MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 9 NUMBER 9 JULY 10, 1995

Cyrix 5x86 Targets Low-End Pentium

"M1sc" Extends Life of 486 Pinout with Improved Performance

by Linley Gwennap

With some sleight of hand, Cyrix has pulled a new chip out of a hat. After mesmerizing the audience with repeated glimpses of the M1, the company instead made a derivative design, code-named the M1sc, appear in the market before the M1 itself. The new chip, now called the 5x86, retains many of the performance enhancements of the M1 but contains only a single pipeline, eliminating the superscalar capabilities of the parent. This change lessens performance but greatly reduces the die size, and thus the cost, of the 5x86.

The new chip outperforms the best 486s but cannot match Pentiums of equivalent clock speed. But at 100 MHz, the first 5x86 chips compare well in performance with entry-level Pentium processors of up to 75 MHz. Cyrix is aggressively pricing this device at $147, 40% less than the least expensive Pentium.

The first 5x86 chips use a 486 pinout, providing compatibility with the wide variety of 486 chip sets, which are less expensive than their Pentium counterparts. This pinout makes the chip an easy upgrade for existing 486 motherboards. Notebook vendors are particularly attracted to the 5x86, as it offers a lower price and lower heat dissipation than Intel's notebook Pentium processors.

The 5x86, expected to achieve volume production in Q3, will render moot Cyrix's problems in boosting the clock speed of its 486 line and provide a bridge to the higher-performance M1, due to appear in systems this fall. Even after the M1 appears, the 5x86 will provide a lower-cost alternative for some time to come. Cyrix plans to deploy multiple versions of the 5x86 at higher clock speeds and with both 32-bit (486) and 64-bit (Pentium) system buses.

Picking the Right Features

Cyrix set up the 5x86 and M1 teams in parallel, tasking the latter with reaching the maximum possible performance while aiming the 5x86 team at an intermediate price point. The M1 *(see **071401.PDF**)* includes superscalar dispatch, an extended pipeline, register renaming, branch prediction, speculative execution, out-of-order execution, memory bypassing, and a 16K unified dual-ported cache. The 5x86 team chose some of these features but rejected others for cost reasons.

The biggest difference is the 5x86's single pipeline. Removing the second pipeline reduces the die size considerably. By eliminating superscalar execution, this decision also greatly simplifies the control logic throughout the chip and reduces the number of ports in the register file. With only a single pipeline, instructions are always executed in order, further simplifying control paths. Halving the peak issue rate also reduces pressure on the instruction fetch unit, which helped Cyrix reduce the instruction buffer from 256 entries to 48.

Speculative execution places a major burden on the M1 design. Before speculatively executing instructions, the M1 must checkpoint important processor state in a set of shadow registers. The M1 supports four levels of checkpointing, quadrupling the storage required for saving processor state. In addition, the eight general integer registers are shadowed within a 32-entry register file using register renaming. The 5x86 eliminates speculative execution, checkpointing, and register renaming to simplify the design, although it can still speculatively fetch instructions.

What remains is a scalar x86 CPU with a six-stage pipeline, branch prediction, and memory bypassing. The key feature of the pipeline, shown in Figure 1, is the AC2 stage for cache accesses, separating this operation from the EX (execute) stage. Arithmetic instructions that reference memory flow smoothly through this pipeline without any delays. In the 486 and Pentium, which combine the cache-access and execute operations, instructions that reference memory stall for at least one cycle.

The problem with extending the pipeline is that mispredicted branches have a four-cycle penalty, versus three for Intel's 486. The 5x86 compensates for this problem by adding a 128-entry branch target buffer (BTB)

similar to Pentium's. It also includes an 8-entry return-address stack. The 486 has neither of these features. The 5x86 should achieve a branch-prediction accuracy of about 85% on SPECint92; the 486, which simply assumes that all branches are not taken, reaches only 40% *(see **090405.PDF**)*.

The 5x86's pipeline combines the two decode stages that the M1 uses, since it decodes only one instruction at a time. Other pipeline stages are also simpler, due to the lack of superscalar overhead and register renaming. Although the initial 5x86 runs at the same speed as the first M1, we expect that the 5x86 will ultimately be able to run faster than the M1 in any given process technology—another benefit of a simpler design.

