Most Significant Bits

Pentium Prices Plunge Precipitously

Intel continues to send shock waves through the PC market by rapidly dropping the prices of its high-end processors. Effective January 30, the company slashed the price of its top of the line, the 100-MHz Pentium, by 28%, indicating the effect of improved yields on that part *(see **090101.PDF**)*. Yields on the 66-MHz version, the fastest 0.8-micron part, have also improved, allowing its price to fall close to that of the 60-MHz speed grade. As the following table shows, low-end Pentium prices have broken the $300 barrier. But Intel saved its biggest moves for its DX4 line, halving the price of 75- and 100-MHz 486s.

	10/31/94	1/30/95	% CHG
Pentium-100	$935	$673	–28%
Pentium-90	$587	$546	–7%
Pentium-75 (PGA)	$535	$301	–44%
Pentium-75 (TCP)	$495	$301	–39%
Pentium-66	$479	$289	–40%
Pentium-60	$383	$273	–29%
DX4-100	$459	$245	–47%
DX4-75	$382	$184	–52%

These prices—reminiscent of the old Crazy Eddie commercial ("so low, they're insane!")—aim to destroy the market for high-end 486 chips from AMD and Cyrix in two ways. First, they motivate system makers to move to low-end Pentiums, forsaking the 486. Those that insist on a 486 will pay reduced prices, preventing AMD and Cyrix from making exorbitant (i.e., Intel-type) profits.

Intel to Reveal Pentium Errata List

Vowing that it has learned from the Pentium FDIV debacle, Intel plans to make public the errata list (bug reports) for its Pentium chips as well as for future x86 processors (i.e., the P6). Although the company continues to work out the details of this plan, the basic concept is that bugs would be made public about 30 days after they are given to OEMs and ISVs under nondisclosure agreements. Intel has also committed to being more thorough in its errata process (the FDIV bug was not added to the bug list until after the public furor arose) and to identify and resolve bugs more quickly.

The company is developing a format that will be useful to the public. The current errata lists, intended for OEMs, contain short, highly technical descriptions of each problem, offering little insight as to how often and under what conditions an end user might encounter it. To be truly useful, some sort of elaboration and classification is needed. On the other hand, Intel would rather let third parties prioritize the bugs, as the company was criticized for classifying the FDIV bug as insignificant.

We applaud Intel's decision to fully disclose the state of its silicon and invite other vendors to do the same. In the meantime, we will offer an analysis of the Pentium errata list in our next issue.

Digital StrongArms Embedded Market

Killing two birds with one throw, Advanced RISC Machines (ARM) and Digital Semiconductor have joined hands to develop a high-end embedded CPU based on the popular ARM instruction set. The new processor, dubbed StrongArm, should solve ARM's lack of a high-end device while helping to fill Digital's fab.

The companies hope to combine Digital's expertise at high-performance processor design, honed on its Alpha line, and ARM's familiarity with compact, low-power devices to produce a low-power embedded processor with at least three times the performance of current ARM7 chips. StrongArm will be significantly faster than even the forthcoming ARM8, making it suitable for set-top boxes, PDAs, and other consumer devices that require maximum performance.

Digital becomes the eighth ARM licensee but the first to codevelop a CPU core. The partners will offer the StrongArm core to other ARM licensees, but Digital's advanced semiconductor processes could give it a lead of 6–12 months in manufacturing the part. If successful, StrongArm chips could help make up for lower-than-expected Alpha volumes, but the tiny part will require relatively few wafers to generate large volumes.

The partnership is a big win for ARM, which has already lost at least one deal due to lack of a high-end part. For Digital, the agreement does little for its flagship Alpha line other than undercut Alpha's embedded plans. The two must now deliver a product; more details about StrongArm are expected later this year.

PowerPC 603+ Announced As 603e

The widely rumored PowerPC 603+ *(see **0808MSB.PDF**)* has been announced as the PowerPC 603e. Both Motorola and IBM will produce the new chip, which is pin-compatible with the 603 but doubles the size of the on-chip caches (to 16K each for instruction and data) and achieves a clock speed of 100 MHz, 25% faster than the standard 603. At this speed, the 603e is estimated to deliver 120 SPECint92 (using the new SPEC rules) *(see **0816MSB.PDF**)* and 105 SPECfp92.

This speed increase is achieved without moving to a new IC process; the 603e remains in the 0.65-micron CMOS-5L process. The Somerset design team modified several critical timing paths to achieve the faster clock rate. Without a process shrink, the bigger caches increase the die size to 98 mm^2, 15% larger than the 603,

and bump the power dissipation to 3.5 W (maximum). The MPR Cost Model estimates that the 603e will cost about $50 to build, an increase of 10% over the 603's cost.

Neither company has released pricing or availability for the new chip. Somerset has received first silicon and expects volume production sometime in 2H95. Given the minimal cost increase, IBM and Motorola should be able to price the part aggressively.

Interestingly, the 603e achieves slightly better performance than the 601 at the same clock speed. Once the disparity in cache capacity was corrected, the 603's more aggressive microarchitecture allows it to surpass its predecessor's performance. In addition, the 603e is smaller than the 601 in the same manufacturing process and is even less expensive to build than the 0.5-micron 601+, according to our cost estimates.

Given these characteristics, Motorola can now promote the 603e for new designs rather than the 601, which it does not manufacture. IBM may also try to move the market from the 601 to the 603e, which improves the overall cost/performance position of PowerPC.

