

K6 to Boost AMD's Position in 1997

Derivative from NexGen 686 Targets Klamath Performance



by Michael Slater

After a year's gestation, NexGen's 686 has been reborn as the AMD-K6. With a Pentium bus interface, MMX-compatible multimedia unit, and large on-chip caches, the K6 is the centerpiece of AMD's x86 strategy for 1997. If the K6 meets its goals, it will be a potent competitor for Intel's P55C (see 101404.PDF), Cyrix's M2 (see page 23), and—AMD hopes—Intel's Klamath.

The Nx686 microarchitecture (see 091401.PDF), first disclosed at last year's Microprocessor Forum, hasn't changed much. As Figure 1 shows, the K6 uses a decoupled dispatch/execute design that can decode two x86 instructions into up to four internal instructions, which AMD calls RISC86 instructions (we stick with the more generic term RISC Ops or ROPs for brevity).

Decoding of x86 instructions is assisted by five pre-decode bits per byte, which are stored in the instruction cache as instructions are fetched from memory. There are four instruction decoders: two "short" decoders that can handle the most common x86 instructions and generate one or two ROPs each; one "long" decoder that can produce up to four ROPs representing a complex instruction; and a "vector" decoder that works with a microinstruction ROM to deliver long sequences of ROPs for instructions such as string moves.

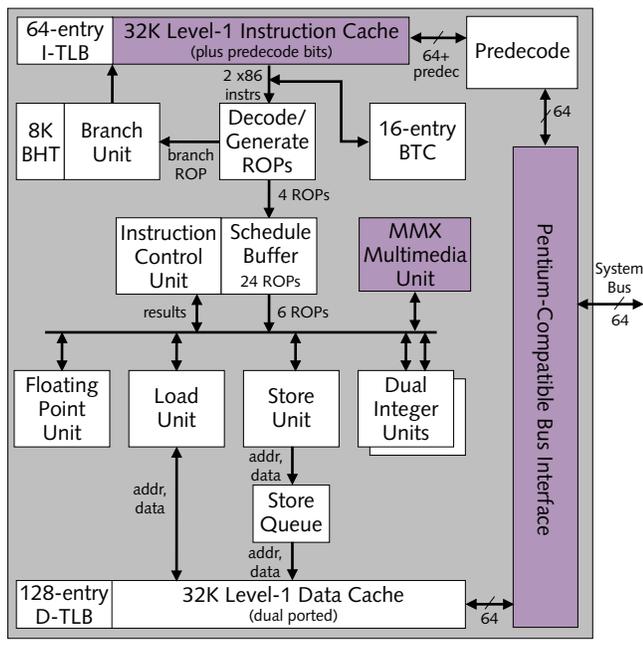


Figure 1. AMD's K6 can decode two instructions per clock cycle, generate four ROPs, and dispatch up to six ROPs from the schedule buffer. Purple blocks are changed from the NexGen 686.

The major changes to the 686 design were increasing the instruction caches to 32K, deleting the L2 cache controller, changing the bus interface and pinout to be Pentium-compatible, and modifying the multimedia unit to be MMX-compatible. The K6's designers also modified NexGen's proprietary System Management Mode (SMM) instructions to match Intel's SMM.

Because the core CPU is largely unmodified from the 686, AMD has been able to test this portion of the design in silicon for more than a year before taping out the K6, minimizing the likelihood of long delays from first silicon to production. The 686 experience is especially valuable because the K6 is entirely different from the K5; unlike Cyrix's M2 and Intel's P55C, the K6 is not a direct derivative of the company's previous-generation chip.

Management, Fabs Shifting

NexGen has been fully absorbed into AMD over the course of this year, from the demise of the Nx586 to shuffling of executives from NexGen into top AMD positions (notably NexGen CEO Atiq Raza to AMD's chief technology officer and NexGen's COO Vinod Dham to VP of AMD's Computation Products Group).

The Austin K5 design team is now focused on the next-generation K7 design. The K6 design is being done by the former NexGen team in San Jose. Another design team in San Jose is already working on the K8.

| | AMD-K5 | AMD-K6 | Intel P55C | Pentium Pro |
|-----------------------|----------------------|-----------------------|-----------------------|---------------------|
| L1 Cache | 16K instr 8K data | 32K instr 32K data | 16K instr 16K data | 8K instr 8K data |
| TLB | 128-entry | 128-entry | 32/64 (I/D) | 32/64 (I/D) |
| MMX? | No | Yes | Yes | Yes |
| Out-of-Order? | Yes | Yes | No | Yes |
| Decode Rate | 1-4 x86 | 2 x86 | 2 x86 | 3 x86 |
| Issue Rate | 4 ROPs | 6 ROPs | 2 x86 | 5 ROPs |
| BTC † | None | 16 entries | None | None |
| BHT Entries | 1,024 | 8,192 | 256 | 512 |
| Return Stack | 16-entry | 16-entry | 4-entry | 4-entry |
| Max Clock | 100 MHz | 180+ MHz | 200 MHz | 200 MHz |
| Voltage | 3.3 V | ~ 2.9 V | 2.8 V | 3.3 V |
| Transistors | 4.3 million | 8.8 million | 4.5 million | 5.5 million |
| IC Process | 0.35µ, 3M | 0.35µ, 5M | 0.28µ, 4M | 0.35µ, 4M |
| Die Size | 181 mm ² | 180 mm ² * | 140 mm ² | 196 mm ² |
| Production | Now | 1H97 | 1H97 | Now |
| Est. Mfg Cost* | \$70 | \$85 | \$60 | \$145† |

