Cahners
MICROPROCESSOR REPORT
www.MPRonline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

TSMC Sets Sights on #1

Biggest Foundry Plans to Be Biggest IC Manufacturer

By Keith Diefendorff {6/5/00-01}

At TSMC's annual technology symposium, its president, F.C. Tseng, pointed to a number of firsts for his company: first foundry to 0.18-micron production, first foundry to offer copper interconnect, first foundry to 0.15-micron production, and first foundry to distribute 0.13-micron design rules to its customers. And the company, which already claims the title of world's largest semiconductor foundry, has its sights set even higher: according to Tseng, TSMC will in 2001 surpass Intel as the world's largest wafer manufacturer of any type.

The company has come a long way in a short time. Three years ago, when TSMC began 0.35-micron production, the company was two full years behind the Semiconductor Industry Association's (SIA) *International Technology Roadmap for Semiconductors* (ITRS). Since then, TSMC has been working to close the gap, and it has succeeded: last year TSMC began volume production of 0.18-micron wafers, exactly on the schedule called for by the SIA roadmap. For 0.13 micron, Shang-Yi Chiang, VP of research and development, says TSMC plans to be more than a year ahead of the SIA roadmap, as Figure 1 shows.

The reason TSMC was once so far behind is that the company was exclusively focused on lowest possible cost, which required very conservative generic processes. Over the years, however, Chiang says the company discovered that whenever it offered a higher-performance option, customers wanted it. The company also recognized that by not offering high-performance processes it was forcing many potential customers to go to IBM for a more competitive process.

These revelations caused TSMC to change strategy and push toward higher performance. While it still intends to deliver the best low-cost process, at 0.18 micron and beyond it also intends to match leaders IBM and Intel in microprocessor (high-performance-logic) processes, Texas Instruments in DSP (low-power) processes, and Lucent in mixed-signal processes.

Success Fuels R&D

This goal is certainly a lofty one, but it is one that TSMC believes it can achieve based on its financial success and its exclusive focus on the foundry business. In 1999, TSMC reported net sales of $2.3 billion—not including revenue from its most recent acquisitions, WSMC and TSMC-Acer (TASMC) (see *MPR 1/31/00-07*, "TSMC Buys Acer Semi,

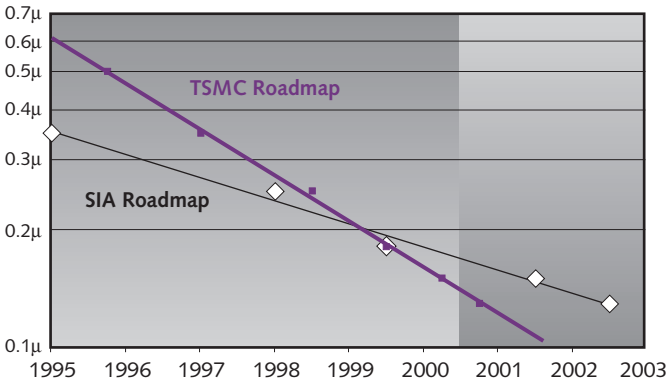

Figure 1. Three years ago, at 0.35 micron, TSMC was two years behind the SIA's technology roadmap. The company caught up at 0.18 micron and intends to pull well ahead in the 0.15- and 0.13-micron generations.

Fab	Wafer Size	Location	Production Start	Capacity (wafers/wk)
Fab 1	150mm	Hsin-Chu	1987	3,950
Fab 2A	150mm	Hsin-Chu	Apr 1990	8,630
Fab 2B	150mm	Hsin-Chu	Mar 1992	9,380
Fab 3	200mm	Hsin-Chu	Aug 1995	9,840
Fab 4	200mm	Hsin-Chu	Feb 1997	8,190
Fab 5	200mm	Hsin-Chu	Oct 1997	7,980
Fab 6	200mm	Tainan	Nov 1999	3,050
Fab 7 (TASMC)	200mm	Hsin-Chu	Dec 1999	7,300
Fab 8 (WSMC)	200mm	Hsin-Chu	1998	7,900
Fab 12	300mm	Hsin-Chu	1Q02	TBD
Fab 14 (was 7)	300mm	Tainan	4Q01	TBD
WaferTech	200mm	Camas, WA	Jun 1998	5,420
Vanguard	200mm	Hsin-Chu	Nov 1999	3,540
SSMC	200mm	Singapore	4Q00	80
Total				75,260

