

Rise Joins x86 Fray With mP6

New Company Offers Superscalar Design for Basic PCs



by Michael Slater

Rise Technology's long-awaited entry into the x86 microprocessor fray debuted at last month's Microprocessor Forum as the mP6, a three-issue superscalar processor designed for entry-level PCs. The chip is now shipping at clock speeds up to 200 MHz, delivering a claimed performance rating of PR266 (equivalent to a 266-MHz Celeron) with a 100-MHz Socket 7 bus. Rise plans to ship first samples of the mP6 II, which adds a 256K on-chip L2 cache, by the end of the year, with production in 1H99.

Unlike other Intel-alternative processors for entry-level PCs, which typically have been weak in their multimedia performance, Rise hopes to distinguish itself by offering low prices and low power consumption without compromising floating-point and MMX speed.

Rise's three-issue superscalar design strives to execute more instructions per cycle than other designs, but it compromises greatly on clock frequency—the fastest version initially offered runs at only 200 MHz in a 0.25-micron process, the slowest top speed in the industry. It also has only 16K of L1 cache, one-fourth that found in its competitors. As a result, Rise will be limited to the most cost-sensitive customers, at least until the mP6 II ships.

Culmination of Five Years of Effort

The mP6 is perhaps most significant as the product that marks the debut of the fifth x86 processor supplier. Rise was founded in late 1993 by David Lin, who previously managed RISC processor marketing at NEC. After five years of effort, Rise has nearly 100 employees at its Santa Clara headquarters and a small group in a sales office in Taiwan. The company has raised about \$30 million from investors that include investment banks BT Alex Brown and Needham & Co., venture-capital firm Draper Fisher Associates, and unnamed PC makers, chip-set makers, and foundries.

Rise has not disclosed the identity of its foundry. The company does say it will use a foundry with an Intel patent

license. This makes IBM Microelectronics, STMicroelectronics, and Texas Instruments good candidates.

In-Order Three-Issue Design

The mP6 core has three instruction decoders, as Figure 1 shows, enabling a peak rate of three x86 instructions per clock cycle. It is a "native" x86 design; it does not translate x86 instructions into an internal format, as do Intel's Pentium II and AMD's K6.

The ability to decode three instructions per clock cycle represents a radically different approach from low-end competitor IDT, whose WinChip family (see MPR 6/1/98, p. 1) uses a streamlined single-issue design. The three-issue design is clearly an advantage over the otherwise similar two-issue design of Pentium/MMX; compared with Cyrix's M II, it remains to be seen whether the third issue slot is

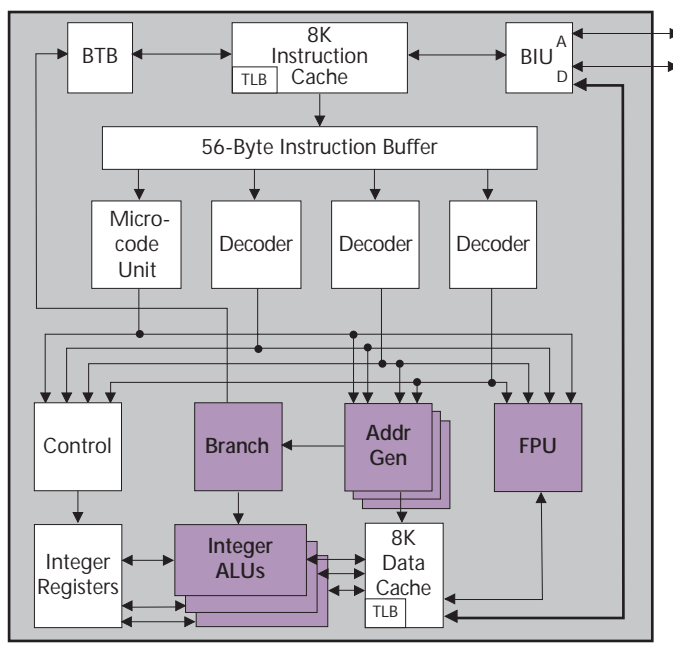


Figure 1. Rise's mP6 is an in-order three-issue design.

more valuable than Cyrix's limited out-of-order capability and larger L1 cache.

