IC Vendors Prepare for 0.25-Micron Leap

Texas Instruments, IBM Have Early Lead Among Microprocessor Makers

by Linley Gwennap

This article updates our look at major microprocessor vendors' IC process technology. See **080504.PDF**, *and* **090905.PDF**, *for details on vendors' 0.5-micron and 0.35-micron processes, many of which are in production today.*

Although pessimists claim the development cycle of new IC process technology is extending to three years, several vendors plan to deploy 0.25-micron microprocessors in 1997, in most cases just a year or two after their 0.35-micron chips first appeared. Once these new processes reach production, computer designers (and users) will see another big leap in price/performance.

As in the past, the new process generation will offer increased clock speeds and greater transistor density. Greater density can either reduce the cost of processors by making them smaller, or it can allow CPU designers to add more complex circuitry—and hopefully performance—to their designs. Typically, processor vendors take both paths, cost-reducing their older chips while introducing a new processor at the high end.

A CPU vendor's IC process technology is at least as important as its design capabilities. When clock speed and transistor-count improvements are combined, gaining a step in the process-technology race can double performance over rivals stuck with an older technology. Since all vendors have access to the same manufacturing tools, getting very far ahead of the pack is tough; with the cost of new fabs skyrocketing, however, falling behind is frighteningly easy.

Cost has not deterred Texas Instruments, which has made massive investments ($1.4 billion in 1995) to push its process technology from a trailing position to the front of the pack. TI is already building 0.29-micron parts and expects to reach the 0.25-micron level in 1Q97. IBM is at 0.27 microns today and, along with partner Motorola, hopes to reach 0.25 microns at about the same time as TI.

These early efforts are hybrid processes, combining 0.25-micron transistors with the metal layers of a 0.35-micron process. Both TI and the PowerPC partners expect to improve their metal layers in 2H97, but by that time most other major microprocessor vendors will have caught up, offering their own 0.25-micron processes. Of these 2H97 processes, TI's appears the best, but it remains to be seen whether all of these vendors will deliver as planned.

Voltage Continues to Plummet

The forthcoming 0.25-micron processes share many characteristics. In addition to reducing the length of the transistor, vendors typically reduce the gate oxide thickness as well. The typical oxide thickness for a 0.25-micron process is about 50 Å, compared with about 70 Å in the 0.35-micron generation. Since an angstrom is roughly the width of a single atom, these oxides are very thin indeed.

The thinner oxides help transistors switch faster. The downside is a reduced voltage tolerance. After moving to 3.3 V for 0.35-micron devices, most vendors are switching to a 2.5-V or 1.8-V supply for their 0.25-micron devices. Some will offer 3.3-V I/O for compatibility with older system designs, but support for traditional 5-V I/O will be rare.

All the 0.25-micron processes discussed here are pure CMOS. The last bastions of BiCMOS, Intel and TI, have given up on bipolar support. Although bipolar transistors provide some performance boost at 3.3 V, the gain is near zero at 2.5 V and below. Jettisoning the extra bipolar layers reduces wafer cost by 10% to 20% and improves yield.

The number of metal layers continues to rise, with five being commonplace among 0.25-micron processes, and a few vendors are pushing forward to six. Most manufacturers have now moved to chemical-mechanical polishing (CMP), which smooths the surface of the chip after each new layer is deposited. With CMP, metal layers can be stacked to virtually any height; beyond five or six, however, additional metal layers typically don't reduce the die size by enough to offset the incremental processing cost. In some cases, the extra metal layers improve performance but not cost.

Most vendors have also succumbed to the difficulties of pushing i-line (mid-ultraviolet) steppers to finer pitches and are putting DUV (deep-ultraviolet) steppers in place for 0.25-micron production. TI is the only one that expects to get 0.25-micron transistors out of its i-line steppers. Putting off the DUV transition has helped TI get its 0.25-micron process running more quickly, but other vendors' investments in DUV should pay off further down the road.

One factor that doesn't appear to be improving is die size. After reaching 300 mm^2 or so last year, no vendors are projecting 400-mm^2 processors in the foreseeable future, despite a historical trend of annual increases of roughly 10% in die size. Although such a large die could conceivably be built, low yields would make it uneconomical. Defect rates would have to drop considerably to enable larger die, but defects are staying constant from generation to generation. Holding die size constant will slightly reduce the growth rate in transistor count, possibly impacting performance increases over time.

