


Consumer ICs: **designing for reliability**

CONSUMERS WANT THE LATEST ELECTRONICS BUT ONLY IF THEY'LL LAST LONGER THAN THE TIME IT TAKES TO TAKE THEM OUT OF THEIR PACKAGES. IC VENDORS MUST TACKLE RELIABILITY ISSUES SO THAT THEIR DEVICES WILL FIND USE IN RELIABLE, LONG-LASTING PRODUCTS.



If you've purchased a Microsoft Xbox 360 or a Sony PlayStation 3, the person who sold you the system most likely recommended adding a cooling fan to go along with the game system. Chances are good that you reluctantly forked over the extra \$30, even though the fan corrects what would appear to be a design flaw that shouldn't be there in the first place. And, if you were an early Xbox 360 customer, you probably received a recall notice from Microsoft offering free replacements of ICs, IC-cooling systems, or both that were prone to system slowdowns or even failures. Even if you didn't yet own a 360, you had probably heard about the recall and the likelihood that 360 may have some design flaws but then went ahead and bought one anyway.

It's an increasingly interesting phenomenon in the industry: Consumers are buying products that they know are prone to early failures. In the consumer-electronics market, the drive for the latest and greatest digital "bling" often overcomes better judgment, so purchases of consumer electronics are quickly becoming emotional ones. Many people buy a new game console every four years when they become available, a new mobile phone and MP3 player every year, and a TV and PC every four to six years.

Although consumers now seem willing to fork over cash for the bling, will they be willing to do so in the future if product failures start occurring before their mobile-phone contract has expired? Even if consumers aren't worried, the makers of consumer products should be, because early defects will sooner or later cause costly recalls and may even turn consumers and OEMs against the defective brands. In the game-console world, consumers have only three choices: the Xbox 360, the PS3, and the Wii. In the TV, cell-phone, and most other consumer-electronics niches, however, consumers and OEMs have many choices—and very long memories.

Product longevity—or the lack thereof—becomes an even more daunting problem when you take into account the ever-increasing complexities of designing and manufacturing the leading-edge IC designs that power consumer devices. The semiconductor industry now focuses largely on ensuring that IC-design fabrication produces sufficient yields, that the ICs pass functional tests so they can go into products in large volumes, and that those products will land on store shelves sooner than those of the competition. But, as IC processes become more advanced and consumer demand for greater performance and system functions increase, IC failures will become more commonplace unless vendors tackle reliability issues.

Providers of military, automotive, and medical ICs have long practiced high-reliability techniques to ensure that devices last. Those designing and manufacturing ICs for the consumer and OEM market also have paid close attention to reliability, and most target an MTBF (mean time between failures) of at least 10 years—longer than most consumers will keep the products. Experts say that reliability will always be a

major concern for semiconductor vendors, but these vendors must overcome many obstacles before they can produce reliable products that meet customers' increasing demands for faster, smaller, and higher performance products. Most consumer-device manufacturers employ reliability-engineering groups that set guidelines for and closely monitor designs through each step of the design, manufacturing, packaging, and burn-in processes. Burn-in is an important step because it puts designs through accelerated-lifetime tests for the best performance under worst-case conditions—high temperature and humidity. As a manufacturer develops each new silicon process, these reliability-engineering groups are constantly on the lookout for both new and re-emerging failure mechanisms (**Figure 1**). Today, they also must monitor trends, such as gate leakage and process variability, that can complicate the manufacture of reliable ICs (see **sidebar** "The shifting sands of silicon" in the Web version of this article at www.edn.com/080306cs).

"There's no such thing as the 'same old, same old' in the reliability world," says Jack Hergenrother, PhD, manager of technology for System Z Test in the IBM systems and technology group. "We've been progressing continuously in our understanding of new failure mechanisms and new ways of looking at potential wear-out and failure mechanisms." According to Hergenrother, this phenomenon is not unique to IBM. "It's an industrywide thing," he says. "In the last 10 years [as Moore's Law has evolved], some new mechanisms have come up, and you need to factor in those [mechanisms] during the qualification and design process. That [need] is true from both a chip- and a system-reliability perspective."



