



Virtual versus vertical

By Geoffrey James -- 12/1/2005

Electronic Business

<http://www.reed-electronics.com/eb-mag/article/CA6287343?pubdate=12%2F1%2F2005>

Sections:

Fab expense

The threat of DFM

Customer partnerships

Industry partnerships

Changing the foundries

There's no question that the pure-play foundry has proven to be one of the most successful innovations in the semiconductor industry. Twenty years ago, these independent foundries didn't exist; in 2004 the combined revenues of the four largest foundries—Taiwan Semiconductor Manufacturing Company (TSMC), United Microelectronics Corp. (UMC), Chartered Semiconductor Manufacturing and Semiconductor Manufacturing International Corp. (SMIC) exceeded \$14 billion. Their success has rippled through the industry, giving rise to fabless design firms and the customer-owned tooling (COT) do-it-yourself design revolution, in which a customer hands a finished chip design to its foundry. "The foundries have done a tremendous job of providing excellent process technology at a relatively low cost to their customers," says James Hines, principal analyst at Gartner.

Despite its obvious success, the basic foundry business model is about to change. The cost of developing and building a state-of-the-art manufacturing process has ballooned to the point where only the very largest foundries can consider the investment. At the same time, the high cost of developing chips for the newest fabs has design firms clamoring for design-for-manufacturing (DFM) tools that break the traditional barrier between design and manufacturing. DFM is forcing the foundries to expose their proprietary manufacturing techniques, their most closely held intellectual property. That's not all. With designs now destined from the start for specific manufacturing processes, designers are losing the ability to jump from foundry to foundry, which was a primary reason COT was economical.

With the cost of fabs skyrocketing, the potential for IP theft growing and designers questioning the financial viability of COT, how will the foundries cope in a world where DFM rules?

### **Fab expense**

In looking at this new world, let's start with the issue of cost. One thing's for certain: Investing in a new fab is hardly for the faint of heart. Some of the dollar figures being thrown around are mind-boggling. For example, in mid-2005 a Taiwan news service announced that TSMC planned to spend \$7.5 billion on a new fab. TSMC vehemently denies this. "We don't know where that came from," says Chuck Byers, TSMC's director of brand management. Of course, in the words of cartoon Dilbert's Dogbert, "a rumor isn't true until it's been denied," but even taking Byers' denial at face value, the official number is only slightly less amazing. The planned state-of-the-art fab, according to Byers, will cost "only" around \$3.5 billion.

Let's put that \$3.5 billion into perspective. The total combined revenue of the third- and fourth-largest foundries (Chartered and SMIC) was less than \$2 billion in 2004. Even No. 2 foundry UMC booked revenues only of a little more than \$4 billion. Let's face it, even for TSMC, with yearly revenues of around \$8 billion, \$3.5 billion is hardly chump change.

"The cost of equipment is the major issue facing the foundries," says Semico Research analyst Joanne Itow. "They must constantly figure out how they're going to provide advanced technology to their customers." Indeed, the cost of equipment may be getting a bit rich for some of the foundries. Despite sharply growing revenues, only TSMC, among the four largest foundries, increased its capital equipment spending in 2005, according to Gartner (see the table "Worldwide Semiconductor Capital Equipment Spending," below). By contrast, UMC and SMIC dropped their capital equipment spending by 61 percent and 46 percent, respectively, and Chartered's spending was so low in both years that the company wasn't among the top 20 spenders in the semiconductor industry, according to Gartner analyst Bob Johnson.

"If the foundries are going to continue to provide advanced processes to their customers, they must continue to invest," says Itow. But coming up with the cash for new equipment may be the least of their worries. Once the new fabs are built, they'll have to be kept busy. And we're talking a whole lot o' silicon to make that happen. According to Byers, that \$3.5 billion fab will have a capacity of 35 thousand 300-millimeter wafers a month, nearly twice the volume of fabs currently under construction throughout the semiconductor industry. And that kind of volume means that there'd better be a whole lot o' design starts flowing into the foundry.