The 5x86's memory bypassing feature allows the CPU to execute instruction sequences such as:

```
ADD [mem], CX; SUB DX, [mem]
```

in successive cycles without any delays. This sequence adds CX to the contents of [mem], then subtracts the result from DX. Most processors, including Pentium, wait for the first instruction to write to the cache, then fetch the operand for the second instruction. On Pentium, this sequence takes four cycles. The 5x86 detects that the memory address for both instructions is the same and feeds the result directly to the next instruction, bypassing the cache and completing in two cycles.

As Figure 2 shows, the 5x86 has an independent floating-point unit. Once an instruction has been issued to the FPU, integer instructions can be processed by the main pipeline while the FPU completes the calculation. The pipeline will stall if a subsequent instruction requires the result from the FPU.

The 5x86 retains the four-entry FP instruction queue from the M1 design. Normally, this queue is disabled, and the pipeline stalls if a second FP instruction is issued before the previous FP operation completes. The queue can be enabled to improve performance, but the 5x86 lacks the M1's ability to checkpoint machine state, preventing it from generating precise FP exceptions, required for full x86 compatibility, in this mode.

The Cyrix chip can queue up to four stores to the cache. This feature allows the CPU to continue processing instructions while these stores wait for access to the cache. Memory bypassing allows pending store data to be used in subsequent calculations without delay.

486-Style Cache and System Bus

The 5x86 retains the 16K unified cache structure of the M1. The M1 uses a multibank cache that allows up to two loads or stores per cycle. The 5x86 can generate only one load or store per cycle, so it uses a single-ported cache. With just a single port, however, the 5x86's cache is blocked from instruction accesses on any cycle in which it is supplying data.

Just like the 486, the 5x86 includes a 48-byte instruction prefetch buffer along with a prefetch engine, helping to alleviate cache conflicts. The prefetch engine attempts to keep the buffer full of instructions, fetching along the predicted path. It can obtain a full cache line (16 bytes) in a single cycle, reducing the number of instruction reads from the cache. Data accesses have priority, so in most cases instruction fetches have no impact on cache performance.

The 5x86, like all other Cyrix processors, uses a write-back cache. Most of Intel's 486 chips use the inferior write-through model, although the company has added a write-back mode to its 486DX2 chips and uses this algorithm in Pentium as well. AMD also added a write-back cache to its latest 486DX4 chips.

The initial 5x86 parts use a 486 system bus and pinout. This pinout allows the part to take advantage of low-cost 486 chip sets and motherboards, many of which handle the write-back signals used by the Cyrix processors, which have been adopted by Intel and AMD. The 5x86 supports both clock doubling and clock tripling;

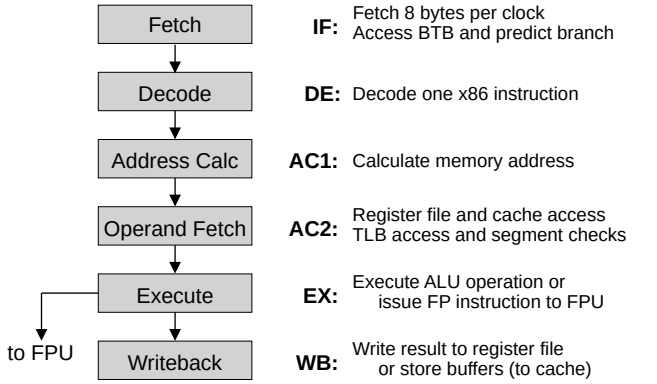

Figure 1. The 5x86's extended pipeline puts the cache access in AC2. Data from the cache is available in time for a calculation in EX, eliminating stalls for instructions that hit in the cache.