Table 1. This table shows the capacity of all 11 of TSMC's fabs, plus its affiliated fabs, WaferTech, Vanguard, and SSMC, as of April 2000. (Source: TSMC)

WSMC")—and it spent $72 million in process R&D (not including capital equipment). In the first quarter of 2000, TSMC reported sales of $920 million, up 126% over 1Q99 and up 19% over 4Q99. And, just as Scotty used to push the Enterprise's warp drive beyond design limits, TSMC operated its fabs at 113% of rated capacity during the first quarter of this year.

Unlike many semiconductor companies, TSMC's only business is its foundry business, and Tseng promises to keep it that way. He says this focus allows TSMC to obtain a very high return on its R&D investments. It also avoids any competitive issues between the company and its customers, which is an issue for foundries like IBM Microelectronics.

Out of its 11 fabs, TSMC produces an enormous number of wafers. In 1999, the company's annualized capacity was 1.8 million 200mm-equivalent wafers. Tseng said that during 2000, TSMC fabs will produce a whopping 3.4 million wafers—a 90% increase over 1999. Affiliated fabs (WaferTech, Vanguard, and SSMC) will produce an additional 500,000 wafers for TSMC in 2000. By 2004, Tseng's plans call for a total annual output of 7.4 million 200mm-equivalent wafers, roughly evenly distributed between 0.13- (and below), 0.15-, 0.18-, 0.25-, and 0.35-micron wafers.

In Fab 6 alone, which at 190,000 square feet already boasts the world's largest cleanroom, TSMC will boost output from 145,000 wafers in 2000 to 602,000 in 2001. Output from Fabs 7 and 8, which are in the equipment-finalization phase now, will go from 652,000 200mm-equivalent wafers in 2000 to 1.4 million in 2001. Table 1 shows the current capacity of all TSMC fabs.

TSMC's capacity-expansion plans include an aggressive move to 300mm (12-inch) wafers. Next summer, TSMC will start up a 300mm pilot line in Fab 6, producing about 4,500 wafers in 2001. Production of 300mm wafers will begin in earnest in TSMC's newest 12-inch-only fab, Fab 12, in 4Q01. In 2002, TSMC expects Fab 12 to turn out about 79,000 300mm wafers. Fab 14 (previously called Fab 7, before the purchase of WSMC and TASMC) will enter production in early 2002, pumping out another 30,000 300mm wafers during that year. (A 300mm wafer yields about 2.3 times as many chips as a 200mm wafer. Once fully operational, a 12-inch fab should produce wafers at about the same rate as an 8-inch fab of similar size.)

For 2000, TSMC projects that about 35% of its capacity will be 0.35-micron; 30%, 0.25 micron; and 5%, 0.18 micron. This distribution indicates that the vast majority of TSMC's customer base is still not using its leading-edge process, even though that process has been in production for more than a year. TSMC believes, however, that this situation will change rapidly, and that by the end of this year, demand for the leading-edge process will significantly increase. To accommodate this general trend, TSMC's strategy is, at each technology node, to first deploy a generic process and then, over time, add optional features and fine-tune that process to offer more application-specific variations.

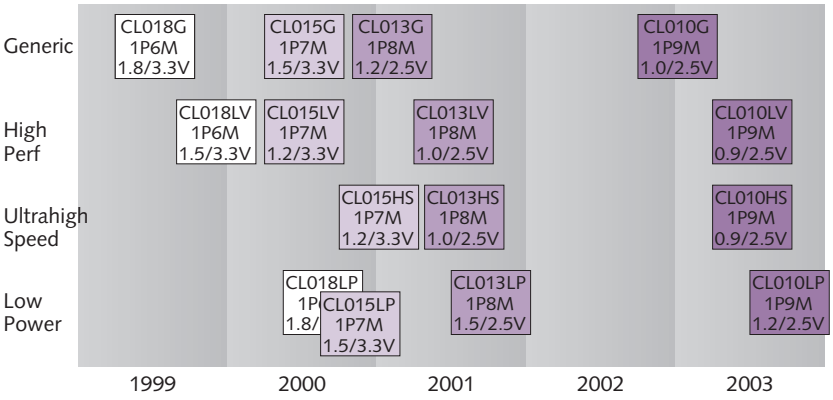

Figure 2. Typically TSMC first places a generic low-cost process into production, then rolls out variations of it with transistors, or other features, optimized for specific application areas. The primary transistor variants are for high performance (low voltage), ultrahigh-speed logic, and low power. Optional modules exist for mixed-signal, embedded DRAM, nonvolatile memory, SRAM, CMOS image sensors, and high voltage. (Source: TSMC)

