

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

AMD Ships Pentium Competitor

5K86 Matches Performance of Pentium-90 at Half the Price

By Linley Gwennap

Aiming to revitalize its microprocessor line, AMD today announced its first Pentium-class product, the AMD5K86. The vendor's first independently designed x86 processor, the 5K86 is based on the K5 core first revealed 18 months ago (see [081401.PDF](#)). The chip is designed to be fully compatible with Intel's P54C Pentium, allowing PC designers to simply drop the 5K86 into existing Pentium sockets. The first members of the family, designated 5K86-P75 and 5K86-P90, will begin shipping this week; AMD plans to ship parts rated as high as P150 before the end of the year.

The company is working on several fronts to achieve this rapid improvement, including a transistor shrink, circuit improvements, and most important, logic redesign. Testing on initial parts identified three key bottlenecks: long-latency instructions, prefetch inefficiencies, and internal bus contention. The next revision of the chip is designed to break these bottlenecks, improving performance significantly.

With 486 volumes waning and prices plummeting, the 5K86 is critical to restoring AMD's microprocessor line to a competitive position. The 5K86-P90 will debut at a list price of \$99, half of Intel's current pricing. Given Intel's projected price cuts, AMD must successfully execute its plan to quickly boost 5K86 performance if it is to return to the realm of three-digit prices.

Roadmap Includes P150 Version

The two new processors use the original K5 design and run at 75 and 90 MHz, respectively. AMD claims these chips deliver slightly better performance than a Pentium of the same clock speed. The P75 and P90 ratings are backed with AMD benchmark data produced in accordance with the recent P-rating specification (see [100202.PDF](#)), but these results have not been independently certified.

With the rollout of the two chips, AMD completed the first phase of its most recent K5 roadmap (see [0913MSB.PDF](#)) on schedule. The vendor had promised "limited" production in 1Q96, with volume parts in the

second quarter. In fact, by shipping a P90 part, AMD is ahead of its plan. The company demonstrated a 100-MHz 5K86-P100 at the recent Cebit conference; it plans to sample this device soon and achieve volume shipments in 3Q96.

Achieving higher performance will require an enhanced version of the K5 core. The company expects this core to match the performance of a Pentium running at a 30% faster clock speed; AMD believes the 90-MHz version of this core will be rated as a P120, and that the 100-MHz part will perform as a P133. With circuit tuning and a process tweak, AMD expects this core to reach at least 120 MHz, resulting in a P150 part before the end of the year.

As Figure 1 shows, Intel plans to rapidly raise the low end of its Pentium line, which it has established at roughly \$100, during 1996. Assuming AMD delivers according to its plan, the 5K86 will stay barely ahead of Intel's rising tide. Of course, AMD can continue selling parts that sink below this waterline, but only for much less than \$100.

We estimate the 5K86 costs about \$70 to build (see [100404.PDF](#)), just under the \$75 list price of the P75 version. This estimate, however, includes the cost of depreciat-

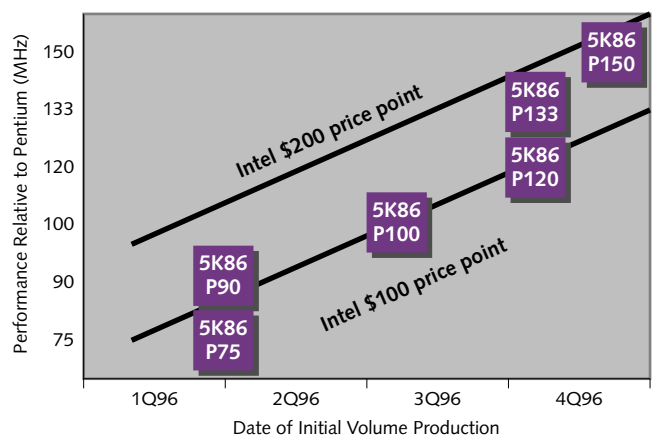


Figure 1. AMD's plan includes processors as fast as a Pentium-150 by year-end, but the company will have problems staying much above Intel's low-end price point. (Source: 5K86 schedule and performance ratings from AMD, Intel price projections by MDR)

ing the fab, in this case AMD's Fab 25. The direct costs (labor and materials) of a 5K86 total only \$45; the other \$25 helps pay for the cost of building and equipping the fab.

With two-digit prices, the 5K86 will generate little profit for the company. But with demand for 486-pinout parts drying up, particularly in the desktop market, AMD has few alternatives for filling Fab 25. Thus, the company may be willing to sell 5K86 parts for less than \$70 if the alternative is to let the fab capacity lie fallow. Accepting a low price simply means there are fewer dollars to pay for the fab.

