VOLUME 11, NUMBER 2
FEBRUARY 17, 1997

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

Klamath Extends P6 Family

New Intel Design Adds MMX, 16K Caches, 16-Bit Speedup

by Linley Gwennap

Two years after introducing the P6 design, Intel has disclosed the chip that will bring that design to the masses. Although the company refers to the new chip only as "a future member of the P6 family with MMX," sources indicate the design will first appear in 2Q97 as the processor code-named Klamath. Some of the features, however, may not be used until the Deschutes processor, a 0.25-micron shrink of Klamath due to appear around the end of this year.

Klamath brings significant performance enhancements while expanding the P6's appeal in the consumer space. In addition, Klamath parts will cost less to build than Pentium Pro, the original P6, due to a less expensive IC process and an inexpensive package. These savings will ultimately bring Klamath prices below Pentium Pro's, although probably not until late this year.

The changes to the original P6 core are modest but effective. To make the part more appealing to consumers, Intel added MMX for multimedia applications and also improved handling of segment loads, a key bottleneck for the 16-bit code found in Windows 95 and older applications. To reduce cost, Klamath exchanges Pentium Pro's expensive multichip module for a plastic LGA package connected to the level-two cache through a small printed-circuit board. The chip itself is built in a pure CMOS process, eliminating the extra bipolar expense of PPro's BiCMOS process.

The major performance benefit comes from improved clock speeds. At the recent ISSCC, Intel demonstrated a supercooled P6 chip running at 400 MHz and presented a performance estimate for a 300-MHz version. While these clock speeds are possible in future versions, we expect to see Klamath products top out at 266 MHz, a 33% boost over the fastest Pentium Pro. Intel did not disclose any meaningful performance data; we expect a 266-MHz Klamath to deliver about 10.5 SPECint95 and 7 SPECfp95 (base). This performance would keep Intel ahead of its x86 competitors, abreast of PowerPC, and behind the high-end RISC workstation processors.

Although Intel has, for the first time, disclosed many technical details about the new P6 processor, it did not reveal specifications such as power or performance. In addition, clock speeds, pricing, availability, and the actual product name have been withheld, presumably until the product introduction set for this spring.

Cache Architecture Redesigned

To reduce costs and allow higher clock speeds, Intel revamped the entire cache design. Pentium Pro *(see* **090202.PDF***)* couples 8K instruction and data caches with a fully pipelined level-two cache that runs at the speed of the CPU, up to 200 MHz. The high bandwidth of this L2 cache compensates for the small L1 caches. To operate at 200 MHz, however, the cache must be built from expensive high-speed SRAM and be physically close to the CPU; Intel combines the PPro CPU and a custom L2 cache chip in a single ceramic package.

For Klamath, Intel needed to push the CPU faster but wanted to reduce system cost. These goals forced the adoption of an external half-speed cache, as Figure 1 shows. For example, a 266-MHz Klamath can use 133-MHz synchronous burst SRAMs for the L2 cache. At this speed, the cache can use low-cost commodity SRAMs that don't need to be in the same package as the CPU. For very low cost systems (and future faster CPUs), the Klamath design supports a one-third-speed cache.

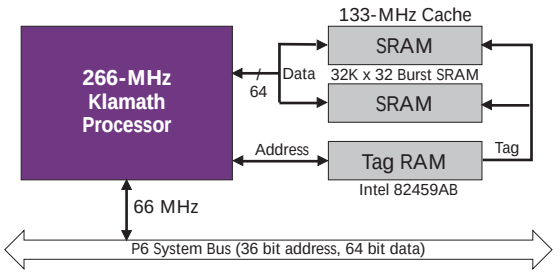

Figure 1. The Klamath CPU, along with external cache tag and data SRAMs, can be placed onto a small daughtercard.

To compensate for the slower L2 cache, Intel increased the L1 caches to 16K each, improving their hit rates and reducing the number of accesses to the external cache. The company had originally hoped the cache redesign would result in similar clock-for-clock performance for Pentium Pro and Klamath, but sources indicate a minor decrease in integer performance and a larger loss on floating-point applications due to the slower cache. Higher clock speeds, however, give Klamath a performance edge.