Table 1. AMD's K6 takes a big step beyond the K5 and should outperform Intel's P55C. †Includes 256K L2 cache. ‡ A branch target cache (BTC) holds instructions at the target of a branch. (Source: vendors, except *MDR estimates)

Manufacturing has shifted away from IBM Microelectronics to AMD's Fab 25—a huge plant waiting for more chips to build. AMD adapted its fab line to handle the five-layer metal and local interconnect to match IBM's process capabilities, and the former NexGen team adjusted its design rules for AMD's tighter metal pitches. AMD has licensed the C4 solder-bump technology from IBM and is adding this capability to its fab.

Boosting AMD's Prospects

The past year hasn't gone well for AMD's x86 microprocessor business. The K5 finally shipped in the spring, but its poor performance limited it to the low end of the Pentium market and made it hard for AMD to gain market share among major players. The K5 recently moved up to the Pentium-133 performance level (see [1013MSB.PDF](#)), but this still limits it to midrange Pentium systems. The K6, AMD hopes, will dramatically improve the company's position in 1997.

The K6 comes just in time for AMD, with Intel raising the Pentium bar with the P55C and luring the high end of the market to Klamath. At the same time, AMD will be competing with Cyrix's M2 for the attention of companies considering non-Intel processors.

The strength of the K6's position depends on two critical factors—price and performance—that have not yet been disclosed. AMD is expecting clock speeds of around 180 MHz, with performance per clock cycle comparable to or better than Klamath's. This combination would put the chip well ahead of Intel's P55C and in the same range as Cyrix's M2. More detailed comparisons must await the disclosure of measured benchmark results.

Table 1 compares AMD's K5 and K6 with Intel's P55C and Pentium Pro. With twice as much on-chip cache as the P55C, more execution resources, and out-of-order capability, the K6 should easily outperform the Intel chip. To make this meaningful, however, AMD will have to deliver on its clock-speed goals much more effectively than it has with the K5.

At 180 MHz, the K6 should be a potent competitor for the 200-MHz P55C. AMD hopes to scale the K6 clock rate to 225 MHz in 1H97 and, after a process shrink, to 300 MHz, while Intel is likely to limit the P55C to 200 MHz.

Intel's answer to performance-oriented users is not P55C at all, of course, but Pentium Pro and Klamath. When compared with these chips, the K6's position is less clear. It seems likely that the K6 will be in the same performance ballpark as Klamath. After the K6 is shrunk to AMD's next-generation 0.25-micron process, its clock speed could be significantly increased, as could its cache size (mitigating the limits of the Pentium pinout).



Lance Smith of AMD explains the advantages of the K6's decoupled design at Microprocessor Forum.

MICHAEL MUSTACCHI

Price & Availability

AMD has not disclosed price or availability for the K6. Samples are promised for late this year, with production in 1H97. For more information, call AMD at 800.222.9323 or 408.749.5703, or access the Web at www.amd.com.

AMD would not disclose the K6's die size, but sources indicate that it is close to the K5 die size—taking advantage of the additional metal layers and tighter physical design to pack more transistors into the same area. As the K6 moves into more advanced process technology, it will quickly shrink well below the K5 size. As a result, if all goes well with the K6 plan, the K5 will be limited to economy-minded markets by the end of 1997. Power consumption is another parameter that has not been disclosed; AMD expects the chip to be close to today's Pentium thermal envelope.

Intel's P55C can issue two MMX instructions at once, while the K6 is limited to one. As with the P55C, the latency for all MMX ALU and shift operations in the K6 is one clock cycle. For MMX multiply operations, the K6 delivers an impressive single-cycle latency, while Intel has single-cycle throughput but a three-cycle latency. The K6's only multicycle MMX operation is multiply-add, with a two-cycle latency; this operation takes two cycles on Cyrix's M2 or three cycles on Intel's P55C.

For floating-point programs, the K6 should do well; its FP latencies are shorter than Pentium's and even Pentium Pro's, although the FPU is not fully pipelined, so its issue rate is lower. The K6 also can issue an FP and an integer instruction (and a second integer instruction and a load and a store as well) in parallel, while the P55C cannot.

Getting Into the P6 Ranks

The K6 should be a strong competitor to Intel's P55C. AMD, of course, would like to position the K6 against Klamath, but as is the case for Cyrix's M2, this will be a tougher battle to fight (see [1014ED.PDF](#)). In the short term, competing with the P55C isn't a bad position for AMD: the P55C price umbrella should be high enough for AMD to make good margins on the K6 while offering users higher performance for the same price. With plenty of fab capacity available, AMD is well positioned for a comeback in 1997—possibly even rekindling the interest of Compaq and other major PC makers.

In 1998, AMD will need to either move beyond the Pentium pinout (presumably with the K7), or show that the K6 can keep up with P6-family performance, to avoid being relegated again to the low-cost segment of the market. □