

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 9 NUMBER 2

FEBRUARY 16, 1995

P6 Underscores Intel's Lead

Sets New x86 Performance Standard, Thrusts into Server Market

by Linley Gwennap

While its competitors struggle to complete their Pentium-class chips, Intel is well on the way to delivering its next-generation P6 processor. Looking under the hood (see page 9), we see that the microarchitecture of the P6 is not that different from the design of AMD's K5 or even NexGen's 586. From the user's perspective, though, the difference is clear: the first P6 should outperform its competitors by about 50% on typical applications, based on initial performance estimates.

Intel estimates that the P6 will deliver 200 SPECint92 at 133 MHz, the target frequency of the initial implementation, although it has not yet announced any specific P6 products. Intel will build the first P6 parts in the same 0.5-micron BiCMOS process (which Intel calls a 0.6-micron process) as current Pentium chips. The company is now testing first silicon, which runs Windows and many application programs. We expect the first P6 systems to begin shipping in volume in 3Q95.

The first P6 processors will be expensive and power-hungry (sound familiar?), making them suited only for high-end servers and very expensive PCs. A shrink to a 0.35-micron process, expected in 1H96, should ease these problems and boost the clock speed by as much as 50%. This revised P6 will appear in premium PCs in 1996 and plunge into the mainstream in 1997. By 1998, a multitude of P6 processors will be available, including a 0.25-micron version that will further reduce costs and increase performance.

The P6 supports a second strategic effort: increasing Intel's penetration of high-end servers. While the company has had much success in the mass market for PC servers, the more expensive systems that are replacing old minicomputers and mainframes are typically based on RISC processors. Intel's new device implements a revamped memory subsystem and a new system bus, improving server performance. The new bus allows up to four P6 processors to be easily connected, reducing the cost of building MP servers.

Accelerating the Generational Pace

The P6 is the first fruit of Intel's effort to double its pace of processor generations. In the past, new x86 processor cores have rolled out every four years or so from what was essentially a single design team at Intel's main facility in Santa Clara (California). The P6 is the first x86 processor core from the company's Hillsboro (Oregon) team, which had previously focused on i960 chips. The P6 project began in late 1991, several months before the first tape out of the Pentium processor.

The decision to start a second x86 design center shows excellent foresight. Without the P6, Intel's high end throughout 1996 would have been limited to 150-MHz Pentium chips. In an increasingly competitive market, Intel could have lost the x86 performance lead. With the P6, the company should retain its hold on the high end in 1996 and probably beyond.

Shortening the generation gap to two years, however, does not accelerate Intel's overall performance pace. Integer performance of the x86 increased at a rate of 50% per year between the first 486 and the first Pentium (see [090103.PDF](#)); from Pentium to the P6, performance has also jumped by 50% per year.

With the P6, Intel may briefly surpass the performance of PA-RISC and MIPS if the first P6 systems ship before the end of this year. But as Figure 1 shows, Digital is about to ship a 330-SPECint92 processor, and all the major RISC vendors expect their next-generation efforts to surpass the performance of the initial P6, in many cases by 50% to 100%.

With Intel's 0.35-micron process, which is currently running test wafers in Hillsboro, a 200-MHz P6 processor is achievable in 1H96, delivering perhaps 280 SPECint92. This product will start to pressure the weaker RISCs. Intel, however, has not committed to delivering a faster P6 before the end of 1996.

Outside of the x86 market, the most important competitor for the P6 will be the PowerPC 620, which also happens to be the slowest of the next-generation RISC

processors. The 620 is slated to deliver 225 SPECint92, a mere 12% better than the P6. Like the P6, the 620 operates at 133 MHz and is expected to appear in systems in 3Q95. Pricing is not yet available for either chip, but the PowerPC vendors will be hard-pressed to deliver on their promise of doubling x86 price/performance with the 620.

Instead, IBM and Motorola hope to counter the P6 by improving the clock speed of their less expensive 604. A planned process shrink could boost performance from 160 SPECint92 to well over 200 while keeping the cost of the chip relatively low. Thus, PowerPC could deliver P6 performance levels at half the price of a P6, but the new Intel processor will make it difficult for PowerPC to double high-end x86 performance in the foreseeable future.

Threatening the Lucrative Server Market

With software barriers continuing to keep RISC processors out of the mainstream PC market, the most significant contention between the two camps will occur in the server arena. The x86 already dominates the market for PC servers, which are typically based on standard PC technology and provide a shared network resource. With the P6, Intel is moving upscale.

From a system-design standpoint, a server is distinct from a PC in memory and I/O capacity, and often uses multiple processors to achieve higher performance on large workloads. In the past, these systems used proprietary CISC processors and were called minicomputers and mainframes. Today, most of these legacy servers are being converted to systems using mainstream RISC processors. These systems support dozens or hundreds of users and carry an average selling price of tens or even hundreds of thousands of dollars, making this a lucrative market when measured in revenue, if not in units.