Process Variants Serve Many Applications

On each baseline logic process, TSMC plans to offer four basic transistor variations: low-power, baseline (ASIC), high-performance (low-voltage), and ultrahigh-speed (CPU class). In general, the lower-power versions trade off transistor delay for low leakage current (I_{off}), whereas the high-performance variants are optimized in the opposite direction.

Figure 2 shows TSMC's process roadmap through the 0.10-micron generation. In this figure, the left edge of each box represents initial "risk production." At TSMC, risk production is defined as the point in time at which the process meets a number of

very specific goals: it must complete all process-qualification criteria; pass a 168-hour product qualification on an internal test vehicle; and demonstrate a manufacturing defect density of no more than 0.12 defects/cm^2 per critical mask layer. (Chiang says a mature process at TSMC normally runs in the range of 0.015 to 0.03 defects/cm^2/critical mask layer.)

Table 2 shows the device characteristics for the baseline process at each node, and Table 3 shows the performance characteristics of the transistors in five variants of the baseline 0.18-micron process, along with the characteristics of the ultrahigh speed 0.15-, 0.13-, and 0.10-micron transistors. Although (according to the MDR FET Performance Metric) these processes don't quite match the performance of the leading-edge microprocessor-logic processes from leaders IBM and Intel, they will be very close. We expect, however, that TSMC's process will be equal to, and in some cases better than, those vendors' processes on static-power consumption (leakage), dynamic power consumption, density, and cost.

To these basic logic processes, TSMC adds a number of optional process modules that allow its customers to further customize a process for their applications. These modules, which generally appear about a year after the baseline process enters production, include such extensions such as mixed-signal features (e.g., metal-to-metal capacitors), embedded-memory cells (DRAM, flash, E^{2}PROM, and SRAM), CMOS image sensors, and high-voltage transistors (CMOS and BiCMOS).

An important component of TSMC's strategy is that all transistor variations and module options use exactly the same design rules as the generic process on which they are based. As a result, customers are free to mix and match options and to grow and evolve their products. This strategy is especially valuable for system-on-a-chip (SOC) designs, which must integrate a wide range of functions.

There are downsides to this strategy, however. To accommodate a wide variety of optional features, the baseline design rules must be somewhat more conservative in order to serve as a lowest common denominator. This suggests that TSMC's ultrahigh-performance option, for example, may never quite match the performance of another manufacturer's process that was designed and tuned for speed without compromise. Chiang claims, however, that in practice very little is sacrificed to this flexibility. For the most part, he says, the various transistors are not that much different, requiring only simple tweaks to parameters such as channel length (L_{eff}) and threshold voltage (V_t), and that process modules are largely independent and have little or no effect on the generic baseline-process definition.

Technology	0.18µm	0.15µm	0.13µm	0.10µm
Risk Production	1Q99	1Q00	4Q00	4Q02
Core Voltage	1.5–1.8V	1.2–1.5V	1.0–1.2V	0.9–1.2V
Poly Half Pitch	0.215µm	0.185µm	0.155µm	0.13µm
L_{gate}	0.16–0.13	0.11–0.12	0.08–0.085	0.05–0.06
T_{ox} (effective)	26, 32Å	20, 26Å	16, 20Å	<12, 20Å
Lithography	248nm	PSM	193nm	PSM
Silicide	Cobalt		Nickel	
Substrate	Bulk		SOI option	
Metal	6, Al or Cu	7, Al or Cu	7–8 Cu	9–10 Cu
Dielectric (k)	FSG (3.7)	FSG (3.7)	<3.0	<2.5
M1 Pitch (cont)	0.46µm	0.39µm	0.34µm	0.26µm
Mn Pitch (cont)	0.56µm	0.48µm	0.4µm	0.32µm
Mtop Pitch	0.9µm	0.9µm	0.9µm	0.9µm
SRAM Cell	4.65µm^2	3.42µm^2	2.43µm^2	<1.4µm^2