Of course, just as the fabs are expensive, so too is the cost of designing chips for them. "If you're working at 90 nanometers or 65 nanometers, it's probably going to cost you 20 to 30 million dollars to get your chip on the market," says Risto Puhakka, vice president at VLSI Research. That's a more-than-geometric jump from the design cost at .13 micron, which is around \$3 million, according to Meyrick Chan, SMIC's director of design service support.

"The smaller foundries are leveraging existing, mature process nodes, by incorporating specialty features to create novel architectures and levels of analog integration not achievable at the larger foundries.

—Boris Petrov, Petrov Group

Why are designs suddenly so pricey? The short answer leads us right back to DFM. As the average size of a component has shrunk, designers are forced to follow design rules so Byzantine that it's not always clear whether problems are design- or manufacturing-related. Back in the good old days, designers worried only about whether the circuitry made sense. They could then "throw the design over the wall" and still expect an acceptable yield from the fab. Not so today. With DFM, designers must account for the peculiarities of specific manufacturing processes, making design chains longer, design tools more complex and design projects more expensive.

### **The threat of DFM**

Because DFM is making design more expensive, the foundries are going to need to make their services as easy to use as possible if they're going to keep those big-ticket fabs operating at capacity. However, because DFM breaks down the traditional barriers between design and manufacturing, the foundries will need to act in ways that may not seem entirely natural to them. "DFM requires a lot of interaction and engineering work between the foundry, third-party IP suppliers and the customer that's designing the chip," says Hines, "With the latest geometries, we're past the point where the classic COT engagement will result in a usable chip."

This may be bad news for the foundries, because DFM may favor the integrated device manufacturers (IDMs), which already have both functions under a single roof. "Companies such as IBM have everything in-house, so the designers can go do some test chips, get experience with the manufacturing process and then use that knowledge in the design flow," says Rich Goldman, vice president for strategic development at Synopsys. Hines concurs: "The situation does seem to favor a more integrated approach."

The decision about whether to go to an ASIC vendor or use COT for a custom chip has always meant balancing cost and risk. Traditionally ASICs implied a high cost per chip but with lower risk, because

the IDM guaranteed yield in return for charging extra for each chip. By contrast, COT traditionally meant a lower cost per chip (because the foundries compete on price) combined with a somewhat higher risk resulting from the do-it-yourself nature of COT. What's important here is that the value proposition for COT is directly dependent on the fabless firm's ability to shop for the least expensive foundry.

DFM changes that value proposition. Because fabless design firms can no longer "throw the design over the wall," the choice of a foundry must be made at the very start of a project. Once that decision is made, DFM makes it difficult and maybe even impossible to change horses midstream. That's not going to be a problem when the relationship between the design firm and the foundry is copasetic, but if the relationship goes south or a process doesn't deliver the expected yield, the designer may be forced to go back to an entirely blank drawing board.

The injection of DFM into the design equation may already be changing the way electronics firms source their custom chips. Although correlation is not causation, it's interesting to note that, according to Gartner, the rising importance of DFM is matched by a leveling off of the decade-long decline in yearly ASIC design starts, a decline that resulted from the growing popularity of COT in the mid-1990s. "DFM, and the increased cost associated with it, is making ASICs more attractive for companies needing custom chips," says Steven Longoria, vice president of semiconductor technology platforms at IBM.

### **Customer partnerships**

The foundries are meeting the DFM challenge head on, by creating closer relationships with their customer base, according to Byers. TSMC, for example, has "moved into an area of collaboration with a great deal of interaction at a business, operations and technical level," he says, citing a new TSMC Web site where designers can obtain detailed information about TSMC's manufacturing processes. UMC also collaborates closely with design firms "early and often throughout the rest of the semiconductor supply chain," according to Lee Chung, UMC's vice president of corporate marketing.

Such cooperation doesn't always come easy, though. Although the foundries are quick to insist that they'll give their customers everything they need in order to achieve yield, DFM may require them to share information they'd prefer to keep private. "The ability of the foundries to compete is tied up with their ability to offer processes that are superior to their competitors'," explains Synopsys' Goldman. "They don't want too much information about those processes to bleed into the rest of the market." This is a sensitive issue for the foundries, as evidenced by TSMC's 2003 lawsuit accusing SMIC of infringing on TSMC's chip-making patents and misappropriating TSMC's trade secrets.

