

# MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 8 NUMBER 16

DECEMBER 5, 1994

## Cyrix M1 Design Tapes Out

### Pentium Competitor Likely to Have Little Impact Until 1996

by Linley Gwennap

A year after first revealing the microarchitecture of its M1 chip (*see 071401.PDF*), Cyrix has taped out the design and is awaiting first silicon of its Pentium-class processor. Having accomplished this task, the company has, for the first time, completely disclosed the physical design of the part. Cyrix did not, however, go so far as to announce pricing.

The M1 is a superscalar x86 processor that supports advanced features, such as register renaming and out-of-order execution, that Intel's Pentium does not. At the same frequency, the M1 should deliver better performance than Pentium, allowing Cyrix to compete with Intel all the way to the top of its product line.

Cyrix reaffirmed its intent to ship 100-MHz M1 processors; these chips should match or exceed the performance of Intel's forthcoming 120-MHz Pentium on binaries that have not been recompiled for Pentium. On SPECint92, a recompiled benchmark, the M1 should be slightly faster than a Pentium of the same clock rate.

The M1 was originally promised for the end of this year, and even recent press announcements claimed a 1Q95 shipment date. Cyrix officials now say the first M1 systems will ship in volume in 2Q95. Even that schedule will require an abbreviated debugging cycle, which the company believes it can accomplish even though it has performed no hardware emulation (e.g., Quickturn), a technique used for most processors of this generation.

Our rule of thumb (*see 0815ED.PDF*) requires 12 months between tape out and first system shipments. By this rule, we expect the first M1 systems to roll out late in 3Q95, putting the M1 on about the same schedule as AMD's K5 and Intel's P6 (*see 0816MSB.PDF*).

#### A Die Size to Die For

We expect few M1 systems to ship in 1995 due to a second glitch revealed in Cyrix's announcement: the die size of the first M1 chips is an astounding 394 mm<sup>2</sup>. This die is easily the largest commercial microprocessor yet

fabricated. If it were a garage, it could fit two cars and a motorcycle, as its area is enough for two P54C die with nearly enough space for a DX4 as well.

This Texas-size die severely diminishes the number of M1 chips that Cyrix's fab partners will be able to produce: our model predicts that the M1 will yield about eight good die per 200-mm wafer. Furthermore, economics will push Cyrix to devote most of its wafer starts to 486 chips. A 200-mm wafer of Cyrix 486DX2 chips yields about 210 good die, or roughly \$21,000 at projected 3Q95 list prices. Assuming that the M1 sells for about \$700, an equivalent wafer would produce just \$5,600 in revenue.

Even if Cyrix can convince its fab partners to triple its current number of wafer starts, the company would be hard-pressed to get more than 10,000 M1 chips per month in 2H95. Intel, in that period, will produce close to 2,000,000 Pentium chips per month, leaving the M1 with less than 1% of the Pentium-class market.

The first M1 design is not optimized for IBM's manufacturing process, as it is intended to be produced by both IBM and SGS-Thomson, Cyrix's other foundry. In IBM terminology, the M1 is built in CMOS-4LC, which combines 0.65-micron transistors with the metal geometries of IBM's 0.8-micron process. The M1 uses only three layers of metal, whereas IBM can support up to five.

#### Quick Process Shrink Will Help

Fortunately, help is on the way. Cyrix and IBM plan to quickly redesign the M1 to use IBM's more aggressive CMOS-5S process (*see 080504.PDF*), the same process as the PowerPC 620. Based on the improved metal geometries and use of five metal layers, we project that the CMOS-5S version of the M1 will be about 200 mm<sup>2</sup>, just over half the size of the initial version. This change will quadruple the number of good chips per wafer and allow Cyrix to achieve a share of the Pentium market similar to its 2–3% share of the 486 market.

The large die size of the M1 is not a total surprise. Cyrix's 486DX2, for example, is 80% larger than Intel's DX2 when both are manufactured in a 0.8-micron pro-

cess. The initial M1, 34% bigger than the original P5 Pentium, comes closer to Intel's mark. Based on our projection, the CMOS-5S M1 will be about 26% larger than Intel's 0.5-micron P54C, although the Cyrix design has the same amount of cache and slightly fewer transistors (3.0 million, versus 3.3 million for the P54C). Intel has a small army of circuit designers devoted to producing compact designs; Cyrix, a \$250-million company, cannot afford this luxury.

