Centaur Improves C6 With No Extra Cost

New C6+ Is Faster on Integer, FP, and MMX Applications; Original C6 Ships

by Linley Gwennap

Even as it begins selling its initial x86 processor, the C6, IDT subsidiary Centaur has already developed a follow-on product with many minor improvements. The C6+, due to reach production in 2Q98, brings MMX and FP performance up to par with the competition while adding other features, such as branch prediction, that raise performance on general PC applications. Centaur remains focused on low cost: these improvements add just 3 mm^2 to the die size of the C6.

Speaking at last month's Microprocessor Forum, Centaur founder Glenn Henry disclosed plans to further boost the performance of the C6, now known as the IDT WinChip family. IC process improvements should allow the part to reach 300 MHz by mid-1998. In 2H98, IDT will use a 0.25-micron process to further increase clock speed while incorporating a 256K on-chip level-two cache with the CPU.

Centaur has independently adopted many of the same techniques being used by AMD and Cyrix to extend their K6 and 6x86MX processors, respectively, including an enhanced FPU, new instructions for 3D graphics, a faster system bus, and an integrated L2 cache. Although the C6 has a die-size advantage over competitive 0.35-micron processors, AMD plans to begin shipping its 0.25-micron K6, with a smaller die than the C6, in 1H98 (see MPR 10/27/97, p. 19).

The C6+ tapeout is imminent but has not yet occurred, so all performance projections are based on simulations. Because of the limited number of changes from the current C6 and the team's demonstrated skills, Henry believes the new part will reach volume production just six months after tapeout, half the time it usually takes for a new processor.

	Centaur C6+		Pentium/MMX	
	Thruput	Latency	Thruput	Latency
FP add	1 cycle	3 cycles	1 cycle	3 cycles
FP store	1 cycle	n/a	2 cycles	n/a
FP multiply (SP)	1 cycle	3 cycles	1 cycle	3 cycles
FP multiply (DP)	2 cycles	4 cycles	1 cycle	3 cycles
FP mul/add*	1 cycle	3 cycles	4 cycles	6 cycles
FP to integer*	1 cycle	3 cycles	6 cycles	6 cycles
FP sq root (SP)*	24 cycles	24 cycles	70 cycles	70 cycles
FP inv sq root (SP)*	24 cycles	24 cycles	109 cycles	109 cycles
MMX add	1 cycle	1 cycle	1 cycle	1 cycle
MMX multiply	1 cycle	1 cycle	1 cycle	3 cycles
MMX mul/add	1 cycle	1 cycle	1 cycle	3 cycles
MMX store	1 cycle	n/a	2 cycles	n/a

Table 1. Compared with the Intel Pentium/MMX, the C6+ has the same or better performance on most MMX and FP instructions except for double-precision FP multiplication. *C6+ uses proprietary instructions; Pentium/MMX uses standard x86 instructions. n/a=not applicable (Source: vendors)

FP and MMX Match Pentium/MMX

The original C6 was designed and put into production in less than two and a half years, a task that takes most vendors three or four years. Given the tight schedule, the designers focused on time to market rather than on maximum performance. One area that didn't get as much attention was the floating-point unit, since the number of PC applications that use FP today is still relatively small.

The C6+ contains a completely redesigned FPU that fits into the same die area as the current unit but is now fully pipelined for most FP operations. The C6+ takes an extra cycle for double-precision multiplies, but this operation is typically used only in technical applications (e.g., CAD) that are not in Centaur's target market.

As Table 1 shows, on most instructions the new FPU is as fast as Pentium/MMX's, and Centaur expects the C6+ to deliver about the same performance on most FP applications as Pentium/MMX at the same clock speed.

Centaur also completely redesigned the C6's MMX unit to improve performance. The C6+ can issue and execute up to two MMX instructions per cycle, with essentially the same restrictions as Pentium/MMX. Except for the MMX unit, however, the C6+ is a scalar processor, so it can't pair an MMX instruction with an integer instruction.