"Companies such as IBM have everything in-house, so the designers can go do some test chips, get experience with the manufacturing process and then use that knowledge in the design flow."  
—Rich Goldman, Synopsys

Although the SMIC case doesn't appear to be DFM-related, Byers admits that IP theft as a result of DFM remains a concern. To keep control of its competitive advantage, TSMC is pursuing a two-pronged strategy, according to Byers. First, TSMC has created its own design libraries that are prequalified to work well with TSMC's manufacturing processes. Because these libraries are already tuned for TSMC's fabs, they can function as black boxes, giving designers no particular incentive to understand the process behind them.

The second prong of TSMC's protection strategy is to haul IP violators into court. "We have very stringent IP protection devices in place and treat our IP very seriously from a legal standpoint," says Byers. Indeed, TSMC showed its seriousness with the SMIC lawsuit, which resulted in SMIC's paying TSMC \$175 million in damages. More important, TSMC sued in the United States, where patent protection is strong, rather than in Taiwan or China, where patent protection is less honored. Still, it's one thing to sue a competitor and quite another thing to sue a big customer, even if that

customer is guilty of passing along proprietary information. One wonders whether TSMC would be so bold if, say, a major customer was found telling tales out of school.

### **Industry partnerships**

Whereas TSMC may plan to aggressively protect its IP rights, most other foundries are taking a lower-key approach, namely working closely with EDA vendors to create tools that make DFM less of a burden, according to Goldman. As an example, he cites new tools that allow designers to insert redundancy into circuits to ensure that at least one of the circuits works properly. "By sharing the process information with EDA vendors, the foundries avoid sharing it with the rest of the world," Goldman explains. The EDA vendors are also working to make third-party IP blocks less process-specific, according to Jan Willis, senior vice president of industry alliances at EDA vendor Cadence Design Systems. "We are making these solutions work as much as possible in the traditional way: by abstracting the design issues from the transistor level into the higher level of the design chain," she says.

But there's a limit to how much the EDA vendors can protect designers from DFM, and Goldman's example of redundant circuitry is a case in point. "You can't do it too much, or you're going to blow up the size of the chip," he admits. With DFM, keeping chips efficient and still achieving acceptable yield will require EDA tools that take advantage of the peculiarities of individual manufacturing processes. "UMC partners early with leading EDA tool companies to develop and validate design flows and tools for technologies such as digital, analog and low-power," says Chung. "We can collect data from real silicon and apply this data into DFM-aware models of lithography and other process steps that designers can use before tape-out to identify problem areas, and information can also feed back to DFM-aware tools to help process data for optical proximity correction after tape-out."

In other words, although the foundries and EDA vendors can make post-DFM COT look as much like pre-DFM COT as possible, chip designers will inevitably lose some flexibility when it comes to selecting a foundry, especially after the design has already started down a particular pathway. One could argue that the best way to lessen the risk of sole sourcing is to establish the closer relationships between designers and foundries that DFM demands. However, even if such a partnership runs smoothly, there's always the possibility that a natural disaster, such as an epidemic of bird flu, might shut down the fab.

With designs tied to specific manufacturing processes, the only way to ensure the availability of sufficient manufacturing capacity is to have multiple fabs running the same process. With DFM it's not enough to build one \$3.5 billion fab; a foundry must now be able to offer its customers access to at least two and probably three (or more!) fabs running that same process technology. And that means major changes for the foundries.

### **Changing the foundries**

There's only one foundry that definitely has the financial clout to build redundant fabs at \$3.5 billion a pop: TSMC. UMC, with half of TSMC's yearly revenue, might also be able to come up with that kind of dough, but for the rest of the foundries, it's just not going to happen. How the other foundries will adapt to the post-DFM reality differs according to their size.

