MICROPROCESSOR REPORT
VOLUME 12, NUMBER 12
SEPTEMBER 14, 1998
THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

The Race to Point One Eight

Microprocessor Vendors Gear Up for 0.18 Micron in 1999

by Keith Diefendorff

Headed for the 100-million-transistor mark, all major microprocessor vendors next year will push their fabs into the 0.18-micron regime—fully two years ahead of the prediction made in the 1994 Semiconductor Industry Association's *National Technology Roadmap for Semiconductors*. Manufacturers succeeded in blowing through the 0.25-micron generation with apparent ease and, since 0.18 micron requires no fundamentally new technologies or equipment, we expect a smooth and rapid transition to 0.18 micron.

IBM is leading the way with its aggressive moves to copper wires (see MPR 8/4/97, p. 14) and silicon-on-insulator (see MPR 8/24/98, p. 8). Motorola, following a similar roadmap, is hot on IBM's heels with both copper, as Figure 1 shows, and SOI. Other manufacturers are pushing lithography as hard but are being more conservative with new materials, choosing to first take the smaller step of implementing low-dielectric-constant (*k*) insulation.

Some manufacturers will go straight into 0.18-micron production with simple 30% shrinks of their aluminum-based 0.25-micron processes. These companies will upgrade their 0.18-micron process over time with copper metallurgy and low-*k* dielectrics to boost speed and gain the manufacturing experience they will need for 0.13 micron, where these new materials will be the ante.

Challenges for 0.18 Micron

Because 0.18-micron circuits can be fabricated using the same 248-nm deep-ultraviolet (DUV) steppers that are currently in use for 0.25 micron, and because there are no other fundamentally new manufacturing procedures or equipment required, the initial transition to 0.18 micron should not be difficult. The half-generation shrink to 0.15-micron features, which we expect a year later, will require some lithography improvements, such as phase-shift masks or new 193-nm DUV steppers. Since 193-nm DUV steppers must be in place for 0.13 micron anyway, we expect manufacturers to upgrade to them during the course of the 0.18-micron generation. At a cost of over $20 million per complete station, this upgrade will involve a substantial capital investment.

This is not to say that the 0.18-micron transition will be easy. Not only does the 30% reduction in feature size make manufacturing more difficult, the physics of smaller devices requires even tighter manufacturing tolerances. Making matters worse, smaller feature sizes make lethal defects out of many particles that were inconsequential at 0.25 micron. Thus, to achieve good yields, improvements must be made in fab cleanliness and the purity of process chemicals, increasing fab costs. Intel is planning to spend $5 billion on capital equipment this year, the majority being targeted to 0.18-micron production facilities.

Originally targeted for 1999 introduction, 300-mm (12") wafer fabs have been put on the back burner for now. The industry appears confident that 0.18-micron demand can be satisfied with existing 200-mm (8") fabs, and it prefers to put off the enormous cost of constructing new 300-mm fabs for a while longer. The 300-mm wafers, which should reduce wafer costs by 30% (per mm^2), are now being moved out to the 0.13-micron generation in 2001 or 2002.

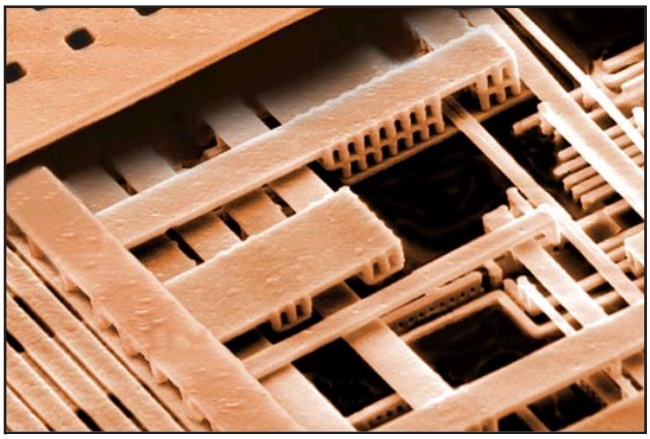

Figure 1. This photomicrograph shows all six copper interconnect layers from Motorola's upcoming 0.18-micron HiPerMOS 6 process. (Source: Motorola)

The most significant change in store for the 0.18-micron generation involves on-chip interconnects, which are rapidly coming to dominate circuit performance. As Figure 2 shows, 0.15- and 0.13-micron circuits could actually be slower without improvements in interconnect technology to reduce resistance and capacitance. This dilemma will force all manufacturers to new metallurgy (for lower resistance) and to new low-*k* dielectrics (for lower capacitance) by the time they reach 0.13 micron.