Intel Focuses on Metal Pitch

Although Intel will probably not be the first to market with a 0.25-micron transistor, it should be among the first with a

Vendor	AMD	Digital	Fujitsu	IBM*	IDT	Intel	TI	TI
Process name	CS-44	CMOS-7	CS-70	CMOS-6X	CEMOS-10+	P856	C07	C07
Example product	K6+	21264+	n/a	PPC 60x+	n/a	Deschutes	n/a	n/a
First production	2H97	1H98	2H97	2H97	1H98	3Q97	3Q97	1Q98
Supply voltage	2.5 V	1.8 V	2.5 V	1.8 V	2.5 V	1.8 V	1.8 V	1.8 V
I/O voltage (max)	3.3 V	3.3 V	3.3 V	3.3 V	3.3 V	2.5 V	3.3 V	3.3 V
Gate length (drawn)	0.25 µm	0.25 µm	0.24 µm	<0.25 µm	0.25 µm	<0.25 µm	0.21 µm	0.18 µm
Channel length (effective)	0.18 µm	0.16 µm	0.18 µm	n/a	0.20 µm	n/a	0.17 µm	0.14 µm
Gate oxide thickness	n/a	45 Å	55 Å	40 Å	65 Å	45 Å	40 Å	36 Å
Number of metal layers	5 metal	6 metal	5 metal	6 metal	4 metal	5 metal	5 metal	5 metal
Local interconnect?	yes	no	yes	yes	no	no	no	no
Stacked vias?	yes	yes	yes	yes	yes	yes	yes	yes
M1 contacted pitch	0.88 µm	0.84 µm	0.9 µm	0.7 µm	0.94 µm	0.64 µm	0.85 µm	0.85 µm
M2 contacted pitch	0.88 µm	0.84 µm	0.9 µm	0.9 µm	1.1 µm	0.93 µm	0.85 µm	0.85 µm
M3 contacted pitch	0.88 µm	1.7 µm	0.9 µm	0.9 µm	1.1 µm	0.93 µm	0.85 µm	0.85 µm
M4 contacted pitch	1.13 µm	1.7 µm	0.9 µm	0.9 µm	1.1 µm	1.6 µm	0.85 µm	0.85 µm
M5 contacted pitch	3.0 µm	1.7 µm	2.7 µm	0.9 µm	1.4 µm	2.6 µm	2.5 µm	2.5 µm
SRAM cell size	n/a	11.5 µm^2	n/a	8.6 µm^2	11.2 µm^2	10.3 µm^2	10.5 µm^2	10.5 µm^2
Routing index	0.60 µm^2	1.1 µm^2	0.62 µm^2	0.53 µm^2	1.0 µm^2	0.67 µm^2	0.56 µm^2	0.56 µm^2
Wafer cost index	$4.0	$3.5	$4.0	$4.7	$3.6	$4.0	$4.1	$4.2

Table 1. Among true 0.25-micron processes, IBM's CMOS-6X and TI's C07 offer superior circuit density, as represented by the routing index, but IBM's process is more costly because of its extra metal layer and local interconnect. See sidebar (page 14) for more information on the parameters and indices (smaller is better). *Motorola's PPC4 is similar to CMOS-6X but may have smaller gates. + indicates shrink version. n/a indicates information not available from vendor. (Source: vendors, except indices by MDR)

contacted metal pitch of less than 1.0 micron, the mark of a "true" quarter-micron process. Metal pitch is an often-ignored parameter that may be more important than the quoted gate length. The metal pitch, particularly for the first few layers, is the critical factor in circuit density. Table 1 attempts to summarize the overall benefit of the metal layers in a single routing index.

Based on this calculation, Intel's 0.25-micron process, which it calls P856, will be one of the densest quarter-micron processes. It features a contacted metal pitch of just 0.64 µm for metal 1 and 0.93 µm for metal 2 and metal 3. The very small metal-1 pitch helps reduce SRAM cell size in the absence of local interconnect.

In addition to packing more chips on a wafer, this focus on metal pitch can also improve performance. On a smaller die, traces are shorter and thus have less resistance and capacitance, speeding signal transmission and reducing power. One countervailing factor is an increase in resistivity when metal pitches are reduced. To combat this, Intel uses metal traces that are thicker than they are wide, resulting in a aspect ratio of roughly 2:1.

P856 is tuned for operation at 1.8 V, lower than most other 0.25-micron processes. This low voltage will reduce the power of the current P6 core by as much as 75% at the same clock speed. Notebook versions of the P6 may operate at even lower voltages, giving up a bit of performance. P856 handles 2.5-V I/O, suitable for the GTL+ interface of the P6, but higher I/O voltages may not be supported.

