Most Significant Bits

ARM Addresses Midrange with ARM8 Core

Advanced RISC Machines (ARM) has announced its second new core in as many months, following StrongArm. The ARM8 core promises to double the performance of the existing ARM7 core, to 80 Dhrystone MIPS, while maintaining many of the traditional ARM values of small size and low power consumption.

The new core gains performance mostly through higher clock rates, which, in turn, are enabled by stretching the pipeline. Like StrongArm, the cooperative venture between Digital Semiconductor and ARM *(see **091504.PDF**)*, the ARM8 core uses a five-stage pipeline. The execution of each instruction is still completed in a single clock cycle (with one exception, noted below), but cache accesses and register writebacks are relegated to individual pipeline stages. The new core is compatible with ARM7, so existing software will be able to migrate to the new chip for a performance boost.

The first ARM8 implementation, ARM810, will wed the new core with an 8K unified cache, write buffer, MMU, and 32-bit external bus. ARM810 chips will appear in the second half of 1996.

The primary virtues of many newer embedded CPU cores is their small die size and uncomplicated pipeline structure. MIPS, ARM, and PowerPC vendors all offer embedded cores with simple three-stage pipelines that can be fabricated in less than 5 mm^2 of silicon. The short pipeline, while advantageous in terms of real estate, holds clock frequencies down because instruction execution is forced into a single clock cycle. ARM is particularly constrained because of its unusual inline shifter placed between the register read port and the ALU.

ARM8 alleviates this critical timing path by streamlining the inline shifter. In an ARM7 design, one register operand may be shifted left or right by any number of bits as part of any other logical, arithmetic, or load/store operation. In contrast, the ARM8 design allows only logical left-shifts and only by 1–3 bits. Larger shift counts are handled by a second, fully featured, barrel shifter—which doubles as part of the integer multiply unit—and require an additional cycle to execute.

Although the capabilities of the inline shifter are severely compromised, ARM estimates that most existing code will be affected very little. Customer surveys indicated that inline shifting is used primarily to index arrays of words or doublewords, so limiting pointer indexing to ×2, ×4, and ×8 operations was not seen as a major drawback.

Instruction prefetching is also modified to increase core performance. The ARM8 includes an eight-instruction prefetch FIFO. As instructions are added to the FIFO, a check is made for branch instructions. Unconditional branches automatically redirect the prefetch unit and are eliminated from the FIFO, essentially hiding the branch latency behind subsequent instructions.

Conditional branches use static branch prediction. Forward branches are assumed to be not taken; backward branches are assumed taken, and the prefetch unit is restarted with the target address. Correctly predicted branches have zero latency, while mispredictions cost 2–3 cycles.

The designers took the unusual measure of defining the cache interface as part of the CPU core design. ARM8 accesses its unified cache much differently than do earlier generations. One 32-bit instruction or data word is transferred during each phase of the internal pipeline clock rather than on each cycle. Using both clock edges doubles cache bandwidth without resorting to a 64-bit bus. This method is not carried to the chip's external bus interface; thus, ARM8 chips will experience a harsher cache-miss penalty than usual.

The sum total of these changes is that the ARM8 will double the Dhrystone MIPS rating of the fastest ARM7 design, according to the company. Much of the advantage accrues to its 75-MHz clock rate—55% faster than the ARM7 currently runs—in a 0.5-micron, three-layer-metal process.

The changes have an attendant cost as well. In equivalent processes, an ARM8 core is slightly more than twice the size of an ARM7 before caches are added. Double the performance for about double the silicon area may be an attractive tradeoff for many ASIC designers, particularly those who couldn't make do with the under-40-MIPS performance of the ARM7.

Between its two most recent core announcements, ARM has squarely addressed concerns that the architecture would be bottled up at the low end, hindered by its pipeline and its power budget. ARM8's promise to deliver competitive integer performance from a relatively conventional 0.5-micron process should allow a number of ARM's licensees to capitalize on the new design and provide a more reassuring roadmap to customers evaluating the low-power core.

Exponential Confirms PowerPC Plan

In its first public statements, Exponential Technology confirmed much of what has been speculated about the company in the past *(see **0817MSB.PDF**)*. The San Jose (Calif.) startup will market a high-end PowerPC processor that it claims will be the fastest in the world. To achieve these speeds, Exponential will use a unique fabrication process that mixes bipolar and CMOS circuits. Unlike traditional BiCMOS designs, most of the logic will be bipolar, with CMOS mainly for the dense mem-

ory structures, a technology strategy similar to MicroUnity's *(see **091402.PDF**)*.

Most other IC vendors have dropped BiCMOS, and its biggest proponent, Intel, plans to move to a pure CMOS process after the 0.35-micron generation *(see **090905.PDF**)*. Conventional wisdom says bipolar devices add little speed below this level. Exponential claims to have licked this problem, as well as the excessive power consumption that has plagued previous bipolar devices. However, the company is not yet offering any details to help evaluate these claims.