IBM Leads With Copper and SOI

IBM is now in production with its copper 0.22-micron CMOS-7S process on the 400-MHz Lonestar processor (see MPR 9/14/98, p. 4); Lonestar is Somerset's PowerPC 750 that IBM has remapped from 0.28-micron CMOS-6S2 to CMOS-7S.

In 1Q99, IBM will upgrade CMOS-7S with silicon-on-insulator (SOI) technology, a move that should boost Lonestar to 500 MHz or more. IBM will be the first microprocessor vendor to go into volume production with SOI.

IBM will enter production with its 0.18-micron CMOS-8S in the middle of next year. CMOS-8S will use copper wires and add a seventh layer of interconnect. It will be IBM's first production process that employs a low-*k* dielectric. IBM plans to use FSG, a fluorine-doped spin-on glass with a dielectric constant of 3.5–3.6, about 10% lower than the silicon dioxide (SiO_2) used in 7S.

CMOS-8S, like 7S, will initially be built on bulk-silicon wafers, then upgraded to SOI about six months later. We anticipate that IBM's Giga processor, targeted to run at more than 1 GHz, will be built using 8S-SOI in 2000.

Motorola Following IBM-like Strategy

Motorola is currently in production with its 0.25-micron HiPerMOS 4 (HIP 4) process, which the company uses for its 350-MHz PowerPC 604e (Mach 5). HIP 4 is roughly equivalent to IBM's CMOS-6X; not a coincidence, as IBM and Motorola have for years kept their technologies synchronized to accommodate Somerset-designed processors.

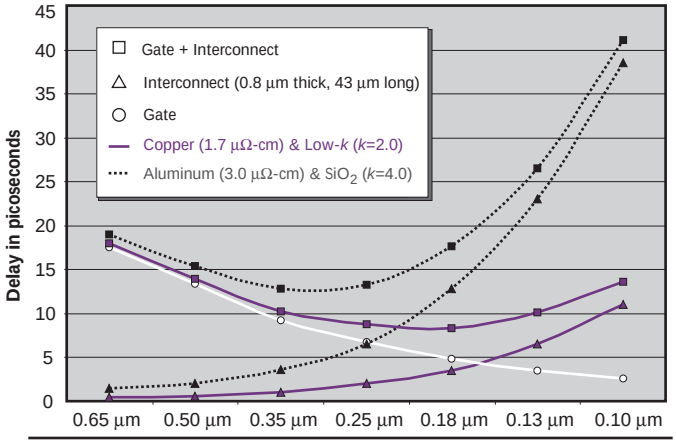

Figure 2. Gate speed continues to improve as geometries shrink, but interconnects slow down. As this example shows, copper and low-*k* dielectrics will become mandatory in future generations. (Source: Mark Bohr, Intel, IEDM 95)

In 4Q98, Motorola will go into production with its copper 0.22-micron HIP 5 process, which is similar to IBM's CMOS-7S. The first HIP 5 processor will be the new G4, which will be disclosed at the Microprocessor Forum next month. In the first half of 1999, Motorola plans to upgrade HIP 5 with a low-*k* dielectric and faster transistors, a move the company expects will increase speed by 20% over processors built with the initial HIP 5.

Now that Motorola and IBM have parted ways on processor development at Somerset, HIP 5 is likely to be the last process similar to IBM's. Working together for many years as they have, however, has set the two companies on a similar course, at least philosophically. This is evident, for example, in Motorola's plan to adopt SOI technology, a move that many other vendors are not contemplating.

Motorola's next-generation 0.18-micron HIP 6 process, shown in Figure 3, should appear in late 1999 and be upgraded to SOI by 2H00. HIP 6 will employ six layers of copper interconnect and use FSG as its low-*k* dielectric.

Motorola did not say what processor it would use as the process driver for HIP 6, but Somerset is working on a follow-on to G4, which is rumored to implement multiple processors on one chip along with a large on-chip L2 cache. With HIP 6, the processor should easily achieve 600 MHz and, after the SOI upgrade, possibly 800 MHz.

