

Mendocino Improves Celeron

New Low-End Chip First With On-Die Level-Two Cache

by Linley Gwennap

Giving a major performance boost to its Celeron line, Intel today announced a new version of the product, code-named Mendocino, that brings the level-two (L2) cache onto the processor chip. This part, Intel's first with an on-die L2 cache, fixes the performance problems of the initial Celeron product, code-named Covington, and should strengthen the acceptance of Intel's new low-end product line.

Mendocino Trounces Covington

The initial versions of Mendocino run at 300 and 333 MHz. This creates an overlap with the Covington parts (see MPR 4/20/98, p. 14), which operate at 266 and 300 MHz. To avoid this overlap, Intel had initially planned not to market a 300-MHz Covington but later changed its mind in the face of the poor performance of the 266-MHz version.

The overlap causes nomenclature problems, as both Covington and Mendocino are marketed under the Celeron name. To distinguish them, Mendocino is sold as Celeron-300A and Celeron-333. The latter part needs no A, since there are no 333-MHz Covingtons, and future Mendocino parts will also do without the suffix.

Intel hopes PC buyers will figure out that A stands for better performance. According to Intel, the Celeron-300A scores 25% better than the Celeron-300 on Winstone98 in a typical low-end system, as Figure 1 shows. On SPECint95, a more CPU-intensive test, Mendocino scores 36% better than Covington at the same clock speed.

Although this gap represents at least three speed grades in Intel's normal lineup, the vendor is pricing the new part as if it were only one speed grade faster: the Celeron-300A lists for \$149, whereas the Celeron-300 lists for \$112.

Covington's poor performance (slower than even a Pentium/MMX-233) stems from its complete lack of L2 cache: any time the CPU needs instructions or data that aren't in its 32K of on-die cache, it must wait for the slower main memory to respond. Thus, the 300-MHz CPU is often reduced to the 66-MHz speed of main memory.

Mendocino solves this problem by adding a 128K L2 cache to the processor die, greatly reducing the number of accesses that go to main memory. Although this cache is only one-quarter the size of Pentium II's off-chip L2 cache, it manages to deliver similar performance. Because the Pentium II cache relies on cheap commodity SRAMs, it operates at just half the CPU speed, whereas Mendocino can run its on-die L2 cache at the full CPU speed. The latency of the on-die L2 cache is also half that of the commodity cache, reducing the number of CPU stall cycles.

The faster cache helps a 300-MHz Mendocino come within 3.5% of a 300-MHz Pentium II on SPECint95, assuming the Pentium II uses the 440LX chip set, which is similar in performance to the 440EX used with Mendocino. The improved 440BX increases Pentium II's advantage to 5%. In either case, Mendocino is about half a speed grade behind Pentium II at the same clock speed.

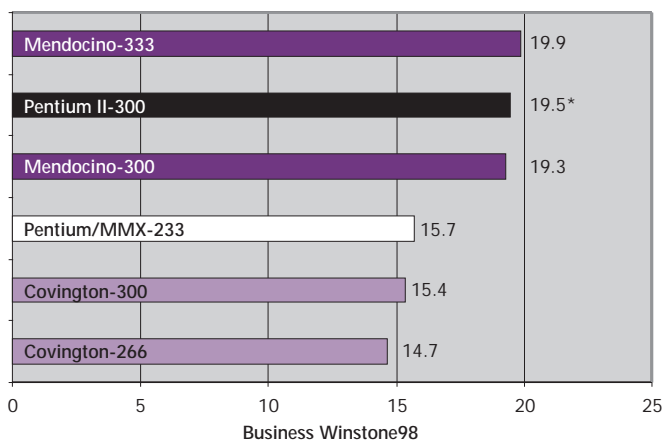


Figure 1. The new Mendocino processors are much faster than the earlier Covingtons and nearly as fast as Pentium II at the same clock speed. All processors tested with 440EX chip set and standard L2 cache, except Pentium/MMX tested with 440TX and 512K L2 cache. All processors tested with 32M SDRAM, Seagate 32112A, and STB Velocity 128 PCI with 4M SGRAM in 1,024 × 768 × 16 mode. (Source: Intel except *MDR estimate)