The MPR Cost Model (*see 081203.PDF*) projects the manufacturing cost of the initial M1 to be around \$340, nearly three times the estimated cost of the P54C Pentium. After the shrink, the cost of building the M1 should be cut in half, but it will still be about 40% greater than the P54C's.

In addition to cutting costs, CMOS-5S will also boost the M1's clock speed to 120 or possibly 133 MHz. By the time this version debuts, however, Intel will have a 150-MHz Pentium, again pacing the M1. In addition, Intel will also be shipping the first P6 processors, keeping it firmly in the performance lead.

### More Advanced Microarchitecture

The latest revelations emphasize the M1's architectural advantages over Pentium. M1 features such as register renaming, memory bypassing, speculative execution, and out-of-order execution have already been discussed (*see 071401.PDF*); none of these features are found in Pentium. The M1 design also has some more subtle advantages.

Cyrix revealed that the initial M1 parts will use a 16K unified on-chip cache, giving it the same total cache capacity as current Pentium devices. The unified design has a better hit rate, in general, than Pentium's split caches. To avoid the bottleneck of a single cache, the M1 includes a fully associative 256-byte instruction buffer.

Both the M1 and Pentium use a banked design to perform two cache accesses per cycle so long as the accesses use different banks. Pentium's data cache has eight banks, interleaving on 32-bit boundaries. The M1 cache has 16 banks, each 16 bits wide, reducing the number of bank conflicts slightly. This effect will be greatest for code that accesses 8- or 16-bit data.

The M1's branch-prediction accuracy should exceed Pentium's. Both use a 256-entry branch history table with two bits per entry, but the M1 supplements this with an eight-entry return address stack. This stack stores an address when a subroutine is called, allowing the return address to be "predicted" by taking it from the stack rather than waiting to retrieve it from memory. This stack will reduce the number of branches that stall.

The M1's unified TLB contains 128 entries, more than Pentium's 96 total TLB entries. To avoid thrashing, the M1 includes an eight-entry victim TLB that holds entries flushed from the main TLB. Like Cyrix's 486

## Price & Availability

Cyrix has not yet announced pricing for the M1. It expects volume production to begin in late 2Q95. For more information, contact Cyrix (Richardson, Texas) at 214.994.8491; fax 214.994.8404.

processors (*see 060501.PDF*), the M1 has configuration registers for the location of SMM space and up to four noncachable regions; Pentium lacks these features.

### Cyrix Avoids Appendix H

From a software standpoint, Cyrix claims that the M1 is fully compatible with Pentium. The only exceptions are the secret "Appendix H" extensions, which Intel has never publicly revealed. Unlike AMD, which has attempted to reverse-engineer these extensions and implement them in its K5 design, Cyrix has chosen to avoid this potential legal minefield. This incompatibility should have little effect, as these extensions are applicable only to operating-system code, and no operating systems are known to be using them—in particular, Microsoft asserts that none of its software takes advantage of Appendix H.

One of the secret extensions allows a single TLB entry to map a large virtual space, a feature that is also included in many RISC processors. Cyrix chose to implement support for 4M pages in the M1 but does it in a way that (presumably) is different from Pentium.

The M1 does support the CPUID instruction introduced in Pentium, which returns the string "Genuine-Intel" when this instruction is executed. The M1 returns the string "CyrixInstead," which happens to contain the same number of characters.

### Plug-Compatible with Pentium

Although Cyrix is rumored to be developing a version of M1 (code-named M9) in a 486DX4 pinout, the initial M1 product will use Pentium's 293-pin PGA package and pinout, allowing system vendors to easily support both processors with a single motherboard. There are some key differences, however, that must be addressed.

Intel has patented the nonlinear address sequence used for certain burst transactions on the Pentium bus. For system-logic chip sets that accept only this order, the M1 will operate in a "noninfringing" compatibility mode. As Table 1 shows, for accesses to the second or fourth words in a cache line, the M1 uses a single access to obtain the needed word first, then issues a burst transaction, starting with the first word, to refill the internal cache. This mode increases bus utilization by about 10% but should degrade CPU performance only slightly.

