MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 8 NUMBER 3 MARCH 7, 1994

Intel Extends 486, Pentium Families

100-MHz DX4, "P54C" Processors Use 0.6-Micron BiCMOS Technology

by Linley Gwennap

Repositioning virtually its entire product line to outflank its touted competitors, Intel has announced a new line of 486 chips, known as the DX4, and faster clock rates for Pentium based on a new design code-named P54C. Both new chips use a new 0.6-micron BiCMOS process to achieve clock frequencies of up to 100 MHz, a 50% improvement over the previous high-end parts. For Pentium in particular, the new process will also significantly decrease the power dissipation and manufacturing cost from the original 0.8-micron parts.

The DX4 reopens the performance gap between Intel's 486 chips and those of its competitors. Intel can now slash the price of its DX2 processors, forcing other x86 vendors to match the new prices and accept lower margins. With the unique DX4 parts, however, Intel can maintain its traditional high margins. For end users, these changes will translate into higher performance at all system price points.

The P54C design is everything that the original P5 Pentium was supposed to be: a 100-MHz processor with reasonable power dissipation and moderate manufacturing cost. It includes a few minor enhancements: an interrupt controller, power management, and a clock multiplier. The new chip is not pin-compatible with the old version and requires 3.3-V system logic. With its 100-SPECint92 performance, the P54C is a potent weapon against PowerPC and other RISC processors. The cost reductions also will allow Intel to cut the price of Pentium and eventually move it into the PC mainstream.

Both of the new parts have been expected for some time *(see **070401.PDF**)*; Intel's development and construction of several 0.6-micron manufacturing lines have been in progress for years, and the benefits of moving both the bread-and-butter 486 family and the next-generation Pentium family to the new process are clear. The new Pentium products are fairly similar to the current chips, but Intel has thrown a few curves with its DX4 announcement.

A 486 by Any Other Name…

The first surprise is the name, officially IntelDX4 (no space). Intel has dropped the "486" designation for these parts; it argues that other vendors have destroyed the meaning of this appellation by selling 486 chips that do not deliver the same level of performance as the original. Left unspoken is the fact that the new name can be trademarked, while "486" cannot. We are confident, however, that other vendors will be quite happy to call their versions 486DX4—when they have such a part.

The DX4 name seems to imply clock-quadrupling, but it does not; the new parts allow the internal clock to run at 2×, 2.5× or 3×the external clock rate but not at 4×, as Table 1 shows (see below).

Another new feature is a 16K on-chip cache, twice that of a 486DX. Since the new process enables a much more compact design, Intel was able to double the cache size and still end up with a smaller chip. Although Pentium implements a write-back cache, the DX4 uses the same, less-efficient write-through policy as previous 486 parts; the company says that adding write-back capability to the part would have delayed its introduction. The larger cache helps offset the extra cycles lost on a cache miss, which result from increasing the CPU clock rate without a corresponding boost in bus speed. In fact, Intel rates its 100-MHz DX4 at 51 SPECint92 using a 33-MHz system bus, more than 50% better than a 66-MHz DX2 with the same system bus.

Like most 0.6-micron processors, the DX4 runs internally at 3.3 V; the smaller transistors cannot tolerate 5-V operation. To connect to existing system-logic chip sets and standard memory chips, the DX4 has "5-V tolerant" I/O pads. The system must provide a 3.3-V supply to the chip, however, so the parts cannot be dropped directly into existing designs. The lower internal voltage keeps the power consumption reasonable even at the higher clock rates; at 100 MHz, the DX4 is rated at 4.3 W (maximum), 28% less than a 66-MHz DX2 (at 5 V).

Otherwise, the DX4 is similar to Intel's other 486

chips. It uses the same pinout and the same packaging options. Like the SL-Enhanced 486 processors *(see **070801.PDF**)*, it includes the Intel power-management features and has a 208-pin SQFP packaging option.

The DX4 includes one other minor addition: the enhanced virtual-8086 mode implemented by Pentium *(see **070402.PDF**)*. Intel is trying to get operating-system vendors to take advantage of this new feature, but has had little success so far. The company could have made the effort to implement a write-back cache instead and improved performance on a much broader range of applications, but Intel appears to be focusing on features that are harder for other x86 vendors to copy.