Klamath does support a full-speed cache, although this mode is unlikely to be used in PCs, as it would require very expensive SRAMs and potentially a multichip package. Such a device could be used in servers and workstations, but we expect a high price tag will prevent this configuration from reaching the PC mainstream.

A potential performance advantage of the external L2 cache comes from increasing its size. Pentium Pro is currently limited to 512K of cache, due to its two-chip package. With the L2 cache on a daughtercard, the Klamath design can support much larger caches, although Intel would not specify an upper bound. A limiting factor for initial Klamath products could be the external tag RAM, an Intel proprietary design (82459AB). This device supplies tag entries for up to 512K of L2 cache. Without a larger tag RAM, Klamath's cache size cannot exceed 512K; we expect to see an improved tag RAM around the end of this year. Intel did not specify the set associativity of the 82459AB.

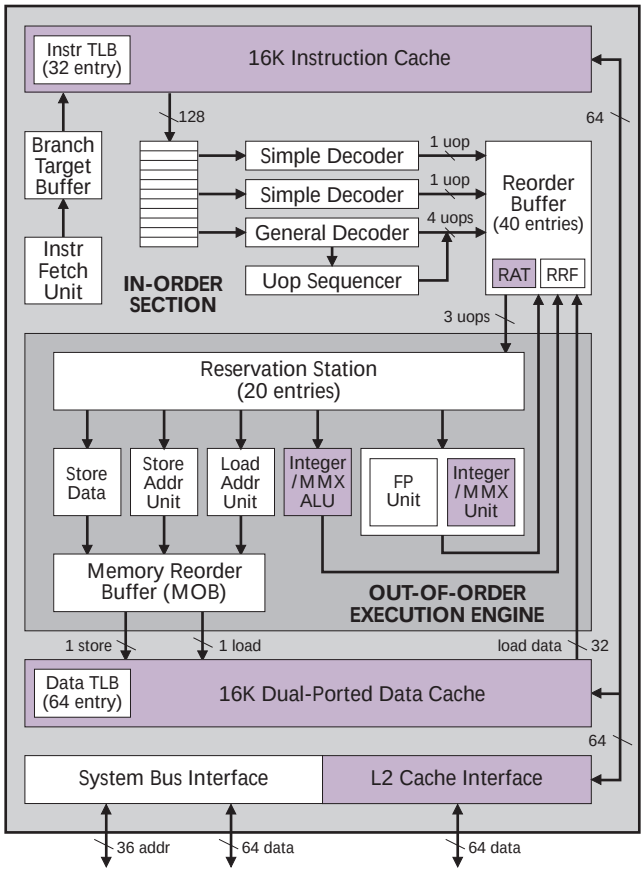

Figure 2. This Klamath block diagram shows the significant modifications over Pentium Pro were doubling the on-chip caches and adding MMX functions to both integer units. The register alias table (RAT) now supports renameable segment registers.

Improvements for Multimedia, 16-Bit Code

Like the P55C, Klamath adds Intel's MMX multimedia extensions *(see* **100301.PDF***)*. These 57 new instructions provide parallel processing for the 8- and 16-bit data often found in multimedia applications. Performing key calculations in parallel increases performance from 40% on motion video to as much as 300% on image processing (e.g., Photoshop). For example, a 233-MHz Klamath will perform full DVD decoding (AC-3 audio and MPEG-2 video) at 30 frames per second entirely in software.

As Figure 2 shows, both integer units in the original P6 core were modified to handle MMX instructions. The main function unit, which handles integer and FP instructions, gains an MMX ALU and multiplier in the Klamath design, as detailed in Figure 3. The second integer unit now includes a second MMX ALU and a shifter. As a result, Klamath can dual-issue almost any combination of MMX instructions except for two multiplies. As in the P55C, Klamath's MMX units all have a latency of one cycle, except for the multiplier, which is fully pipelined with a three-cycle latency.