Motorola is already working on its 0.15-micron, 1.2-V HIP 7 process, which the company expects will allow it to surpass the one gigahertz mark in 2001. One innovation that Motorola is hoping to introduce with HIP 7 is a high-dielectric-constant gate insulation material. The purpose of a high-*k* gate insulator is to counteract the short-channel effect whereby the gate loses control over the channel, robbing the transistor of gain and reducing its performance.

The traditional solution to this problem has been to reduce the gate-oxide thickness (T_{ox}). Unfortunately, this necessitates lowering the supply voltage, which also lowers performance. The solution Motorola, among others, is pursuing is a high-*k* gate insulating material that would increase the coupling of the gate to the channel without having to thin the gate oxide.

TI Leading With Low-k

Texas Instruments, determined to keep its own DSPs and Sun's UltraSparc processors competitive, has been investing heavily in process technology. The company is gearing up this quarter to go into production with its 18C07 process. This process, which the company labels 0.18 micron, is similar in most respects to the 7S and HIP 5 processes that IBM and Motorola call 0.22 micron.

Instead of copper wires, however, TI has implemented low-*k* intrametal dielectrics to reduce capacitive coupling, thereby improving on-chip interconnect performance. In 18C07, TI employs a spin-on glass called HSQ (hydrogen

silsesqui-oxane, $HSiO_{1.5}$) as an intrametal dielectric. HSQ has a *k* of 3.2, about 20% lower than the SiO_2 used by most other vendors and still used by TI between metal layers.

TI will deploy its so-called 0.15-micron 15C05 process late next year, about the same time that IBM and Motorola deploy their comparable 8S and HIP 6 processes. Sun expects to use 15C05 on a 750-MHz UltraSparc III processor in 1H99.

Although 15C05 will initially use HSQ and aluminum wires, TI will upgrade it a year later to all-copper interconnects, a move Sun will exploit in UltraSparc IV to break through 1 GHz. Copper on all layers is a change in plan for 15C05, which was to use copper for the upper two metal layers and aluminum for the lower four. When it moves to copper, TI will dispense with HSQ and, like IBM, go to FSG. While the *k* of FSG is not as low as that of HSQ, FSG can be used between layers as well as between traces, yielding a simpler, homogeneous process with a net overall lower *k*.

Last year, TI announced that it is developing a low-*k* dielectric material called xerogel, which is a highly porous silicon-dioxide material. Because air in the pores has a dielectric constant of 1.0, xerogels can deliver dielectric constants as low as 2.0. More work, however, remains to be done on xerogels, and at this time it is not clear when TI will integrate this technology.

TI recently divulged some details of its next generation 0.12-micron 12C35 process, which it intends to prototype in 2000 and put into production in 2001. The 1.2-V process will have a gate length of 0.10 micron and include a full 30% shrink of the metal pitches, giving it a metal-1 contacted pitch of 0.35 micron. The process will employ 7 or 8 levels of copper interconnect and use a dielectric material with *k* lower than 3.0. The process will probably use a nitrided gate oxide thinner than 30 Å. (Nitride in the gate oxide increases the *k* and improves reliability by preventing boron in the polysilicon gate from migrating into the channel.)

Another feature being planned for the 12C35 process is dual threshold voltages. With this feature, low-V_t devices can be used where high speed is required, while high-V_t devices are used elsewhere to keep overall leakage-current low. TI claims this feature gives some of the speed advantages of SOI but at lower cost and with much less yield impact. But IBM also offers low-V_t devices on 7S and gets a 7–10% performance boost from low-V_t devices in both SOI and bulk.

While TI continues to investigate SOI, at this time the company feels it is not economically justifiable, believing that it adds at least 20% to the processed wafer cost (IBM claims less than 10%). TI has no current plans to use SOI in mainstream DSPs, but we expect that Sun, in light of IBM's announcement, is pushing TI to re-evaluate.

A major objective of the 12C35 process is to integrate analog circuits on the same die with high-speed digital circuits. TI's goal is to produce system-on-a-chip DSPs operating at more than 1 GHz but dissipating only 1–3 watts. TI is also planning to use 12C35 on the newly revealed 1.5-GHz UltraSparc V that Sun is developing for 2001.