Table 2. On most parameters, TSMC's 0.18-micron process stacks up favorably against both IBM's 0.18-micron CMOS-8S and Intel's 0.18-micron P858 microprocessor-logic processes (see *MPR 1/25/99-06*, "Intel Raises the Ante With P858"). Future processes will also be very competitive. (Source: TSMC)

The performance parameters TSMC has published for the processes on its roadmap indicate that Chiang may be correct. Nevertheless, we expect highly optimized microprocessor-logic processes from leaders like IBM and Intel to continue outperforming TSMC's best, although the gap could quickly be reduced to as little as 5% if those companies relax at all.

A True Believer in Copper

Two years ago, IBM shook up TSMC by announcing it would introduce copper into its 0.22-micron microprocessor-logic process. Immediately following that announcement, TSMC launched a crash program to catch up. Within a year, TSMC

Feature	CL018G	CL018LP	CL018ULP	CL018LV	CL018HS	CL015HS	CL013HS	CL010HS
Voltage (V_{dd})	1.8V	1.8V	1.8V	1.5V	2V	1.5V	1.2V	1V
T_{ox} (Effective)	42Å	42Å	42Å	35Å	42Å	29Å	24Å	21Å
L_{gate}	0.16µm	0.16µm	0.18µm	0.15µm	0.13µm	0.11µm	0.08µm	0.05µm
I_{DSat} (n/p) (µA/µm)	600/260	500/180	320/130	700/320	780/360	860/370	920/400	780/350
I_{off} (leakage) (pA/µm)	20	1.60	0.15	80	300	1,800	13,000	62,000
V_t (n)	0.42V	0.63V	0.73V	0.42V	0.4V	0.29V	0.25V	0.21V
Ring Osc Delay	26.7ps	34.0ps	49.1ps	23.2ps	20.2ps	13.8ps	10.2ps	7.2ps
FET Perf* (GHz)	30.1	21.9	14.0	39.8	43.1	52.3	79.5	110.6

Table 3. This table describes the performance of various TSMC processes (G = generic, LP = low-power, ULP = ultralow-power, LV = high-performance/low-voltage, HS = ultrahigh-speed). According to the MDR FET Performance Metric, TSMC's generic 0.18-micron transistor (CL018G) is about 35% slower than IBM's CMOS-8S and Intel's P858 transistors, but TSMC's ultrahigh-speed version (CL018HS) is only about 6–10% slower. (Source: TSMC, except *MDR estimate)

had successfully integrated copper into its 0.25-micron process, and it now offers both two-layer and all-layer copper options in its production 0.18-micron process—actually beating IBM to the punch in offering copper to its foundry customers.

This success astounded Chiang, who had fully expected Murphy's Law to inject problems into volume production. But to Chiang's surprise, no problems have developed; Chiang says copper yields are every bit as good as aluminum yields, and process reliability has surpassed his most optimistic expectations. He attributes TSMC's remarkably smooth integration experience to the copper-tool vendors, which, he says, made the transition a simple one.

TSMC has become a true believer in copper, even though most of its customers are not yet taking advantage of it. Customer reluctance is due to the fact that, at 0.18-micron, the higher speed of the interconnect layers shows up only if the customer is willing to undertake significant redesign to exploit copper's lower resistance and capacitance. Since copper is still a more expensive option (because new tool costs have not yet been amortized), customers are slow to make the leap. But at 0.15 micron, the benefits of copper become more compelling, and TSMC expects to see many more of its customers opt for it. At 0.13 micron, aluminum-wire delays become untenable, and TSMC says its 0.13-micron process will be copper only.

Along with interconnect metallurgy comes the issue of intrametal and interlayer dielectric material. TSMC's 0.18-micron process currently uses fluorosilicate glass (FSG), which has a dielectric constant (k) of about 3.7, only about 5–10% better than pure silicon dioxide. At its technology symposium, the company indicated it plans to move to a material with a much lower k ($k < 3.0$) during its 0.13-micron generation but has not yet decided which material to use.

