

WATCH YOUR STEP

IC technology and tools face economic hurdles

THE SEMICONDUCTOR INDUSTRY is unique in having a blueprint for development that has proved accurate for 39 years.

In 1965, *Electronics* magazine asked Gordon Moore, PhD, a

founder of Intel Corp and then director of the research-and-development laboratories at Fairchild Semiconductor to write an article discussing the future of semiconductor technology (Reference 1). In it, Moore predicted an exponential growth in the number of transistors packaged in a single IC—a prediction later termed Moore's Law.

Despite its name, Moore's Law is not a scientific law. It is a principle that describes the opportunity for exponential improvements by advances in semicon-

Illustration by Randy Lyhus



**SEMICONDUCTOR FABRICATION HAS REACHED THE
65-NM PROCESS NODE, BUT THE COST OF DEVELOPING
A PRODUCT WITH THIS TECHNOLOGY COULD BE SO HIGH
THAT ONLY A FEW COMPANIES WILL BE ABLE TO AFFORD IT.**

ductor technology. The real importance of the law is not what it predicts, but that it is a powerful industry motivator that continues to justify new inventions. The challenge to stay on the development path that Moore's Law predicts has become so formidable that the semiconductor industry is more and more conducting its research efforts, including consortia and collaboration with suppliers, in a precompetitive environment.

Although Moore's Law is expressed in number of transistors per IC, the industry uses a different metric to identify the manufacturing stages corresponding to Moore's Law steps. The unit of measure-

AT A GLANCE

- ▶ Gordon Moore predicted that the number of transistors on a chip would double every 18 months.

- ▶ Fabricating features smaller than visible light is challenging.

- ▶ Each new process node increases the difficulty for design engineers.

- ▶ Project managers find that projects become more complex and more expensive as they move to the next process node.

ment is the half-pitch between two traces on the die—in other words, half the distance between two features diffused on a substrate. In 1999, the half-pitch dimension was 180 nm. This value is significant, because, for the first time, the critical dimension became smaller than the wavelength of the light used to expose the material before etching.

Today's lithography equipment, which will likely remain in use for the rest of this decade, uses a wavelength of 193 nm. In 2001, semiconductor manufacturers reduced the half-pitch to 130 nm, and in 2003, a half-pitch of 90 nm became the leading-edge process. A process capable

THE REAL WORLD MEETS 65-NM TECHNOLOGY

By Jon Stoner, AMI Semiconductor

With a few companies announcing products in 90 nm, the industry has reached another milestone on the road to infinitely small devices. Yet, so many industry pundits have started calls for the next generation—65 nm. Mixed-signal products make up three-quarters of AMI Semiconductor business with a heavy content of precision analog. A crucial difference in AMI's SOCs (systems on chips) is their link outside the digital world.

The digital world, or the world of logic design, is known as the "virtual world." The "real world" is analog. AMI customers develop SOCs that connect to electronics in the real world, such as motors and sensors. At first glance, industry observers might say that the real world does not need 65-nm technology. But the industry has found a way to follow Moore's Law, and engineers have continued to design and produce mixed-signal products in shrinking technologies (albeit two to three generations behind the leading edge). Could 65-nm SOCs eventually dominate real-world mixed-signal applications?

The general case for applying increasingly smaller digital technologies to mixed-signal SOCs has been the argument that more digital computational

power compensates for analog performance. As long as a signal can be digitized, DSP techniques provide FFTs and significantly improve filter performance. Because DSPs are primarily a series of multipliers, accumulators, and registers, they scale with each technology. Other merits of this architecture include the stability and test coverage a primarily digital design achieves. All in all, an ADC-DSP-DAC single-chip architecture has much to recommend it *if* (and it's a big if) the onboard converters can process the signal's amplitude and bandwidth and it can achieve the appropriate input-signal conditioning and output drive.

Small digital technologies are good for digital design but miserable for analog circuits. Analog circuits don't scale like digital ones, because the bad analog features, mostly the power-supply scaling, outweigh all the good analog features of the small-generation technologies. An often-cited study shows that the real "sweet spot" from a minimum power perspective is 0.25 or 0.35 microns, or 250 to 350 nm (Reference A). Using smaller features increases the power requirements for the analog circuitry—just the opposite of

digital circuits. Gate leakage is also a problem that becomes severe at less than 90 nm. It creates another mismatch effect as device sizes increase—exactly the method used to offset small-geometry device variations and inherent noise (Reference B).

