

# Microprocessors Lead the Way to 0.35 Microns

## NEC, IBM, Intel Are First to Reach This Level—Others Follow Soon

by Linley Gwennap

*This article updates our look at microprocessor vendors' 0.5-micron IC manufacturing processes (see 080504.PDF) with the latest advances.*

Microprocessors are now leading advances in IC manufacturing, taking the baton from the traditional process driver, DRAM. Even as memory vendors are struggling to bring 64-Mbit DRAMs to volume production in 0.35-micron processes, companies such as NEC, IBM, and Intel are already producing microprocessors with similar geometries. CPU vendors expect to deploy 0.25-micron processes in 1996 or 1997, well before DRAM vendors reach production of 256-Mbit parts.

These microprocessor processes are, in fact, more complex than DRAM processes of the same geometry. Although both may have the same gate length, CPU vendors have emphasized thinning the gate oxide to reduce transistor switching time. They have also led the way to four and five layers of metal, improving the density of their more complex designs.

While giants such as Intel forge ahead with their own processes, other CPU vendors have banded together in teams—including AMD/HP and Hitachi/VLSI—to share the cost of process development. NEC, IBM, and Texas Instruments continue to use the traditional model of using memory chips for early development work.

IC process technology is a key determinant of both CPU performance and cost. By reaching the 0.35-micron level ahead of the competition, NEC, IBM, and Intel hold a trump card that allows them to deliver higher performance. Even as other vendors reach this level, the early entrants will have had more time to shake out the bugs in their processes, giving them a cost advantage.

But not all 0.35-micron processes are equal. Some vendors have emphasized high performance, others low cost, still others compatibility with previous processes. Intel's vast pocketbook has helped it overcome previous manufacturing lags and vault into the lead in IC technology, making it even harder for its competitors to keep up. Examining the details of CPU vendors' manufacturing technology helps determine the likely schedule, clock speed, voltage, and cost of future products.

### Moving to Lower Voltages

In many ways, the typical 0.35-micron process is much like its predecessors, except that the devices are smaller. But this generation has caused a few new issues to arise, which vendors are handling in different ways.

For chip users, the most significant issue is that of the power supply. The move from 0.8- to 0.5-micron technology caused most vendors to reduce power supplies from the age-old 5-V standard to roughly 3.3 V. In the move to 0.35-micron designs, IBM and Digital are leaping to a 2.5-V supply, while most others are trying to keep to 3.3 V for one more generation. Intel is in both camps with various products and is attempting a compromise at 2.9 V with the initial P6.

The power-supply change is caused by continued shrinkage in gate-oxide thickness (see 070705.PDF). In general, thinning the gate oxide improves the switching speed of the transistor. But if the supply voltage is too high, these thin oxides will disintegrate over time, causing the transistor to fail. To avoid long-term reliability problems, chip vendors have adjusted the voltage downward as oxide thicknesses have plunged well below 100 Å. Although circuits typically run slower at lower voltages, the speed improvement from thinning the oxide layer more than makes up for this change. An additional benefit is lower power dissipation.

Changing voltages, however, forces system vendors to redesign their products and may cause incompatibility with memory and interface chips. Vendors that wish to maintain 3.3-V compatibility are holding oxide thickness at 70–80 Å. Going below this value requires a shift to a lower supply voltage. Those vendors that shift to a 2.5-V supply can thin the gate oxide to about 60 Å, giving their chips a clock-speed advantage over similar designs built with thicker oxides.

While vendors have traditionally improved performance by speeding the transistors, metallization is becoming increasingly important. A gate that drives a short trace may switch very quickly, but critical signals must often be driven across a large die. In this case, the switching time is limited by the resistance and capacitance of the metal trace.

Metal resistance is proportional to the width and thickness of the trace; as the metal lines get smaller, the resistance goes up, reducing gate speed. CPU vendors are highly motivated to shrink metal lines, however, as this is the biggest factor in reducing die size and thus cost. To resolve this catch-22, vendors with the narrowest metal traces, such as Intel, are keeping the thickness (depth) of the traces constant, reducing their resistance.

Capacitance, on the other hand, diminishes as trace width shrinks. To further reduce capacitance, TI and other vendors are investigating new materials to insulate the metal layers. These materials have a smaller

dielectric constant, reducing capacitance. In addition to speeding signals, lower capacitance also curtails power dissipation, a crucial issue as transistor counts continue to skyrocket.