Performance Enhancements to the K5 Core

Achieving the higher performance points requires enhancements to the basic K5 core. The initial 5K86 parts use the original K5 design with all features fully functional. This version was code-named SSA/5 at one point, but that name has been dropped. According to AMD, this design meets its goal of 30% better performance than Pentium on compiled benchmarks such as SPECint92, although the company has not published optimized SPEC results.

Unfortunately, the chip does not perform as well on typical PC applications. As Figure 2 shows, AMD's Winstone 96 results show the 5K86 is just slightly faster than a Pentium of the same clock speed, despite extensive presilicon performance simulations. To the company's chagrin, the instruction mixes of these simulations did not correlate well with those of large PC applications. Thus, AMD did not find the performance shortfall until it received first silicon last year.

Analysis of the problem led to the discovery of three major performance bottlenecks. The most significant is that certain multicycle instructions, such as far CALL and REP MOVSB, take much longer to execute on the K5 than on Pentium. AMD's designers did not realize these instructions occur fairly frequently in many 16-bit PC applications and thus spent their efforts optimizing other areas. Upon discov-

ering this situation, the designers revised the chip to slash the execution time of these instructions. For example, far CALL is reduced from 15 cycles to 9, versus 4–5 cycles in Pentium.

A second issue concerns prefetching. On an instruction cache miss, the cache is filled one instruction at a time, due to the K5's predecode logic. If the program branches before the end of the line, the processor aborts this slow cache fill and begins fetching from the new target. If the program later returns to the original code stream (as in a subroutine call), an additional bus transaction is required to reload the instructions. In a processor without predecoding, such as Pentium, this transaction isn't needed, since the entire cache line is filled the first time through.

AMD's simulations did not expose this hazard. To fix it, the enhanced K5 adds a small prefetch cache that holds cache lines before they are predecoded. In the situation described above, when the subroutine returns, the original cache line is likely to be in the prefetch cache, eliminating the need for the additional bus transaction.

The third issue is contention for the internal bus that connects the instruction cache, data cache, load/store unit, external bus interface, and other function blocks. This bus is required for accesses between the load/store unit and the data cache as well as for any transactions between the external bus and the caches, causing frequent conflicts that delay program execution. AMD eased this problem by improving the arbitration algorithm and by adding a small side bus for a few critical signals, such as the MESI state bits used for cache coherency.

AMD expects these enhancements, along with other minor changes, to improve Winstone 96 performance to the goal of matching a Pentium of 30% faster clock speed. The company would not comment on the state of the enhanced core, however, nor does it yet have any benchmark data to support its performance claims. Although it is unusual for a chip to improve in performance by such a significant amount with such minor changes, it is certainly possible if the original design contains critical bottlenecks. AMD says it is "confident" the new design will meet its targets.

The die size of the initial 5K86 is 181 mm², slightly larger than originally planned. The company expects that the new changes will not further increase the die size; compacting other areas of the design compensates for the minor additions. In fact, the die may shrink a bit due to this compaction effort.

New Process Raises Yield, Drops Voltage

The initial 5K86 chips are built in AMD's 0.35-micron CS-34 process (see [090905.PDF](#)) and use a 3.3-V supply voltage, the same as Intel's Pentium chips. The transistors in CS-34 are similar to those in Intel's 0.35-micron process P854, but the metal layers are fewer and not as tight, resulting in significantly lower circuit density and a larger die.

AMD is preparing to move the 5K86 design to a modified process called CS-34E. This process reduces the channel

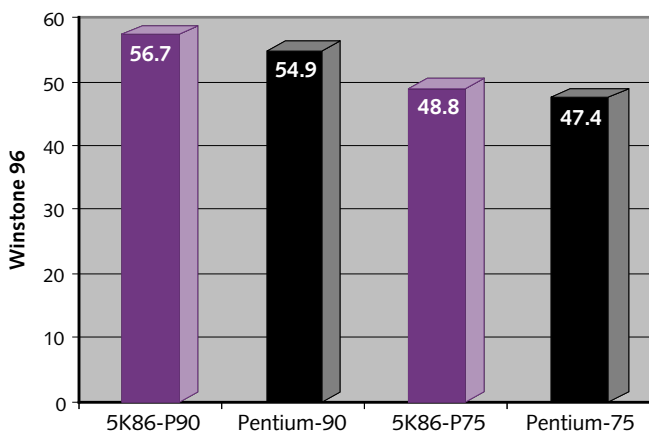


Figure 2. Using Winstone 96 under Windows 95, AMD measured its 5K86 to be slightly faster than a Pentium of the same clock speed. Both chips ran in the same system with a 256K synchronous cache, 16M of EDO DRAM, a VIA Apollo Master chip set, Diamond Stealth 64 graphics card in 640 x 480 x 8 mode, and 1.2G Quantum Fireball disk. (Source: AMD)

length and gate-oxide thickness, making it comparable to Intel's 0.28-micron process, but it retains the same looser metal layers as CS-34. To permit the thinner oxides, AMD plans to reduce the supply voltage to less than 3 V, although the exact value has not been set. Intel plans a similar move to a 2.5-V supply for its 0.28-micron parts, such as the forthcoming P55C Pentium.