The Centaur chip is faster than Pentium/MMX on several MMX operations. As Table 1 shows, the C6+ has better latency on MMX multiply or multiply-add, but most applications depend on throughput, not latency, for performance. Stores take a single cycle, half as long as on the Intel chip. As a result, the C6+ should slightly outperform Pentium/MMX on most MMX-based benchmarks, according to Centaur.

New Instructions Aid 3D Graphics

Like AMD and Cyrix (see MPR 10/27/97, p. 22), Centaur has added new instructions to its next part to speed 3D geometry calculations. Centaur has taken a more radical approach, however, that is not compatible with either the AMD 3D instructions or Cyrix's MMXFP. Instead of simply increasing throughput by pairing single-precision FP values and operating on them in parallel, the C6+ implements a new set of FP registers and instructions that use them.

The new part adds 22 directly addressable 80-bit floating-point registers to the eight-entry FP stack defined by the x86 architecture. This greatly increases the number of values that can be stored in the chip, reducing time-wasting memory

Price and Availability

IDT is now shipping its Winchip (C6) at clock speeds of 180 and 200 MHz. The "suggested retail price" is $90 and $135, respectively. The company did not reveal its 1,000-piece price, which is presumably lower. IDT expects to sample 225- and 240-MHz parts this month, with production shipments in 1Q98. The C6+ is scheduled to ship in 2Q98. For more information, try *www.winchip.com*.

accesses. The designers then added 53 new instructions (using 12 x86 opcodes) to operate on these new registers. These are fully IEEE compliant and handle all precisions.

In addition to the usual arithmetic and load/store operations, the new instructions include a fully pipelined floating-point multiply-accumulate. The x86 instruction set has no such operation, and issuing a multiply and a dependent add takes four cycles on a Pentium/MMX. The new instructions also include fast square root and inverse square root operations, as Table 1 shows. These operations are frequently used in lighting calculations and other 3D geometry algorithms.

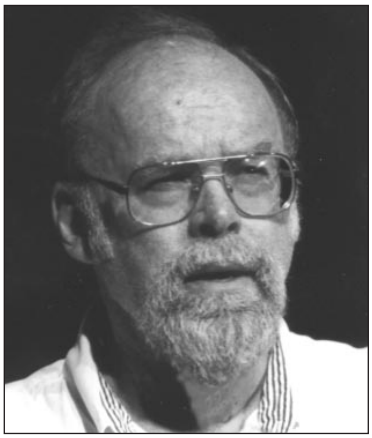

MICHAEL MUSTACCHI

Centaur founder Glenn Henry describes the new features of the C6+ at the Microprocessor Forum.

Speed Gains Require Direct3D

As a result, Henry claims the C6+ will have significantly better performance on 3D games than Pentium/MMX at the same clock speed. Yet the die size impact of the new instructions is less than 1 mm^2, mostly for the larger register file. Decoding the new opcodes is simple, and most of the data paths were already in place, supporting microcode primitives to execute existing x86 instructions.

These improvements are moot, however, unless the new instructions are used by software, a struggle with which Centaur, along with AMD and Cyrix, must contend (see MPR 10/27/97, p. 35). Centaur is developing its own version of Microsoft's immediate- and retained-mode library for Direct3D. If Microsoft agrees to distribute this code as part of DirectX 6 and Windows 98, any 3D application that uses this API could take advantage of the new instructions transparently.

Microsoft has supplied its code to Centaur but hasn't committed to distributing Centaur's version. Instead, Microsoft would prefer that the x86 vendors agree on a single set of 3D extensions, and Henry has volunteered to adapt his chip to one of the others. At the Microprocessor Forum, AMD CEO Jerry Sanders publicly offered to license his company's AMD 3D extensions, and the K6 3D will ship before Cyrix's Cayenne, making AMD the logical choice. Henry couldn't confirm whether he will use AMD's or Cyrix's extensions.