### New Manufacturing Techniques Seen

These new IC processes are more likely to utilize chemical-mechanical polishing (*see 080504.PDF*). CMP adds cost to the manufacturing process but allows more than three metal layers. IBM pioneered this technique; Intel and AMD have used CMP to some degree in previous process generations. With more vendors adding a fourth metal layer to compact their designs, CMP is becoming nearly standard in the industry. For example, Digital and NEC have added CMP to their latest processes; Fujitsu and TI are phasing it in over time.

Intel, IBM, and Digital are leading the move from i-line steppers, which use ultraviolet light, to deep ultraviolet (DUV) light. Steppers use a light source to expose the die through a mask, printing the resist with a pattern that is then etched into the chip itself (*see 070705.PDF*). These vendors have reached the point where the wavelength of the light used by the i-line steppers is approaching the size of the features being patterned, causing undesirable diffraction. DUV has a shorter wavelength, eliminating this problem.

Other vendors are sticking with i-line steppers at 0.35-micron, but even they admit that a move to DUV is needed to take the next step. Those vendors that make the move now gain experience for the next process generation. But DUV requires new stepper machines, adding to the cost of fabs and thus increasing wafer cost. Fortunately, the older i-line steppers can still be used for less critical layers, such as the upper metal layers.

DUV also requires a different chemistry, as the old resists are not sensitive to DUV light. Developing and testing these new resists may increase the time needed to deploy the new process. IDT and others are investigating using phase shifting instead to reduce diffraction, extending the life of i-line steppers and current resists. Late adopters are willing to let others do the work of developing these techniques.

### Intel Seizes Density Lead

Contrary to our recent item (*see 090402.PDF*), Intel was not the first microprocessor vendor to put a 0.35-micron process into production. But the microprocessor giant was hot on the heels of NEC and IBM, and Pentium has reached this milestone months before any other x86 processor. Furthermore, Intel's 0.35-micron process, known as P854, carries tighter metal pitches than any of the others shown in Table 1 (*see below*); we believe this allows the company to pack circuits more densely than any other CPU vendor, giving Intel a cost advantage and allowing greater output from its fabs.

The 0.35-micron process, which is currently in production, is designed for compatibility with Intel's 0.5-micron process, called P852. (Intel labels P852 as a 0.6-micron process; for consistency, we refer to all processes by their drawn gate length.) The 0.35-micron P854 maintains compatibility with the 3.3-V supply and BiCMOS layers used by the P54C Pentium design. These features allowed the P54C design to be quickly shrunk to P854, creating a design known as P54CS.

Intel has developed a second version of P854 that modifies the transistor characteristics while maintaining the same metal layers. This version uses a 0.28-micron gate and a 2.5-V supply, allowing a reduction in the gate-oxide thickness. The thinner oxide helps reduce the effective gate length, speeding the transistor switching time. At this level, the presence of bipolar gates adds little performance while increasing wafer cost significantly, so Intel left them out of the 2.5-V P854. Eliminating the four bipolar layers reduces the total number of photomask layers by 20%. This move to pure CMOS reduces wafer cost even as the gate size shrinks.

The 2.5-V process will be used for a redesigned version of Pentium code-named P55C, which is due in 1H96. Although the shrunk P54C will probably top out at about 150 MHz, we expect the P55C to achieve 166 MHz and possibly 180 MHz, due to the enhanced transistor design. The P6 will also move from a 2.9-V version of P852 to the 2.5-V P854; we expect this version to reach at least 200 MHz and ship around mid-1996. The lower voltage levels have an additional benefit of maintaining reasonable power dissipation as clock speeds increase.

Both versions of P854 use the same metal pitches. Intel leads the industry with a contacted metal pitch of 0.88 microns for metal-1, allowing an SRAM cell size of just 21  $\mu\text{m}^2$ . Metal-2 and metal-3, which help determine the density of logic areas, are also ahead of the pack at 1.16  $\mu\text{m}^2$ . This combination lets Intel pack more features into a given die size than any other CPU vendor.

### IBM Pushes Performance

IBM is one of the few microprocessor vendors that could keep up with Intel in IC manufacturing investment, although lately it has fallen off the pace, spending \$600 million on research and new fabs, compared with Intel's massive \$3 billion budget. Despite this deficit, Big Blue beat Intel to the 0.35-micron level by a few months, putting its 100-MHz PowerPC 601 into production in October of last year. (IBM had previously classified its CMOS-5X as a 0.5-micron process but recently disclosed that its transistors physically measure 0.33 microns.)