The DX4 die, shown in Figure 1, is 77 mm^2, just 5% smaller than the DX2. The larger cache takes nearly all the area saved by the smaller geometries of the 0.6-micron process. According to the MPR Cost Model *(see **071004.PDF**)*, the DX4 will be more expensive to build than the DX2 due to the higher wafer cost of the new process (see "0.6-micron" sidebar); the model estimates that the DX4 costs about $50 to build, compared with $40 for the DX2. As the new process matures and the defect density declines, the manufacturing cost of the DX4 will approach that of the DX2.

Sorting Out the Options

While much of the early speculation on the DX4 focused on its role on the desktop, the new processor will also play a significant role in notebook systems. The new chips include the same power-management features that notebook vendors have been using, and the 3.3-V supply keeps the power consumption of the new parts within reason for a portable system.

The plethora of frequency options is better understood by categorizing chips for either the notebook or desktop markets. Many notebook vendors have stayed with a 25-MHz bus (including DX2-50 systems) even as most desktop vendors moved to a 33-MHz bus. Thus, even though a 75/25-MHz DX4 will have performance similar to a 66/33-MHz DX2, these notebook vendors will prefer the former part while desktop systems use the latter.

CPU	CPU Clock	External Clock	1Q94 Price	2Q94 Price	4Q94 Price**	Volume Availability
Pentium*	100 MHz	66 MHz	$995	$995	n/a	now
Pentium*	100 MHz	50 MHz				(limited)
Pentium*	90 MHz	60 MHz	$849	$849	n/a	now
Pentium	66 MHz	66 MHz	$871	$750	$595	now
Pentium	60 MHz	60 MHz	$793	$675	$550	now
DX4	100 MHz	50 MHz	$649	$580	$450	now
DX4	100 MHz	33 MHz				
DX4	83 MHz	33 MHz	—	—	n/a	2H94
DX4	75 MHz	25 MHz	$499	$475	$375	now
DX2	66 MHz	33 MHz	$440	$360	$260	now

Table 1. Intel's new Pentium and DX4 chips are available with several CPU and system bus speeds. Prices are 1,000-unit list prices. *3.3-V (P54C) processors. (Source: Intel except **PC Week)

Price and Availability

The price and availability of the new IntelDX4 and Pentium processors are shown in Table 1. For more information, contact your local Intel sales office or call 800.548.4725.

Intel is not offering a 100/25-MHz version for notebooks; to reach 100 MHz, system designers must switch to a 33-MHz bus. The mismatch in clock speeds of a 100/25-MHz processor would keep its performance from significantly exceeding that of a 75/25-MHz part; Intel would rather see notebook vendors move to a 33-MHz system bus to improve the performance of the DX4.

On the desktop, most system vendors probably will support the DX2-66 and the DX4-100 using the same 33-MHz system bus. While the DX4-100 also supports an external clock of 50 MHz, few PC vendors want to cope with a system bus at that speed, although some of the larger vendors may use it as a differentiator. The faster bus is more appealing for servers (one of the few areas where the 50-MHz 486DX has been accepted), but these systems are more likely to move to Pentium. Thus, the 100/50-MHz DX4 is likely to see little usage initially, although it will become more popular as chip set vendors eventually begin to support 50-MHz devices.

Later this year, Intel plans to offer an 83-MHz DX4 using a 33-MHz external clock. (The company is still working on the "clock two-and-a-halfing" circuit.) This chip will fill the gap between the 66-MHz DX2 and the 100-MHz DX4 for desktop systems; Intel believes that system vendors will use all three chips to deliver a product line with 20–25% jumps in performance between models. Because these options can all be supported with a single 33-MHz motherboard, it will be fairly simple for system vendors to supply all of them. Intel's strategy avoids leaving any gaps for its competitors to fill.

System Upgrade Possibilities

The DX4 would appear to be a convenient end-user upgrade for DX2 systems, a sort of over-Overdrive part. One problem is the requirement for a 3.3-V power supply, which is not present in the DX2 upgrade socket. This socket is instead intended for the forthcoming "P24T" Pentium processor, which runs at 5 V.