Proponents of 65-nm, mixed-signal SOCs would correctly argue that these device and circuit problems do not matter if you can achieve the appropriate converters. Engineers can create good converters in small digital technologies; it just takes time. An equivalent figure of merit for ADCs shows that their performance doubles about every five years; digital figures of merit (such as MIPS) double every 1.5 years (Reference C).

However, none of these problems is as significant as the issue of connecting to the real world. Transistors produced at 65 nm are delicate. They can't handle voltage transients, their gates are too thin, and their junctions are too shallow. And they must live in a virtual world of a few hundred millivolts. The real world provides input signals on common-mode voltage levels of 10 to 50V with signal swings of 100 mV to 5V.

The real world requires mixed-signal SOCs—to directly

drive relays, stepper motors, and displays—and robust, not 600-mV, transistors. The requirement for transistors in AMI Semiconductor's mixed-signal SOCs with 5 to 10V, and even 50V operating capability is the biggest deterrent to migrating to smaller transistor geometry. Ultimately, higher voltage transistors dramatically raise the cost of 65-nm technology. At the 65-nm level, the transistors, junctions, and wells are so different from higher voltage transistors that little of the circuitry can be shared. Eventually, it all comes down to cost versus function. If the faster computational capabilities of digital logic at 65 nm can offset the cost of integrating analog circuitry drawn with different rules, 65-nm ICs will make it in the real world.

REFERENCES

- A. Annema, Anne-Johan, *IEEE Trans on Circuits and Systems*, Volume 46, No. 6.
- B. Annema, Anne-Johan, et al, *Planetanalog*, Feb 18, 2004.
- C. Boser, B, University of California—Berkeley presentation, 2004.

AUTHOR BIOGRAPHY

Jon Stoner is chief technology officer at AMI Semiconductor.

of a 65-nm half-pitch is in development and expected to be available for volume production in 2007.

ENGINEERING CHALLENGES

As semiconductor technology has progressed from 2 microns, to today's 90-nm processes, and on to 65 nm, the number of design elements that have gone from being fixed or "given" to being variable has been increasing. In many cases, this situation requires a trade-off among speed, area, power, and yield. Designers make these trade-offs following design rules that the foundry provides. As the half-pitch decreases, the number of rules increases by several factors. When working with 90-nm geometries, an engineer must consider more than 500 rules when making a design decision.

Reference 2 discusses the issues involved in DFM (design for manufacturing).

Engineers need circuit-design expertise to make correct decisions when dealing with such multidimensional problems. Unfortunately, the industry has convinced the US education system that logic designers need to understand little, if any, physics and electronics theory to develop good designs. In fact, some EDA companies are even advertising that software engineers can create good electronic circuits. Such an assertion is credible only for simple circuits that designers implement on FPGAs or structured ASICs. Engineers cannot confront problems at 90 nm and below without understanding circuit design, and design teams often require an engineering expert in semiconductor-manufacturing issues.

Many problems stem from the need to use lithography equipment that uses a 193-nm wavelength as the illuminating source when the half-pitch of the resulting geometries is 65 nm. To produce a good circuit using the 90-nm process, manufacturers must use RET (reticle-enhancement technology) and OPC (optical phase correction). Both techniques modify the pattern of light to expose the photoresist layer, with dimensions almost half the size of the wavelength of the light source.

What designers see when they look at the prefabrication layout of a chip is no longer what they get. **Figure 1** shows an




Figure 1

RET/OPC techniques at 90 nm improve the results of optical exposure on the die.

intended pattern in green, an actual pattern without RET/OPC corrections in purple, and a couple of actual geometries using different combinations of corrections. The circuitry in purple would result in nonworking silicon. Even with the best feasible corrective measures, designers can only approximate the desired shapes. The choice of various corrective measures impacts speed, power, and yield, and the amount of OPC impacts the area. Making the correct decision is generally the critical contribution to product profitability. Stone Pillar Technologies offers products that interconnect the process and mask details with subsequent electrical-test or yield data to provide engineers with insight into the likely cause of failure.

Iroc Technology maintains that reliability is becoming the fifth element to consider in the design process. Soft errors are the dominant causes of reliability failures. Cosmic rays impacting the silicon circuitry cause most of these failures. Iroc has established that the average soft-error FIT (failure-in-time) rate on 130 nm or less is around 500 per megabyte of memory. This value is almost 100 times the classic reliability numbers and 10 times the general market requirements. To ensure continued proper functioning, engineers must design error-correction circuitry into each design.