Experts say that the industry has been able to adequately and quickly address reliability issues at all phases of development. According to John Chen, vice president of technology and foundry operations at graphics-processor vendor Nvidia, the industry will be able to overcome these problems in the next few years. “Designers need to be aware of these issues, so they can take full advantage of advanced technology and avoid pitfalls,” he says. Both Nvidia and Xilinx are in the forefront when it comes to creating designs employing new logic processes, so they and their foundry partners must be aware of potential failures, according to Glenn O’Rourke, senior director of product-development engineering in the advanced-products group at Xilinx (see sidebar “Fab or fables: Reliability is still top goal”). “We more than double the [number of] transistors in our designs every 18 months because graphics engines require massive processing power,” says Chen.

Nvidia’s co-founder, Chris Malachowsky, in 1996 designed the company’s first chip, a 1 million-transistor design that was massive for its time. In comparison, the company’s latest graphics processor, which the company built using 65-nm technology, exceeds 1 billion transistors. “We can always use smaller, faster, and higher performing transistors, unlike some applications that are pad-limited and are not scalable,” says Chen. “We have a great opportunity for riding the wave of Moore’s Law; we are always at the forefront of the technolo-

AT A GLANCE

- Most companies target their devices for a 10-year life expectancy and do several stress tests to ensure that their ICs will meet or exceed that goal.
- Soft errors have become bigger problems for logic-based design over the last two technology nodes.
- Companies such as Nvidia and Xilinx are typically the first to use new processes and begin working with foundry partners.
- IBM takes a holistic approach to reliability, examining potential reliability issues at the technology, device, package, and system levels.

gy. However, new challenges come with being one of the first companies to use a new technology.”

IC-FAILURE MECHANISMS

For the 130-, 90-, 65-, and 45-nm process nodes, IC-reliability groups have paid the most attention to failure mechanisms such as NBTI (negative-bias-temperature instability), hot-carrier effects, EM (electromigration), gate-oxide integrity, and SERs (soft-error rates). NBTI and hot-carrier effects are two commonly monitored failure mechanisms, both leading to a loss of gate control (references 1 and 2). NBTI is a key reliability issue that is of immediate concern in CMOS devices enduring stress from negative-gate voltages. Hot-carrier effects occur when an electron, or

“hole,” gains sufficient kinetic energy to overcome a potential barrier, becoming a “hot carrier,” and then migrates to a different area of the device. In both NBTI and hot-carrier effects, the driving current to a transistor becomes smaller, degrading or locking up the timing of the gate, potentially causing failures.

The issue of NBTI became a problem at the 90-nm node, but manufacturers quickly addressed the issue. The initial studies of NBTI typically focused on devices running on always-on dc current, in which the problem is worse, according to Li-Pen Yuan, group R&D director for extraction- and power-integrity products at Synopsys. Devices running on ac have less of a problem with NBTI because the current is discontinuous, thus it does not overstress the transistors. NBTI remains an issue that reliability and design groups must monitor, however, especially if their designs target dc-system applications, such as mobile computing or handheld devices.

NBTI hasn’t disappeared but has gone into the background, says IBM’s Hergenrother. “A few years ago, it caused some problems,” he says. “You don’t hear about it too much anymore because we have figured out how to deal with it. Today, you hear more about PBTI [positive-bias-temperature instability], which is similar to NBTI, except that it happens on a PFET rather than an NFET. The physics of PBTI are different enough that it will become a problem at later technology nodes. This time, the industry will likely be more ready for it.”