Figure 1. Next-generation processors (shown by stars) from RISC vendors should surpass the performance of Intel's P6, although the P6 may briefly exceed the integer performance of the leading MIPS and PA-RISC chips. (Source: vendors, MDR estimates)

The relatively small unit volume of this market has made it unattractive to Intel, which has concentrated on its bread and butter, the PC market. As a result, the company holds just 14% of the server market when measured by system revenue (according to Dataquest). But like a skilled despot, Intel has realized that it must crush all opposition; tolerating the existence of RISC servers gives competitors a base from which to attack Intel's stronghold on the desktop.

The P6 is the next step in the company's plan to encircle its RISC competitors, forcing them to wage war on all fronts. The first step was the inclusion of a much-improved floating-point unit in Pentium, allowing that processor to compete with RISC systems in the volume portion of the workstation market. By tuning the P6 for servers and retaining a fast FPU design, Intel leaves its RISC competitors nowhere to hide.

Sun, for example, sold \$4.6 billion of workstations and servers in 1994; the company reinvests a portion of this money to develop and promote SPARC. With this revenue stream threatened by Pentium and P6 systems, it will be more difficult for Sun to make incursions into the mainstream desktop market. It is too early to tell how this strategy is playing out, but companies such as Sun and Digital appear vulnerable. PowerPC's base of support from Apple and IBM is more resistant to harm.

The key to the P6's ability to compete for server wins is its new system bus, which provides 528 Mbytes/s of sustainable bandwidth and the ability to connect up to four processors to a single system-logic chip set without any glue logic. These features match those of next-generation RISCs such as HP's PA-8000 and the MIPS R10000. Only the PowerPC 620, with its 128-bit system bus, offers significantly better bus bandwidth than the P6. Intel hopes that the lower cost of support logic, plus compatibility with x86 software, will push server vendors toward the P6.

Improving PC Standards

While the server market may feel the earliest impact from P6, Intel's primary focus remains the PC market. In that arena, the P6 must hold the fort against x86 processors from a variety of vendors. The P6 is crucial because a number of these vendors will have Pentium-class processors shipping by the end of this year, forcing Intel to move to a new level to stay ahead.

The P6 does just that. Figure 2 shows the projected performance of the new Intel chip compared with that of expected high-end x86 processors from AMD, Cyrix, and NexGen. With its higher clock speeds and fast level-two cache, the P6 looks as if it will remain well ahead of most of the x86 competition throughout 1996.

Of these competitors, NexGen has the most ambitious plans. It hopes to ship its 686 device by the end of this year, offering performance similar to that of the ini-

tial P6. NexGen has established credibility by being the first competitor to bring a Pentium-class product to market and by matching the performance of Intel's current top of the line. But even if NexGen can deliver on its 686 promises, the tiny startup lacks the stature and system compatibility to take significant business from Intel.

AMD is gearing up a new fab to propel its K5 processor into the Pentium-class market. But the company does not expect its K6 design to reach volume production until 1H97 (see **090101.PDF**). At least until then, the K86 family will lag the P6 in performance. Likewise, Cyrix's M1 will also fall well behind the P6. Cyrix has not disclosed plans for a P6-class processor.

P6 Prices Start High, Go Low

This performance advantage will allow Intel to charge premium prices for P6 products. As the K5 and M1 reach their stride in 1996, however, they will begin to threaten Intel's share of the Pentium-class market. In fact, AMD's planned 150-MHz K5, due in 1H96, should offer better performance than Intel's fastest Pentium. At that time, Intel will be forced to choose between maintaining high P6 margins or dropping prices to compete with the souped-up K5.

In keeping with its history, as well as its "crush the opposition" strategy, Intel is likely to take the latter course. The 0.35-micron P6 will enable the company to drop P6 prices without losing too much margin. A 166-MHz P6 will offer a significant performance advantage over a 150-MHz K5, leaving room on the price curve for a premium 200-MHz P6. These lower prices will make P6 the preferred choice for PC power users by the end of 1996. As prices continue to fall in 1997, the P6 will head toward the mainstream, reaching annual sales of about 25 million in 1998 (according to CI InfoCorp projections).

Figure 3 gives a closer look at our projections for P6 price and availability. Intel intends, at least initially, to market the P6 CPU with an integrated 256K level-two cache, making it more difficult to compare the price of the two-chip P6 set with standalone Pentium CPUs. For comparison, the figure shows the approximate price of a Pentium CPU with 256K of fast secondary cache.

We estimate that Intel will introduce the 133-MHz P6 at a 1,000-unit list price of roughly \$1,400. Intel will also market a lower-speed version; given the narrow frequency yield of Intel's other BiCMOS products, this version will probably be 120 MHz. We expect this product to debut at about \$1,100. The relatively large difference represents the price premium Intel often charges for its fastest chip (e.g., 100-MHz Pentium) and provides a buffer in case yields are low at the top speed.