Whirlwind Development Effort

Unlike most Intel processors, which are planned well in advance of their launch, Mendocino was conceived and executed very quickly. Intel revamped its entire low-end strategy last fall, after watching the sub-\$1,000 PC surge from nearly nothing to 40% of the U.S. retail market by August 1997. In response, the company began simultaneous development efforts for Covington and Mendocino. The two products would anchor a new brand, Celeron, designed to make Intel more competitive in the emerging sub-\$1,000 PC market.

Since Covington required only removing the external L2 cache from the Pentium II module, that product was ready first, with shipments starting in June. Mendocino, however, required revising the Deschutes die. Although the modifications were relatively minor, any silicon rework requires extensive design and verification efforts.

But Mendocino was completed in record time. Originally expected to ship late this year, the product was pulled into an August release. Although most processors take two or three tries before they can be shipped, the first silicon of Mendocino was so clean that Intel is using it to begin volume shipments.

To accelerate first shipments, Intel kept redesign work to a minimum. From the CPU's standpoint, the on-die L2 cache appears identical to the off-die cache, just faster. The Deschutes cache controller is already designed for a full-speed cache, as used in the Xeon line. Bringing the L2 on die, however, greatly reduces its latency. The interval between the time the L1 miss is signaled and the time the first (critical) word of data is received is just 8 cycles on Mendocino, compared with 14 cycles for Xeon and 18 cycles for a standard Pentium II.

The design team made one significant change: Mendocino uses two unidirectional data buses to replace the single bidirectional data bus in Pentium II. The data buses remain 64 bits in width. The split eliminates dead cycles needed to turn the bus around and allows reads and writes to overlap in some cases, slightly improving performance. As Mendocino's




Figure 2. Although Mendocino's on-die L2 cache increases the cost of the die, eliminating the external L2 cache results in an overall cost savings of about \$5. (Source: MDR Cost Model)

Price & Availability

In quantities of 1,000, the Celeron-300A and -333 are priced at \$149 and \$192, respectively. Both are available in volume now. For more information, access the Web at www.intel.com/celeron/index.htm.

L2-cache connections are entirely on the die, there is no pin-count penalty for adding the extra bus.

A more aggressive design could have taken further advantage of the L2's proximity by increasing both the width of the data transfers and the number of sets in the cache. Wider transfers could refill an entire L1 cache line in a single cycle, speeding back-to-back misses, and greater associativity would increase the cache's hit rate. These features are likely to appear in future Intel processors that use on-die L2 caches.

Reducing Cost for Celeron Line

A key reason for creating the new Celeron processors was to reduce manufacturing costs compared with the Pentium II products. The MDR Cost Model estimates that a Mendocino module costs about \$65 to manufacture, while Pentium II costs about \$70. As Figure 2 shows, the Mendocino die itself actually costs more to build than a Deschutes die, as the on-die L2 cache increases the die size from 131 mm² to 154 mm².

The added area is cache, as Figure 3 shows, and this area is protected by redundant columns. Redundancy prevents most defects in this area from ruining the chip, so Mendocino yields should be nearly the same as those for Deschutes. The die cost is still higher, however, since fewer die fit on each wafer.

This extra cost is made up for by savings in the module. The Celeron module is less expensive, not only because it lacks SRAMs but also because the PC board has fewer layers, fewer discrete components, and no black plastic case.

The initial Mendocino chips are packaged in the same 528-pin LGA used for Deschutes, although more than 100 pins are unused, due to the lack of an external L2 cache. Intel could have cut costs by moving to a smaller package, but instead it will eventually offer Mendocino in a 370-pin plastic PGA, eliminating the module entirely. This packaging option will be available in 1Q99 and should reduce manufacturing costs to about \$55, as Figure 2 shows. Intel has not yet quoted pricing for the moduleless Mendocino, but we doubt the vendor will pass its cost savings on to its customers.