IC manufacturers squeeze more speed from transistors and minimize leakage power by using strain engineering—a technique for enhancing performance by modulating strain, or stress, in the transistor channel. Modulating strain enhances electron mobility and, thus, conductivity through the channel. One of the side effects of the technique is that it can introduce hot-electron effects into the design. These effects can shift the voltage threshold and reduce the lifetime of an IC. “Intuitively, if you use strain engineering, you make the transistor faster and higher power and may cause more hot-electron, or hot-carrier, effects,” says Chen. He explains that strain engineering induces a higher electrical field near the drain side of the transistor and causes the electrons in

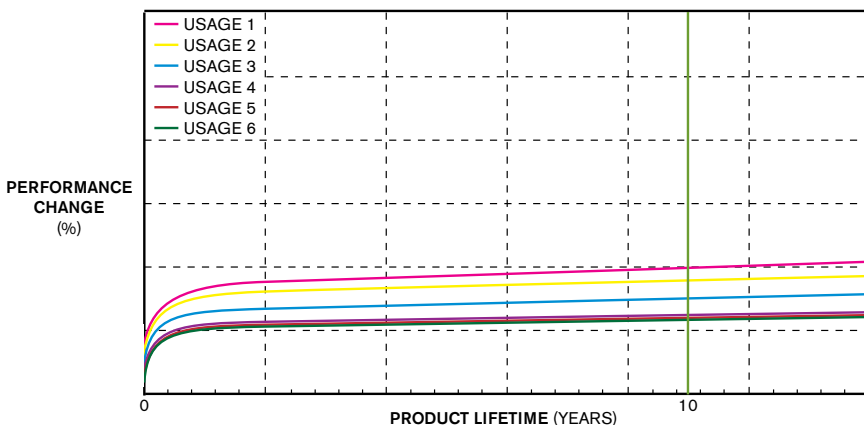


Figure 1 Xilinx devices find use in multiple applications, which in turn target long lifetimes. Xilinx uses lifetime-performance analysis, which runs full characterization during various points in the burn-in process.

the N channel to quickly reach velocity saturation. Electrons must move as fast as they can because doing so provides current. “[The moving electron] will hit other electron-hole pairs and generate other electrons,” he says. “It’s an avalanche effect—impact ionization—that creates more electrons, and, when they get too much energy, they jump into the MOS-gate dielectric and get trapped there, causing a threshold shift and ultimately device failures. But manufacturers have figured out how to increase the barrier entering the gate dielectric. In that regard, it helps: It increases the hot electrons but creates a barrier to stop electrons from getting trapped in the dielectric. The net effect is equal or fewer hot-carrier effects.”

The most diligently monitored failure mechanism, EM, occurs when too much current passes through thin metal traces connecting transistors. When two thin

traces are close together and carrying current or switching at once, one can splinter, causing an open. This splinter can then touch the adjacent trace, causing a short circuit, which can lead to a device failure. EM usually occurs over time, leading to failures long after the chips have left testing. Both the semiconductor and the EDA industries have been aware of EM for many years. “EDA vendors offer analysis tools to detect areas of a design that are susceptible to EM,” says Synopsys’ Yuan. As new processes emerge, EM has grown but not excessively. “A typical design 10 years ago would have a few areas that were sensitive to EM,” says Yuan. “Today, a design may have just 10 [areas]. It isn’t like the problem is exploding.” As EM continues to be a problem, however, tools for preventing it will likely become more common in the mainstream designer’s toolbox.

Another failure mechanism is gate-oxide breakdown or integrity, in which current causes a slow breakdown of the gate dielectric, which can lead to failures. Chen notes that new materials, such as high-k-metal gates will help improve reliability in this area. Intel pioneered these materials, and the rest of the silicon manufacturers will soon follow. Chen notes that some 45-nm and, more likely, 32-nm designs will likely use high-k metal dielectrics composed of hafnium oxide instead of the more traditional gate oxide. Manufacturers grow gate oxide on the silicon during the manufacturing process, and doing so creates a remarkably smooth surface. But in high-k fabrication, manufacturers deposit the hafnium oxide on the silicon in composite layers. “If you use one type of layer, it usually doesn’t work,” says Chen. Using multilayer, high-k dielectric usually means hav-