One key improvement over the P55C, according to sources, is that the context-switching time between MMX and FP modes has been slashed. Few applications today switch modes frequently, but some 3D and other software may require fast switching in the future.

The MMX changes do not affect the overall microarchitecture of the P6. The CPU still decodes up to three x86 instructions per cycle, for example, and translates them into micro-ops. The processor still executes up to five micro-ops per cycle with dynamic reordering. To simplify the register-alias table (RAT), MMX values are physically stored in the FP registers rather than in a separate register file, as in the P55C.

The only other significant change to the core is in the handling of writes to the segment registers. In the original P6 design, these writes cannot execute speculatively; before the

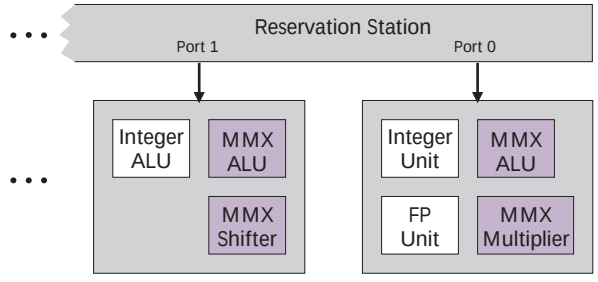

Figure 3. A detailed diagram of the two main function units shows that both have an MMX-capable ALU, but only port 0 can accept MMX multiply instructions, and only port 1 can accept MMX shift instructions.

write can occur, the entire pipeline must be drained to ensure that any instructions referring to the old segment value are completed. The segment write itself, performed via microcode, takes several more cycles, resulting in a total delay of 20–30 cycles.

The Klamath design renames segment registers when they are written, and the time needed to update the value has been reduced. This design allows segment writes to execute speculatively; the new value can be discarded if necessary. In addition, instructions referring to the old and new values can coexist in the out-of-order core, avoiding the need to drain the pipeline and execute sequentially. As a result, Klamath can execute segment writes with essentially no delay.

Although the segment registers are never written in most 32-bit applications, this event is common in older 16-bit code *(see* **091001.PDF***)*. Pentium Pro fares relatively poorly on 16-bit applications and even on Windows 95, due to the amount of 16-bit code in that OS. This problem has made the chip much more popular for Windows NT than Win95, a positioning that has worked well in the business community. Few consumers, however, use NT, and many still use 16-bit software, so the improvement will make Klamath more attractive for home users. Intel estimates the segment-register fix alone increases 16-bit performance by 8–10%.

Circuit Designers Make Impact

Klamath will be built in the same 0.28-micron four-layer-metal CMOS process as the P55C. This process is a variation of Intel's P854 *(see* **090905.PDF***)* and features the same metal layers as the 0.35-micron BiCMOS process used for Pentium and Pentium Pro. For overall clock speed, the smaller transistors of the new process roughly compensate for the lack of fast bipolar transistors. Klamath's clock speed improvement over PPro is mainly due to improved circuit design. In addition, PPro's clock speed is partly limited by the speed of the cache chip; with a faster cache, the PPro CPU might have reached 233 MHz.

The circuit designers also did an excellent job on die compaction, producing a 7.5-million-transistor die that measures 203 mm^2. Despite adding 16K of cache and two MMX units while compensating for the loss of bipolar transistors, the CPU size increased by just 7 mm^2 from Pentium Pro to Klamath. In contrast, when the P55C designers made very similar changes, that chip ballooned by 50 mm^2. In fairness, Intel's method is to ship an initial core design and compact it later. The Pentium core was already quite compressed when the P55C designers got hold of it; the P6 core clearly was not as compact before Klamath.