Intel plans to release its first 0.25-micron processors in 3Q97, a bit later than some CPU vendors. The process is designed as a linear shrink from the CMOS version of Intel's 0.28-micron P854 process; the lead product will be the Deschutes processor, an optical shrink of the P854 Klamath design. We expect this P6-family device to achieve clock speeds of up to 333 MHz once P856 is fully mature.

TI Attempts to Grab Lead in Transistor Size

A year ago, TI laid out an aggressive plan to be the first microprocessor vendor to market with a 0.25-micron product. So far, these efforts are on track. A key milestone was the recent delivery of UltraSparc-2 processors, the first chips in TI's C10 process. As Table 2 shows, these first C10 products use a 0.29-micron drawn gate, but TI plans to reach the 0.25-micron level with a minor adjustment to the C10 transistor while keeping the same metal layers. The company hopes to deploy the 0.25-micron version of C10 in 1Q97. We expect this process to boost UltraSparc-2 clock speeds beyond 300 MHz in that time frame.

Although the company previously thought C10 would require DUV steppers, it now plans to stick with i-line technology. The company had also planned to move away from tungsten plugs in C10, instead using force-filled aluminum to fill vias. Both of these changes have been delayed until the next process generation, speeding the company's move to 0.25-micron CMOS.

TI's 0.25-micron plan appeared aggressive when it was first revealed; its new 0.18-micron plan seems incredible. The next step after C10 is a process called C07. The company expects to reduce the contacted metal pitches from 1.2 to 0.85 microns; according to our routing index, this will double the number of transistors that can fit into the same die area. The initial products, which TI foresees in production by 3Q97, will use a 0.21-micron gate length. As with C10, a gate shrink is planned to occur about six months later, drop-

ping the gate length to 0.18 microns in 1Q98 while trimming the gate-oxide thickness to just 36 Å.

Although this plan, if successfully executed, would make TI the first to reach 0.18-micron gates, the metal layers of the C07 process are similar to those in most 0.25-micron processes. Thus, from a density standpoint, TI's 0.18-micron process will have little advantage over these 0.25-micron processes. Furthermore, even the faster transistors may not have a significant impact because of delays due to the interconnect.

In C07, TI is addressing interconnect delays by using new materials to fill the gaps between metal layers. Reducing the dielectric constant of these materials has a corresponding effect on capacitance. Lower capacitance, in turn, speeds signal transmission and reduces power consumption. Other vendors are working on similar materials, although Intel, for one, would not comment.

TI offers the option of dual oxide thicknesses in its processes. For chips that require 3.3-V or (in C10) even 5-V I/O, the I/O drivers can use special transistors with a thicker gate oxide that is more reliable at the higher voltages. If used, this option adds a small cost but gives designers more flexibility in supporting legacy system interfaces.

Vendor	AMD	IBM*	IDT	TI
Process name	CS-34EX	CMOS-6S2	CEMOS-9+	C10
Example product	K6	PPC 60x+	R4700+	UltraSparc-2
First production	1H97	1H97	4Q96	3Q96/1Q97
Supply voltage	2.7 V	2.5 V	3.3 V	2.5 V
I/O voltage (max)	3.3 V	5.0 V	5.0 V	5.0 V
Gate length (drawn)	0.30 µm	0.25 µm	0.28 µm	0.29/0.25 µm
Channel length (effective)	0.22 µm	0.18 µm	0.23 µm	0.25/0.21 µm
Gate oxide thickness	70 Å	50 Å	75 Å	57 Å
Number of metal layers	5 metal	5 metal	4 metal	5 metal
Local interconnect?	yes	yes	no	no
Stacked vias?	yes	yes	optional	yes
M1 contacted pitch	1.4 µm	1.0 µm	1.17 µm	1.2 µm
M2 contacted pitch	1.4 µm	1.25 µm	1.4 µm	1.2 µm
M3 contacted pitch	1.4 µm	1.25 µm	1.4 µm	1.2 µm
M4 contacted pitch	1.8 µm	1.25 µm	1.8 µm	1.2 µm
M5 contacted pitch	4.8 µm	1.25 µm	—	2.5 µm
SRAM cell size	n/a	16 µm^2	17.5 µm^2	22 µm^2
Routing index	1.5 µm^2	1.1 µm^2	1.7 µm^2	1.1 µm^2
Wafer cost index	$3.1	$3.6	$2.8	$3.2/$3.4

Table 2. These hybrid processes combine sub-0.35-micron gates with the metal layers of a 0.35-micron process. *Motorola's PPC3 is similar to CMOS-6S2 but has a 0.23-µm gate. — indicates not applicable. n/a indicates information not available from vendor. (Source: vendors, except indices by MDR)