Intel has not announced details of the P24T, but with its narrower, slower system bus and 66-MHz clock, that chip will not match the performance of a 66-MHz Pentium and might not exceed the integer performance of a 100-MHz DX4, particularly on code that has not been recompiled specifically for Pentium. Furthermore, the expected manufacturing cost of the P24T is much, much higher than that of the DX4. The DX4 is available now,

well before the P24T, and it alleviates the heat problems plaguing some systems with the P24T socket.

Thus, there may be an opportunity to offer the DX4 as an end-user upgrade. To work around the 3.3-V problem, an upgrade would need to include a voltage converter. One option would be to combine a DX4 in the small SQFP package with a discrete converter on a tiny PC board; such a device would be no larger than the existing PGA socket and could be manufactured and distributed by a third party. If this market appears lucrative enough, Intel could use a PGA package with an integrated voltage converter to achieve the same result, but this would undercut the market for the P24T.

The DX4 itself also has an upgrade path. Intel says that it will provide a Pentium-class product as an end-user upgrade for DX4 systems. This product, codenamed P24CT, will be similar to the P24T but will use the 0.6-micron process, resulting in a higher clock rate and a 3.3-V supply. The company would not comment on the availability of this upgrade, but it will be no sooner than the 5-V P24T, due late this year.

If the market for the P24T is stunted by third-party (or Intel) DX4 upgrade modules, a version of P24CT could be developed quickly with 5-V tolerant I/O and a voltage converter for the 3.3-V supply in the package. Such a chip could plug into existing P24T sockets and, running at 100 MHz, deliver higher performance than either a DX4 or a P24T. On the other hand, Intel may decide that the market for such a device is too small to justify a special part.

P54C: The True Pentium

When Intel first announced the Pentium processor, it looked like a wonderful product. But upon lifting the hood, we saw that the chip is difficult and expensive to build, uses nearly three times as much power as a 486, and yet has a clock speed (on a good day) no better than a 486DX2's. Any of these problems could have prevented Pentium from ever becoming a volume desktop CPU.

Fortunately, Intel has found that the new IC process resolves these issues. The original P5 design, using the 0.8-micron BiCMOS process, has a die size of 294 mm^2, yielding an estimated manufacturing cost of $320, according to the MPR Cost Model. Figure 2 shows that the new 0.6-micron process shrinks the P54C to 163 mm^2, a reduction of 45%. More important, the estimated cost of building Pentium is reduced by more than half, to about $150. This cost will drop by another 25% or so as the new process matures.

The new design will also help Intel greatly increase the number of Pentium processors that it produces. We estimate that the P5 design yields five good die per wafer; for the P54C, the yield should improve to about 32 chips per wafer. Three factors create this sixfold improvement. A single wafer can hold more of the smaller dice, and these smaller dice are also less likely to contain defects. Finally, the 0.6-micron factories use 200-mm wafers with 80% more area than the 150-mm wafers used in the older fabs. These increases in yield are somewhat offset by a higher wafer cost, but the overall cost is still much improved.

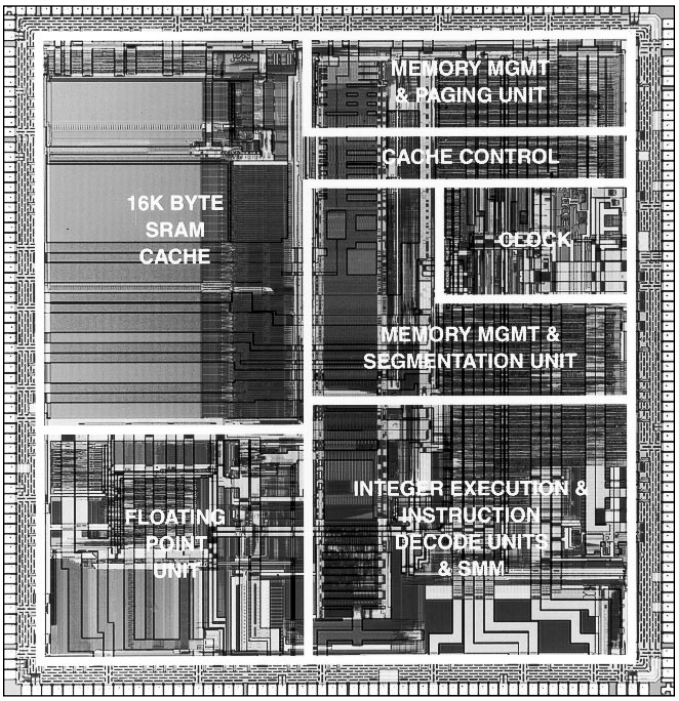

Figure 1. The IntelDX4 measures 8.6 ×8.9 mm (339 ×351 mils) and contains 1.6 million transistors, compared with 1.2 million in a 486DX.