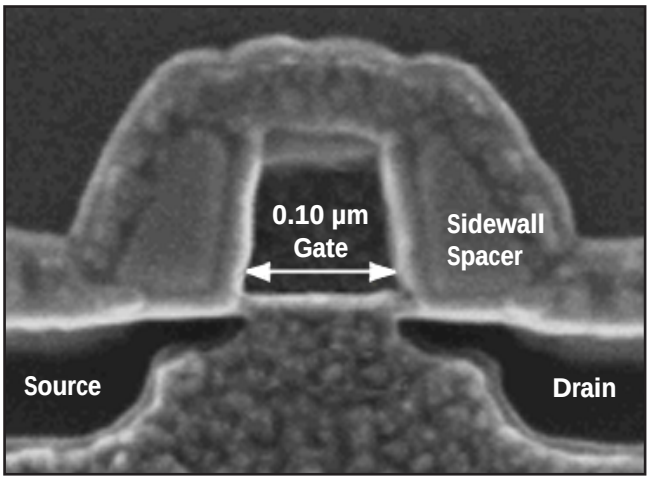

Figure 3. Transistor gate length is often pushed well below a process's typical minimum feature size, as shown in this photomicrograph of a transistor from Motorola's 0.18-micron HiPerMOS 6 process. (Source: Motorola)

Intel Keeping Its 0.18-Micron Plans Under Wraps

Intel was one of the first to deploy 0.25-micron technology when it put P856 into production in 3Q97 with its Tillamook mobile Pentium/MMX and later with its Deschutes Pentium II. Intel is now in production with P856.5, a 5% shrink of the original P856. The P856.5 process is used for Mendocino and for the 450-MHz Deschutes—which is 50% faster than the original 300-MHz Deschutes—and will be used for a 500-MHz Katmai early next year. According to our FET Performance Metric, the transistors in the latest version of P856.5 are 45% faster than those in the initial P856.

For higher frequencies, Intel will probably rely on its next-generation 0.18-micron P858 process. Intel will not disclose details on P858 until 1Q99, but said it expects P858 to begin production in 2Q99 with shipments in 3Q99. We expect the first P858-based processor to be Coppermine, a shrink of Katmai with a 256K or larger on-chip L2 cache that should run up to 700 MHz (Coppermine does not use copper interconnects despite its name.)

Intel is developing P858 in its D1B development fab in Hillsboro (Oregon). We expect that by the end of 2000, Fab 12 (Chandler, Ariz.), Fab 14 (Leixlip, Ireland), Fab 17 (Hudson, Mass.), and Fab 18 (Kiryat Gat, Israel) will all be converted to 0.18-micron production with an aggregate capacity of 28,000 wafers per week (15×10^{15} transistors per year!).

Intel did say that P858 will operate at 1.3–1.5 V and employ a low-*k* dielectric, but not copper wires. While Intel acknowledges copper's lower resistivity, it sees no frequency gain at the 0.18-micron level. Intel wants to avoid copper at this time because the manufacturing tools are not yet widely available, at least not enough for its enormous needs. As a result, Intel will wait until the 0.13-micron generation to introduce copper. Intel is investigating a copper-etch process—as opposed to IBM's inlaid dual-damascene process—that might simplify copper's integration into existing fabs.