The PowerPC chip is nearing tapeout, and Exponential expects it to appear in systems in early 1997. The company's fab partner remains unnamed. Apple is one of the company's backers and has provided some assistance with the design. Both IBM and Motorola are said to look favorably upon the effort—perhaps realizing that anything done to advance PowerPC performance will benefit all PowerPC vendors. However, neither has provided any specific technology or even a formal license to the instruction set, although the latter is under negotiation.

If Exponential delivers on its claims, its chip could fill the high-end gap left by the disappointment caused by the PowerPC 620. Exponential's chip is a 32-bit implementation designed for desktop systems, leaving the 620 to address the server market. The new device will be compatible with the PowerPC 60x bus used in today's desktop systems, making it simple for Apple and others to upgrade their boxes. Many have claimed to be developing the world's fastest processor but failed to deliver; it will be interesting to see if this ambitious startup can beat the odds.

SGS-Thomson Speeds DX4 to 120 MHz

Scant weeks after announcing its 100-MHz 486DX4 chip *(see **0915MSB.PDF**)*, SGS-Thomson has increased the chip's maximum clock rate to 120 MHz. Like the 100-MHz part, the new DX4-120 has an 8K write-back cache and an Intel-compatible pinout. The faster part is housed in a plastic pin grid array (PPGA), which the company claims is only half as expensive as a conventional ceramic PGA package.

By releasing a faster version of the 486DX4, SGS-Thomson pulls ahead of Texas Instruments in 486 performance. Although both SGS-Thomson and TI share the same Cyrix design, the French company has utilized a more aggressive 0.35-micron (effective) process, reducing the chip's manufacturing cost and enabling the company to achieve a higher clock rate. SGS-Thomson is now second only to AMD in the 486 clock-speed regatta, falling behind the curiously named 133-MHz Am5x86-P75 *(see **0915MSB.PDF**)*. The new DX4-120 will sell for around $72 in 1,000-piece quantities when it enters production in 1Q96; the price of the 100-MHz part has been reduced to $40.

Digital Does MPEG-1 Encoder/Decoder

Digital Semiconductor has cut the cost of real-time MPEG encoding with its 21230 PCI video codec. The 230 supports real-time compression of video to MPEG-1 IBP- or IBBP-frame and I-frame standards. For videoconferencing, the chip can also support the H.320 standard (including H.261) for ISDN lines and H.324 (including H.263) for transmission over conventional copper wiring.

Resolutions at NTSC, PAL, and CIF scales are all available, with interpolation performed as required to maintain a 4:3 aspect ratio. The chip does not perform audio compression and decompression internally but relies on a host CPU to provide G.711, G.722, or G.728 processing. With a 90-MHz Pentium, the 230 is said to process MPEG-1 audio and video streams in real time. This performance allows the 230 to be used as the basis of a low-cost videoconferencing system. Pricing for the chip has been set at $75 in quantities of 10,000.

Although a complete PCI videoconferencing board requires substantial additional logic (graphics controller, frame buffer, RAMDAC, video decoder, audio decoder, clock, etc.), Digital estimates the complete bill of materials costs less than $200. While other vendors have delivered less expensive MPEG-1 decoders, Digital's 21230 is the least expensive solution for encoding MPEG-1 in real time, undercutting competitors' prices.

Toshiba Shows MPEG-2 Decoder

At Comdex, Toshiba revealed its newest single-chip decoder for MPEG-2. The TC81211 includes a generic host bus interface, video port, audio DAC port, and DRAM controller. The chip is capable of decoding an MPEG-2 video stream at CCIR-601 resolution in real time with 2M of memory. Unlike the earlier TC81200, the 211 decodes a synchronized MPEG-1 audio stream.

The chip is housed in a 208-lead PQFP and operates from a 3.3-V supply. In sample quantities, the device is priced at $150, and is available now.

Erratum: Nx586-133 Performance

In our previous issue, we erroneously stated that NexGen's 586-133 would be slower than a Pentium-133 system, even if the Pentium had an asynchronous cache. In fact, Pentium-133 systems with asynchronous caches are much slower than those with synchronous caches, and NexGen's claim that its 586-133 is faster than a Pentium-133 with an asynchronous cache is correct.

Tests by *PC Magazine* of 40 Pentium-133 systems show that those with asynchronous caches range from 193 to 237 on CPUmark32; NexGen's 586-133 delivers a CPUmark32 score of 248, according to NexGen. Users seeking peak performance will do better with a Pentium-133 with synchronous cache, however; such systems rated as high as 300 in the *PC Magazine* testing, with most of them above 270. ♦