The P54C also addresses the power consumption issue. Like the DX4, the new Pentium chip runs at 3.3 V internally. This reduces the maximum power dissipation to 10 W at 100 MHz, compared with 16 W for the 66-MHz P5. To further reduce power, the new design automatically stops the clock, on a cycle-by-cycle basis, to the cache or to the floating-point unit when those circuits are not being used, reducing power with no performance penalty. Intel claims that these features reduce the average power dissipation of the new Pentium to less than 4 W in typical applications.

The P54C design also includes the full SL Enhanced

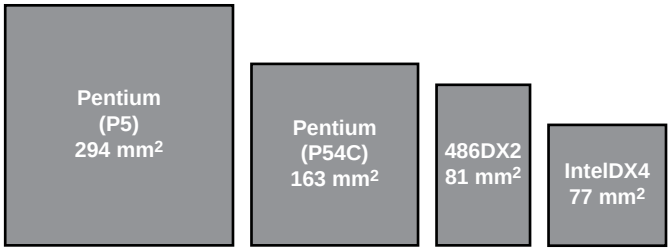

Figure 2. The new manufacturing process shrinks the Pentium processor by 45%. The reduction in size of the DX4 is less dramatic due to the added cache memory.

Intel Moves to 0.6-Micron

The DX4 and P54C are the first products to use Intel's new 0.6-micron fabrication process, now on line at Fab D2 in Santa Clara (Calif.) and at Fab 10 in Leixlip, Ireland. Intel plans to begin 0.6-micron production at Fab 11 in Albuquerque (N.M.) by the end of the year. These fabs, which use 200-mm wafers instead of 150-mm ones, will greatly increase Intel's manufacturing capacity and remedy its current inability to meet demand for its processor chips.

The new process shrinks the drawn transistor size from 0.8 micron to 0.6 micron, reducing circuit area by more than 30%. A fourth layer of metal, which is used to route power and clocks while the other three metal layers carry signals, reduces area by another 10%. The BiCMOS process incorporates both CMOS and bipolar transistors, which can be combined in BiNMOS drivers that speed the transmission of heavily loaded signals.

Although these features improve die area and performance, they also increase wafer cost by about 30% over a three-metal CMOS process with a similar transistor size. Taking into account the smaller geometries and larger wafer size, total wafer cost for the new process is more than twice that of Intel's 0.8-micron BiCMOS process. Given Intel's high production volume, the company is willing to trade reductions in die area against increases in wafer cost.

Intel says it has no plans to move the standard 486DX and DX2 chips to the 0.6-micron fab even though it ultimately would reduce the manufacturing cost of those parts. The company has only recently completed moving all 486DX production to 0.8-micron lines, and moving additional products to the new fabs would result in unused capacity in the older 0.8-micron facilities. By the time the new process has improved to the point that costs drop below those of the older fabs, the volumes of the DX and DX2 parts will be dwindling and it probably will not be a wise investment to move these designs to the new process.

Intel is continuing to invest in more advanced fabrication methods and plans to accelerate its IC process development cycle from three years to two. The company is building a new factory in Albuquerque that will use a 0.4-micron process that is reportedly CMOS only (not BiCMOS). Intel expects that this new process will be in production by the end of 1995; it will be used to build P6 and Pentium chips.

feature set used in Intel's current 486 family. These features, designed to allow hardware and software power management to further reduce power consumption, include system-management mode (SMM), the ability to stop and quickly restart the processor clock, and a automatic power-down mode *(see **070801.PDF**)*. While the P5 includes SMM, the "stop clock" and "auto halt" features are new in the P54C.

These power-reduction features make the P54C processor more suitable than the P5 for notebook systems by greatly extending battery life, which depends on typical power consumption. Cooling must still be provided for the worst-case heat dissipation, but in a notebook, an active power-management system can monitor the temperature of the CPU and slow the clock if the chip is getting too hot. Pentium notebook systems should begin rolling out by Fall Comdex and become widespread in 1995.