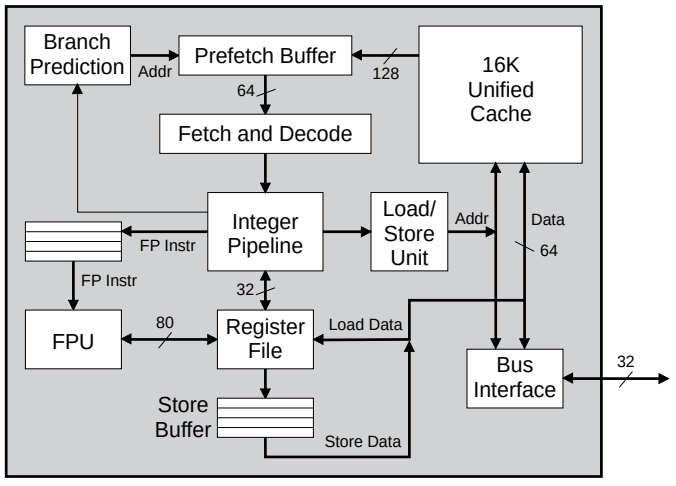

Figure 2. The 5x86's block diagram is similar to the M1's except that a single execution pipeline replaces the M1's dual pipeline. The 486 also uses a single pipeline, unified cache, and prefetch buffer.

Cyrix's performance estimates assume a 33-MHz bus. The company plans to deliver in 1H96 a version of the 5x86 with a Pentium bus, which Cyrix estimates will improve performance by 5–10%.

At 100 MHz, the 3.3-V 5x86 dissipates a maximum of 4.3 W, with a typical dissipation of about 3 W. This puts it about on par with a 100-MHz DX4. Intel's new VRT notebook Pentiums *(see **090702.PDF**)* also dissipate 3 W or less (typical) but carry a maximum rating of up to 6.5 W. Thus, portable systems using a 5x86 will have about the same battery life as either DX4 or Pentium systems. Pentium notebooks, however, need better thermal designs in case the CPU reaches its maximum power dissipation.

Seeking Performance with Efficiency

To meet the 5x86's aggressive price point, Cyrix had to give it a much lower manufacturing cost than the battleship-sized M1. The initial M1 measures 394 mm^2 in a 0.65-micron process and carries an estimated cost of $260 *(see **081601.PDF**)*. Amazingly, the 5x86, in the same IC process, measures just 144 mm^2 and has an estimated cost of only $50. As Cyrix shrinks both parts using more advanced processes, the 5x86 will retain a cost advantage, although the gap will shrink.

As Table 1 shows, the 5x86 contains 1.95 million transistors, less than two-thirds as many as the M1. More important, the number of logic (noncache) transistors is about half that of the M1. Because logic transistors consume most of the die area, this halving of logic transistor count, combined with better die layout and a smaller number of pads, resulted in the smaller die size. Figure 3 shows the 5x86 die.

The biggest reduction in transistor count is in the instruction decoder. The 5x86's single decoder consumes 40,000 transistors, but a two-way superscalar processor requires extensive logic to check for dependencies and issue constraints, pushing the M1 decode logic to 210,000 transistors. Similarly, most of the M1's function areas require more than twice the transistor count of the 5x86's; along with doubling the number of pipelines, each area of the M1 must contain additional coordination logic. The branch unit was reduced primarily by halving the size of the BTB.

Some areas had less change. The 5x86 uses the same FPU as the M1. Both chips use 16K of cache, but the 5x86 eliminates overhead due the M1's dual-ported multibank structure and also removes 80% of the instruction prefetch buffer. As Figure 3 shows, the 16K cache array consumes less than 10% of the 5x86 die; moving to an 8K cache would have significantly reduced performance with only a small change in die area.

These tradeoffs give the 5x86 about a third less performance than the M1 at the same clock rate, but a much lower cost. As both parts move to advanced processes, the Pentium-pinout version of the 5x86 will replace low-end M1s. For example, a 150-MHz 5x86, achievable in IBM's 0.35-micron CMOS-5X process, should match the performance of a 100-MHz M1 at a fraction of its manufacturing cost.

Function Area	5x86	M1
Instruction decode	40,000	210,000
Branch unit / BTB	126,000	250,000
Address calculation	120,000	270,000
Execution unit(s)	160,000	380,000
Floating-point unit	160,000	160,000
MMU / load/store unit	76,000	120,000
Bus interface	66,000	70,000
Total logic transistors	748,000	1,460,000
Cache / buffers	1,200,000	1,550,000
Total transistors	1,948,000	3,010,000

Table 1. By eliminating superscalar execution, the 5x86 pares half or more transistors from most function areas. (Source: Cyrix)

But Is It Really a 586?