IBM focuses on the effective channel length to improve clock speed; at 0.25 microns, the CMOS-5X effective length is smaller than that of Intel's current P854 process. But with its focus on performance, IBM does not match the circuit density of Intel's process. The metal-2

Vendor Process name	AMD† CS-34	Intel P854	Intel P854	IBM‡ CMOS-5S	IBM CMOS-5X	IBM CMOS-6S	TI EPIC-3	TI EPIC-3	TI EPIC-4	Hitachi§ "0.35µ"
Example product	486, K5	P54CS	P55C	PPC 620	PPC 601+	PPC 604+	UltraSparc	486DX4	U'Sparc-2	SH-4
First production	4Q95	1Q95	1Q96	4Q94	4Q94	2Q96	3Q95	4Q95	1H96	2Q96
Supply voltage	3.3 V	3.3 V	2.5 V	3.3 V	2.5 V	2.5 V	3.3 V	3.3 V	2.5 V	3.3 V
BiCMOS?	no	yes	no	no	no	no	no	no	no	no
Gate length (drawn)	0.35 µm	0.35 µm	0.28 µm	0.44 µm	0.33 µm	0.27 µm	0.47 µm	0.42 µm	0.29 µm	0.38 µm
Channel length (effective)	0.25 µm	0.30 µm	0.22 µm	0.39 µm	0.25 µm	0.20 µm	0.37 µm	0.35 µm	0.22 µm	0.35 µm
Gate oxide thickness	70 Å	70 Å	60 Å	80 Å	70 Å	55 Å	80 Å	80 Å	57 Å	80 Å*
Number of metal layers	3-4 metal	4 metal	4 metal	4-5 metal	5 metal	5 metal	3-4 metal	3-4 metal	4-5 metal	3-5 metal
Local interconnect?	no	no	no	yes	yes	yes	no	no	no	no
Stacked vias?	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
M1 contacted pitch	1.4 µm	0.9 µm	0.9 µm	1.4 µm	1.2 µm	1.1 µm	1.8 µm	1.5 µm	1.2 µm	1.4 µm
M2 contacted pitch	1.4 µm	1.2 µm	1.2 µm	1.8 µm	1.8 µm	1.4 µm	1.8 µm	1.5 µm	1.2 µm	1.4 µm
M3 contacted pitch	1.4 µm	1.2 µm	1.2 µm	1.8 µm	1.8 µm	1.4 µm	1.8 µm	1.5 µm	1.2 µm	1.4 µm
M4 contacted pitch	2.4 µm	3.0 µm	3.0 µm	1.8 µm	1.8 µm	1.4 µm	4.0 µm	3.3 µm	1.2 µm	1.4 µm
Routing index (min metal)	1.9 µm <sup>2</sup>	—	—	2.4 µm <sup>2</sup>	—	—	3.1 µm <sup>2</sup>	2.1 µm <sup>2</sup>	1.2 µm <sup>2</sup>	1.9 µm <sup>2</sup>
Wafer cost index (min)	2.3	—	—	2.5	—	—	2.0	2.1	2.9	2.3
Routing index (max metal)	1.7 µm <sup>2</sup>	1.2 µm <sup>2</sup>	1.2 µm <sup>2</sup>	2.2 µm <sup>2</sup>	2.0 µm <sup>2</sup>	1.3 µm <sup>2</sup>	2.8 µm <sup>2</sup>	1.9 µm <sup>2</sup>	1.1 µm <sup>2</sup>	1.5 µm <sup>2</sup>
Wafer cost index (max)	2.7	3.1	2.9	2.7	3.2	3.5	2.3	2.5	3.4	3.0

Table 1. IBM, Intel, and NEC are shipping 0.35-micron microprocessors today, and others will soon follow suit. See sidebar (page 20) for more information on key parameters and indices (smaller is better). n/a indicates not announced. †HP has rights to AMD's 0.35-micron process. ‡Motorola has rights to IBM's processes. §VLSI has rights to Hitachi's processes. (Source: vendors except \*MDR estimates)

and metal-3 pitches are the same as in IBM's older CMOS-5S process and are 50% larger than Intel's metal pitches. The routing index indicates that IBM's chips could be 40% smaller if redesigned for Intel's process.