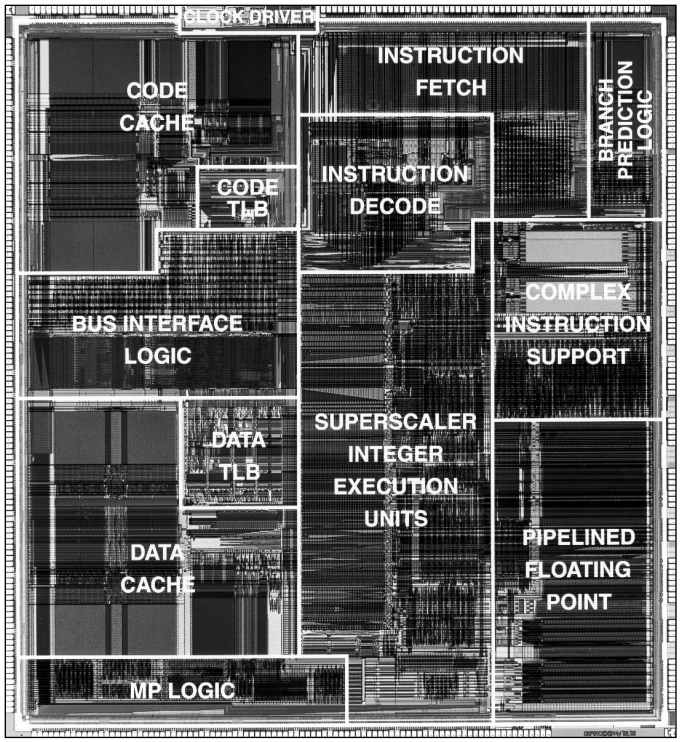

Figure 3. The P54C measures 13.3 ×12.3 mm and uses 3.3 million transistors, slightly more transistors than the original Pentium.

Achieving Maximum Performance

The 0.6-micron process enables the P54C to run faster as well. The first products using this new design reach speeds of 90 and 100 MHz, as Table 1 shows. Intel had initially planned for a 75-MHz version, but yields at 90 MHz have been high enough to eliminate the need for this speed grade. The yields on 100-MHz parts are still low, however, and Intel says that the fastest parts will be in limited production until late this year. Thus, the 100-MHz parts will appear in servers initially, while volume desktop systems use the 90-MHz version.

At the ISSCC conference, the company displayed a 150-MHz Pentium processor. This chip was hand-picked from the production line and operated with a special cooling unit, but Intel hinted that the 150-MHz Pentium may eventually become a product. In 1991, the company demonstrated a 100-MHz 486 at that year's ISSCC and is now able to ship such a part using the next generation of manufacturing technology. Once Intel has progressed

to a 0.4-micron process (see "0.6-micron" sidebar), faster Pentiums will be possible—and probably inevitable.

Intel rates the 100-MHz Pentium at 100 SPECint92 and 80 SPECfp92; Table 2 shows the complete list of SPEC benchmark values. Thus, the new parts retain the ability of the slower Pentium chips to deliver 1.0 SPECint92 per MHz, even though memory speeds are not increasing with the processor clock rate. Although most of the performance increase is due to the faster processor clock, it also includes some compiler improvements. Note that typical PCs will be 10–20% slower than the optimized designs shown in Table 2, which use fast caches and memory systems.

The DX4 and P54C promise a significant increase in performance for every type of PC. Low-end systems can jump to the DX2 if Intel continues to drop the price of that part. Midrange boxes will move from the DX2 to the DX4 while the high end migrates from today's 60-MHz Pentium to faster P54C parts. Thus, by the end of the year, most PC users will see a 50% performance increase for the same system price.

P54C Supports Glueless MP

While the P54C is nearly identical to the P5 design, Intel has taken the opportunity to make a few functional improvements. In addition to the new SL Enhanced features, the design now incorporates Intel's advanced priority interrupt controller, or APIC (see "APIC" sidebar), making the new chip suitable for a glueless dual-processor configuration in which two CPUs share the same cache, as Figure 4 shows.

This dual-processor configuration would not deliver the same performance boost as a traditional MP design with separate caches for each processor; depending on the application, the gain would range from 30% to 70%, according to Intel. This design would be much less expensive than a traditional multiprocessor configuration, as the only cost of adding the second processor would be the CPU chip itself.