Unlike Pentium II, which also can decode three instructions per clock cycle, the mP6 is strictly an in-order design. Furthermore, it has only 8K L1 instruction and data caches (16K total). As a result, it may spend many cycles stalled, waiting for memory to respond.

Unique BTB Drives Autonomous Prefetch

The mP6 implements a unique branch predictor that completely decouples instruction fetch from the execution pipeline. Central to the design is a 512-entry branch-target buffer (BTB) that is indexed with the target address (B in Figure 2) of the last predicted-taken branch (at A). Each BTB entry holds the target address (D) of the next predicted-taken branch (at C), along with an 8-bit byte offset (C-B). This offset locates the address (C) in code-stream #2, where instruction fetch stops; it resumes at address D in stream #3.

This branch-prediction scheme allows the instruction fetch engine to intelligently prefetch instructions down the predicted instruction path (B→C→D→E→F→) independently of instruction execution. Decoupled in this manner, instruction fetch can proceed far ahead of execution, hiding some instruction-cache miss latency. Once the fetch engine is four instruction streams (i.e., four predicted-taken branches) ahead, the mP6 shuts down the BTB to conserve power.

As Figure 3 shows, the instruction-fetch engine occupies stages 1 and 2 of the mP6's eight-stage pipeline. As long as branches are correctly predicted and I-cache miss latency is not too long, the instruction buffer will contain a linear stream of instruction bytes that it can deliver to the instruction decoders in stage 3 of the pipeline.

Unlike most other contemporary x86 designs, the mP6 aligns and decodes x86 instructions in this one pipeline stage. This is an enormous amount of work to perform in a single cycle and could account for the mP6's slow clock rate.

In stage 4 of the pipeline, instructions are paired or tripled, based on instruction types and data dependencies. Stage 5, which is a no-op for reg-op-reg instructions, is used

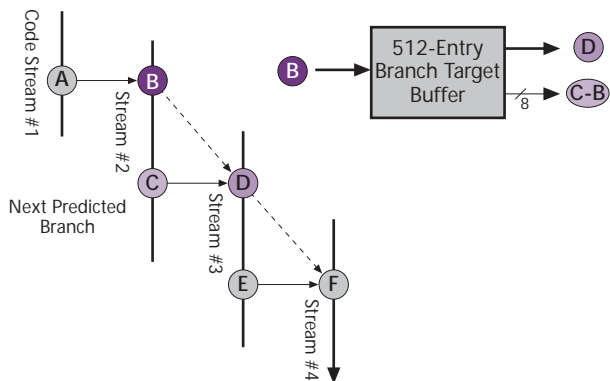


Figure 2. The mP6's unique branch prediction mechanism allows instruction fetch to get far ahead of instruction execution, even across multiple control flow boundaries.

to compute addresses for memory operands. Register and memory operands are read and instructions are issued to the execution units in stage 6.

With reg-op-reg instructions delayed an extra cycle in stage 5, and with the data cache being read a cycle ahead of execution, the mP6 pipeline imposes no load-use penalty. This unusual structure has the benefit of allowing reg-op-reg instructions to be issued in the same cycle as loads on which they are dependent. The only load-use penalty occurs when a load or store is dependent on the previous load or load-op for its address; in these cases, the load-use penalty is two or three cycles, respectively.

In the mP6's scheme, a mispredicted branch causes the instruction buffer to be flushed and the pipeline restarted at stage 1. With the instruction buffer empty, the first instructions out of the instruction-cache in stage 2 bypass the buffer, reducing by one cycle the time required to refill the pipeline. The misprediction penalty for an unconditional jump is five cycles, while a conditional branch incurs a six-cycle penalty behind a load and a six- or seven-cycle penalty behind an ALU operation, depending on pairing.

Three-Input ALU Collapses Dependencies

The mP6 has three integer ALUs, as Figure 4 shows, each with special characteristics. Two ALU operations, along with one move or jump, can be executed in parallel. Each ALU can (with some restrictions) be assigned to whichever pipeline needs it, providing more flexibility than designs such as Pentium and the M II, where execution resources are permanently associated with particular pipelines and thus with instruction positions within an issue packet.