FAB OR FABLESS: RELIABILITY IS STILL TOP GOAL

IC reliability is a top concern for both companies that own their own fabs and those that don’t. For example, IBM creates its products in house, allowing the company to address reliability concerns and even design trade-offs at all levels of product development:

technology and transistor development, circuit-level design, chip design, package design, and system implementation. “Fab-less” companies, such as Nvidia and Xilinx, on the other hand, must rely on external foundries to manufacture products. And because the two companies

are in highly competitive markets, they tend to be the first to jump to a new foundry’s process when it becomes available. But before they begin designing with that process, they make sure it passes rigorous qualifications, such as ISO (International Organization for Standardization) 9000x, ISO 14000, and OCEA (Office of the China Economic Area) standards.

Glenn O’Rourke, senior director of product-development engineering in the advanced-products group at Xilinx, says that the company uses both UMC (United Microelectronics Corp) and Toshiba as suppliers, so Xilinx must ensure that both sources can make comparable versions of Xilinx’s devices. To achieve this goal, Xilinx develops a refer-

ence model of its designs and ensures that both suppliers can meet their objectives. “We develop a silicon reference model upfront, and ... both fabs drive to those goals,” says O’Rourke. He also notes that, because Xilinx’s chips find use in a variety of applications, the company analyzes its designs’ lifetime performances (Figure A). “We do an accelerated burn-in to mimic the lifetime of the product,” he says. “We do full characterization across temperature and voltage to see how the product’s performance is changing over its lifetime. We leverage that data to cover the usage and the specifications for the lifetime of the product.” Xilinx then makes the results of those reports available to customers.

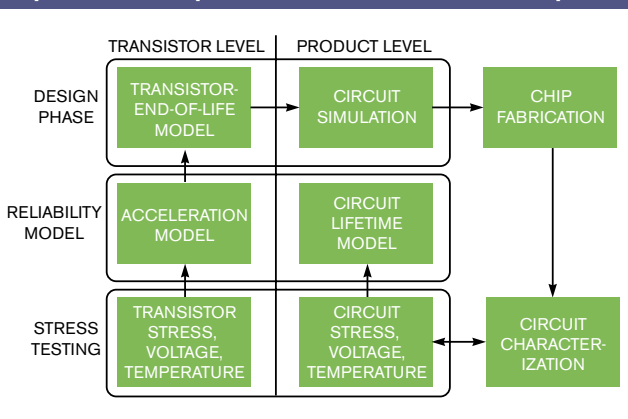


Figure A Xilinx and other large companies examine reliability at multiple phases in IC development.



MORE AT EDN.COM

- ⊕ For another look on design challenges at 45 nm, go to www.edn.com/article/CA6475006.
- ⊕ To read about chip design after the 90-nm process node, go to www.edn.com/article/CA6347251.
- ⊕ For help in choosing system-on-chip processes, visit www.edn.com/article/CA6518679.
- ⊕ Go to www.edn.com/080306cs and click on Feedback Loop to post a comment on this article.

ing fewer pin holes because it's harder to align pin holes for multiple layers. Using high-k dielectric materials usually improves time-dependent-dielectric-breakdown performance. However, unlike silicon dioxide, composite layers have more traps, and more traps can cause electron or N- or P-channel hole trapping, which can cause soft breakdowns, he says. Those things degrade mobility and, in the long term, can create threshold instability. The manufacturers have come up with various process tricks to overcome this issue. "One way is to put a silicon-dioxide layer between the high-k-metal layer and the silicon," says Chen.