Vendor Process	Intel P856.5	IBM CMOS-7S	Motorola HIP 5	AMD CS44E	National CMOS 8	VLSI VSC 9	VLSI VSC 10	TI 18C07
Process Generation **Example Product**	0.25 µm Mendocino	0.22 µm Lonestar	0.22 µm G4	0.25 µm K6	0.25 µm M II	0.25 µm n/a	0.20 µm n/a	0.18 µm Ultra-2
First Production	3Q98‡	3Q98	4Q98	4Q97	2Q98	1Q98	3Q98	3Q98
Supply Voltage	2.0 V	1.8 V	1.8 V	2.2 V	2.5 V	2.5 V	1.8 V	1.8 V
I/O Voltage (max)	2.5 V	3.3 V	3.3 V	3.3 V	3.6 V	3.3 V	3.3 V	3.3 V
Poly Half-Pitch	0.32 µm	0.27 µm	0.27 µm	0.31 µm	0.30 µm	0.30 µm	0.28 µm	0.30 µm
Gate Length (L_{gate}) **Gate Oxide (T_{ox})**	0.20 µm 41 Å	0.16 µm 35 Å	0.15 µm 35 Å	0.18 µm 37 Å	0.24 µm 50 Å	0.25 µm 55 Å	0.20 µm 40 Å	0.18 µm 40 Å
Substrate **Metal Layers**	Bulk Si 5 Al	Bulk Si 6 Cu	Bulk Si 6 Cu	Bulk Si 5 Al	Bulk Si 5 Al	Bulk Si 6 Al	Bulk Si 6 Al	Bulk Si 6 Al
M1 Contacted Pitch	0.61 µm	0.63 µm	0.63 µm	0.88 µm	0.64 µm	0.84 µm	0.72 µm	0.70 µm
M2 Contacted Pitch	0.88 µm	0.81 µm	0.81 µm	0.88 µm	0.80 µm	0.90 µm	0.80 µm	0.85 µm
M3 Contacted Pitch	0.88 µm	0.81 µm	0.81 µm	0.88 µm	0.80 µm	0.90 µm	0.80 µm	0.85 µm
M4 Contacted Pitch	1.73 µm	0.81 µm	0.81 µm	1.13 µm	0.80 µm	1.19 µm	1.19 µm	0.85 µm
M5 Contacted Pitch	2.43 µm	0.81 µm	0.81 µm	3.00 µm	0.90 µm	1.19 µm	1.19 µm	1.90 µm
M6 Contacted Pitch	–	0.81 µm	1.60 µm	–	–	4.00 µm	4.00 µm	1.90 µm
M7 Contacted Pitch	–	–	–	–	–	–	–	–
Local Interconnect	–	Tungsten	Tungsten	Tungsten	–	–	–	–
Intrametal Dielectric (*k*)	SiO_2 (3.9)	SiO_2 (3.9)	SiO_2 (3.9)	SiO_2 (3.9)	SiO_2 (3.9)	SiO_2 (3.9)	SiO_2 (3.9)	HSQ (3.2)
SRAM Cell Size	9.3 µm²	6.8 µm²	7.6 µm²	n/a	13 µm²	14 µm²	8.7 µm²	7 µm²
Ring Oscillator Stage	22 ps	22 ps	n/a	n/a	37 ps	32 ps	32 ps	27 ps
Routing Index* (µm²)	0.50	0.43	0.43	0.59	0.44	0.57	0.44	0.48
Wafer Cost Index* ($)	4.2	5.0	5.0	4.2	4.1	4.2	4.7	4.7
FET Performance* (GHz)	29.9	27.1	29†	25†	19.5	20.0	21.7	22.8

Table 1. Of this year's processes, IBM and Motorola have the most aggressive interconnects, and IBM, Intel, and Motorola have the fastest transistors. See sidebar on page 6 for more information on parameters. ‡ P856.5 is a 5% shrink of the original P856 that Intel put into production in 3Q97. (Source: vendors, except *MDR, † estimate.)

In the meantime, Intel is leading the industry in deploying organic flip-chip packages, which are now in millions of Intel processors. Organic substrates, like the FR4 used in printed circuit boards, allow copper interconnects at the package level. With flip-chip technology, this is like having copper on the upper routing layers of the chip itself.

Intel is negative, but not dogmatic, on SOI. In the near term, the company sees SOI as too expensive. In the long term, it believes junction capacitance will become a less significant contributor to the overall delay and that SOI's floating-body problem makes scaling more difficult. This opinion is in clear opposition to IBM's and Motorola's.

Intel's philosophy is that nothing is more effective than simple scaling. Therefore it prefers to avoid complex new materials and push scaling as hard as possible. While this philosophy is simple and elegant, it will run into a roadblock once interconnect delays dominate circuit performance. Intel doesn't see this as a problem until 0.13 micron.

AMD Teams With Motorola

Three years ago, AMD quickly upgraded its 0.35-micron CS34 process to accept NexGen's design (now the K6)—which was designed for IBM technology—but then struggled getting it into volume production. The company also had problems getting its 0.25-micron CS44E process into high-volume production in its large Fab 25 in Austin (Texas). Those problems have now been fixed, and 350-MHz CS44E-based K6-2 chips are in production, with 400 MHz expected in fourth quarter. The K7 will also be produced on CS44E in 1H99.