Cyrix's choice of the 5x86 name has already caused controversy among industry watchers and particularly its competitors. Cyrix points out that its 5x86 has several features found in Pentium that are not in the 486. On the other hand, the Cyrix chip is certainly missing key features from Pentium and similar processors.

Table 2 compares the 5x86 feature set to that of Intel's high-end 486, the DX4, as well as Pentium. The Cyrix chip can be thought of as a 486 enhanced with branch prediction, a deeper pipeline, and memory bypassing. Or it could be a Pentium that's missing superscalar execution, split caches, and a dual-ported data

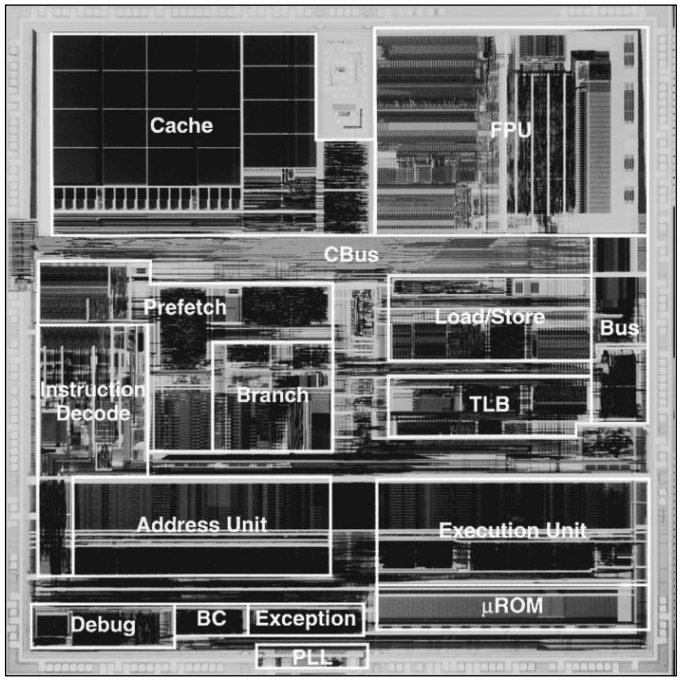

Figure 3. The 5x86, with 1.9 million transistors, measures 12 x 12 mm using a 0.65-micron three-layer-metal CMOS process.

Price & Availability

The 100-MHz 5x86 is currently sampling; Cyrix expects volume production in 3Q95. The part is available in either a 168-pin PGA or 208-pin PQFP at a price of $147 in 1,000-unit quantities. Contact Cyrix (Richardson, Texas) at 800.462.9749 or 214.968.8388; fax 214.968.8404.

cache, and is a bit short on instruction-buffer entries, BTB entries, and TLB entries—and uses a 486-style bus. You make the call.

But these feature-set comparisons are ultimately meaningless. The 5x86 includes some features that even Pentium doesn't have, such as a return-address stack. The bottom line is performance and cost. The 486, Pentium, and M1 should all reach 120 MHz in a 0.5-micron process. (Intel's fastest DX4 runs at 100 MHz, but we believe that this is a marketing constraint, not a design issue.) We expect the 5x86, however, to reach at least 133 MHz when it shrinks to that level. Thus, the 5x86 design clearly delivers better performance than the DX4, but it doesn't come close to Pentium performance on a clock-for-clock basis.

On a cost basis, again adjusting to comparable 0.5-micron processes, the 5x86 will be more expensive to manufacture than a DX4 but significantly less than a Pentium chip. Thus, on the crucial metrics of performance and cost, the 5x86 lies somewhere in between the two camps. In an ideal world for end users, it would be called a DX5.

Cyrix, however, is in the business of selling chips, not idealism. The company got a lot of mileage out of its 486SLC processors despite the fact that they used a 386 pinout and delivered about 20% less performance than an Intel 486 of the same clock speed. It shouldn't be a surprise, then, that Cyrix has chosen a 586-like label for its latest processor, even though it uses a 486 pinout and appears to be about 25% slower, clock for clock, than a Pentium. Cyrix has been referring to the M1 as a "sixth-generation" processor in recent releases, setting the stage to roll out that chip as the 6x86.