One ALU has a multiply/divide/shift unit. The second ALU is a three-input design that allows some dependencies to be eliminated. Consider, for example, the instruction sequence:

$$AX \leftarrow AX + BX$$

$$DX \leftarrow DX + AX$$

Because the second instruction depends on the result of the first, the two cannot be executed in the same clock cycle. Using the three-input ALU, however, this series can be transformed into a parallel instruction pair:

$$AX \leftarrow AX + BX$$




Figure 3. The mP6's eight-stage pipeline has no load-use penalty (for cache hits), allowing reg-op-reg instructions to be paired with loads on which they are dependent.

$$DX \leftarrow DX + AX + BX$$

Thus, a pair of instructions that would have been forced to issue sequentially can now be issued together.

The third ALU has only one input; it passes the data for a move operation and checks condition codes for jumps.

Integer multiply completes in two cycles for byte operands, three cycles for 16-bit operands, and five cycles for 32-bit operands. Divide latency is 6, 10, or 18 cycles, depending on the operand size. Neither operation is pipelined.

High-Performance MMX, FP Units

Figure 5 shows the mP6's MMX and FP units. Up to three MMX instructions per clock cycle can be executed by the seven execution units. All MMX instructions execute in a single cycle, except for multiply and multiply/add, which are fully pipelined but take two cycles to complete. The MMX units are also able to collapse some dependencies.

This MMX unit allows more parallelism than other designs, in many cases, but its performance will be highly dependent on the instruction mix. Compared with the benchmark results of other Socket 7 processors, the mP6's are stronger (at a given speed) for multimedia and imaging applications than for business applications. It is slower than Celeron-266 on these benchmarks, however, possibly because of its small L1 caches and lack of a fast L2 cache.

The four-stage FP pipeline can execute one instruction per clock, with the now-common optimization of pairing an FXCH instruction with an FP operation. This puts it on a par with Intel's Pentium II and ahead of the K6 and M II, neither of which has a pipelined FPU.

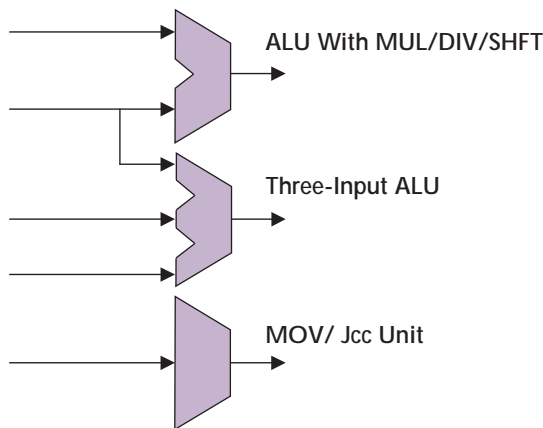


Figure 4. The mP6 has three integer ALUs. One ALU has three inputs, enabling some dependencies to be collapsed.



MICHAEL MUSTACCHI

Ken Munsen, Rise's chief architect, describes the company's new mP6 processor at the Forum.

Costs and Benefits of Complex Design

Rise's three-issue design may deliver good performance at a given speed, but the additional complexity has a cost in both reduced clock speed and larger die size. Initially, Rise is offering three speed grades of the mP6, with performance ratings of 266, 233, and 166 and clock speeds of 200, 190, and 166 MHz. As with Cyrix's M II, Rise is using "PR" performance ratings in place of clock speed to specify the chips. For example, the mP6-266 has a 200-MHz clock, but Rise rates it as delivering performance equivalent to a Pentium/MMX-266 (for mobile) or Celeron-266 (for desktops). In each case, the bus runs at half the processor speed: 100, 95, or 83 MHz. The 166-MHz part is aimed not at PCs but at products such as Windows terminals and set-top boxes.