Although Mentor's early understanding of the lithography problem and decision to dedicate engineering resources to work with foundries before the situation became critical has given it a commanding lead in the RET/OPC market, both Cadence and Synopsys are investing significant resources in the area. Synopsys has obtained much-needed knowledge of the required technology following its acquisition of Numerical Technology, and Cadence is working closely with MaskTools, an ASML company.

Designers can no longer limit themselves to understanding the problem and lingo of digital design. When working with 65-nm processes, they will need to work closely with and understand mask designers, manufacturing engineers, and even process-development technologists. It will no longer be sufficient to understand the terminology; team members will need to appreciate the nature and seriousness of each problem they encounter.

Andrzej Strojwas, chief technologist at PDF Solutions contends, "DFM rules must complement design rules in nanotechnology," noting that DFM rules intrinsically differ from design rules but not necessarily sufficiently to achieve good yields. "For example, typical 90-nm DFM rules recommend doubling vias as well as spreading wires to minimize critical area. However, adding the metal needed for doubling vias will increase the critical area of the metal. For technologies with low-k dielectrics, this [step] may lead to an increase in stress and eventually cracking of the dielectric, causing the yield to drop. With so many rules, designers are facing contradictory directives and cannot be sure that they have [made] the correct choice until the ASIC is produced. At this point, correcting the design [becomes] quite expensive."

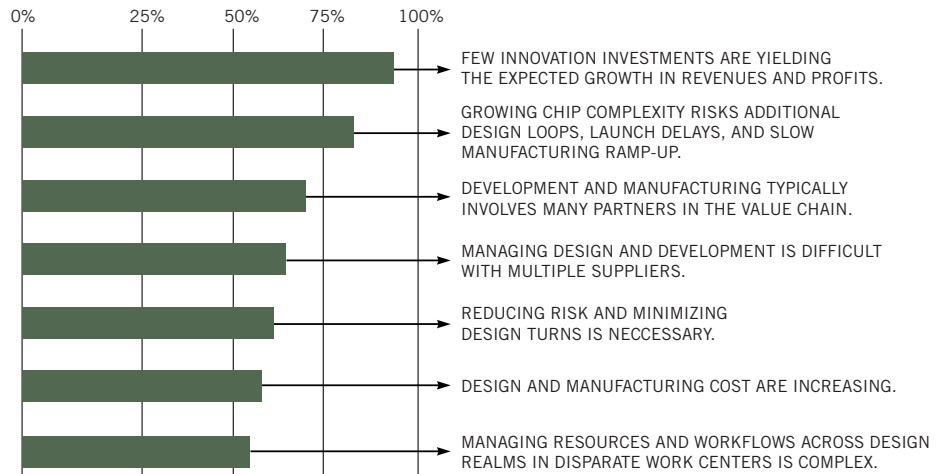
MANAGEMENT CHALLENGES

Although the engineering challenges are severe, managers also face a serious increase in job complexity (**Figure 2**). In a presentation at September's Chartered Semiconductor Technology Forum in San Jose, CA, Walter Lange, PhD, a field

executive in the Systems Solution department of IBM, stated, “The power to manage one’s destiny becomes increasingly dependent on collaboration with partners and wise risk management.”

At the 180-nm node, only two parties are involved in the successful development and production of an ASIC: the company developing the ASIC and the foundry that manufactures it. The design company decides which EDA tools to use independently of the manufacturer’s choice of photo-mask-production software. In addition, the development costs are lower and more predictable. Using the 130-nm process is more complex. At this process node, the choice of an EDA methodology and tool suite impacts both the design team and the manufacturer. Because physical effects, such as signal integrity and timing, are just as important as correct logic design, manufacturers prefer to certify a suite of tools to ensure that both the development and the verification tools can handle the physical problems inherent in the process.

