Most Significant Bits

Cyrix Lays Out M1 Production Plan

Cyrix is now shipping samples of its M1 processor and expects to begin production in June. Initial production will use an enormous 394-mm^2 die *(see **081601.PDF**)*, fabricated in a 0.65-micron, three-level-metal process. We estimate the yield of this chip at about 15 chips per (200-mm) wafer, severely limiting production volume.

The initial parts will build a market for the M1 while Cyrix completes two die shrinks. First, a redesign to use five layers of metal is expected to produce a 225-mm^2 device, with production planned for 4Q95. Then an optical shrink to a 0.5-micron process will bring the die size down to about 170 mm^2—less than half the size of the initial design—in 1Q96. At this size, which is still 15% larger than the latest 0.6-micron Pentium die, Cyrix should be able to deliver considerable volume.

Cyrix expects the initial design to reach 100 MHz. The company expects the first shrink to yield chips running at 120 MHz, and the second shrink to push clock speeds beyond 133 MHz. Even with these increases, the M1 will lag Pentium in clock speed, since 120-MHz Pentiums are shipping now *(see **090402.PDF**)*, with both 133-MHz and 150-MHz chips due later this year.

The M1 does not need to match Pentium's clock to exceed its performance, however. Cyrix is now claiming that the M1 is more than twice the speed of Pentium at the same clock rate—on certain small benchmarks. The company has not revealed any specific benchmark results, but it has indicated that it expects an advantage over Pentium (at the same clock rate) of about 30% on applications. Over-hyping the chip's performance by making the 2× claim could ultimately hurt Cyrix's credibility, since the real performance will become apparent when systems begin shipping.

Cyrix has also begun pitching the M1 as "sixth generation technology in a Pentium socket." The M1 uses a considerably simpler microarchitecture than Intel's P6, but Cyrix believes that it will offer comparable performance at the same clock rate. Apparently stung by criticism of its dual-pipeline design as being old-fashioned, the company has cited a list of features—including superpipelining, register renaming, out-of-order completion, speculative execution, and multiple branch prediction—that are common to the M1 and P6 designs.

What the M1 lacks is the decoupled dispatch and execution of the P6, K5, and NexGen designs, which should enable more out-of-order execution and extract more parallelism from the instruction stream. Cyrix has critiqued the competing approach as being "non-native" and incurring extra overhead for translating x86 instructions for the RISC core, but it is far from clear that this criticism has any substance. Despite its sophistication, however, NexGen's Nx586 delivers only about 7% better application performance than Pentium at the same clock rate. Until benchmark results are available for the K5 and the M1, it is impossible to say whether the apparent advantages of the K5 design will result in higher performance.

It is difficult for an outsider to evaluate the performance of the various microarchitectures, especially since the companies' public disclosures have omitted numerous details. The M1 design is clearly superior to Pentium but appears to fall short of the P6. Ultimately, delivered performance—at whatever frequency each company is able to ship in volume—is what counts, and we expect the M1 to fall closer to Pentium than to the P6.

Indeed, since Intel should be shipping its 150-MHz P55C *(see **090402.PDF**)* by the time Cyrix ships 120-MHz parts, it isn't clear that the M1 will be faster than the fastest Pentium. What the M1 should deliver, in 1996 if not sooner, is very good performance for a lower price than Intel's P55C or P6; this may be an excellent market position, but it is different than the one Cyrix is claiming to have.

Intel Continues P7 Program(s)

Despite a rash of recent reports—first that the P7 will use the new Intel/HP architecture, then that it won't—Intel officially maintains that it hasn't made a decision yet as to the direction of this project. The P8, which already has a small staff in Oregon beginning development, will definitely implement the new architecture *(see **080801.PDF**)*, which we have dubbed P86. The P6, in contrast, is a pure x86 implementation. The P7, due in late 1997, is in the gray zone.

The subject arose when Intel COO Craig Barrett, speaking before an audience of financial analysts and reporters, stated that the P7 would implement the new architecture. As Intel spokespeople attempted to explain Barrett's statement, the situation became more muddled. One thing is clear: the company cannot present a unified story as to the status of the P7.

Sources indicate that there is internal competition between two projects: the original P7, a pure x86 implementation, and the new P86 design, being developed with assistance from HP. If the P86 chip delivers better performance than the P7 and is expected to ship at a similar time, Intel will probably cancel the original P7 and assign the P7 name to the P86 chip. On the other hand, if the P86 program is delayed, Intel may need to forge ahead with the original P7 to remain competitive.