Jack Y.-C. Sun, senior director of advanced-technology development, indicated that the company is having some difficulties integrating spin-on materials, such as Dow Chemical's SiLK, which IBM has already committed to (see *MPR 5/1/00-01*, "IBM Paving the Way to 0.10 Micron"), and that TSMC is still considering chemical-vapor deposited (CVD) materials, because of their superior hardness and adhesion. The problem with those materials, however, is that there is no known path to future porous materials with k values below 2.5, so we expect TSMC to eventually settle on SiLK, or something similar. TSMC said it still has much reliability testing to do, however, before making that commitment.

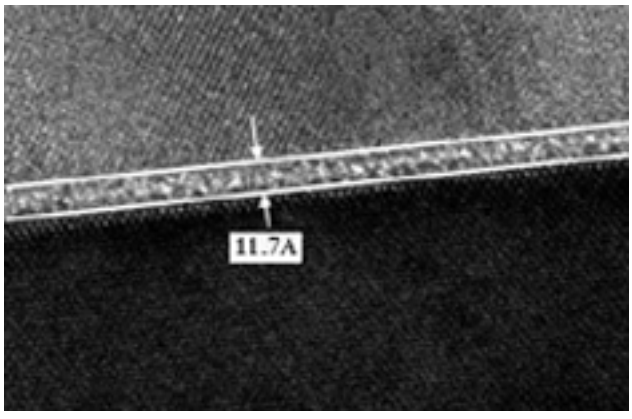

Figure 3. This scanning-electron micrograph shows a sub-12Å gate oxide TSMC is developing for deployment in its 0.10-micron-generation transistor. (Source: TSMC)

If the company appears somewhat uncertain on its low-k plans, it openly confesses to being bewildered about next-generation lithography (NGL). The company is now using 248nm deep ultraviolet (DUV) with resolution-enhancement techniques (RET), such as phase-shift masks (PSM) and optical-proximity correction (OPC), for 0.18- and 0.15-micron production. For 0.13-micron, TSMC will shift to 193nm optical steppers and use RETs to get to 0.10 micron. The company says it strongly favors 157nm optical steppers for the 0.07-micron generation.

Beyond 0.07 micron, however, things become less clear. The company admits to being torn between the extreme-ultraviolet (EUV) approach favored by Intel, AMD, and Motorola and the electron-beam projection (EPL) approach favored by IBM (with PREVAIL) and Lucent (with SCALPEL). Chiang said he has just hired an optical lithography expert from Yorktown Heights (New York), presumably from IBM Research, to resolve this issue.

In the meantime, TSMC is conducting advanced research of its own in other areas. Sun says the company is developing, among other things, a very small stacked-capacitor/transistor embedded-DRAM cell; silicon-on-insulator technology (see *MPR 8/24/98-02*, "SOI to Rescue Moore's Law"); a silicon-germanium (SiGe) BiCMOS transistor for mixed-signal applications; ultralow-k ($k < 2.2$) dielectrics; sub-12Å (physical) gate oxides; high-k gate oxides; and nickel salicide for low-resistance poly and diffusions. Figure 3 shows a scanning-electron micrograph of an 11.7Å gate oxide TSMC is developing for its 0.10-micron process.

This Changes the Whole Picture

As recently as three years ago, the processes deployed by the major Taiwanese foundries (TSMC, UMC, and Chartered) were more than a full process generation behind those used by the major U.S. and Japanese independent semiconductor device manufacturers (IDMs). This gap meant that fabless semiconductor companies had little realistic hope of building a performance-oriented product that could compete with products from the IDMs, and, if they tried, the product was highly vulnerable. The only viable option for these fabless companies was to partner with a big IDM—which often meant selling their soul to the devil—or to resort to IBM's leading-edge foundry. The wafer costs in IBM's

foundry, however, were notoriously high, because IBM's leading-edge capacity is limited, and it is in high demand. Life was not easy for a fledgling high-performance fabless semiconductor business.