At 90 nm, RET/OPC techniques become critical. Therefore, vendors of semiconductor-manufacturing equipment now must work with EDA vendors to ensure that the reticle-production software can successfully modify the layout file, avoid electrical problems, and maintain logic functions. Obviously, managing a project that involves four parties is more complex than dealing with three entities. And when the 65-nm process becomes available for manufacturing, that four-party team expands to five with the addition of the IP (intellectual-property)-core provider. The scenario can become even more complex if the design team is using IP cores from more than one vendor. The additional partner is necessary, because RET/OPC tools will manipulate the IP in the same way as the rest of the logic through the photo-mask process and thus may significantly modify the physical characteristics of the IP, depending on the location of the core on the die. Applying OPC techniques for a location on the die can impact the adjacent features, also requiring changes. The result is that IP cores considered constant from a functionality point of view are



SOURCE: IBM/MERCER PLM SURVEY FINDINGS, APRIL 2001

Figure 2 The number of factors managers must consider in managing schedules and costs increase as the half-pitch of the process decreases.

now variables, because their electrical characteristics have changed.

Second-sourcing has been an industry practice since the early 1970s. Systems companies have traditionally used two or three semiconductor foundries for IC manufacturing, which provides both technical and business flexibility. The practice continued unchanged through the 180-nm process node. The two foundries used for second sourcing, when working at the 130-nm node, need to support compatible EDA tools and methodologies as well as equivalent design rules and manufacturing processes. The two foundries must have a common design flow and process technology at 90 nm and below. Therefore, it will become difficult to establish a second source, and both foundry and customer must enter into a true partnership to succeed. IBM and Chartered Semiconductor forged an agreement two years ago that assures design teams that both companies are running the same process at each other’s fabs.

Managers face an additional challenge. The cost of developing an ASIC from architectural design to good die is increasing by an order of magnitude. This increase is the result not only of the greater complexity of a larger design, but of the cost to develop, verify, and correct the set of masks required for manufacturing. In the last few years, systems companies have pointed to the increasing cost of photo-masks as a significant hurdle to profitability. Gary Smith, principal ana-

lyst for EDA at Gartner Dataquest, has shown that the mask cost per gate diffused has been decreasing, because the cost of a single mask set does not double to produce twice as many transistors when migrating from one process node to the next. The observation is useful in demonstrating the efficiency of the technological progress within the industry but serves no purpose to managers to whom the absolute cost of a project is the determining factor. Cost per gate is relevant only when comparing the number of functions that a designer can integrate into an IC, but total development cost is a deciding factor in product profitability.

Although mask costs are relevant, the principal reason many managers hesitate to move to the next process node is the increased uncertainty in arriving at a definitive development-cost value. For example, moving from 130 to 90 nm has injected a great deal of indeterminism into the cost equation. Lack of experience prevents the industry from predicting how many mask turns are necessary to achieve acceptable yields for a correctly functioning device or even the cost of finding and fixing a problem once a circuit fails verification tests. Managers must ensure that projects finish on time and stay within budget. But when a mask set costs more than \$1 million and the time required to fix a problem is indeterminate because the nature of the difficulty is not well-understood, managing cost and schedules becomes difficult. Add in the

cost of being late to the market, and one project can easily reach a cost variation of tens of millions of dollars.

Because time to market is the most important factor in determining product profitability, inexperience at a new process node makes it impossible to predict how long it would take to fix a problem. Therefore, managers tend to stay at a process node that is "good enough." They make trade-offs among the number of functions one device integrates, the increase in computational speed, or even the final product form factor, in favor of a more deterministic development cost and schedule prediction. According to Felicia James, vice president and general manager of the Cadence Virtuoso Custom Design Platform, "[Cadence] is engaged in a number of projects providing tools and expertise helping customers migrate designs back from the 90-nm to the 130-nm process." As the sidebar, "The real world meets 65-nm technology" indicates, many applications require analog circuitry. Manufacturing and integrating such circuitry present difficult engineering challenges at dimensions much smaller than 250 nm, significantly increasing the cost and scheduling uncertainty of a project.

FEASIBLE ALTERNATIVES

Semiconductor companies use two types of preferred devices with regular structures to qualify new processes and technologies: memories and FPGAs. Often, these device types provide the first




Figure 3

New product families have improved the volume-versus-cost trade-off when designers consider alternatives to traditional ASIC development.

commercial application of a new process, which occurs when a foundry and a fab-less semiconductor company have, or a company that owns both the design and the fabrication functions has, achieved an acceptable yield level producing either memories or FPGA devices. Because the number of ASIC design starts is decreasing yearly, some semiconductor companies are offering a new configurable fabric called a structured ASIC. Esmat Hamdy, PhD, senior vice president of technology and operations at Actel Corp, observes, "Structured ASICs are proof that ASIC vendors see traditional ASICs falling by the wayside."