AMD expects to deploy its 0.18-micron CS50 process in Fab 25 with the K7 in 2H99. The process looks competitive with other 0.18-micron processes in most respects but will use only conventional aluminum interconnect and SiO_2 dielectric. AMD will upgrade CS50 with a low-*k* dielectric but has no current plans for adding copper in Fab 25. The supply voltage for CS50 was set lower than initially planned to accommodate the high power consumption of the K7, indicating that part is probably being targeted as a high-performance processor against the top of Intel's line.

Although AMD had been working on its own copper technology, it has now decided to join forces with Motorola to develop HIP 6 (see MPR 8/3/98, p. 10). AMD will implement this process in 2000 in its new Fab 30 in Dresden (Germany). With this process, AMD expects 1-GHz K7s.

AMD, like Intel and TI, is not enamored of SOI. AMD says it is working on SOI in the lab but has no plans to deploy it at 0.18 micron. Motorola, however, is planning to integrate SOI into HIP 6 so, presumably, Motorola and AMD will eventually align on this issue—one way or the other.

VLSI Enables Hot ASICs

While VLSI Technology may not have to compete with Intel for high-end microprocessors, it does have a number of ARM-based embedded processors and a large ASIC business that benefit immensely from dense, high-performance semiconductor processes. As Table 1 shows, we found VLSI's process plans to be quite competitive with other processes.

VLSI is in production with its 0.25-micron VSC9 process and plans to go into production with its 0.20-micron

Vendor Process	IBM 7S-SOI	IBM CMOS-8S	Motorola HIP 6	AMD CS50	National CMOS 9	VLSI VSC 11	TI 15C05	TI 15C05 Cu
Process Generation	0.22 µm	0.18 µm	0.18 µm	0.18 µm	0.18 µm	0.15 µm	0.15 µm	0.15 µm
Example Product	Lonestar+	n/a	G5	K7	M3D	n/a	Ultra-3	Ultra-4
First Production	1H99	2H99	1H00	2H99	4Q98	2H99	4Q99	4Q00
Supply Voltage	1.8 V	1.5 V	1.5 V	1.5 V	1.8 V	1.5 V	1.5 V	1.5 V
I/O Voltage (max)	3.3 V	2.5 V	2.5 V	2.5 V	3.3 V	3.3 V	3.3 V	3.3 V
Poly Half-Pitch	0.27 µm	0.21 µm	0.21 µm	0.23 µm	n/a	0.21 µm	0.23 µm	0.23 µm
Gate Length (L_{gate})	0.16 µm	<0.13 µm	0.13 µm	0.14 µm	0.18 µm	0.15 µm	0.13 µm	0.11 µm
Gate Oxide (T_{ox})	35 Å	<30 Å	30 Å	25 Å	40 Å	32 Å	29 Å	27 Å
Substrate	SOI	Bulk (SOI 1H00)	Bulk (SOI 2H00)	Bulk Si	Bulk Si	Bulk Si	Bulk Si	Bulk Si
Metal Layers	6 Cu	7 Cu	6 Cu	6 Cu	5 Cu	4 Al, 2 Cu	6-7 Cu	6-7 Cu
M1 Contacted Pitch	0.63 µm	0.49 µm	0.49 µm	0.63 µm	0.56 µm	0.46 µm	0.49 µm	0.49 µm
M2 Contacted Pitch	0.81 µm	0.63 µm	0.63 µm	0.63 µm	0.72 µm	0.52 µm	0.59 µm	0.59 µm
M3 Contacted Pitch	0.81 µm	0.63 µm	0.63 µm	0.63 µm	0.72 µm	0.52 µm	0.59 µm	0.59 µm
M4 Contacted Pitch	0.81 µm	0.63 µm	0.63 µm	0.90 µm	0.72 µm	0.80 µm	0.59 µm	0.59 µm
M5 Contacted Pitch	0.81 µm	0.63 µm	0.63 µm	0.90 µm	0.80 µm	0.80 µm	1.36 µm	1.36 µm
M6 Contacted Pitch	0.81 µm	1.26 µm	1.24 µm	2.16 µm	–	2.00 µm	1.40 µm	1.40 µm
M7 Contacted Pitch	–	1.26 µm	–	–	–	–	–	–
Local Interconnect	Tungsten	Tungsten	Tungsten	Tungsten	–	–	–	–
Intrametal Dielectric (*k*)	SiO_2 (3.9)	FSG (3.6)	FSG (3.6)	SiO_2 (3.9)	SiO_2 (3.9)	FSG (3.6)	HSQ (3.2)	FSG (3.6)
SRAM Cell Size	6.8 µm²	4.2 µm²	4.5 µm²	n/a	8.4 µm²	<3.7 µm²	3.5 µm²	3.5 µm²
Ring Oscillator Stage	17 ps	18 ps	n/a	n/a	26 ps	21 ps	23 ps	20 ps
Routing Index* (µm²)	0.43	0.25	0.26	0.29	0.35	0.19	0.23	0.23
Wafer Cost Index* ($)	5.0	6.0	5.5	5.3	4.5	5.5	5.5	5.7
FET Performance* (GHz)	30.9	33.0	33†	29†	22.4	32.0	26.9	33.4

