

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 7 NUMBER 11

AUGUST 23, 1993

Cyrix Readies 486DX-Compatible CPU

M7 Combines 486-Like Core with 8K Write-Back Cache, FPU, 486 Bus

By Michael Slater

Cyrix is preparing to announce this fall its first 486DX-compatible processor, code-named M7 and not yet officially christened, joining AMD in competing with Intel's mainstream product line. The M7 uses a core CPU similar to that in earlier Cyrix chips, but it has an 8K cache, an FPU, and a 486 pinout. It also offers a clock-doubling option (as in the 486DX2), but the maximum internal clock rate initially will be 50 MHz. A 66-MHz version is planned for early 1994.

Cyrix's CPU efforts began with the 386SX-pin-compatible Cx486SLC, and the company has extended the line with the 486DLC (386DX-pin-compatible) and 486S (486SX-pin-compatible) versions (*see 060501.PDF, 060701.PDF, and 070704.PDF*). The first two parts have a 1K write-through cache; the 486S doubles the cache to 2K and adds a write-back mode. The new M7 brings Cyrix's cache size up to the Intel-standard 8K, while also bringing on-chip Cyrix's enhanced 387-compatible FPU.

The M7 includes an optional clock doubler, but Cyrix plans to promote it primarily for 50-MHz, non-clock-doubled systems using a VESA local bus. This configuration brings it closest in performance to Intel's popular 486DX2-66, which has a faster CPU speed but a slower (33 MHz) local bus. The M7's 50-MHz local bus could improve graphics performance somewhat, assuming a display controller that can keep up with that rate, but it remains to be seen how significant this is. Other than the higher-speed local bus, performance is similar to an Intel 486DX2-50; M7 is 5–10% slower on integer programs, and about 10% faster on floating-point.

Near-486 Core

The CPU in the M7 is similar to that in earlier Cyrix chips, but it has been enhanced with a more efficient coprocessor interface, additional write buffers, improved prefetching, and faster non-aligned accesses. It implements the full 486 integer instruction set but has a dif-

ferent microarchitecture that lacks a dedicated address adder. As a result, although Cyrix's core matches the performance of Intel's on most register operations, it is one cycle slower on instructions that involve an address computation, including all instructions with memory-based operands, all jumps, and all calls.

In the M7, the slower speed of the core is partially offset by the write-back cache, which Intel's chips lack. To provide fast cache line flushes, the M7 extends Intel's 486 bus protocol with burst writes. The M7's overall performance also benefits from Cyrix's faster FPU design and eight-level write buffer—twice as deep as Intel's.

Cyrix's FPU is much larger than Intel's, however, which is partially responsible for the portly 148 mm² (228K mils²) die size of the M7—78% larger than Intel's 0.8-micron 486DX. Other factors are that the 0.8-micron CMOS process has only two layers of metal, compared to Intel's three, and that Cyrix's layout is less dense. The chips are currently being fabricated for Cyrix by SGS-Thomson, and TI may serve as an alternate foundry.

Cyrix has not announced any pricing information, but based on the company's past strategy, the M7 probably will be priced well below Intel's 486DX. Thanks to Cyrix's relatively low-overhead operation and the fat margins of the 486 business, this should be possible despite the large die size.

The chip's static design allows the clock to be slowed or stopped to reduce power. A unique feature of the clock doubler is that it does not use a PLL and thereby allows the clock speed to be changed rapidly. Clock doubling is pin-selectable, but Cyrix may choose to market "DX" and "DX2" versions separately.

Performance Lags on Integer, Leads on FP

Performance (based on PowerMeter MIPS and Whetstones) is claimed to be about 9% slower than Intel's 486DX (at the same clock rate) on integer code and 10% faster on floating-point software. Because these benchmarks fit in the on-chip cache, they yield the same results for a clock-doubled 25/50 system or a full 50-MHz

design. They are also independent of the second-level cache and memory system.