Rise's speed grades represent the bottom of the desktop PC market today; Rise's top grade, PR266, is the slowest grade AMD and Cyrix offer, and Intel has all but abandoned it. In notebooks, which lag far behind desktops both in clock speed and in shifting away from Socket 7, Rise's chips are more competitive; mobile processors today top out at 300 MHz, while Rise reaches PR266 (but note that this rating is relative to Pentium/MMX, not Pentium II).

The 3.6-million-transistor mP6 die is surprisingly large for a chip with only 16K of cache: at 107 mm² in 0.25-micron five-layer-metal process technology, it is larger than all other Socket 7 processors and nearly twice the size of IDT's WinChip 2. Rise asserts that its costs will nonetheless be competitive because of its close relationship with the foundry and its low-cost packaging.

Power Suitable for Notebooks

For notebooks, the mP6 includes Intel-compatible system-management mode (SMM), as well as the other standard power-management states. Rise's designers also took pains to minimize power consumption, using techniques that are common today but that Rise believes it has taken further than prior designs. For example, selective clocking shuts off




Figure 5. The mP6's MMX unit can issue three MMX instructions per cycle to the seven execution units. There is only one FP pipeline, but FXCH instructions can be paired with FP operations.

Price & Availability

The mP6 is in production now. The mP6-266, -233, and -166 are priced at \$70, \$60, and \$50, in thousands. The mP6 II will be sampling to early customers by year-end, with volume production on 1H99. Prices and clock frequencies have not been disclosed. Contact Rise at 408.330.8800, fax 408.330.8900, or www.rise.com.

blocks, such as the FPU, not being used in a given cycle; operands at the input of the multiplier do not switch unless the multiplier is actually being used; and the BTB is not accessed on cycles when its results are not needed.

The 200-MHz (PR266) mP6 dissipates 8.4 W maximum, similar to AMD's mobile K6-266 and IDT's WinChip 2 and in the acceptable range for notebooks, unlike Cyrix's M II. The CPU core runs at 2.7 V, rather high for a 0.25-micron process, with 3.3-V I/O.

Rise packages the mP6 in a thermally enhanced 387-contact "Turbo Thermal" BGA that uses a pin arrangement similar to that of the traditional Socket 7 PGA to simplify board redesigns. Rise also offers the mP6 in a PGA format that is actually a BGA device mounted on an interposer board with PGA pins.

Like other Intel-compatible processor makers, Rise has omitted features not used in the PC environment. The mP6 does not implement Intel's APIC (multiprocessor interrupt controller) or master/checker (redundant pair) capability.

Rise has tested its processor with chip sets from Intel, ALi, SiS, Utron, and VIA. Both AMI and Award offer BIOS

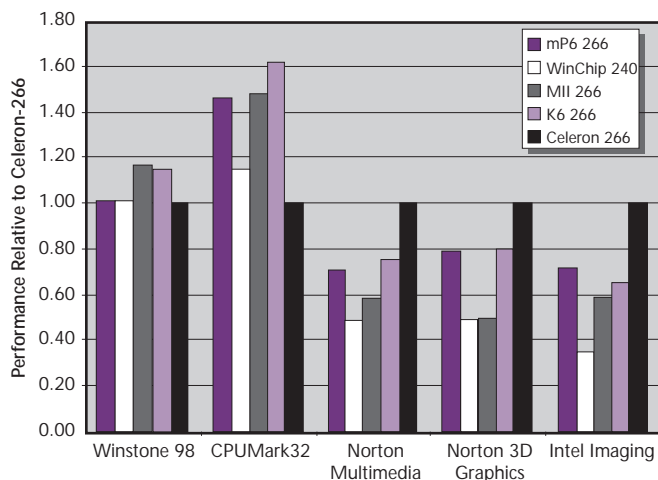


Figure 6. The mP6 fares well against Celeron on Winstone, while falling below Celeron but ahead of AMD and Cyrix on multimedia and FP benchmarks. All the Socket 7 systems have a 1M L2 cache, while the Celeron-266 has none. All tests were run with 64M of SDRAM, a Quantum Fireball 2.1G UDMA hard drive, and ATI Rage-Pro AGP video card. The Socket 7 systems use a VIA MVP3 chip set, while the Celeron system used an Intel 440EX. (Source: Rise)

support. XXCAL, an independent testing lab, has certified the chip's x86 compatibility. In response to the CPUID instruction, the mP6 reports a manufacturer code of "RiseRiseRise" (instead of "GenuineIntel") and a family ID of 5 (same as P55C).