We expect prices to ease gently until 2Q96, when both AMD's K5 and the P6 move to 0.35-micron processes. As noted, this combination of events will motivate Intel to price the P6 more aggressively. The new




Figure 2. Intel should retain its performance lead over x86 competitors due from AMD, Cyrix, and NexGen. (Source: MDR estimates)

process should allow 166-MHz and 200-MHz speed grades, with the price of the former dropping quickly while the latter retains a premium price.

As a side effect, this pricing policy will ultimately eliminate the 0.5-micron P6 from the market, since it will have few PC design wins at that point. As with Pentium, the second-edition P6 will have a lower manufacturing cost than the original version, motivating Intel to move quickly to the new process. By the end of 1996, low-end P6 prices could drop below \$700.

Integrated Secondary Cache

The P6 CPU and its second-level cache are combined in a single PGA package. Initially, the P6 will ship with a 256K cache, but the company is already working on a 512K cache for the part, which will be built using




Figure 3. MicroDesign Resources projects that the P6 will carry a high initial price, but steep cuts should occur in 1H96. For comparison, the dashed line shows the price of a Pentium plus 256K of cache; the P6 prices include 256K or 512K of cache.

Price & Availability

Intel has not yet announced price and availability for the P6, which we expect to begin shipping in 3Q95. For more information, contact your local Intel sales office or access the Worldwide Web at <http://www.intel.com>.

the 0.35-micron process. This version is likely to be used with the faster, 0.35-micron P6, but it is possible that Intel could mix and match. A slower P6 with the larger cache would be attractive for servers, while a faster CPU with the small cache would be less expensive for PCs.

For price-sensitive PCs, Intel may ultimately deploy a version of the P6 without the cache chip. System vendors could then add a secondary cache using external SRAM, as in today's PCs. This configuration, however, would have lower performance than the integrated cache, due to the extra cycles needed to access the external SRAM. Furthermore, this system may not be any less expensive to build than one with an integrated cache.

A more interesting possibility is for Intel to expand the first-level caches on the CPU chip, which would be possible with the 0.35-micron process, and eliminate the secondary cache bus, as in the MIPS R4400PC. This configuration would also have lower performance but would clearly be less expensive than the integrated cache.

Intel is motivated to retain the integrated cache, which boosts the P6 price (and thus Intel's revenues) and helps to fill its enormous fabs. For system vendors, however, this configuration eliminates the ability to differentiate their products based on the size and speed of the L2 cache. It also continues the trend of giving more and more of the system revenue to Intel. System vendors must also design new motherboards for the P6 using, at least initially, Intel system-logic chip sets.

In early 1998, at about the time the first fruits of its "P86" alliance with HP appear, Intel should move the P6 to a 0.25-micron process, further reducing cost and increasing clock speeds, perhaps to as high as 300 MHz. Even at this speed, the P6 will be outrun by the first P86 chips, but the 0.25-micron P6 will play a much bigger

role initially. For Intel to remain competitive against AMD and others, this version must bring P6 fully into the PC mainstream and push Pentium into the ashcan of history.

P6 May Adopt Pentium Name

We expect Intel to announce P6 products this fall, when P6 systems are ready to ship. At that time, the company will reveal the exact speed grades, along with their initial price and availability. More detailed performance measurements should then become available. Intel will also announce details of its first system-logic chip set for the P6.

Another unresolved issue is the new product's name. P6 is simply the internal code name given to the project. When the P5 was originally announced as Pentium, many joked that the P6 would be marketed as "Sexium." Since then, Intel and its OEMs have spent hundreds of millions of dollars to promote the Pentium brand name, both in print and through incessant television ads. In a perverse way, the \$475 million FDIV bug was the final act needed to cement Pentium into the consumer consciousness. Intel hopes that these consumers will remember that the company stands behind its Pentium products and forget the earlier waffling.

Thus, it would not be surprising if Intel applies the Pentium name to the P6 as well. To avoid confusion between the P6 and the original Pentium (which will overlap in clock speeds) and between the P6 and the Pentium competitors, Intel would need some additional appellation for P6 processors: perhaps Pentium Plus, Super Pentium, or in the grand Intel tradition, Pentium DX.

By any name, the P6 will still provide Intel with the weapon it needs to prevent AMD and others from encroaching too far onto its hallowed turf. By the time these competitors can grab a significant share of the Pentium market, Intel will already be moving the PC community to the P6. But to accomplish this task, Intel will have to price its new processor aggressively, forgoing the high margins it has traditionally held. The integrated cache is an interesting and innovative strategy to increase revenue in the face of falling prices; we must wait to see how well it is accepted by system vendors. ♦