According to Cyrix, the BAPCo benchmark, which better reflects application-level performance, shows that the M7 provides performance equal to Intel's 486DX2-50 in a cacheless system design, where the on-chip write-back cache is especially valuable. In a system with a 256K second-level cache, the M7's BAPCo performance falls short of Intel's by 4% in clock-doubled mode and by 7% for full 50-MHz operation.

In a cacheless system, the write-back cache is particularly valuable. Cyrix's measurements show that while Intel's 486 is faster in a system with a second-level cache, the M7 matches the 486 in a cacheless system with fast (3-2-2-2 access pattern) DRAM and outperforms the Intel chip by 5-10% with slower DRAM.

Power Management

Typical current drain is 680 mA (1050 mA max) at 50 MHz (with a 5V supply). A 3.3V version will also be available, but only at 25 and 33 MHz; typical current at 33 MHz and 3.3V is 320 mA.

The wide spread between typical and maximum values is due, in part, to the fact that the FPU is powered-down when no FP instructions are being executed. The maximum rating is measured while repeatedly executing the FCOS instruction, while the typical value is measured while running Whetstone.

Like most of Cyrix's previous microprocessors, the M7 includes a system-management mode (SMM) and a suspend mode. SMM is entered by asserting the system-management interrupt (SMI#) or executing the SMINT instruction. Suspend mode is entered by asserting the SUSP# input or executing the HALT instruction.

When in suspend mode, the processor asserts SUSPA#, telling external logic that the clock can be stopped. Typical power consumption in suspend mode (with the clock running) is 75 mW at 5V and 50 MHz, or 27 mW at 3.3V and 33 MHz. Stopping the clock drops power dissipation below 2.5 mW.

Write-Back Cache and Bus Design

Cyrix currently makes the only 486 microprocessors

Function	Signal Name	Description
Configuration	CLKMODE	Clock-doubler enable
	WM_RESET	Warm reset
	UP#	Float
System Management Mode	SMADS#	SMM address strobe
	SMI#	System management interrupt
Cache Control	HITM#	Hit on modified data
	INVAL	Invalidate
	RPLSET1, RPLSET0	Replacement set
	RPLVAL#	RPL valid strobe
Suspend Mode	SUSP#	Request suspend mode
	SUSPA#	Suspend mode acknowledge

Table 1. New signals implemented by Cyrix's M7.

Price & Availability

Sampling of the M7 has already begun, and production is scheduled for the fall. Price and availability details will be revealed at the formal product announcement later this year.

Cyrix, 2703 N. Central Expwy., Richardson, TX 75080; 800/848-2979 or 214/234-8388; fax 214/699-9857.

with write-back caches, although both Intel and AMD are expected to introduce new 486s with this feature in the coming months.

The 8K cache is four-way set-associative and uses a least-recently used (LRU) replacement algorithm. The line size is 16 bytes; there is one valid bit per line and one dirty bit per 4 bytes. When a dirty cache line must be re-allocated, it is not necessary to write the entire line back to memory; only the modified 32-bit words must be written. Cache lines are not allocated on writes.

Table 1 lists the additions to Intel's 486 signals. Because the write-back cache can contain data that is more recent than the version in main memory, the cache must be checked when an external bus master (such as a DMA controller or another processor) accesses memory. This is implemented with cache inquiry ("snooping") cycles, which allow an external bus master to poll the CPU's on-chip cache. In Intel's 486, cache inquiry cycles affect only the setting of the valid bit; since the 486 has a write-through cache, there is never any dirty data. The M7, on the other hand, must supply data from the cache when a read snoop hit occurs on dirty data.

Two new signals implement the basic control for the write-back cache. HITM# is asserted by the processor when it snoops external bus activity and a hit occurs on dirty cache data (i.e., when a cache inquiry cycle is in progress and the cache holds modified data for that address). The INVAL input is sampled during cache inquiry cycles to let external logic control whether or not the cache line's valid bit is cleared.

When a cache hit occurs on dirty data during a cache inquiry cycle, the M7 implements an "abort and retry" protocol by asserting HITM# and writing the dirty cache line (or the dirty words within the line) to memory. The other bus master then reads the data from memory. This protocol is very similar to that implemented by Pentium.