Alternatively, Intel could have left the APIC off of the processor and assumed that it would be in the system logic. Chip-set vendors balked at the added cost of the APIC, however, and any chip sets that include the logic to support dual processors will be more expensive than standard uniprocessor chip sets. Intel believes that this cost differential, in the highly competitive PC market, would have led to a dearth of MP-capable systems. By including the APIC, which uses less than 5% of the die area, on the CPU, system vendors can offer dual-processor capability for the cost of a second CPU socket, seeding the market with lots of these systems.

Of course, to take advantage of the second processor at all, a multiprocessor operating system is required. And unless an application is multithreaded, the second CPU is active only when two or more tasks are running.

System	Intel XPRESS	Siemens PCE-5S/66	Micronics M4P PCI	Compaq Deskpro
Processor	Intel Pentium	Intel Pentium	Intel DX4	Intel 486DX2
Clock Rate	100/66 MHz	66 MHz	100/33 MHz	66/33 MHz
Cache (on/off-chip)	16K/512K	16K/256K	16K/256K	8K/256K
espresso	96.6	62.5	48.2	30.6
li	134.4	88.8	71.7	49.0
equtott	103.8	60.1	52.4	29.7
compress	54.4	42.9	31.6	20.1
sc	145.7	102.0	73.9	51.9
gcc	93.8	64.2	43.7	23.9
SPECint92	100.0	67.4	51.4	32.2
spice	64.9	51.6	34.9	19.2
doduc	79.1	52.1	25.6	15.3
mdljdp2	95.2	65.0	28.2	16.8
wave5	55.8	41.5	17.6	5.9
tomcatv	77.7	73.6	27.4	16.3
ora	93.7	64.1	33.0	22.7
alvinn	170.5	122.3	45.4	27.5
ear	210.7	160.0	47.2	28.8
mdljsp2	44.9	30.2	13.7	9.0
swm256	45.1	42.6	15.6	10.3
su2cor	56.9	49.4	28.1	19.5
hydro2d	83.0	57.6	25.1	16.3
nasa7	60.7	52.7	26.4	15.4
fpppp	117.5	85.7	25.9	18.1
SPECfp92	80.6	61.5	26.6	16.0

Table 2. At 100 MHz, the new Pentium and DX4 processors deliver about 50% better performance than their 66-MHz predecessors.

Given that neither DOS nor Windows (including the future Chicago version) can handle multiple processors, Intel expects that the dual-processor Pentium will be used primarily for high-end desktops or servers running UNIX or Windows NT.

In these high-end markets, the dual-processor mode can be used as an upgrade strategy for the P54C. For the majority of users, however, Intel will provide a traditional upgrade chip that usurps control of the system from the original CPU. The company will not discuss any specifics about this upgrade part but expects it to be available in 1996. Thus, it is possible that the upgrade will take advantage of the P6 processor core.

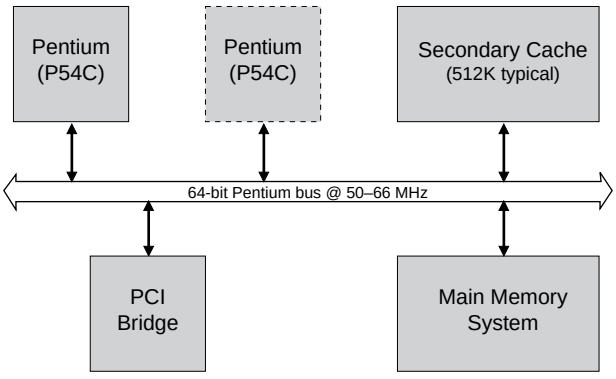

Figure 4. With integrated interrupt controllers, two P54C processors can share a single cache in a low-cost multiprocessor system.

Pentium Goes "APIC"

Intel's advanced priority interrupt controller (APIC) architecture was first announced more than a year ago *(see **0615MSB.PDF**)* in its first instantiation, the 82498DX. The APIC architecture replaces the old 8259 interrupt controller originally designed for the 8080 and inherited by all PCs since then. With a more flexible priority scheme and faster response time, the APIC has some benefit for uniprocessor systems, but its major advantage is in supporting multiprocessor systems, which is not possible with the simple 8259 design.