SER, another failure mechanism that has long been a concern in the military and aerospace-IC and memory markets, is now becoming a greater concern in logic devices (Reference 3). Alpha particles in packaging materials or neutron strikes that occur naturally in the environment are the typical causes of soft errors. Essentially, an alpha particle or neutron can strike a device, generate noise, and flip bits in memory devices or even flip latches in your circuit. "It is getting to be a bigger challenge with each technology generation," says IBM's Hergenrother. "At the active areas of the devices, the volume of the critical-area devices keep going down, which means that you have to deposit smaller and smaller amounts of charge to create an upset in your transistor." It's difficult to remove alpha particles from packaging materials, so you must build immunity to both cosmic and alpha particles into your system. You can address soft errors at many levels. "[IBM] looks

at SER at the technology level to make transistors soft-error-tolerant, at the circuit level ... to arrange transistors into latches and flip-flops so that it is robust even if one of the transistors does flip," he says. "Then, we look at the chip level for robust error-detection and -correction mechanisms, so, even if there is an error, we catch it and correct it before it propagates any undesirable data. On top of those mechanisms, we have system-level protection, which is another layer of error detection and correction."

Several failure mechanisms can lead to reliability issues. The semiconductor industry has been diligent about identifying and thus correcting failure mechanisms before they ever reach consumers. As devices move closer to the limits of physics and CMOS, however, you may wonder whether reliability will become a worse problem to deal with. **EDN**

REFERENCES

- 1 Peters, Laura, "NBTI: A Growing Threat to Device Reliability," *Semiconductor International*, March 1, 2004, www.semiconductor.net/article/CA386329.
- 2 Peters, Laura, "Strained Silicon: Essential for 45 nm," *Semiconductor International*, March 1, 2007, www.semiconductor.net/article/CA6418539.
- 3 Santarini, Michael, "Cosmic radiation comes to ASIC and SOC design," *EDN*, May 12, 2005, pg 46, www.edn.com/article/CA529381.

FOR MORE INFORMATION

Apache Design Solutions
www.apache-da.com

IBM
www.ibm.com

Intel
www.intel.com

Microsoft Corp
www.microsoft.com

Nintendo
www.nintendo.com

Nvidia Corp
www.nvidia.com

Sony
www.sony.com


Synopsys
www.synopsys.com

Toshiba
www.toshiba.com

UMC
www.umc.com

Xilinx
www.xilinx.com

You can reach
Senior Editor
Michael Santarini
at 1-408-345-4424
and michael.santarini@reedbusiness.com.



THE SHIFTING SANDS OF SILICON

Although the semiconductor industry understands and addresses failure mechanisms, a few trends in silicon manufacturing may seemingly complicate designing reliable ICs. Transistor leakage, multimode- and multivoltage-design techniques, and IC-process variability can all affect reliability. Perhaps the greatest concern is transistor leakage. Starting at the 130-nm node, shrinking transistors leak power when running at peak performance. This leakage creates heat, and the heat in turn creates more leakage (Reference A). This issue has been of primary importance to companies designing mobile devices because leakage wastes an increasingly greater percentage of overall power and, thus, battery life (Reference B). In high-performance microprocessors and graphics processors, however, leakage and the heat it emits drove many processor companies to turn to multicore architectures rather than crank up the clock rates as they had traditionally done with single-processor architectures.

Leakage is one of the reasons that many consumer devices, especially gaming systems, require elaborate heat-sink and fan sys-

tems and that vendors often recommend that consumers purchase extra cooling systems. And, as designs move to 45-nm processes, leakage will cause the loss of more than half a design's power. The semiconductor industry does not currently consider transistor leakage a reliability concern because the industry is seeking new design methods to lower power, and foundries are using new materials, such as high-k-metal gates, to reduce power. Design engineers employ several techniques to optimize their design for power. One of those techniques is multivoltage, multimode design: Designers group functions of a design by performance and power requirements. This approach could somewhat complicate reliability testing. Traditionally, designers send designs to tapeout, run wafer-level tests, place the designs in packages, and then test them at peak speed and performance under worst-case temperature and humidity conditions. However, ICs employing multimode design undergo a certain amount of stress in switching operating voltages and turning on and off. A cell phone, for example, has a peak operating mode when a

user is receiving a call, a low-power mode when the user replies to e-mail, and a standby mode when the user is not employing the device. Testing the reliability of devices that employ multimode ICs can, therefore, differ from normal lifetime testing.