PowerPC Vendors Make Rapid Progress

IBM has already deployed its sixth-generation process. The first version for microprocessors, CMOS-6S, has a physical gate length of 0.27 microns, although IBM sometimes refers to it as a 0.29-micron process. This process is currently in production with the P2SC processor *(see* **101104.PDF***)*. To reach a 0.25-micron gate, IBM plans a linear 10% shrink, resulting in a process called CMOS-6S2. This latter version should be in production in 1H97. Because of the linear shrink, it will be easy to move products from CMOS-6S to 6S2.

The metal pitches of 6S2 are comparable to those in Intel's 0.35-micron process, however. According to the routing index, IBM's 6S2 processors will be about 60% larger than Intel's 0.25-micron chips, assuming the basic designs are of similar complexity. This increased die size puts a cost burden on the PowerPC processors and Cyrix chips that IBM builds.

IBM plans to remedy this problem with its CMOS-6X process, due to begin production in 3Q97, about the same time as Intel's P856. This process tightens the metal pitches while adding a sixth metal layer. When using all six metal layers, IBM's processors may actually be smaller than comparable Intel chips. IBM has historically taken much longer than Intel to move its processors into a new process and ramp its volume; we will see whether it can pump out processors in its 6X technology at the same rate that Intel ramps P856.

Last year, Motorola invested more money ($2.2 billion) to improve its IC manufacturing capabilities than any other semiconductor company except Intel. This investment has paid off, as the company has gone from a laggard with technology best suited for the embedded market to near-parity with IBM. Although the two partners develop their process technology separately, they must support essentially the same design rules, allowing them to produce the same PowerPC chips.

Motorola's plans include a PPC3 process that matches IBM's CMOS-6S2, except Motorola's gate length is 10% smaller. In theory, this should allow Motorola to produce faster clock speeds than IBM for PowerPC chips. When IBM deploys CMOS-6X, Motorola intends to match it with a similar process called PPC4.

Digital's High Clock Speeds Constrain Density

With a much smaller CPU revenue stream, Digital is unable to match the investment of larger microprocessor players. Even with a smaller budget, however, the company has managed to deliver the manufacturing technology needed for Alpha processors to reach their eye-popping clock speeds. Since both Samsung and Mitsubishi have licensed Digital's process technology, the U.S. vendor is also driving the logic processes of these Asian vendors, although neither has gotten Digital's 0.35-micron process into production yet.

Digital's 0.25-micron CMOS-7 process is not due until 1H98, a year or so behind the leaders. Most of the equipment in its Hudson (Mass.) fab is ready for 0.25-micron production, including the DUV steppers needed for the critical layers, but the company still needs to fine-tune the production parameters. Although the process is a bit late, it may be worth the wait: to improve switching time, the transistors use a 45-Å oxide. To maintain reliability at this thinness, the

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 wafer size, drawn gate length, number of metal and poly layers, metal pitches, and local interconnect. The wafer costs are expressed as a ratio, with a 0.5-micron three-layer-metal 150-mm wafer as the base value (1.0). 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, and some vendors have chosen not to quote this parameter at all.

oxide is built from "treated" SiO_2, and the supply voltage is constrained to 1.8 V.

In addition to 0.25-micron transistors, CMOS-7 features 0.84-micron contacted metal pitches for layers 1–3. These are the tightest pitches for metal 2 and 3 of any 0.25-micron process described here, giving Digital a superior routing capability. The company also plans to support up to four "coarse" metal layers for a total of seven, more than in any other vendor's plans. This works out to a routing index of 0.53, the best of any process discussed here.

Unfortunately, none of Digital's upcoming processors is likely to take advantage of this density. Both the high-end Alpha chips and the embedded StrongArm processors run at relatively high clock speeds: up to 500 MHz in the case of Alpha. To achieve these speeds, Digital cannot use narrow traces to route most global signals; there is not enough time in the cycle to drive a signal across the chip through a long, skinny trace. For this reason, all StrongArm and Alpha processors to date, including the forthcoming 21264, are designed with a coarse metal-3 layer. In this configuration, shown in Table 1, the CMOS-7 routing index balloons to 1.1, preventing Digital's processors from achieving circuit density that is competitive with that of other 0.25-micron processors.