The Bottom Line: Performance

Rise claims the mP6 core is more than 15% faster than Pentium II's CPU core (ignoring cache and bus effects) at a given clock speed, but this claim is based on a single, synthetic benchmark (WinTune). Because it fits in the L1 cache, it eliminates the effects of L2 cache and bus performance. It is also likely that the instruction mix is not representative of real applications. In any case, since the mP6's top clock speed is less than half of Pentium II's top speed, it is far behind in delivered performance.

Figure 6 shows Rise's benchmark results for the mP6. This graph compares 266-MHz or PR266 versions of each processor, except WinChip, which tops out at 240 MHz. On Winstone 98, the mP6 matches the performance of WinChip and Celeron but falls about 15% behind the K6 and M II. On multimedia and 3D benchmarks, it delivers only 70–80% of Celeron's scores, is comparable to the K6, and consistently leads the M II and WinChip.

mP6 II Adds On-Chip L2

The mP6's small L1 cache size presumably was dictated by the need to keep the die size down, given a fairly large core CPU. To remedy this weakness, Rise will follow quickly with the mP6 II, which adds a 256K on-chip L2 cache. Rise's Ken Munson said the mP6 II's L2 cache latency will be only four cycles for instructions and three cycles for data, half the latency of Mendocino's L2 cache.

In the same 0.25-micron technology as the mP6, we estimate the mP6 II would be about 170 mm²—very large for the low-end market. Rise plans to move this design to a 0.18-micron process in mid-1999; it is in this process that the chip may have its chance to shine.

Rise expects the mP6 II to dissipate about the same amount of power as the mP6; although it has the additional power consumption from the L2 cache, it saves power by reducing off-chip accesses. Thus, the mP6 II should be suitable for notebook computers. In a 0.18-micron process, it should be able to achieve substantially higher clock rates, and by running the core at 2.0 V, Rise expects to keep the power dissipation below 10 W.

Battle for the Entry Level

With its initial mP6, Rise aims to offer a higher-performance alternative to IDT's WinChip 2 at the low end of the PC market. With most of the desktop PC market focused on speed grades of 300 MHz and up, Rise is starting out in a domain below most of the market served by AMD, Cyrix, and Intel. Only in the notebook arena, where clock speeds are lower, does Rise's first chip offer much appeal—but even there, it

competes with Pentium/MMX and AMD's K6, not with Pentium II.

The mP6 II, if tuned for speed and built in a faster process, could elevate Rise's position. Rise hopes to compete against Intel's Mendocino with the mP6 II. With an on-chip L2 cache, the core will not be held back so much by the small L1 caches. Whether this chip actually boosts Rise's relative market position or just enables it to keep pace remains to be seen. We expect nearly all of Intel's processors to have on-chip L2 caches and run at speeds of 400 MHz and up in 1H00, placing the bar much higher than it is today.

The primary competition for Rise will not be Intel's processors, however, but those from AMD, Cyrix, and IDT. Table 1 compares the various Socket 7 processors that are scheduled to ship by 1Q99. Rise has staked its positioning on the mP6's ability to decode and execute three x86 instructions per clock cycle—a feat none of the other Socket 7 processors matches. Only Intel's Pentium II does this today. It is unclear, however, that this distinction will mean much. The Rise design makes a number of other tradeoffs, such as in-order execution and small L1 caches, that may negate the benefit of the three-issue design. Furthermore, it is difficult to evaluate how Rise's unusual pipeline design will perform on a range of real applications.