Centaur has gone a step further than any other vendor, even Intel, by adding state to the machine. Unless Microsoft agrees to modify Windows 98, which seems unlikely, the state of the 22 new registers will not be saved and restored on a context switch. If the new instructions are used only for 3D, and only one 3D application is running at a time, there should be no data corruption.

A further concern is the potential incompatibility with future instructions. If Intel were to use one of the 12 new opcodes for a different purpose, say an instruction in its forthcoming MMX2 extensions, it would be difficult for Centaur to implement both its instructions and the new Intel instructions. Unless Intel reveals its MMX2 encodings soon, it could put Centaur over a barrel.

Standard PC Applications Gain About 6%

The C6+ includes several minor improvements to speed standard integer applications. Some couldn't be implemented in the C6 due to its tight schedule; others became apparent later when analyzing code traces for performance bottlenecks. Taken together, these changes improve performance over the C6 by about 6% on the Winstone 97 Business benchmark, a collection of popular PC applications.

For example, the C6+ improves the timing of several instructions. Integer multiplies take 6 cycles (compared with 10 for Pentium/MMX). Arithmetic instructions that write their results to memory take two cycles, one fewer than on the Intel chip. The otherwise scalar execution core can pair two PUSH or two POP instructions, emulating the two-way superscalar Pentium core.

Cache improvements also contribute to better integer performance. The new data cache is four-way set associative, twice the number of sets as in the C6. The new cache also supports a write-allocate option, reducing the number of writes to the L2 cache when this mode is enabled.

Highly Accurate Branch Prediction

Perhaps the most important change is adding branch prediction. The C6, like the 486, has no branch prediction and takes a three-cycle penalty on every taken branch (except subroutine returns, which are handled by a return-address stack). The C6+ has a 4,096 × 1-bit branch history table (BHT) that predicts the direction of each branch but not the target address, as Pentium's branch target buffer (BTB) does.

Thus, correctly predicted taken branches create a single-cycle "bubble" in the execution stream. Because the C6 uses an instruction queue to decouple the fetch stream from the execution units (see MPR 6/2/97, p. 1), this bubble reduces performance only when the queue is empty, which is about 15% of the time. By not caching target addresses, the C6+ is

able to track the history of eight times as many branches as Pentium/MMX, improving its prediction accuracy.

Accuracy is further improved by two advances in the prediction algorithm. The C6+ uses a two-level indexing method known as Gshare and a new encoding method called "agrees" (see MPR 11/17/97, p. 22); Intel has not provided the details of Pentium/MMX's two-level BTB. As a result of these changes and the larger BHT, the C6+ correctly predicts 93% of the branches in Winstone 97 Business, compared with 82% for Pentium/MMX, according to Centaur. Yet the C6+ BHT consumes less than one-eighth of the space of Intel's BTB.

Roadmap Includes On-Chip L2 Cache

IDT is currently selling the C6, which is built in a 0.28-micron hybrid process (the metal layers are similar to those of a 0.35-micron process), at speeds of 180 and 200 MHz. The part is currently sampling at speeds of 225 and 240 MHz, and the company plans to ship these speed grades in 1Q98.

Since the C6+ is essentially the same size as its predecessor, IDT plans to offer it at the same price, providing more performance for free. In addition, circuit rework should boost its performance to 266 MHz in the same process. IDT also has a version of the 0.28-micron process that reduces the supply voltage from 3.3 V to 2.5 V; this reduction allows thinner gate oxides and thus faster transistors. In this process, the C6+ is projected to reach 300 MHz by mid-1998, as Figure 1 shows.

In 2H98, IDT plans to shrink the C6+ to a true 0.25-micron process. This process should reduce the die size to less than 60 mm^2, smaller than any competitive x86 processor. Running at 2.5 V, the chip should have moderate power dissipation despite higher clock speeds, making this new part suitable for notebooks or desktops. Regarding clock speed, Henry says only that the 0.25-micron C6+ should operate in excess of 300 MHz.