The chip, shown in Figure 4, uses a new package to fit the signals required by the L2 cache interface. Whereas the Pentium Pro package has 387 pins, Klamath requires 528. To fit these connections into a compact form factor, Intel chose a land-grid array (LGA) package, which is similar to a ball-grid array (BGA) with small pads but no solder balls. The

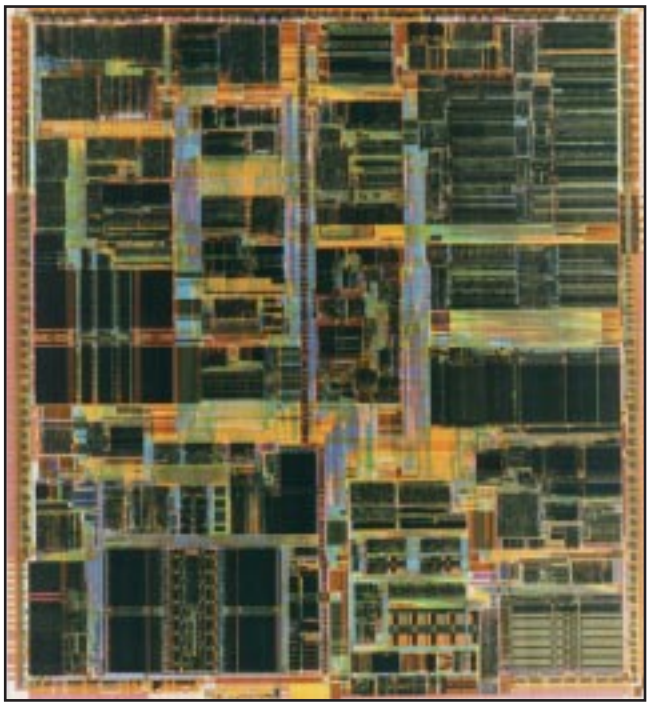

Figure 4. Intel's Klamath has 7.5 million transistors and measures 14.9 x 13.7 mm in a 0.28-micron four-layer-metal CMOS process.

Klamath LGA measures 42 mm on a side, just 36% of the area of the bulky Pentium Pro PGA. The MDR Cost Model estimates the manufacturing cost of Klamath to be about $80. Since Intel plans to sell Klamath mainly on daughtercards, few vendors will deal with this package directly.

Except for a metal heat spreader, the LGA is plastic, lowering cost. For adequate heat dissipation, the spreader can be attached directly to a heat sink or, in the case of Intel's single-edge cartridge (SEC) daughtercard *(see* **1016MSB.PDF***)*, to the protective metal case. Intel would not disclose the power dissipation of Klamath, but the 2.8-V core should provide some relief from the 3.3-V Pentium Pro; we estimate a Klamath-266 will dissipate about 25 W maximum.

Intel revised the Klamath core to reduce operating power by stopping the clock to unused function units and similar techniques. Thus, the typical power dissipation should be much lower than the maximum figure. The vendor also added unspecified features for software power-management, presumably STPCLK and other features consistent with the mobile 486 and Pentium processors.

Preparing for 0.25-Micron Shrink

Although Klamath is unlikely to appear in notebooks, these power-management features will be useful in the next-generation part known as Deschutes. We expect this processor to be essentially a relayout of Klamath into Intel's P856, a 0.25-micron five-layer-metal CMOS process *(see* **101203.PDF***)*. Since P856 is designed to operate at 1.8 V, Deschutes will have a maximum power dissipation of well under 10 W, making it well suited for notebook PCs.

For More Information

Intel has not yet announced pricing, availability, or the actual product name of the processor known as Klamath. For information on Intel's ISSCC paper, contact ISSCC at *www.isscc.org.*

The smaller transistors will allow clock speeds of at least 333 MHz while reducing die size considerably; based on our analysis of routing density, we estimate the Deschutes die will measure about 115 mm^2, shaving manufacturing cost to about $55, roughly the cost of a P55C. We expect Deschutes to be available around the end of this year at 300 MHz, with clock speeds increasing in 1998.

Deschutes may also offer system-bus speeds faster than 66 MHz, the current limit. Intel confirmed that the high-speed Klamath demonstrated at ISSCC ran its system bus at more than 90 MHz, but of course this was merely a demo. For Deschutes, however, the process shrink will enable the CPU to clock the system bus faster, and, by also upgrading its system-logic chip sets, Intel could support bus speeds of up to 100 MHz, probably by mid-1998.