Multimode designs are also vulnerable to ESD (electrostatic-discharge) failures. "Unlike EM [electromigration] that causes fatigue on metal layers over time, ESD failures result from a sudden surge of current that causes the meltdown of metal or vias," says Andrew Yang, chief executive officer of Apache Design Solutions. "ESD failures are getting worse at lower process nodes as the device density increases and metals get thinner. Also, advanced low-power technologies, such as multiple-voltage-domain techniques make designs more susceptible to ESD failures." He says that ESD causes 20 to 30% of IC failures. Designers add ESD-protection circuits and need accurate analysis tools to determine the effectiveness of the protection circuit.

The use of high-k metals is especially promising for leakage issues. High-k metals help out in two re-

gards, according to John Chen, vice president of technology and foundry operations at Nvidia. "The higher k material allows you to do thicker gate dielectric; therefore, you have less leakage through the gate. But leakage doesn't just come from just the gate; it comes from the source and drain and junctions. ... [Using high-k metals also] allows us to further enhance the transistors," he says. "When you turn the transistor off, you still draw a tiny current. In the old days, we didn't worry about this current because we didn't have that many transistors on a chip; now, it's a big problem."

In addition to new materials, companies can also employ cooler packages or cooling systems, such as heat sinks or fans, to further dissipate heat of newer devices. If Nvidia continues to more than double gate density every 18 months, the follow-up to the company's billion-plus-transistor graphics processor will have a few billion transistors. No one knows how that increase will affect the devices' thermal profiles. Chen notes that Nvidia does extensive research and qualification with its foundry partners years before releasing a design and does comprehensive testing of its devices after packaging and during burn-in testing.

Another wild card that can affect reliability is process variation. As process geometries continue to shrink, IC manufacturers must also deal with extreme calibration and modeling of their new processes and manufacturing equipment. And, with every new process reduction, a small change in a fabrication line can produce ICs that vary greatly from targeted performance and power goals. "Process variation causes electrical variation in devices," says Chen. "If you vary the gate-dielectric thickness or transistor length, the current varies accordingly. When the electrical current varies, it can cause a different lifetime of a product, and that [difference] is a reliability concern."

Jack Hergenrother, PhD, manager of technology for System Z Test in the IBM systems and technology group, agrees that variability is key. "It becomes increasingly more difficult to control all the dimensions to the same level of relative accuracy when you move from one technology to the next," he says. "It's easier to print a 300-nm line than a 30-nm line with 1% precision." He says that the company has had to address such fundamental issues as dopant fluctuations.

Some variations can cause errors in binning. Microprocessor

and graphics-processor vendors have traditionally drawn the greatest margins from their highest performing devices. Today, however, vendors using the latest processes must discard devices that exceed their top-performance goals because the leakage and, thus, heat profiles can also increase and lead to early product failures. "After we go through manufacturing, we measure our ICs to get the real distribution of performance for devices," says Chen. That performance ranges from 330 to 450 MHz, and leakage ranges from 100 mA to 1A. The highest performing devices tend to have the highest leakage; thus, the devices that operate faster than 450 MHz typically exceed the 1A leakage target. "The transistors run fast but are also too leaky," says Chen. "We can't sell these parts; we have to throw them away."

REFERENCES

- Santarini, Michael, "Thermal integrity: a must for low-power-IC digital design," *EDN*, Sept 15, 2005, pg 37, www.edn.com/article/CA6255052.
Santarini, Michael, "Taking a bite out of power: techniques for low-power-ASIC design," *EDN*, May 24, 2007, pg 46, www.edn.com/article/CA6442429.