To allow the M7 to be used in systems that don't support snooping, a control bit (BARB) enables a mode in which all dirty cache contents are written to memory whenever HOLD is asserted (such as when a DMA controller requests the bus) and before the pro-




Figure 1. Die photo of Cyrix's M7, which measures 12.2 × 12.1 mm (480 × 476 mils) and incorporates 900,000 transistors.

cessor asserts HLDA. External logic also can force all dirty data in the cache to be written to memory by asserting the FLUSH# input. This would be used prior to entering suspend mode, for example, since no snooping occurs in that mode.

Burst writes, if enabled via a control register, are performed whenever a cache line must be replaced or flushed and all four 32-bit words of the line are dirty. Burst writes use the same BRDY# control signal to pace the transfer as Intel-standard burst reads. Assuming a three-clock initial transfer and single-cycle transfers within the burst (i.e., 3-1-1-1), burst mode cuts the line write time in half.

The RPL signals identify which one of the four cache lines in a set is replaced. Intel's 486 lacks these signals, and as a result, it is impossible for a second-level cache to precisely track the contents of the on-chip cache.

The WM_RESET signal resets the processor but does not flush the cache, and it leaves configuration registers intact. This feature is provided for compatibility with older software that resets the processor to switch from protected to real mode.

System-logic chip sets require a small amount of additional logic to handle the write-back cache signals and burst writes. Cyrix says that OPTi, SiS, UMC, and several other chip-set vendors will have chip sets designed to work with the M7 (supporting the 50-MHz bus and write-back cache) in the fourth quarter. Chip-set makers have been developing similar logic for Pentium and P24T chip sets, as well as for anticipated 486s with write-back caches from Intel and AMD.

Facilitating 50-MHz Design

By promoting the 50-MHz, non-clock-doubled version of its CPU, Cyrix is attempting to reinvigorate a speed that Intel has de-emphasized. Since the emergence of the 486DX2-66, the DX-50 has faded in popularity because of the DX2's simpler system design (due to the 33-MHz bus) and higher core CPU performance. Cyrix hopes to make 50-MHz systems more appealing by promoting the 50-MHz VL-Bus (see sidebar below).

To provide the needed 50-MHz components, Cyrix is working with numerous suppliers, including Tseng, S3, Cirrus, Weitek, IIT, and ATI for graphics controllers; Adaptec and Promise for hard-disk controllers; and Cypress and Micron for cache SRAMs.

Cyrix says that its chip is easier to design with than Intel's 50-MHz 486 because of the way the timing is specified. Intel's timings are specified with no capacitive loading, and delays must be derated for the appropriate loading. Cyrix's specifications, on the other hand, include a 50-pF load.

PLL-Less Clock Doubler

Unlike Intel's 486DX2, the M7 does not use a phase-locked-loop (PLL) for clock doubling. Instead, it uses a delay circuit that generates a series of pulses after each clock edge. This has the advantage of having no limit on how fast the frequency can be changed. In a chip with a PLL clock generator, such as Intel's 486DX2, the frequency cannot be changed rapidly because the PLL will not stay locked onto the changing frequency. This limits the degree to which power-management circuitry can dynamically slow the clock to save power.

Cyrix's clock-doubler circuit is triggered only on the rising edge of the clock. In response to this edge, it generates a series of four pulses, with the time between pulses set by an on-chip delay line. Each pulse toggles a flip-flop, which creates the frequency-doubled output. The delay time between pulses is set so that at the maximum clock frequency, the fourth pulse arrives one delay time before the next rising edge on the clock input. As the clock input is slowed, the spacing of the four pulses remains constant, so only the last half-cycle of every alternate clock cycle is stretched. This stretching does not bother the logic, however, and this circuit allows the clock frequency to be changed dynamically without restriction.

Seeking Niches

Unlike AMD's frontal assault on Intel, Cyrix is focusing on niches in which it can provide some benefit that Intel's chips do not offer. Cyrix has relatively limited access to production capacity and has decided not to pitch the M7 as an alternate source for Intel's 486DX-33 or AMD's 486DX-40, but to focus on the DX-50 version.