The M1 can also perform burst transactions using a

simple linear addressing scheme, also shown in Table 1. Some Pentium system-logic chips, such as Opti's Viper chip set (see *0816MSB.PDF*), support the linear addressing mode in addition to the Intel burst mode; these chip sets allow the M1 to achieve slightly higher performance. As the cost of supporting this second mode is fairly minor, Cyrix expects that most forthcoming Pentium chip sets will also support the linear burst mode.

Cyrix's chip does not include the advanced priority interrupt controller (APIC) that is part of Intel's P54C Pentium. The APIC (see *080302.PDF*) is intended for use in multiprocessor systems but is protected by Intel system-level patents. For MP systems, Cyrix instead is promoting its SLIC design (see *0808ED.PDF*), which requires no special logic in the processor. Only Via's Apollo chip set currently supports SLIC.

Although the P54C operates its system bus at two-thirds the speed of the CPU, the M1 uses a half-speed system bus. A 100-MHz M1, for example, has a 50-MHz bus. This choice simplifies the synchronization of the CPU with the system, but it decreases the available bandwidth by 25%. This decrease will mainly affect applications that frequently miss the on-chip cache.

Like the P54C, the M1 operates from a 3.3-V supply. Unlike the P54C, the M1 is compatible with 5-V chip sets and cache-memory chips. This feature will allow the M1 to be used with components that are more widely available and, in some cases, less expensive. Cyrix expects the M1, at 3.3 V and 100 MHz, to consume a maximum of 10 W, slightly more than the 100-MHz P54C.

### Competing Against K5

Cyrix likes to position the M1 against Pentium, but its real competitor is AMD's K5 (see *081401.PDF*). Once a system vendor decides to look at Intel alternatives, the M1 and K5 are the obvious contenders, with NexGen's 586 as yet another possibility.

The K5 uses a more radical microarchitecture than the Cyrix chip, breaking down x86 instructions into simpler "RISC operations" that can be executed speculatively and out of order. The M1 supports speculative execution but only very limited out-of-order execution. Both chips implement register renaming to avoid the

First Address	Pentium Burst Order	Cyrix Noninfringing	Cyrix Linear Mode
00	00,08,10,18	00,08,10,18	00,08,10,18
08	08,00,18,10	08—00,08,10,18	08,10,18,00
10	10,18,00,08	10,18,00,08	10,18,00,08
18	18,10,08,00	18—00,08,10,18	18,00,08,10

Table 1. The M1 system bus supports two burst orders, neither of which appears to infringe upon Intel's patented Pentium order.

register bottleneck in the Pentium design. The K5 can execute a maximum of four x86 instructions per cycle but is more likely to achieve two or three; the M1 is limited to two x86 instructions per cycle. The K5 has a 16K instruction cache plus an 8K data cache, which should deliver a higher hit rate than Cyrix's single 16K cache.

Based on projections by the two vendors, the K5 and M1 will have similar performance, although AMD's impressive feature set may give it an edge. A critical factor will be whether the companies actually deliver 100-MHz devices at first release.

If these parts deliver similar performance, the battle will come down to price and the ability to supply parts. The K5 will debut in a true half-micron process and is likely to have a significant manufacturing cost advantage over M1. Cyrix will have to accept lower margins to match AMD on price, but Cyrix's business model is designed for lower margins. Once AMD's Fab 25 goes into production in mid-1995, AMD will have an advantage in accounts that require a large supply of parts, particularly before Cyrix is able to deploy its shrink version.

The initial version of M1 is technically impressive but won't sell many chips, even by Cyrix's standards. The shrink version looks stronger. It will carry a higher manufacturing cost than Pentium, but with no fab and a small design team, Cyrix can tolerate lower margins, particularly compared with Intel's towering profits.

The M1 should give Cyrix a small slice of the Pentium market by 1996, just in time for Cyrix to avoid getting trapped in the collapse of the 486 market. Although the M1 will not have much impact on the overall Pentium market, this market's sheer size will allow Cyrix to do quite nicely with just a small share. ♦