Midsized foundries, such as Chartered and SMIC, will need to build alliances with other semiconductor companies, including the IDMs, to build fabs that share the same process technology. For example, Chartered has been working with IBM, Samsung and Infineon to develop new chip manufacturing processes for the fabs of all four firms. Under this agreement, in the event of a disaster or an unexpected demand for a given chip, Chartered customers can buy excess manufacturing capacity from IBM or Samsung. (Infineon is keeping the process for in-house use only). "We now have three very large companies supporting customers in terms of manufacturing their designs," says Kevin Meyer, Chartered's vice president of worldwide marketing and platform alliances.

"If you're working at 90 nanometers or 65 nanometers, it's probably going to cost you 20-30 million dollars to get your chip on the market."

—Risto Puhakka, VLSI Research

The alliance also helps design firms manage the risk of custom chip design, according to IBM's Longoria. He envisions electronics firms coming to IBM for an ASIC design but eventually moving mass manufacturing over to Chartered, where it can be accomplished at a lower unit cost. Longoria also sees a reverse scenario in which a fabless firm encountering problems with a COT project turns to IBM for help in converting to an ASIC effort. "As you merge the ASIC and the foundry market, you can drive a substantial share of design and manufacturing work to the companies that are best suited to perform it," he explains.

The smaller foundries, however, aren't likely to be partnering with giants such as IBM and Samsung, much less coughing up beaucoup billions for multiple new fabs. Because of this, DFM will inevitably force the smaller foundries into niche markets in which state-of-the-art manufacturing processes aren't crucial, according to Boris Petrov, managing partner at the Petrov Group, a strategy consulting firm. "Some of the emerging foundry players are pursuing the analog, RF and high-voltage markets, reflecting the emergence of power management, wireless and analog startups," he says, citing the example of Jazz Semiconductor, a privately owned foundry specializing in TV tuners, consumer-segment high-voltage/power management and MEMS/cellular integration. "The smaller foundries are leveraging existing, mature process nodes, by incorporating specialty features to create novel architectures and levels of analog integration not achievable at the larger foundries," he adds.

Specializing in these niches makes sense, because building effective circuits and chips isn't as dependent on miniaturization as are the high-volume digital applications that require the latest process technology. Thus, rather than trying to compete head-to-head against TSMC, the smaller foundries are building infrastructures and business models that take advantage of the relative stability and low cost inherent in older process technology. For example, "the smaller foundries are expanding their ability to generate prototypes by using multiproject wafers and multireticle masks to help their customers get a chip to market more quickly," says Petrov.

DFM, in short, will change the entire structure of the foundry sector of the semiconductor industry. At the top, competing with the big IDMs, will be TSMC and probably UMC. Chartered and SMIC will be forced to forge alliances with the IDMs and gradually turn from pure-play foundries into hybrids servicing both COT and ASIC design methodologies. The smaller foundries will be left behind, participating in niche markets that, over time, will slowly decline as SoC functionality moves onto the increasingly complex chips manufactured with the latest processes.

There is the possibility, of course, that a disruptive innovation in manufacturing—such as a smaller wavelength for photolithography—might change everything, making DFM unnecessary and returning the industry to the good old days when COT ruled. Barring that, there's little question that the pure-play foundry market of tomorrow will only vaguely resemble the foundry market of today.

How are you planning on solving DFM challenges? Send your thoughts to [feedback@eb.reedbusiness.com](mailto:feedback@eb.reedbusiness.com).

#### THE FOUR LARGEST FOUNDRIES AT A GLANCE (2004) Revenue Profit Growth Net margin

TSMC 8,128 2917 36.2 36%

UMC 4082 1006 45.1 25%

Chartered 932 6.6 68.9 1%

SMIC 975 89.7 166.5 9%

OURSCE: HOOVERS

#### WORLDWIDE SEMICONDUCTOR CAPITAL EQUIPMENT SPENDING 2004 2005 (est.)

TSMC 3,843.0 5,900.0 54%

UMC 2,670.3 1,034.0 -61%  
SMIC 1,838.8 1,000.0 -46%  
Chartered 700 750 7%  
SOURCE: GARTNER

Geoffrey James writes Electronic Business' EDA column.

© 2005, Reed Business Information, a division of Reed Elsevier Inc. All Rights Reserved.