Today, this picture has changed significantly. In fact, TSMC believes it has now closed the performance gap completely, or soon will. Although we believe this claim is overly optimistic, there is little question that TSMC has gained considerable ground. The ongoing research and development work at many IDMs, especially IBM, will probably allow those companies to maintain a performance lead over TSMC for the foreseeable future. TSMC doubts, however, that IBM's strategy to stay ahead of the industry (see *MPR 5/1/00-01*, "IBM Paving the Way to 0.10 Micron") can last forever. The problem with IBM's strategy, says Chiang, is that it lacks a volume component that can sustain the required investment long term. His point is well taken. The power behind the semiconductor industry, after all, has proved to be massive volume; volume drives the learning curve and brings with it an economy of scale that no amount of high-margin low-volume manufacturing can match.

Even if IDMs such as IBM, Intel, Motorola, and TI can stay ahead of TSMC in performance, their lead will surely be much smaller in the future than it has been in the past. Moreover, TSMC is very likely to match IDM processes on density and power, and may beat them on cost. And TSMC is not alone in this game. Both UMC and Chartered have similarly aggressive roadmaps and serious capacity-growth plans.

Whether TSMC is ahead of or behind these other foundry competitors at any given moment is hard to call. The race between TSMC and UMC is especially tight; nearly every week, one or the other stakes some new claim of "first company to something." TSMC claims to be the first company to deliver production 0.15 micron, and claims superior performance. UMC also claims 0.15-micron production, and although TSMC's process does appear somewhat faster, both are nearly identical on density parameters. UMC's Web site (*www.umc.com*) claims the company is the first foundry to use 0.13 micron, having demonstrated on May 2 a fully functional 2MB SRAM with a tiny 2.28μm^2 cell. But TSMC isn't far behind and says it intends to widen its performance gap over UMC in the 0.13-micron generation. Chartered (*www.charteredsemi.com*) doesn't yet offer 0.15- or 0.13-micron processes, but it has agreements with Lucent and with Motorola (see *MPR 3/8/99-msb*, "Chartered, HP Go HiPer (MOS)") to get them.

It is notable that both UMC and Chartered have chosen a different strategy than TSMC to obtain advanced technology. While Chartered has partnered with Lucent and Motorola, and UMC has partnered with IBM (see *MPR 2/14/00-02*, "IBM, Infineon, UMC Gang Up On 0.13"), TSMC has decided to go it alone. TSMC hopes its size will generate enough revenue to adequately fund its advanced technology work, and it prefers to control its own destiny rather than depend on partners. On the other hand, advanced process development is enormously expensive, and UMC's and Chartered's strategy to share costs with other companies has some merit.

The availability of high-performance processes from TSMC, UMC, and Chartered will undoubtedly have profound implications for the whole electronics industry. The most obvious effect will be on the fabless semiconductor companies. Not only do these companies gain easy access to high-performance, low-cost processes, but the availability of large intellectual-property (IP) libraries from the foundries will allow small companies to design complex parts that previously required hundred-person design teams. As a result, the fabless-semiconductor business model becomes significantly more attractive.

The losers in this high-stakes game are the IDMs. Those companies will no longer enjoy the protection from small fabless competitors that their advanced technology, proprietary fabs, and large design staffs once provided. As a result, IDMs could see what was once their largest asset turn into an albatross, causing them to rethink their level of investment.

Furthermore, no IDM, including even the mighty Intel, can afford to put in enough reserve fab capacity to meet peak demand. As a result, IDMs (Motorola, for example) have begun to think seriously about outsourcing a portion of their manufacturing (see *MPR 1/25/99-02*, "Integration, Diversification Key in 1998"). The availability of nearly equivalent processes from the foundries makes outsourcing an attractive alternative to new multibillion-dollar fabs—fabs that run the very real risk of sitting idle in an economic downturn. As IDMs take more advantage of outsourcing, however, they become more motivated to make their processes compatible with those of the foundries. This trend could destroy the last remnants of process differentiation for the IDMs.

Exactly how this scenario will play out is unclear. But it is clear that TSMC is enabling a whole new generation of small startup companies to compete head-on with the Goliaths of the semiconductor industry. Over time, these small companies could siphon considerable power away from large monolithic IDMs. While it is worrisome to see more process technology and fab capacity migrate outside the U.S., the increase in competition could ultimately be good for consumers and the industry. ◇