The performance and cost of the 603e undoubtedly could be improved further by moving it to a true 0.5-micron process (such as CMOS-5X). While IBM has demonstrated its ability to build 0.5-micron chips, Motorola has not and, other than the low-volume PowerPC 620, has no announced plans to do so. If Motorola cannot duplicate IBM's 0.5-micron processes, it could delay the entire PowerPC effort or force IBM to forge ahead alone.

Intel Launches Triton PCI Chip Set

Intel's new Triton chip set for PCI-based Pentium PCs promises to raise system performance to new heights through a combination of advanced memory support, bus-master IDE, and PCI data streaming. The chip set supports the 3.3-V Pentium family at up to 100 MHz.

The four-chip set connects the Pentium bus to the second-level cache, DRAM, and PCI and ISA buses. It supports standard, burst, or pipelined-burst SRAMs for the L2 cache (which can be either 0, 256K, or 512K), and page-mode or EDO (extended data out) DRAMs for main memory. The 64-bit memory interface supports up to 128M of DRAM using either 3.3-V or 5-V chips.

Unlike Intel's earlier high-end chip sets, the L2 cache is optional; Intel says that systems using the Triton chip set with EDO DRAM and no L2 cache can match the performance of today's best systems with an L2 cache. This cost-saving option is likely to be especially popular for 75-MHz Pentium systems, which Intel hopes to push quickly into the mainstream market.

The PCI-to-ISA bridge chip also provides a bus-master IDE controller, which moves data from the disk to the PCI bus at up to 22 Mbytes/s and reduces bus utilization during disk transfers from 20% (using programmed I/O) to 1%. With Windows 95, which supports multitasking, having the disk transfer run in the background while using little bus bandwidth is a significant benefit. This chip also provides a plug-and-play port for connecting an audio codec or other motherboard peripheral.

Deep FIFOs and "snoop-ahead" logic seek to prevent stalls, allowing data to move between a PCI device and DRAM at more than 100 Mbytes/s. Intel claims that this performance is twice that of existing PCI chip sets and is important for native signal processing (NSP) applications, which require efficient bus and system designs to move large data streams through the processor. Intel's assertion implicitly concedes that existing Pentium systems will not be able to deliver on the NSP promise.

The Triton chip set is priced at $41.95 in quantities of 10,000. Production is planned for this month. Intel is using its own fabs, as well as outside foundries, to ramp Triton production rapidly; most of its other chip sets have been built solely by foundries. With most processor production moving to 0.5-micron fabs, plenty of 0.8-micron capacity is becoming available for chip sets. Intel's aggressive chip-set stance could make life even tougher for other chip-set makers.

TI Takes Floating-Point DSP to New Low

Texas Instruments has introduced the TMS320C32 DSP chip, cutting the entry price of its popular 'C3x family of floating-point DSPs by more than 50%. The new chip lists for $25 in quantities of 5,000, and TI expects it to sell for less than $10 in high-volume applications (more than 250,000 units). The 'C32 further reduces system cost by allowing an 8- or 16-bit-wide memory system; other 'C3x parts require a memory width of 32 bits.

Despite the lower cost, the new design carries the same performance rating as the more expensive 'C31: 20 native MIPS at 40 MHz. The new chip also adds a second on-chip DMA controller. Its two banks of on-chip RAM, however, are only 256×32 bits each, one-fourth the size of the 'C31's on-chip memory. Like other family members, it is built in 0.72-micron CMOS and runs at 5 V. All 'C32 versions use a 144-pin PQFP. Volume production at speeds up to 60 MHz is set for 3Q95.

The company hopes that this low-end design will move floating-point DSPs into high-volume consumer products that currently rely on inexpensive fixed-point DSPs. Although the 'C32 takes a big step toward leveling the playing field, these applications are very price sensitive, and most will still find adequate performance from fixed-point devices at a lower cost. But the aggressive pricing of the new chip will make it attractive for a range of medium- and high-volume designs, broadening the horizons for TI's floating-point DSPs.

Cirrus to Ship Rambus Graphics Chips

Becoming the first standard-product vendor to back the

technology, Cirrus Logic will incorporate a Rambus interface into new graphics chips expected to debut later this year. The company did not reveal any details about these devices, but we expect the first Rambus chips to be high-end graphics accelerators that will begin appearing in systems by the fall.

By switching to a Rambus interface, Cirrus can achieve a bandwidth of 500 Mbytes/s using an 8-bit bus *(see **070304.PDF**)*. In contrast, VRAM or synchronous DRAM require a more expensive 64-bit interface to approach this performance level. Furthermore, the narrow bus reduces the minimum number of memory chips, allowing more flexibility when moving to 16-Mbit devices. As a result, both controller pin count and overall subsystem size can be reduced.

Several other companies have lined up behind Rambus, but all are DRAM or infrastructure vendors, including Samsung, the latest vendor to announce plans to ship Rambus DRAMs. Cirrus, one of the leading makers of graphics accelerators, gives Rambus a big credibility boost. Other wins in the graphics area are likely to follow. Rambus claims that eight additional companies have licensed its technology but are not ready to publicly reveal product plans. One of these is presumably Silicon Graphics, which, according to RDRAM-supplier Toshiba, is developing a Rambus-based graphics subsystem. ♦