The voltage change will require OEMs to modify their designs to supply the lower voltage to these 5K86 parts along with a 3.3-V supply for the pad ring. Because Intel is moving in the same direction with the P55C, a voltage change should not be a problem for AMD as long as it picks the same core voltage as Intel. This may be the reason that AMD has not yet revealed its new voltage level.

The advantage of the modified process is a speedup of the basic transistor, which should improve speed yields at the higher frequencies. AMD expects to produce 100-MHz P100 parts using CS-34 but may need the new process to get adequate yields. The enhanced core includes circuit tuning intended to speed the part as well; thus, it should yield well at 100 MHz even in the older process. By combining the new process with the enhanced core, AMD hopes to reach the 120-MHz level, creating a P150-class product.

Ready for Pentium Sockets

The 5K86 is designed to be fully compatible with all PC software and has passed AMD's extensive compatibility testing. The chip implements all standard x86 instructions, including SMM enhancements and even, according to AMD, the secret Pentium-specific instructions from Appendix H. The chip does not execute the recently announced MMX extensions (see [100301.PDF](#)); AMD plans to include these to its forthcoming K6 processor shortly after Intel delivers its first MMX-based CPUs.

To simplify PC designs, the 5K86 is pin-compatible with the P54C socket in virtually all Pentium PCs today. Unlike the Cyrix 6x86, AMD's part even supports Intel's patented burst order, delivering optimum performance with Intel chip sets. One exception is AMD's lack of support for Intel's APIC pins; because this feature is used only in multi-processor systems, its absence should not affect the 5K86.

AMD is targeting its part solely at desktop PCs, at least for 1996. The 90-MHz 5K86 has a maximum power dissipation of 8.5 W and a typical dissipation of 4.5 W. These figures are about 25% higher than those of a 90-MHz Pentium, although lower than those of Cyrix's 6x86. Even the 5K86-P75, at 3.7 W typical, is 20% hotter than Intel's notebook Pentium chips. AMD has no immediate plans to offer the 5K86 in a TAB or other low-profile package, as Intel does with its Mobile Pentiums. Instead, the 5K86 is available only in a 296-pin PGA.

The 486, including AMD's so-called 5x86, will be a viable notebook processor throughout 1996, although after

Price & Availability

The 5K86 is currently shipping in 75-MHz (P75) and 90-MHz (P90) versions. In 1,000-unit quantities, the P75 lists for \$75, and the P90 lists for \$99. For more information, phone AMD at 800.222.9323 or 408.749.5703, or access the Web at www.amd.com.

midyear it will be restricted to a small portion of the market. We expect AMD to produce a notebook version of the 5K86 around the end of this year to protect its share in the notebook market. The CS-34E version, with its smaller transistors and lower supply voltage, should reduce power dissipation to within the Mobile Pentium range, at least for the lower-speed parts. Some notebook vendors continue to use PGA-packaged processors, but AMD may need to develop an alternative package to succeed in this market.

Even without notebook sales, the desktop market is large enough to absorb AMD's 5K86 parts for at least the next year. AMD says it has the capacity to ship 3 million 5K86 parts by the end of 1996, which would represent about 6% of the Pentium-class market. We expect 5K86 capacity will approach 2 million units per quarter in 1997. Although this figure would be 12% of the Pentium market, it would comprise about half of the low-cost (sub-\$200) desktop portion, a share that AMD would find difficult to achieve.

Thus, for AMD to sell as many chips as it can build in 1997, the company must find a way to break out of the low end of the market; a notebook part would also help. Even with enhancements, the K5 core will not match the performance of Intel's midrange chips. The 5K86's clock speed appears to be 40% lower than Pentium's in comparable processes, offsetting its per-clock performance advantage, and the 5K86 die is about 75% more costly than Pentium. AMD expects its K6 processor, due in 1H97, to address these higher performance points.

Until then, AMD must sell enough 5K86 chips to keep Fab 25 running at a reasonable load. The initial 5K86-P90 pricing is aggressive: \$99 is half of Intel's current Pentium-90 price and 30% less than its projected 4/30 price. Cyrix, in contrast, is currently capacity constrained and is pricing its Pentium-class parts roughly equivalently to Intel's. By offering lower prices, AMD is attempting to rapidly attract enough business to fill its new fab.

We believe the 5K86 has a good chance of meeting its shipment targets. The remaining issues are demonstrating performance comparable to Intel's, avoiding compatibility problems, and ramping production as planned. The ease of dropping AMD's part into existing motherboards and the price advantage over Intel should convince enough PC vendors to adopt the new part to meet AMD's goals. ■