Startup Faces Many Challenges

Rise pitches the mP6's faster MMX performance as a key advantage over Cyrix's M II, AMD's K6, and IDT's WinChip, potentially enabling more effective soft modems or allowing soft DVD to be implemented in less-expensive systems. This advantage is likely to be fleeting, however. Next spring, Cyrix will introduce the M II+ (code-named Jedi), a Socket 7 version of its Cayenne CPU core that will upgrade the M II line with faster FP and MMX. IDT's WinChip 2 already offers a big boost over the original WinChip, and AMD's K6-2 improves on the K6's MMX performance. The mP6 is far behind even the slowest Celeron on multimedia benchmarks.

The mP6 lacks 3DNow, which could weaken its position in the performance-focused game market—but the chip doesn't have the speed for that market anyway. Rise plans to skip 3DNow and implement Katmai New Instructions in future processors.

Rise hopes to leverage the mP6's modest power consumption into success in the notebook market. Most of the notebook business is controlled by a few large OEMs that are unaccustomed to non-Intel processors, but they may be eager to escape Intel's grip. There is considerable interest in

	Rise mP6	Cyrix M II	AMD		IDT WinChip 2
			K6	K6-2	
Max Issue Rate (x86 Instructions)	3 instr	2 instr	2 instr	2 instr	1 instr
MMX Issue Rate	3 MMX	1 MMX	1 MMX	2 MMX	2 MMX
Pipeline Stages	8 stages	7 stages	6 stages	6 stages	6 stages
L1 Cache (I/D)	8K / 8K	64K unified	32K / 32K	32K / 32K	32K / 32K
3D Extensions	None	None	None	3DNow	3DNow
Pipelined FP	Yes	No	No	No	Yes
Max Clock Speed	200 MHz (PR266)	250 MHz (PR333)	300 MHz	350 MHz	266 MHz
Core Voltage	2.7 V	1.8 V	2.2 V	2.2 V	3.3 V
Power at PR266	8.2 W	24 W	9.8 W	14.7 W	12.5 W
Transistors	3.6 million	6.5 million	8.8 million	9.3 million	6 million
Die Size	107 mm ²	88 mm ²	68 mm ²	81 mm ²	58 mm ²
Process	0.25μ ?M	0.25μ 5M	0.25μ 5M	0.25μ 5M	0.25μ 4M
Est Mfg Cost*	\$45	\$35	\$30	\$35	\$25
Price (1000s)	\$70 (PR266)	\$79 (PR266)	\$159 (mobile 266)	\$136 (333 MHz)	\$51 (240 MHz)

Table 1. The mP6 will compete primarily with other Socket 7 processors from AMD, Cyrix, and IDT. It is a wider-issue machine but has a larger die size and lower clock rate. (Source: vendors, except *MDR estimates)

driving down notebook prices, however, and Rise's chip could be attractive in sub-\$1,000 notebooks.

Rise will share Cyrix's burden in positioning its processors on the basis of performance rather than clock speed. While Cyrix has demonstrated that this is possible, it isn't easy—and Cyrix is striving to make sure that its next-generation Jalapeno core (see MPR 11/16/98, p. 24) will deliver the raw clock speed to get it out of the PR game. The PC market is fixated on clock speed, despite its lack of accuracy as a performance metric. Clock speed, at least, is a simple, objective measurement; PR ratings leave a lot open to interpretation. Ultimately, it is delivered performance that should matter, but it is harder to sell that performance without the clock speed to match.

Rise also takes Cyrix's former position as the only fabless x86 supplier. The lack of an intimate coupling with a foundry, and the corresponding lack of leading-edge process technology, made it difficult for Cyrix (or NexGen, in its day) to compete. The natural desire for fab independence tends to create large chips that don't fully exploit the potential of any particular process. Rise will have to develop better fab resources and execute better than its predecessors if it is going to thrive as a fabless supplier. If Rise is successful, it will become an attractive acquisition target for a company that owns a fab and an Intel patent license.

With the mP6, Rise has gained entry into the PC processor business. Now it must follow quickly with the mP6 II and increase its clock speeds dramatically if it is to gain much ground. If Rise succeeds only in holding its position relative to the rest of the processor landscape, it will find the going tough, with many competitors and thin profit margins. The company's future depends on its ability to improve its performance faster than its competitors improve theirs—a daunting challenge in today's competitive x86 business. 