Table 2. For next year's 0.18-micron processes, IBM is still in the lead but other manufacturers are closing the gap. The similarities among all these processes are more striking than the differences due to the fact that most of the vendors use the same or similar manufacturing tool sets. Only National is significantly behind, but its process enters production six months earlier than anyone else's. Note that the FET Performance metric does not capture the speed benefit from SOI's floating body; in practice 7S-SOI should be slightly faster than 8S, and 8S-SOI (not shown) faster yet. n/a: not available. (Source: vendors except *MDR, † estimate.)

VSC10 in 4Q98. The company's 0.18-micron VSC11 process, scheduled for production in 2H99, will begin life with six layers of aluminum interconnect and FSG dielectric. As Table 2 shows, VSC11's tight pitches give it outstanding density. In the first half of 2000, VLSI will upgrade the process by replacing the upper two layers of interconnect with copper—similar to the plan TI had originally contemplated but then abandoned.

Like VLSI's previous processes, VSC11 has an option for dual gate-oxide thicknesses. This feature allows VLSI to support high I/O voltages while at the same time having fast transistors in the core. Some manufacturers, such as TI, have also used this approach, while others, like IBM, prefer to use cleverly designed I/O drivers and receivers, which have less process complexity and lower cost.

Looking into the future, VLSI is experimenting with cobalt silicide and nickel silicide to reduce gate resistance in very short channel devices. The company is also looking into a carbon-rich SiO_2 dielectric with a *k* of 3.0, as well as parylene, a spin-on dielectric material with a *k* as low as 2.7. VLSI is not currently looking at SOI, even though the company admits it could be useful for reducing substrate coupling and noise in RF applications.

An important market for VLSI is battery-powered devices. To support these applications, the company is working on 1-V operation, multiple threshold and multiple supply voltages, and ultra-low-leakage transistors. The company is also working on embedded flash and DRAM. Embedded DRAM is an attractive option for filling space and eliminating pins in ASICs that are becoming increasingly pad-limited. VLSI says it is looking into a possible solution to the problem of compromised logic performance that has plagued embedded DRAM.

National Stresses System On a Chip

Believing that the industry lacks a performance-hungry application to drive demand for high-performance processors, National is focusing its processes on low cost and low power to support its system-on-a-chip strategy.

After years as a technology laggard, National is now in production with its 0.25-micron CMOS 8 process, which the company uses to build Cyrix's M II and MediaGX. Production will begin later this year on 0.18-micron CMOS 9, which the company will use to shrink the M II and to implement the MXi (see MPR 10/27/97, p. 22) and M3D. The M3D, which we expect in late 1999, is based on Cyrix's forthcoming Jalapeno core and, like the MXi, integrates 3D graphics.

CMOS 9 has transistor specifications similar to IBM's 7S but with tighter metal pitches; in fact, CMOS 9's metal pitches are tighter than any of the other process that will be deployed in 1998. National did not disclose details of its 0.15-micron CMOS 10 but did indicate that it would use copper wires and go into production in 2Q99. The company characterizes its SOI effort so far as "dabbling" but has plans to get more serious about it in 4Q99.

National has expressed concern about the return on investment for advanced process development. The company

New Process Metrics

L_{gate}: In previous years (see MPR 8/4/97, p. 14 and MPR 9/16/96, p. 11) we specified L_{drawn} as the physical gate length. But, especially for recent processes, this is a misnomer, since the physical gate is often smaller than the drawn gate after mask and process biases have been applied. This year, we are specifying L_{gate} to clarify that we mean the actual physical gate length on the finished die.

