Cahners
MICROPROCESSOR REPORT
www.MPRonline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

TI Releases Details on First 'C64xx

By Jennifer Eyre and Jeff Bier, BDTI {6/26/00-06}

At **Embedded Processor Forum 2000**, Texas Instruments revealed new details about the first implementation of its recently announced high-performance DSP architecture, the TMS320C64xx. The TMS320C64xx is the successor to TI's well-known TMS320C62xx architecture and bumps up the earlier processor's performance with significant architectural enhancements and a much higher clock speed (see *MPR 3/6/00-01*, "TI Cores Accelerate DSP Arms Race"). Without going so far as to announce a specific chip, TI lifted the veil just enough to disclose new information about the memory structure and peripherals that will be included in the first 'C64xx product.

Because the 'C64xx is capable of very high levels of parallelism (for example, execution of four 16-bit multiplications in parallel with four 16-bit additions), getting data in and out of the processor core is a key concern—particularly given its high initial target clock rate of 600–800MHz. DSP processor architects have historically shied away from using caches, because they detract from the programmer's ability to predict code execution times, but the drive for higher clock speeds is leading DSP architects down some of the same paths previously trod by high-performance CPU designers. One of these is cache use. TI had already begun to explore this path with the TMS320C6211, the cache-based low-budget member of its 'C62xx family, and continues with the 'C64xx. For DSP processors with caches, the availability of a cycle-accurate instruction-set simulator that correctly models cache effects is a key component in the success of the product; DSP software developers need this tool to guarantee real-time performance and effectively optimize their software.

The first TMS320C64xx product, like TI's TMS320C6211, will use a two-level cache configuration. The sizes of the caches are larger on the 'C64xx than on the 'C6211, since the newer architecture is capable of much higher levels of parallelism and requires commensurately more data to operate at peak efficiency.

The 'C64xx chip will contain two level-one (L1) caches, one for data and one for instructions. Each of the L1 caches contains 16K of memory (four times that of the L1 caches on the 'C6211). The L1 program cache is direct mapped; the L1 data cache is two-way set-associative. The L1 caches are fed by a unified level-two (L2) cache, which contains four 32K memory banks totaling 128K (double the L2 of the 'C6211). The L2 cache can be configured as non-cached RAM, as a set-associative cache, or as a partitioned combination of the two. For example, the L2 can be configured as one bank of RAM, with the remaining three banks set up as a three-way set-associative cache. The L2 cache is fed via what TI calls an enhanced DMA controller, or EDMA. The EDMA supports 32 channels, and TI claims it can support more than 2.6GB/s of bandwidth. The chip also will include three multi-channel buffered serial ports, three timers, and three off-chip interfaces: a 64-bit interface for connection to memory, a 16-bit interface for I/O, and a 32-bit host-port interface.

The processor's cache configuration will be an important factor in its performance, given that the width of the external memory interface, combined with the chip's high projected clock speed, means that there will be a substantial performance penalty for applications that make extensive use of off-chip memory. This performance penalty was sometimes problematic for the 'C62xx and is likely to

remain so on the new chip. Although TI doubled the width of the external memory interface relative to that of 'C62xx chips, this increase won't be enough to compensate for the increased bandwidth demands caused by the new device's massively increased clock speeds. The 'C62xx, at 250MHz, is hard pressed to get data from off-chip memory fast enough to allow the core to operate at peak efficiency; at 1.1GHz, the 'C64xx won't be able to get anywhere near realizing its full potential when accessing off-chip memory. Programmers will need to ensure that their code and data fit mostly in on-chip memory if they expect the processor to perform as advertised.

When Texas Instruments introduced the 'C62xx several years ago, the processor drew widespread attention because of its huge speed advantage relative to other DSP chips of the time. However, its success in the market has been somewhat hampered by its high energy consumption and poor code density. With the 'C64xx, TI has taken steps to address these issues, but in the process it has created some new tradeoffs—particularly in the use of cache memory. The problem of execution-time predictability is mitigated somewhat, because L2 cache can be configured as RAM, but most other (noncached) DSP processors don't have this problem at all. Then again, most other DSP processors won't be running at 1.1GHz. It remains to be seen whether TI's newest firebreather will find acceptance in the broad market or will mainly serve to give TI performance bragging rights while it goes about selling millions of its low-power 'C5xxx chips.