IBM plans to narrow this gap next year when it deploys CMOS-6S. Taking advantage of the local-interconnect layer and tight metal-1, CMOS-6S has an SRAM cell size of 20 µm<sup>2</sup>, slightly smaller than in Intel's P854 process. This process significantly reduces the pitch of the metal-2 and metal-3 layers, but it still doesn't match Intel's density. As a result, IBM can match Intel in cache density but not in the logic areas that make up most of a typical CPU, putting IBM's microprocessors at a higher cost than Intel's.

IBM's process sequence shows how microprocessors are driving the company's process advancement. The company developed its CMOS-6 process, which uses a 0.35-micron gate, for 64-Mbit DRAMs. The CMOS-5X process, however, uses a similar feature size and reached production at about the same time as IBM's 64-Mbit DRAMs. The company has fabricated a 256-Mbit DRAM using a 0.25-micron process (as have other memory vendors), but these parts are not expected to ship in volume until 1999; CMOS-6S will deliver products with similar transistors next year.

Even as its microprocessor processes pull ahead of its DRAM processes, IBM believes that it gains some value from its DRAM experience. The memory chips help flesh out the basic process steps and improvements to design tools. But the logic chips generally use more metal layers, thinner gate oxides, and ultimately smaller transistors than the memory devices, which are not pushing the limits on these fundamental parameters.

Although IBM has produced 0.35-micron processors at about the same time as Intel, this parity is not as significant as it appears. Both Cyrix and NexGen compete with Intel using x86 processors built by IBM, but neither of these vendors has been able to move its designs into CMOS-5X in a timely fashion. Cyrix's 5x86 (*see 090901.PDF*) and M1, as well as NexGen's Nx586, are being built in 0.65-micron processes that are two generations behind CMOS-5X. Both x86 vendors plan to move their chips to CMOS-5X in 1H96, bringing them to process parity with Intel's 0.35-micron Pentium.

On the PowerPC side, only the 601 has reached the 0.35-micron process. IBM plans to move the 603, 604, and 620 to CMOS-5X in 1H96, more than a year after that process first reached production. It isn't clear why these chips haven't reached IBM's most advanced process already; the company denies having significant capacity constraints on its 0.35-micron process.

Motorola—IBM's partner on the 603, 604, and 620 processors—has access to all of IBM's process technology but has not yet brought even the 0.45-micron CMOS-5S process on line, although this milestone is planned for later this year. Motorola has not disclosed its plans to move to 0.35-micron technology, but it appears the company will not achieve that goal until sometime in 1996, well behind the rest of the industry.

### TI Produces Aggressive 0.25-Micron Plan

Texas Instruments is known for low-cost manufacturing, not high-performance IC processes. The company recently moved its 486DX2 into its 0.47-micron EPIC-3 process and is now selling the part for as little as \$66, undercutting the competition. We believe TI's defect

Digital CMOS-6	Digital CMOS-6	Fujitsu CS-55	Fujitsu CS-60ALE	IDT CEMOS 8+	NEC† "0.35μ"
21164A 1Q96	n/a n/a	Hal CPU 2Q95	n/a 1Q96	R4400-200 1Q95	R4400-200 4Q94
2.5 V no	2.5 V no	3.3 V no	3.3 V no	3.3 V no	3.3 V no
0.33 μm 0.25 μm	0.33 μm 0.25 μm	0.40 μm 0.35 μm	0.35 μm 0.28 μm	0.30 μm 0.25 μm	0.35 μm 0.28 μm
65 Å 3-4 metal	65 Å 4-5 metal	80 Å 3-4 metal	80 Å 3-5 metal	90 Å 3 metal	85 Å 3-5 metal
optional yes	optional yes	no no	no yes	no no	yes yes
1.2 μm 1.2 μm	1.2 μm 1.2 μm	2.1 μm 2.1 μm	1.3 μm 1.3 μm	1.5 μm 1.8 μm	1.5 μm 1.5 μm
2.5 μm 2.5 μm	1.2 μm 2.5 μm	2.1 μm 210 μm	1.3 μm 1.8 μm	2.3 μm —	1.5 μm 1.5 μm
2.6 μm <sup>2</sup> 2.8	1.2 μm <sup>2</sup> 3.2	4.4 μm <sup>2</sup> 2.4	1.5 μm <sup>2</sup> 2.6	4.1 μm <sup>2</sup> 2.9	2.1 μm <sup>2</sup> 2.5
1.8 μm <sup>2</sup> 3.0	1.2 μm <sup>2</sup> 3.4	4.0 μm <sup>2</sup> 2.8	1.2 μm <sup>2</sup> 3.3	— —	1.8 μm <sup>2</sup> 3.3

Table 1 (cont.) †Toshiba uses a process similar to NEC's.

rates are among the best in the industry, allowing it to achieve low manufacturing costs. Now, the company is looking to deliver high performance as well.