Positioning Klamath Is Tricky

Technically, Klamath is the third version of the P6 core, as the original Pentium Pro started life in 0.5-micron BiCMOS before moving to 0.35-micron BiCMOS. In Intel's product life cycle, however, it is similar to the second-generation Pentium (P54C) in many ways, reducing cost while breaking socket compatibility with its predecessor. PC makers have not invested heavily in Pentium Pro on the desktop, in part because of the upcoming board changes. Once Klamath appears, the P6 will be ready for launch.

For home PC buyers, Klamath offers significantly better CPU performance than a Pentium system as well as the hot new feature, MMX. With higher clock speeds and improved 16-bit performance, Klamath will outperform P55C even on Windows 95 and older applications. By moving to the P6 family now, PC buyers can easily upgrade to future faster processors. Initially, Klamath will be available mainly in high-end PCs costing $2,500 or more, but system prices should ease a bit by Christmas. Intel will use Pentium with MMX (P55C) to fill the lower price points.

For business users, MMX is less interesting, but Klamath still offers compelling performance, particularly on Windows NT, which is becoming popular in the corporate world. But neither Klamath nor NT addresses the needs of mobile users, so businesses must still rely on P55C and Win95 for notebook systems, a dichotomy that won't be resolved before 1998. Klamath's improved prowess on Win95 may speed the acceptance of that OS in the corporate world.

Klamath will not be an immediate hit in two areas: workstations and servers. We expect the half-speed cache will cause significant performance degradation for software that frequently accesses the L2 cache. Many technical applications fit into this category. Server performance is also very dependent on cache and memory bandwidth. In addition, Klamath will not support the four-CPU configuration popular in servers. For these types of systems, PPro will continue to be widely used. But once Klamath (or Deschutes) reaches speeds of 300 MHz and cache sizes in excess of 1M, PPro should quickly disappear.

Still the Fastest x86 Chip

Klamath should keep Intel in the lead in x86 performance. AMD's K6 could be a tough competitor; we expect the chip to deliver clock-for-clock performance similar to Klamath's. The K6 is unlikely to exceed 233 MHz in its 0.35-micron version, however, whereas Klamath should reach 266 MHz, leaving the K6 behind in performance. The two vendors are racing to deploy 0.25-micron versions of their chips; if AMD reaches this mark before Intel, it could grab the performance lead late this year. Whether either chip reaches its speed and performance targets, however, remains to be seen.

Cyrix's best offering against Klamath will be the M2, which should appear at speeds up to 225 MHz. The M2 core is architecturally inferior to the P6 core, however, delivering less performance at the same clock speed. Thus, we expect the M2 to trail the K6 and fall well below Klamath. Both competitive chips will be even further behind on applications that use floating-point or MMX operations, as Intel has placed more emphasis on these areas than AMD or Cyrix. In addition, even with the half-speed cache, Klamath offers superior cache and memory bandwidth, outperforming the competition on applications that stress these areas. AMD and Cyrix counter with more L1 cache, but this improvement is effective only on software with small working sets.

The new Intel design fares less well compared with its RISC brethren. Processors such as the 21164, PA-8200, and 0.25-micron R10000—all due by mid-1997—should offer about 30% better integer performance and three times the floating-point performance of Klamath. But these chips have a higher manufacturing cost and are likely to appear only in expensive servers and workstations. Where price (and x86 compatibility) is not an issue, RISC will still be the preferred solution, but Klamath should appeal to price-conscious buyers who need strong performance.

Within the PC world, Klamath will initially appear as a high-end extension to Pentium Pro, aimed mainly at business users. With its lower cost structure, however, Klamath prices should drop below Pentium Pro's around the end of this year, broadening the P6 family's appeal and drawing in consumers as well as business users. By the end of 1998, Klamath and its direct descendant, Deschutes, will sweep all vestiges of Pentium from the PC market, establishing the P6 as the PC processor of choice. M