At this time, the company has enough resources to keep both projects going, particularly with HP's help on the P86 design. By the end of this year, however, the P7

 © 1995 MicroDesign Resources

will be about a year away from tape out and will require extensive staffing. This resource demand will probably force Intel to make its P7 decision late this year. Until then, the company may continue both projects as a hedge. Barrett's comment may indicate his preference for moving to the new architecture as soon as possible.

The P86 architecture specification is nearing completion, with the two companies devoting many engineers to the project. If this project goes well, we expect Intel to shift the P7 to the new architecture. The P86 design should provide superior native performance and allow Intel to differentiate its products from x86 chips designed by AMD and Cyrix. These benefits should convince Intel to make the switch sooner rather than later.

HP Tapes Out PA-8000

The PA-8000, likely to be HP's final PA-RISC core, taped out on March 6; the company expects the new chip to appear in systems in 1Q96. At Compcon, HP announced performance estimates of at least 360 SPECint92 and 550 SPECfp92, within the range that we predicted when the part debuted *(see **081501.PDF**)*.

The company would not discuss the clock speed associated with these performance ratings. Sources indicate that the company hopes to reach 200 MHz, but the quoted numbers assume a speed "slightly less" than that. Until the first silicon is received and tested, we won't know whether the chip will actually achieve these aggressive performance targets, which are greater than those of any other announced processor.

SPEC95 Benchmarks Nearly Complete

The System Performance Evaluation Committee (SPEC) is nearing completion of the SPEC95 benchmark suites. These suites will replace the aging SPEC92 benchmarks, which are becoming increasingly inadequate to represent current application performance and are also overly susceptible to compiler tuning. SPEC has a list of candidate benchmarks for SPECint95 and SPECfp95; the members will meet on April 4–8 to decide which will be included in the final suites. If all goes well, the first SPEC95 results will be published in June.

The new benchmarks will appear not a moment too soon. Intel recently pushed the SPECint92 score for the 100-MHz Pentium to 122; the same part was rated at just 100 SPECint92 last fall. Other x86 processor vendors have given up on SPECint92, apparently because they cannot match Intel's compiler-tuning expertise.

The committee is considering retaining the baseline and peak measurements of the current system *(see **0803MSB.PDF**)* but making optional the publication of baseline results, which exclude most compiler tuning. Unfortunately, this change would destroy the usefulness of baseline results, making it even harder to find scores for a broad enough range of processors to make comparisons. Most vendors would continue to rely on peak (optimized) measurements, which don't reflect the performance achieved by many users.

We advocate using SPEC95 as an opportunity to eliminate excessive compiler tuning by defining only baseline run rules for the new suites. This change would make the numbers more representative while making the suite easier to use. For more information on the SPEC95 suite or to give your opinion to SPEC, contact your company's SPEC representative or send e-mail to *walter.bays@eng.sun.com*.

VLSI to Sell NexGen Chip Set

Boosting NexGen's credibility, VLSI Technology has announced plans to market a PCI system-logic chip set for the Nx586 processor. The prospects for this Pentium-class chip have been hindered by its incompatibility with Pentium system logic. Currently, only a single chip set is available for the 586: NexGen's own VL-Bus design.

The processor vendor has been working on a PCI chip set, which is due to sample soon. The companies announced plans to "codevelop" a PCI chip set for shipment by midyear; this "new" chip set is actually the existing NexGen design with little or no change. Under the agreement, the PCI chip set will be available solely from VLSI and not NexGen.

While VLSI's backing enhances NexGen's image, simply taking over the PCI chip set from the CPU company does nothing to expand the number of chip-set options for Nx586 customers. VLSI may develop its own 586-compatible chip sets in the future, but only if NexGen garners more volume. Having a partner with VLSI's stature may help NexGen in this regard.

TI Gains 486DX2 Design

Completing its legal settlement with Cyrix *(see **0817MSB.PDF**)*, Texas Instruments acquired rights to that company's 486DX2 design on March 1. TI immediately began adapting the design to its own manufacturing process and plans to sample the DX2 later this year, with volume production commencing in 1Q96. The company did not announce pricing for the part, which it expects to deliver in 66- and 80-MHz speed grades.

TI today is positioned at the low end of the 486 market with its SXL processors. These parts sell mainly to makers of low-cost portable systems and to developing countries. Although the DX2 appears to upgrade TI's line, in fact the market will have changed by the time the DX2 reaches it. By 1Q96, the SXL may be a tough sell, even overseas, and the faster DX2 parts will be required for TI to address even the low end of the x86 market.

The company will not receive the M1 nor any other x86 processor designs from Cyrix; TI is working on its own Pentium-class processor core that will eventually offer a growth path from the 486DX2. ♦

 © 1995 MicroDesign Resources