The APIC is physically divided between the processors and the system logic. The "I/O APIC," typically part of the system logic, accepts system interrupts much like the 8259. Unlike the older part, however, the I/O APIC can pass interrupts to multiple processors, each of which must have its own local APIC module. The various APIC modules are connected via a private interrupt bus, allowing interrupts to be communicated without using the normal system bus structure.

The 82498 has not been widely used. Most vendors with multiprocessor x86 systems had already defined their own MP interrupt protocol and saw no reason to change, although a few have adopted the APIC. Desktop systems have not incorporated the 82498 due to its cost ($26), and system-logic vendors have seen no reason to incorporate a complete APIC (both I/O and local modules) in their chip sets.

The P54C Pentium integrates the local APIC module and communicates to the I/O APIC (and other P54C processors) using a three-wire bus. Although the P54C implementation is register-compatible with the 82498, the three-wire bus is not compatible with the 82498's five-wire protocol. Intel will include the I/O APIC logic in its P54C chip sets and is licensing it to other system-logic vendors. The I/O APIC is relatively small, and Intel expects that most vendors will include it in their basic chip sets.

For compatibility with software that does not include APIC code, the P54C can disable its on-board APIC and use an external 8259-type controller. Today, few software vendors support the APIC, although a special HAL for Windows NT is available. By increasing the installed base of APIC hardware, Intel hopes to spur other MP operating systems to support it.

The other new feature of the P54C design is a phase-locked loop (PLL) that lets the CPU run at 1.5×or 2×the system bus frequency. This keeps the system bus between 50 and 66 MHz while the CPU runs as fast as 100 MHz. Although the 2×ratio allows for a 100/50-MHz system, this configuration will not significantly outperform a 90/60-MHz design, so vendors may try for a 100/66-MHz arrangement to maximize performance. As in the DX4, the P54C clock multiplier is pin-selectable; there is a single 100-MHz version of the chip that supports both bus frequencies.

Unlike the DX4, the P54C supports only 3.3-V I/O signals, forcing system designers to use low-voltage cache memory and chip sets. Intel plans to release a 3.3-V version of its 82430 chip set for new designs, and expects that similar chips sets will be available from Opti, VLSI, and others by the end of the year. By forcing a move to 3.3 V at this time, Intel is preparing system vendors to support future Pentium chips and Pentium upgrades, all of which will use 3.3-V I/O.

The new processor uses a 296-pin PGA, 23 more pins than the P5 package. Three of the new pins are used for the APIC, and most of the rest are left as no-connects to allow for "future functional enhancements." Because of the new package and the 3.3-V I/O, it is impossible to simply drop the P54 processor into an existing Pentium motherboard; in fact, upgrading to P54C will require a significant redesign.

5-V Pentium Line Will Continue

Despite the vast manufacturing cost improvement of the P54C, Intel says that it will continue to build 5-V Pentium chips using the P5 design for the foreseeable future. Because of the extensive redesign required for the new version, many vendors will continue to ship systems using 60- and 66-MHz Pentium chips for some time. The lower frequencies also provide additional performance points for Pentium systems, although the 100-MHz DX4 overlaps the 60-MHz Pentium on some applications.

Continuing the 5-V line also maximizes the number of Pentium chips that Intel can produce in 1994. By the end of the year, the company plans to have three factories building the P54C design along with the two currently making the older P5 version. With all five fabs continuing to crank out chips, Intel expects to build about six million Pentium processors this year. In fact, given the amount of fab capacity coming on line, the company is obligated to build a large number of chips simply to defray the costs of building the new fabs.

Thus, Intel will aggressively price Pentium chips, cutting prices on the 0.8-micron versions even though they are more expensive to build than the higher-speed chips. Intel has already announced that it will slash prices on 66-MHz Pentiums from $871 to $750 in the second quarter and, according to a document obtained by *PC Week*, that price could fall to $595 by 4Q94, as shown in Table 1. These price cuts make room for the higher-speed parts, which will start at $995 for a 100-MHz Pentium and $849 for the 90-MHz version.

These price cuts will reduce the margins on the older P5 parts below the high margins of Intel's other processors, but the company will still make a significant profit on these parts. Furthermore, by increasing the penetration of Pentium in the PC market, Intel devalues the product offerings of its 486-based competitors.