While other vendors are moving to 0.35-micron production this year, TI has a more modest plan of shrinking its current process by 10%, achieving a 0.42-micron gate with 1.25-micron metal pitches. This process should allow TI to build 100-MHz 486DX4s late this year. As a further step, the company may produce another shrink, bringing the gate size to 0.35 microns with the same metal pitches; this version would be in production in early 1996. All of these versions of EPIC-3 retain the same 80-Å gate oxide to maintain a 3.3-V supply voltage.

TI's next big step will be to EPIC-4, which moves to a 2.5-V supply and an oxide thickness of just 57 Å. This process compresses the gate size to 0.29 microns and the metal pitches to 1.0 micron. The company hopes to have this process in production in 1H96. But EPIC-4 requires a move to DUV steppers and CMP planarization, which may delay its implementation.

The second-generation UltraSparc will probably use EPIC-4, pushing clock speeds to 250–300 MHz and reducing die size below 160 mm<sup>2</sup>. From the 0.29-micron level, TI believes it can move quickly to 0.25 microns; the company expects to do so by the end of 1996, making it one of the first to reach the quarter-micron level.

TI expects to move away from tungsten plugs in the EPIC-4 generation. Instead, a technique called force-filled aluminum will be used to fill holes created by vias, forming a solid metal bond between layers. This process reduces the opportunity for defects and, properly implemented, lowers cost by eliminating the tungsten steps. This technique also improves reliability by simplifying the structure of the via.

## AMD Pursues Intel with Fab 25

AMD's first processor megafab, Fab 25 in Austin (Texas), is designed to match the capabilities of Intel's best fabs. AMD is a bit behind in the 0.35-micron race; the company plans to ship 0.35-micron 486 processors by the end of the year, about three quarters after Intel's first 0.35-micron shipments. K5 processors at this level are expected in 1Q96.

AMD's 0.35-micron process has significantly worse metal pitches than Intel's, making AMD's chips about 40% larger than they would be if built by Intel. Still, the new process will allow AMD to boost 486 clock speeds to at least 150 MHz and perhaps as high as 180 MHz. This move will also reduce die size and allow AMD to double the on-chip cache with little incremental cost.

AMD's initial 0.35-micron process runs at 3.3 V. The company hinted that it will quickly move to a 2.5-V supply and thinner gate oxides; this refinement could be in production as early as mid-1996.

AMD's challenge is to bring the new process on line quickly in its new fab. Its 0.35-micron technology was initially developed in partnership with Hewlett-Packard at that company's Palo Alto (Calif.) facility, then transferred to AMD's Submicron Development Center (SDC) early in 1994. The company is now replicating the process at Fab 25, but a switch from 150-mm to 200-mm wafers forced some equipment to change in this transfer. AMD says that the same sequence was successfully used to transfer its latest flash-memory process to a new Japanese fab, and the company expects few problems in bringing the 0.35-micron process on line in Fab 25.

## Digital Seeks to Shed Cost

Digital's CMOS-5 process has a number of extra mask layers that boost performance while adding cost (see [080504.PDF](#)). In CMOS-6, the company has kept its eye on performance while shedding some of these layers. For example, Digital engineers figured out how to produce precision resistors, needed for ECL compatibility, without extra process steps. CMOS-6 also eliminates the low-resistance layer over the pad ring without significantly increasing pad resistance.

While the new process retains the local interconnect of CMOS-5, initial designs may not use it to further reduce cost. CMOS-6 adds stacked vias, which helps circuit density. Digital is using DUV for a few critical layers and expects more widespread use in future processes. The company is currently running test wafers for CMOS-6 and expects products in this process to sample in 3Q95 and ship in 1Q96.

The initial products include a faster version of the Alpha 21164, which uses four metal layers, and the first StrongArm chips, which will use only three. This process could push Alpha clock speeds well beyond 400 MHz.