ASIC providers recognize that ASSPs (application-specific standard parts) and FPGAs will continue to maintain a stronghold in the market, and, as a result,

they are attempting to stay in the game with structured ASICs. Of course, ASIC designs will not totally disappear. A few systems companies will still find markets that offer enough volume and reasonable price levels to justify the investment necessary to transform an architectural concept into an IC manufactured at 65 nm. A significant decrease in ASIC-design starts will impact EDA vendors whose revenue relies greatly on back-end-tool licenses. Tom Kingsley, product marketing manager for lithography verification at Synopsys notes, "The decrease in the number of customers will be in part compensated by the increase in the price of a license, since back-end tools will need to increase in sophistication. A leading-edge RET/OPC tool may become as valuable as a wafer stepper." Such equipment sells at list prices of around \$30 million and brings up the question of whether system companies will be willing to pay \$1 million for a single yearly license. Managers might, if the tool also allows design corrections resulting from process-modeling functions.

FPGA devices, when compared with an equivalent ASIC implementation, require more silicon area and power but provide less operating speed. Their unit price is also greater, and, even taking into account the greater development cost of an ASIC, FPGAs become more expensive at around 50,000 units.

Both Xilinx and Altera now ship devices manufactured with 90-nm processes. Altera offers the Stratix II family, which supports as much as 9 Mbytes of memory, DSP blocks with frequency as high as to 370 MHz, and the Nios proces-

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in EDN.

Actel

www.actel.com

ChipX

www.chipx.com

Mentor Graphics

www.mentor.com

Altera

www.altera.com

eASIC

www.easic.com

NEC

www.necusa.com

AMI Semiconductor

www.amis.com

IBM

www.ibm.com

PDF Solutions

www.pdf.com

Anadigm

www.anadigm.com

Iroc Technologies

www.iroctech.com

Stone Pillar Technologies

www.stonepillar.com

Cadence Design Systems

www.cadence.com

LSI Logic

www.lsillogic.com

Synopsys

www.synopsys.com

Chartered Semiconductor

www.charteredsemi.com

MaskTools

www.masktools.com

Xilinx

www.xilinx.com

sor. Altera also provides ARM cores, giving designers a powerful inventory of building blocks for their SOC (system-on-chip) designs using FPGA devices. Altera also offers a structured ASIC. Customers who want to produce an FPGA in large volume can “harden” the circuit using Altera’s Hardcopy and significantly decrease the unit cost of the device.

Xilinx offers the Virtex 4 family of devices to implement SOC devices on an FPGA. Customers can use PowerPC cores and choose from three application-specific platforms: one for DSP applications, another for high-speed serial I/O, and a third for digital-logic integration. Xilinx does not use a classic structured-ASIC approach to lowering unit costs. Instead, it bases EasyPath on the observation that failing to test unused portions of the device will increase the yield and thus lower the cost. **Figure 3** shows

the cost/volume trade-off among traditional ASIC, FPGA, structured-ASIC, and yield-enhancement techniques, such as EasyPath.

Digital structured-ASIC options are available from ChipX (formerly Chip Express), NEC, LSI Logic, eASIC, and others. Anadigm has just introduced a structured-ASIC option for analog designers. Structured ASICs provide a significant advantage to designers over hard-IP cores, because they provide silicon-constant cores; a hard-IP core is hard only at the GDSII level.

The ASIC design market will continue at a reasonable volume if both designers and manufacturers find a viable alternative to the current methods. The most promising combines platform-based design with RTL sign-off. Platform-based design provides a proven circuit designed for a specific application market

but allows users to implement some amount of proprietary circuitry in the same die. RTL sign-off separates the functions of logic design from the methods of generating a gate-level netlist and placing and routing the design. Given the additional area, speed, and lower power requirements, designers may finally be willing to relinquish their control over layout to the manufacturing team, which will be able to apply the correct geometrical and electrical modification to the netlist to achieve an economical acceptable yield level. □



REFERENCES

1. Moore, Gordon, *Electronics*, Volume 38, No. 8, April 1965.
2. Moretti, Gabe “Mano a mano with manufacturing,” *EDN*, Oct 14, 2004, pg 34.

TALK TO US

Post comments via TalkBack at the online version of this article at www.edn.com.