Poly Half-Pitch: Because poly lines (gates) are often made narrower than poly spaces (to speed up the transistors), L_{gate} is not a good indicator of feature sizes. Therefore, this year we are including the poly-line half-pitch, which is related to feature size and yet not subject to variations in gate length within a process generation.

L_{eff}: As process geometries have shrunk, the effective electrical channel length has become an ambiguous indicator of process speed, being more sensitive to the algorithm used to measure it than to anything else. As a result, we no longer report this parameter.

FET Performance Metric: In place of L_{eff}, we have developed an FET performance metric to indicate the relative transistor speed. This metric takes into account gate capacitance, junction capacitance, supply voltage, and p- and n-FET drive currents. Larger values are faster.

Wire Performance Metric: We are also working on an indicator of the relative delay of an interconnect wire, based on resistance and capacitance. Unfortunately, the data for this metric was not ready in time for this article.

has observed that while processor manufacturing costs are currently distributed roughly 80% to die cost, 10% to package cost, and 10% to test cost, its products may migrate toward a 30-45-25 distribution. In light of the skyrocketing costs of reticles, steppers, testers, and the like, National is questioning whether more R&D dollars should be spent on package and test technology and fewer on process development.

This would be a risky strategy for any company to try alone. Should National, for example, go in that direction, it could easily fall behind in process technology—a mistake the company made once before, with nearly disastrous consequences. It could create a problem, for example, if Intel suddenly managed to reinvigorate the demand for processor performance and catch National flat-footed. National says it is not pursuing this strategy now, nor does it plan to in the near future. But it is questioning whether the industry as a whole is headed in the right direction.

IBM in Front by a Length

Manufacturers optimize processes to meet different business objectives. In addition, technologists disagree on the best ways to optimize a process to meet a given objective. These factors make comparing process technologies difficult and ultimately unfair. It is not fair, for example, to say that one company's aggressive (and expensive) process technology is "better" than another's more conservative, and therefore higher-yield, process. No valid metric for "better" exists.

Of the companies we surveyed, however, IBM gets our vote as the company with the most aggressive, most innovative process roadmap. IBM has everything: fast transistors, more interconnect layers, and tight routing pitches. The company, although slow to adopt low-*k* dielectrics, is pushing many other innovations, such as copper and SOI, into production sooner than anyone else. This is not unexpected, as IBM has a solid reputation as a leader in semiconductor technology innovation.

Motorola is following a roadmap similar to IBM's but delayed by six months to a year. The lag is not surprising, as Motorola has always been conservative when bringing new processes into production in order to maximize yields. We'll have to wait and see whether Motorola—whose process technology had fallen behind the the industry prior to the formation of the PowerPC alliance with IBM—can keep up the pace. Its new partner, AMD, should help in this effort. Likewise, the partnership should help AMD keep pace with its wealthy archrival, Intel.

To minimize costs, other companies have adopted low-*k* dielectrics first, putting off the development expense and high cost of copper-deposition equipment. The problem with this approach is that when they do switch to copper, they must also switch to a dielectric material that is compatible with copper. IBM and Motorola, by biting the copper bullet first, avoid a second dielectric transition.

Technologically, Texas Instruments is close behind IBM, being similar in most respects but less aggressive on copper and SOI. TI succeeded in bringing its 0.25-micron process in line with the leaders and it is now well situated to hold its position in 0.18 micron. This is good news for Sun, whose UltraSparc will be faced with a formidable challenge from Intel's Merced in 2000.

Considering that it is not competing for high-end PC microprocessor sockets, VLSI surprised us with a very strong process roadmap. The company has made excellent progress since its 0.25-micron VSC 9 process, which was late and not quite up to par.

Although little is known about Intel's next-generation 0.18-micron process, we don't expect it to be quite as dense or as fast as IBM's CMOS-8S. This is not because Intel is incapable of creating such a process but because it is under a whole different set of constraints (see MPR 9/14/98, p. 3). Intel's capacity requirements, and the costs and logistics of upgrading megafabs, dictate a process that emphasizes high volume more than ultimate performance. Still, based on the company's remarkable performance in 0.25-micron technology, we expect Intel to field a process that will make the race to point one eight a close one indeed. M