## Routing and Cost Indices

The routing index shown in the table attempts to capture the circuit density of a process. If a design in one process is reworked to take full advantage of a smaller process, the die area should change by roughly the ratio of the routing indices of the two processes. Because most global routing is done with metal-2 and metal-3, we calculate the index as the product of these two pitches, with small adjustments for stacked vias and additional routing layers. For processes with a local interconnect, some global routing can be done with metal-1; in these cases, the index multiplies metal-2 by the mean of metal-1 and metal-3.

The wafer cost index is based on the MDR Cost Model (*see 081203.PDF*), which computes wafer costs based on the wafer size, drawn gate length, number of metal and poly layers, local interconnect, and bipolar layers. The wafer costs are expressed here as a ratio, with a 0.5-micron three-layer-metal process as the base value (1.0). Fujitsu, the only vendor in our survey still using 150-mm wafers, is assessed a 10% cost penalty. Although the Cost Model assesses volume discounts, for this comparison all costs are computed with equivalent volumes.

Other parameters in the table are simpler. Drawn gate length is the physical width of the polysilicon traces. Metal pitches are measured from contact to contact. One caution: the method of measuring effective channel length varies among manufacturers.

Digital has also developed a five-layer-metal version of CMOS-6 but has no immediate product plans for it.

### MIPS Vendors Emphasize Gate Length

NEC became the first vendor to ship 0.35-micron processors with its 200-MHz R4400 chips last fall. The company recently announced a 250-MHz version of that processor using this process. Both NEC and Toshiba will build the upcoming R10000 using this 0.35-micron process; this chip will use four metal layers, while the R4400 uses only two. IDT currently has the smallest gate length of all the vendors surveyed, producing its CEMOS 8+ process with 0.3-micron gates.

Both the NEC and IDT processes use relatively thick gate oxides, yet the narrow poly traces help these processes to deliver competitive speed. Another drawback of the MIPS vendors' processes is poor metal pitch. NEC helps compensate for this shortfall with a local-interconnect layer, and it plans to make a significant improvement in metal pitches in its 0.25-micron process. CEMOS 8+ has no local interconnect and even wider metal-3 than NEC, giving IDT inferior circuit density.

IDT plans to move to 0.25-micron gates by early 1996, which would make it one of the first to reach that level. This CEMOS 9 process will shrink the metal lay-

ers by 20% and add a fourth metal layer. IDT currently supports two poly layers and plans to add a third in CEMOS 9 to further reduce SRAM cell size.

### Other CPU Makers Move Later

Fujitsu, which builds MicroSparc, HyperSparc, and Hal's Sparc64 chips, continues to advance its process technology. Fujitsu is currently building its SPARC portfolio in its 0.4-micron CS-55 process. This process offers a 30–40% speedup from the 0.5-micron CS-50 but uses the same metal pitches as its predecessor, so it does not reduce die size.

Fujitsu's next step is CS-60ALE, a true 0.35-micron process with much smaller metal pitches, putting it among the leaders in density. Products moving to this process could have their die size cut in half and gain 20% or so in clock speed. Fujitsu expects the first products in this process to reach production in 1Q96.

VLSI Technology and Hitachi have continued their long-standing partnership by developing a 0.38-micron process. Both companies will manufacture this process and can second-source ASICs. Each will also use it for its own processors. For Hitachi, this means faster versions of the current SH-3 chips as well as the future SH-4 design. VLSI is likely to build ARM7 and ARM8 devices in this process. The companies don't expect to see products from the new process until 2Q96, putting them behind most others in reaching 0.35 microns.

Even when this process reaches production, it will not be fully competitive. Its slightly larger effective channel length and relatively thick gate oxide will limit performance. At 1.4 microns, the metal layers are competitive but not outstanding. This process was clearly designed for low cost, not peak performance.

### No End in Sight for Process Development

None of the companies interviewed sees any immediate slowing of process development. In fact, with IDT, IBM, and TI planning 0.25-micron processes in 1996, there may even be a slight increase in the pace, which has averaged about two years per generation over the past few cycles. CMP and DUV will add cost to next-generation facilities, making it more expensive to stay in the game, but for those companies that ante up, the next round will come quickly.

Other vendors may find the cost too steep. Hewlett-Packard, which trailed in the 0.5-micron race by nearly two years, may turn to new partner Intel for production of its future processors. The Hitachi/VLSI partnership, trailing in 0.35-micron implementation, stays competitive only because neither company builds processors with leading-edge performance, aiming instead at the embedded market. For those CPU vendors that want to stay in the performance race, heavy investments will be required, with Intel and IBM pacing this spending. ♦