What makes good business sense for Cyrix, however, could cause end-user confusion in the long run *(see **0909ED.PDF**)*. To avoid this problem, Cyrix plans to use some sort of performance-based metric similar to NexGen's "P" ratings, allowing the company to better communicate the positioning of its parts to end users.

Attacking Pentium's Pricing

The 5x86 gives Cyrix a weapon to exploit a chink in Intel's pricing strategy. Intel has cut off its DX4 line at 100 MHz, choosing instead to offer OEMs aggressively priced Pentium processors for mainstream (soon to become low-end) desktops. This strategy would work fine in the old single-vendor x86 market, but it has a shortcoming: Intel's primary processor for this price point, the 75-MHz Pentium, carries a significantly higher manufacturing cost than a high-end 486.

Over time, Intel has improved the yield curve of its 0.5-micron parts enough that nearly all yield is at 100 MHz and above. In effect, the company is selling 90- and even 100-MHz chips as 75-MHz parts, garnering less revenue than if these parts were sold at their true speed.

Cyrix will instead target this performance point with the 5x86. As Table 3 shows, the 100-MHz 5x86 has a much lower list price than the 75-MHz Pentium, due in part to its manufacturing cost, which we estimate to be just over half of Pentium's cost. Amazingly, the Cyrix part has slightly better margins than the Intel product, and Cyrix's position will improve as it shrinks the 5x86 to a 0.5-micron five-layer-metal process early next year. In short, Cyrix has more margin to cut prices as the action heats up in this high-volume market segment.

As Figure 4 shows, AMD will compete in this space with 120-MHz and faster 486 chips *(see **0908MSB.PDF**)*. These chips will provide integer performance comparable to that of low-end Pentiums and have an even lower manufacturing cost than Cyrix's 5x86. For some applications, AMD's chips will be hampered by their 8K cache, half the cache on the 5x86 and Pentium. AMD may double the on-chip cache in future high-end 486s; in a 0.35-micron process, such a change would not significantly increase their manufac-

	Intel DX4	5x86	Pentium	Cyrix M1
Pipeline depth	5 stages	6 stages	5 stages	7 stages
Peak execution rate	1 instr	1 instr	2 instr	2 instr
Total on-chip cache	16K	16K	16K	16K
Split instruction/data caches?	no	no	yes	no
Instruction buffer	32 bytes	48 bytes	128 bytes	256 bytes
Instruction fetch width	128 bits	128 bits	256 bits	256 bits
Dual-access cache?	no	no	yes	yes
Avoids stalls on cache access?	no	yes	no	yes
Memory bypassing	no	yes	no	yes
Branch target buffer	none	128 entries	256 entries	256 entries
Return address stack	none	8 entries	none	8 entries
TLB entries	32 entries	32 entries	96 entries	128 entries
System bus width	32 bits	32 bits†	64 bits	64 bits
Die size (current process)	77 mm^2	144 mm^2	148 mm^2	394 mm^2
Die size (0.5-micron process)	77 mm^2	95 mm^2*	148 mm^2	225 mm^2
Clock-for-clock performance‡	0.6	0.75	1.0	1.2*

Table 2. Cyrix's 5x86 comes in between a high-end 486 and Pentium in both feature set and performance and, projecting to equivalent 0.5-micron processes, in die size as well. ‡on typical (unrecompiled) integer PC applications, based on vendor claims. †64-bit bus in future (Source: vendors except *MDR estimates)

turing cost.

There is a bigger difference among these competitors in floating-point performance. A standard 486, such as AMD's, is quite slow on programs that make significant use of floating-point math. Cyrix's 5x86 has a faster FP unit but can't match the speed of Pentium's pipelined FPU. For example, 80-bit FP multiplies take 16 cycles on a 486, 4–9 cycles on a 5x86, and just 3 cycles, with a single-cycle issue rate, on a Pentium. This gap could hurt AMD, and Cyrix to a lesser extent, on 3D-graphics applications and on other programs that use FP.