The 21264 uses a total of six metal layers, including four coarse layers. These extra layers help manage other problems caused by fast clock speeds, such as minimizing clock skew across a large die and ensuring that V_{CC} stays close to 1.8 V even when large numbers of signals are switching at once. The wider, thicker metal layers greatly reduce resistivity. Digital also maintains thick (up to 2 μm) dielectric layers to reduce capacitance.

By addressing these RC issues, Digital can minimize the twin phenomena of AC and DC "droop" across large die at high clock speeds. Given that the current CMOS-6 chips are reaching 500 MHz, we expect CMOS-7 to produce clock speeds up to 750 MHz when it appears in 1H98.

AMD Simulates IBM's Process for K6

AMD is currently producing its 5x86 and 5K86 processors in its 0.35-micron CS-34 process. Later this quarter, the company will begin shipping 5K86 chips using CS-34E, which has the same metal layers but reduces the transistor length to 0.30 microns. Although this gate shrink will not affect die size, AMD hopes to achieve a 25–30% speed boost from the enhanced process and circuit-design improvements.

To produce its next processor, the K6, AMD needed to further modify its process technology. The K6 was originally designed at NexGen as the Nx686, to be fabbed in IBM's 0.33-micron CMOS-5X process. That design took advantage of IBM's five metal layers, local interconnect, and C4 (flip-chip) attachment. To avoid bloating the K6 design, AMD added these features to its process portfolio. The new process, CS-34EX, combines the 0.30-micron transistor with these new features.

AMD is planning a big leap when it moves to its 0.25-micron process, CS-44. This process reduces the metal pitches to 0.9 microns for the first three layers while retaining the five metal layers and local interconnect of CS-34EX, resulting in a routing index of 0.60. This ranks among the best 0.25-micron processes and could result in as much as a 60% reduction in the die size of the initial K6. The company plans to put CS-44 into production in 2H97, about the same time as Intel's P856.

Other Vendors Also Have 0.25-Micron Plans

As usual, the MIPS processor vendors have aggressive plans to reach the next process level ahead of most other vendors. NEC, Toshiba, and IDT have historically emphasized smaller transistors, but their metal pitches have lagged. Both NEC and Toshiba plan to begin shipping a 0.25-micron version of the R10000 in 1Q97, reaching that level at about the same time as TI. The companies would not provide details about the process for this part, but we expect it to be a hybrid process with 0.35-micron-class metal layers.

IDT plans to put a 0.28-micron version of the R4700 into production late this year, using its CEMOS-9+ process. Because it is aimed at low-cost products, this process supports just three metal layers, with an optional fourth. Even with four layers, the routing index is 1.7, which is worse than for most 0.35-micron processes.

The SRAM vendor has developed a 0.25-micron process, CEMOS-10, for its memory chips, but since it is not producing the R10000, IDT has no immediate plans to convert this into a quarter-micron logic process. The logic version, CEMOS-10+, is expected to appear by 1H98, when it makes sense for IDT to move its embedded processors to the next level. This process reduces the metal pitches by 20% and adds a fifth layer of metal, pulling the routing index down to 0.93, still well behind that of most 0.25-micron processes. The SRAM cell size is very competitive, however, due to the use of two polysilicon layers instead of the usual one.

Fujitsu, which builds the HyperSparc and Hal processors, plans to jump to a 0.25-micron process in 2H97, along with the rest of the pack. As Table 1 shows, this CS-70 process is similar in most ways to the leading 0.25-micron processes. It includes a local interconnect layer to keep SRAM cell size small. This process should keep Fujitsu's SPARC processors competitive.

Although HP has not officially acknowledged it, the company appears to have given up the pursuit of ever-smaller transistors. HP has yet to put even a 0.35-micron process into production. The company admits it will not fab Merced, the first Intel/HP processor, and it seems likely that Intel will be the sole source of subsequent parts as well. HP has discussed a 0.35-micron shrink of the PA-8000, but it isn't clear whether the company will build that part or seek an outside foundry.

Forging Ahead to 0.18-Micron Gates

Aside from TI, which plans to reach 0.18-micron transistors in 1Q98, other vendors would not discuss their 0.18-micron plans. We don't expect most vendors to reach this level until 1999, about two years after reaching 0.25 microns. In 1998, some vendors will probably hit an intermediate point, such as 0.22 microns, while others (including TI) will offer smaller transistors but without the tight metal pitches of a true 0.18-micron process.

Although some issues remain with 0.18-micron production, all the vendors surveyed believe these issues can and will be solved. After that, no one is quite sure what will happen. Naysayers have always seen a wall after the next couple of generations, but somehow the industry has always found a way to get through. We see no reason why this barrier, too, cannot be surmounted. MR