To take advantage of the tiny CPU core as well as IDT's experience as a leading SRAM vendor, the company plans to pack a 256K level-two cache onto the same die as the CPU, providing an integrated product along with the nonintegrated version. This cache will be eight-way set associative and operate at the full processor speed, offering much better performance than an external cache. In fact, this design provides many of the benefits of Intel's Pentium II dual-bus architecture without breaking compatibility with Socket 7. The integrated part will be particularly good for notebooks, where the power savings should be substantial.

This part appears to be the ultimate goal of IDT's x86 strategy. By adding a modest processor core to its existing 2-Mbit SRAM, IDT can increase the value of that chip by an order of magnitude. Instead of simply providing the cache for a PC, the company can now provide the CPU/cache subsystem, greatly improving its profit margins.

Performance Matches Other Low-End Chips

According to Centaur, today's 200-MHz C6 delivers about the same performance as a 200-MHz Pentium/MMX on the Winstone Business 97 benchmark in a low-cost system configuration. With its integer improvements, the C6+ should outperform the Intel chip by about 7% and match or exceed the performance of AMD's 200-MHz K6 and Cyrix's 6x86MX-PR200.

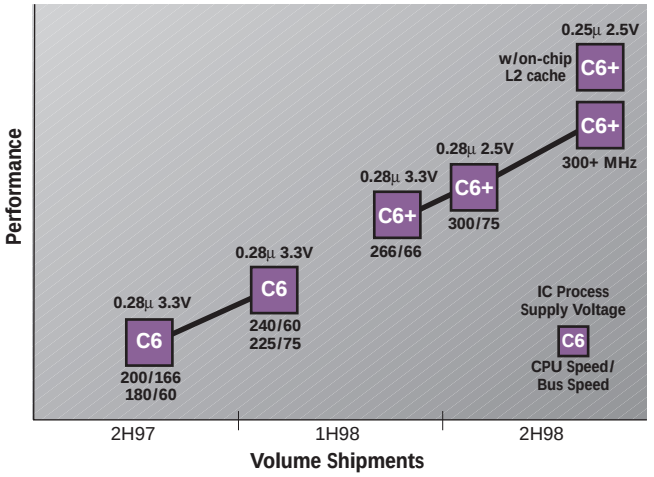

Figure 1. IDT's C6 roadmap shows faster clock speeds in 1Q98, the improved C6+ in 2Q98, and a 0.25-micron shrink in 2H98 that allows an optional on-chip 256K L2 cache. (Source: IDT/Centaur)

The improvements in MMX and floating-point performance are fairly well matched by those in AMD's K6 3D and Cyrix's Cayenne parts. The K6 3D is due in 1H98, about the same time as the C6+, but Cayenne is not expected until 2H98. The new C6+ instructions are aimed at speeding the same 3D applications as the AMD 3D and Cyrix MMXFP instructions, but it is too early to tell if any of the three has a performance advantage over the others.

AMD plans to stay with Socket 7 through the end of 1998 and also expects to add a 256K L2 cache to its K6 3D part in 2H98. This part should be similar to IDT's integrated offering, although IDT probably can't match the 400-MHz clock speeds in AMD's roadmap. Cyrix hasn't clarified its socket strategy for Cayenne but intimated that it may add a backside bus to Socket 7 or move to a new interface in 2H98.

In summary, the advances in the C6 roadmap should allow it to keep pace with the low-end offerings from AMD and Cyrix. Intel will be moving its Pentium II into the low end of its line by 2H98 (see MPR 11/17/97, p. 4). This transition will put pressure on all competitors still relying on Pentium/MMX's Socket 7 at that time.

Given the relatively modest cost of developing the C6, IDT's goal is to gain only a small (1–2%) share of the x86 market. To do so, it must underprice the competition, namely AMD and Cyrix. IDT would reveal only the single-unit prices for the C6-180 and C6-200, which are $90 and $135 respectively. In comparison, AMD sells a K6-200 for $160 in 1,000-piece quantities. With the C6's estimated manufacturing cost of $40, the margins are small but appreciable. For an SRAM vendor, it must seem like a fine way to make a living. M