To build momentum for Pentium, Intel must convince PC makers to move beyond the 5-V, 33-MHz system bus that they are comfortable with. Today, many companies are buying complete Pentium motherboards directly from Intel, which by some counts is the largest vendor of Pentium boards today, because of the difficulty of designing with the 60-MHz system bus. Unless Intel plans to continue growing its motherboard business, which upsets those vendors that actually have the resources to design their own products, it must provide simple design kits and chip sets that can handle the faster Pentium processors.

Pressure on the Competition

The new parts open a wide gap between Intel and other x86 CPU vendors. AMD, Cyrix, and TI are limited to 66 MHz for their 486DX2-type chips, and both Cyrix and TI use a slower core than Intel. None of these vendors is likely to achieve higher clock rates this year. Thus, the DX4 will hold a 40–50% performance advantage in the 486 class for some time.

The only other vendor likely to reach 100 MHz in 1994 is IBM. Currently, its fastest part is the 75/25-MHz Blue Lightning, which has only slightly better performance than Intel's 66-MHz DX2. Blue Lightning is hampered by its narrow bus and lack of a floating-point unit, and it cannot be sold as a standalone chip.

Other than Intel, no vendor has delivered, or even announced, a Pentium-class product. AMD's K5 and Cyrix's M1 are about a year away from shipments, and IBM's 586 is still a rumor. The first Pentium competitor may emerge from NexGen, which has been sampling its 586 processor *(see **0714MSB.PDF**)*. Because of its small size and fabless status, NexGen is unlikely to obtain either the fab capacity or the customer interest to make much of a dent in the Pentium market.

Thus, for at least the rest of this year, Intel will be able to wield the DX4 and P54C largely unopposed in the midrange and high-end markets, leaving its competitors to fight over the low-margin bottom end. Intel will not abandon the low end; profits from its flagship products will subsidize heavy discounting of 486DX2 and other low-end chips.

	Pentium (P5)	Pentium (P54C)	PowerPC 601	PowerPC 603
CPU Clock	66 MHz	100 MHz	80 MHz	80 MHz
Bus Clock	66 MHz	66 MHz	40 MHz	40 MHz
On-Chip Cache	8K I, 8K D	8K I, 8K D	32K unified	8K I, 8K D
Bus Width	64 bits	64 bits	64 bits	64 bits
Max SPECint92	67 int	100 int	~85 int	~75 int
Max SPECfp92	62 fp	80 fp	~105 fp	~85 fp
IC Process	BiCMOS 0.8 µ, 3M	BiCMOS 0.6 µ, 4M	CMOS 0.6 µ, 5M	CMOS 0.5 µ, 4M
Die Area	294 mm^2	163 mm^2	121 mm^2	85 mm^2
Est. Mfg. Cost*	$350	$160	$100	$55
List Price (1K)	$871	$995	$520*	n/a
Availability	volume	limited	volume	sampling

Table 3. The P54C significantly improves both the cost and performance of Pentium but still lags the price/performance of PowerPC. (Source: vendors except *MPR estimates)

Intel also hopes to head off any incursion from PowerPC and other RISC processors. As Table 3 shows, the lower cost of the P54C design brings it much closer to the cost of the PowerPC 601 and 603, although the RISC chips still hold a significant advantage. The list price of the P54C is nearly twice that of the 601, which reflects Intel's higher margins as much as its higher manufacturing cost. The 603 will have an even bigger price advantage over the P54C but can't match its performance.

The new 100-MHz Pentium will outrun either the 601 or the 603 on integer code. IBM and Motorola expect to gain a performance advantage when the PPC 604 begins to ship around the end of the year, but that chip is likely to be more expensive to build than the 601. The prospect of even faster Pentiums, and ultimately the P6, looms in the future.

On the eve of Apple's first PowerPC announcement, Intel has shown that it will not relinquish its performance leadership easily. Given the ability of PowerPC to deliver similar performance at a lower CPU price, Apple (and other system vendors) may be able to translate this advantage into a system-level price/performance advantage, but that remains to be seen. Until PowerPC can improve its position, it will not threaten Intel. And in the meantime, the new x86 chips should allow Intel to continue its dominance of the x86 market while generating its traditional enormous profits. ♦