AMD may also be hampered by offering a 486 product while Intel's tremendous marketing resources continue to tell both OEMs and end users that Pentium is the solution. Cyrix hopes to avoid the 486 backlash by positioning its 5x86 as a Pentium-class device.

The 486 has more life left in the notebook market than on the desktop. While Cyrix markets its new chip as a 586-class device, its 486 pinout and moderate heat dissipation make it well suited as an upgrade to existing 486 notebook systems. The 5x86 should find its greatest success in this area, although it will certainly be used in desktop systems as well.

Penetrating the Top Tiers

Despite its upscale marketing position, the 5x86 may lose sales to AMD's fast 486s due to that company's existing customer base. AMD's current 486 chips are already used by major PC vendors such as Compaq, Digital, and Hewlett-Packard; these companies are likely to upgrade to the faster parts when they are available later this year. Among the top PC vendors, only Epson has so far indicated it will adopt the 5x86; other system makers endorsing the new chip are smaller players, such as ASE, Chicony, Dataexpert, and Veridata.

Cyrix is the only major x86 processor vendor that has not inked a deal with the world's largest PC maker, Compaq. Given Compaq's public disputes with Intel, this omission is surprising. So far, however, Compaq has adopted non-Intel processors only in its consumer products, which currently use AMD chips. This strategy leaves little room for both NexGen, which already has a deal with Compaq, and Cyrix.

	Am486DX4	Intel DX4	Cyrix 5x86	Pentium
Clock speed	120 MHz	100 MHz	100 MHz	75 MHz
Cache size	8K	16K	16K	16K
Typical power	3.2 W	3.5 W	3.0 W	2.4 W
Max power	4.0 W	4.3 W	4.3 W	5.2 W
IC process	0.5µm, 3M	0.5µm, 4M	0.65µm, 3M	0.5µm, 4M
Die size	56 mm^2	77 mm^2	144 mm^2	148 mm^2
Est mfg cost	$30*	$35*	$50*	$95*
List price†	$165	$170*	$147	$220*
Performance‡	0.7*	0.6	0.75	0.75

Table 3. Cyrix's first 5x86 is priced much lower than comparable Intel chips but will compete more closely with AMD's high-end 486. †List price in 1,000-unit lots for 3Q95. ‡Performance on typical (unrecompiled) integer applications relative to 100-MHz Pentium. (Source: vendors except *MDR estimates)

Cyrix, which expects to generate about $400 million in revenue this year, needs only a small portion of the roughly $10 billion x86 CPU market to prosper. The 5x86 is a strong contender against Intel's low-end Pentium chips and fills a gap in Cyrix's product line. The new chip will solidify the company's hold on its current customers and should help attract new ones. ♦

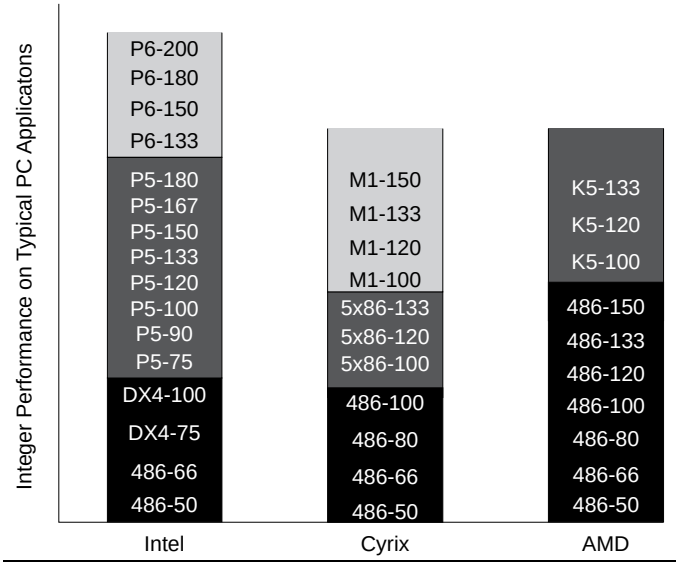

Figure 4. A map of products expected though 1996 shows how Cyrix's 5x86 and AMD's 486 overlap the low end of Intel's Pentium line. (Source: MDR estimates)