Faster Mendocinos Coming Next Year

With four speed grades now in the Celeron line, Intel's pricing strategy for that line is becoming evident. As we predicted when the brand was first announced (see MPR 3/30/98, p. 1), the high end of the Celeron line is settling in at just under \$200; the Celeron-333 was announced at a list price of \$192. At the bottom end, the Celeron-266 now lists for \$86.

Pentium II Hits 450 MHz

Intel also extended its desktop Pentium II line with a new 450-MHz part. The new speed grade is the result of Intel's growing experience with its 0.25-micron process, known as P856, and the Deschutes CPU. The process has now been in production for nearly a year, and Deschutes has been shipping since January. After locating and correcting key speed limiters, Intel can now build faster parts. The 450-MHz processor is available immediately at a list price of \$669.

With the latest announcement, the desktop Deschutes is now running 50% faster than its predecessor, the 0.28-micron Klamath. We do not expect faster versions of Deschutes to emerge, as 50% is generally the limit for speedups due to process shrinks. Katmai, due in 1Q99, will be the next step up for the desktop line.

Although Intel is currently selling Pentium II processors for as little as \$159, we expect Pentium II products to stay above the \$200 mark in the future. This will provide a clear delineation in price between Celeron and Pentium II.

The distinction in performance will be fuzzier. Since Mendocino is based on the same Deschutes CPU core as the current Pentium II parts, there is no technical reason that Mendocino won't yield at 400 MHz or even 450 MHz. To maintain proper separation between the two lines, however, Intel can dole out new clock speeds for Mendocino only after Pentium II gains new performance points.

For example, Intel plans to introduce a 500-MHz version of Katmai in 1Q99. Around the same time, we expect Mendocino to jump to 366 MHz. Based on the Pentium II line, one might expect the next step to be 350 MHz, but a 350-MHz Mendocino wouldn't be compatible with current 440EX-based Celeron systems, which don't support a 100-MHz system bus. We don't expect Intel to support a 100-MHz bus in Celeron systems until the release of the Whitney chip set (see MPR 4/20/98, p. 18), slated for mid-1999.

In 2H99, Intel will deliver a 0.18-micron version of Katmai code-named Coppermine. This part will allow the Pentium II line to reach clock speeds in excess of 600 MHz. These faster parts will allow Intel to boost Mendocino clock speeds again. Using Whitney, the company is likely to deploy 400- and 450-MHz versions of Mendocino at that time.

While Intel has so far avoided serious competition in the Pentium II space, the Celeron line competes directly with inexpensive products from AMD, Cyrix, and other vendors. Thus, Intel may not be able to control the pricing and timing of its products as well as it has in the past. If another vendor offers processors at 400 MHz or faster in 1H99, Intel may accelerate the launch of the faster Mendocinos, since there is no technical reason that it couldn't.




Figure 3. The lower part of Intel's Mendocino die is similar to the Deschutes CPU. The L2 cache, at the top, consumes about 23% of the die. Mendocino contains 19 million transistors and measures 10.4 mm × 14.8 mm in Intel's 0.25-micron four-layer-metal process.

L2 Caches Move Onto the Processor

Mendocino is Intel's first microprocessor with an on-die L2 cache, but it won't be the last. Intel is already preparing a slightly modified version, code-named Dixon, that will include 256K of L2 cache. Dixon should appear around the end of this year. Even as Mendocino takes over the Celeron line, its big brother will push the nonintegrated Pentium II processors out of the mobile line.

Discrete L2 caches will persist for a longer time in the Pentium II line. Sources indicate that Katmai does not have an on-die L2 cache, but once Intel moves to the Coppermine version, moving the L2 onto the processor makes economic sense. By mid-2000, discrete L2 caches are likely to be found only in the high-end Xeon line, as these products need massive amounts of cache.

This shift will practically eliminate the market for PC cache SRAMs and will require Intel to increase its own wafer capacity. But as shown by Mendocino, the change will result in better performance at lower cost—a clear win for Intel. 