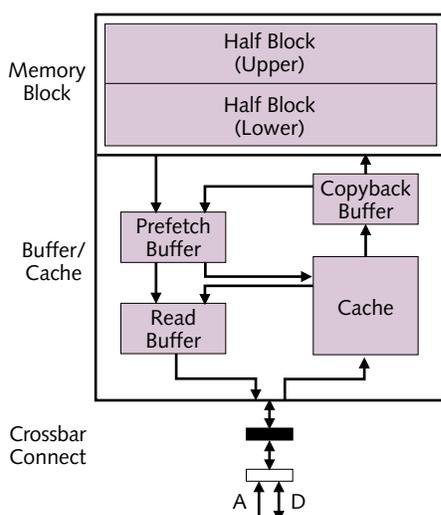
Figure 2. The eDRAM's bandwidth and latency are significantly improved with the use of four-way set-associative 16K caches.

memory system commands, the cache also supports locking on a per-way basis to avoid replacing specific cache contents.

Refreshing News

As with any DRAM, the TS201's eDRAM requires periodic refresh that automatically occurs every 32ms per subarray at 85 degrees Celsius, although the chip supports the programming of higher-frequency refresh rates. The good news is that refresh-associated stalls have a minimal negative effect on performance, thanks to the integrated cache and low leakage. The better news is that eDRAM cells have virtually zero leakage current, especially when compared with SRAM. However, eDRAMs are more difficult to power down than SRAMs (i.e., eDRAMs are not static devices). Furthermore, the TS201S's buses are tied together through the crossbar switch without the use of a memory controller to manage the power down of individual blocks.

Will Customers Take the Bait?

Despite the low leakage current of the eDRAM, many other factors contribute to the chip's power consumption. The TS201S processor has separate power supply connections for internal logic (VDD), analog circuits (VDD_A), I/O buffer (VDD_IO), and internal DRAM (VDD_DRAM) power supply. Combined, these power sources add up to 3.4W at 500MHz and 1.0V operation, including internal peripheral activity but not I/O power. (The core consumes 2.39A at 1V, and the eDRAM adds another 0.67A at 1.5V when using the external regulator.)

From a performance perspective, ADI claims the TS201 supports an "all software solution" for processing both the chip and symbol-rate functions of a 3G base station. Although we do not doubt this, it's questionable whether this is the approach OEMs want to take, as it deviates from the traditional base stations that still use hardwired functions for symbol-rate processing. However, we believe some of the resistance to software-defined radio may be coming from

Price & Availability

ADI will offer three pin-compatible TigerSHARC processors with different memory configurations. The ADSP-TS201 is offered at 500- and 600MHz with 3MB of eDRAM; the 500- and 600MHz versions sell for $207 and $299, respectively, in 10,000-unit quantities. The ADSP-TS202 (1.5MB) and ADSP-TS203 (0.5MB) are offered at 500MHz and sell for $149 and $34.95, respectively.

OEMs that employ a large number of ASIC designers. Furthermore, it should be no surprise that both ADI and Texas Instruments are going after the base station business. ADI's toughest challenge will be to get customers to jump into SHARC-inhabited waters and to steer them away from TI, whose DSP chips currently reside in systems from eight of the top-ten base station manufacturers.

In the future, ADI will take advantage of the Tiger-SHARC's SoC bus and integrate more on-chip peripherals. In the meantime, the TS201 is geared toward multiprocessing but is weak on peripheral support. It would benefit by including simple serial ports and other interfaces, such as Utopia; currently, OEMs will be required to build custom ASICs or use FPGAs to support these external functions. ADI offers Verilog and VHDL code that allows designers to implement a link-port interface using Xilinx FPGAs.

When it comes to performance, the TS201 is a winner. With its ability to crank out eight 16- × 16-bit MACs every clock cycle, it doubles the performance of TI's C64x. Moreover, with its large on-chip memory combined with the unique cache feature, the TS201 has plenty of bandwidth to keep up with its processing prowess. At $299 for a 600MHz version, the TS201 is pricey, but ADI offers lower-cost versions with less memory to help the TigerSHARC angle for the fisherman. ◇