VL-Bus 2.0 Adds Write-Back Caches and 50-MHz Operation

The Video Electronics Standards Association (VESA) is finalizing Version 2.0 of the VL-Bus specification. The current proposal contains several extensions to the original VL-Bus (see *060902.PDF*), including 64-bit transactions, 50-MHz timing, and write-back cache support. The proposal currently is being reviewed by all 183 voting members; VESA expects no major changes before the final Version 2.0 is approved.

The 64-bit extensions give VL-Bus the same growth potential as PCI. The PCI specification defines a 64-bit version, but this width is not yet being implemented. The 64-bit version of VL uses the existing connector by multiplexing the extra data bits onto the address bus. Two new signals indicate when a 64-bit transfer is in progress; systems that support 64-bit transactions must also support 32-bit and 16-bit data as well, so existing 32-bit cards should work in future 64-bit systems. The 64-bit extensions are an optional feature in VL-Bus 2.0.

The original specification allowed 50-MHz operation on the motherboard but not with VL-Bus add-in cards. The new revision allows one or two add-in boards in a 50-MHz system, depending on the motherboard loading. The clock skew for 50-MHz systems is cut from 2 ns to 1 ns to further ease high-speed designs. Several vendors now have prototype systems running at 50 MHz using the new specifications.

VL-Bus 2.0 defines the *WBACK#* signal for write-back caches. All new VL bus masters must support this signal, which is used for snooping. A processor with a write-back cache must snoop all transactions originating on the VL-Bus to see if they access data that has been modified in the on-chip cache. The *WBACK#* signal is asserted when there is

such a snoop hit, requiring the VL bus master to abort its transaction and wait for the CPU to write the modified data to main memory. This protocol allows VL-Bus to be used with processors such as Pentium and Cyrix 486 chips that include write-back caches.

The new specification clarifies the implementation of burst transactions and makes them more like the 486 burst mode. Other minor changes and clarifications are also included.

A 50-MHz, 64-bit VL-Bus has a peak bandwidth of 400 Mbytes/s, compared to 266 Mbytes/s for a 33-MHz, 64-bit PCI bus, the fastest version of that bus. In practice, however, average bandwidth is considerably lower, and VL's technical lead in peak bandwidth may or may not result in any genuine benefit.

The VL-Bus is best-suited for 486 systems, which will stick with a 32-bit local bus. It is easy to synthesize VL from the 486 processor bus, allowing for low-cost interfaces. Support for 50-MHz buses and write-back caches will extend the lifetime of VL through the next generation of 486 chips from Intel and others.

For Pentium and other non-486 processors, the cost advantage of VL is much smaller because both VL and PCI require complex bridges. Furthermore, a 32-bit PCI bus can be implemented in just 45 pins, about half of the number needed for a 32-bit VL-Bus. These factors led to the adoption of PCI by Digital for its Alpha processors and by Apple for future PowerPC-based Macintoshes. This wider market adoption, and the eventual eclipse of the 486 by Pentium, give PCI a long-term edge, but VESA's bus will continue to be popular as long as the 486 is in use.

—LG

Cyrix is the first vendor to put a write-back cache in a 486 pinout. The need to support the write-back cache to gain the chip's full performance, however, means that system makers cannot simply substitute the Cyrix chip in an existing system design without losing performance; this may slow its adoption.

By focusing its efforts on 50-MHz systems with the 50-MHz VL-Bus, Cyrix hopes to add another price/performance point—one that Intel has largely abandoned—to the spectrum. Such systems will fall between 486DX2-50 and 486DX2-66 systems in perfor-

mance. Presumably, what will attract buyers is that the systems typically should cost less than 486DX-33 and DX2-50 systems while providing higher performance, and they will offer the additional allure of the fast VL-Bus.

It remains to be seen if the marketplace will accept this alternative. Larger system makers are unlikely to be interested in adding another configuration to their line, and many are focusing on PCI, rather than VL-Bus, for next-generation products. For smaller system makers, however, the M7 may be attractive. ♦