

TRANSISTOR WARS

RIVAL ARCHITECTURES FACE OFF IN A
BID TO KEEP MOORE'S LAW ALIVE

BY KHALED AHMED & KLAUS SCHUEGRAF



In May, Intel announced the most dramatic change to the architecture of the transistor since the device was invented. The company will henceforth build its transistors in three dimensions, a shift that—if all goes well—should add at least a half dozen years to the life of Moore's Law, the biennial doubling in transistor density that has driven the chip industry for decades.

But Intel's big announcement was notable for another reason: It signaled the start of a growing schism among chipmakers. Despite all the great advantages of going 3-D, a simpler alternative design is also nearing production. Although it's not yet clear which device architecture will win out, what is certain is that the complementary metal-oxide semiconductor (CMOS) field-effect transistor (FET)—the centerpiece of computer processors since the 1980s—will get an entirely new look. And the change is more than cosmetic; these designs will help open up a new world of low-power mobile electronics with fantastic capabilities.

There's a simple reason everyone's contemplating a redesign: The smaller you make a CMOS transistor, the more current it leaks when it's switched off. This leakage arises from the device's geometry. A standard CMOS transistor has four parts: a source, a drain, a channel that connects the two, and a gate on top to control the channel. When the gate is turned on, it creates a conductive path that allows electrons or holes to move from the source to the drain. When the gate is switched off, this conductive path is supposed to disappear. But as engineers have shrunk the distance between the source and drain, the gate's control over the transistor channel has gotten weaker. Current sneaks through the part of the channel that's farthest from the gate and also through the underlying silicon substrate. The only way to cut down on leaks is to find a way to remove all that excess silicon.

Over the past few decades, two very different solutions to this problem have emerged. One approach is to make the silicon channel of the traditional planar transistor as thin as possible, by eliminating the silicon substrate and instead building the channel on top of insulating material. The other scheme is to turn this channel on its side, popping it out of the transistor plane to create a 3-D device. Each approach comes with its own set of merits and manufacturing challenges, and chipmakers are now working out the best way to catch up with Intel's leap forward. The next few years will see dramatic upheaval in an already fast-moving industry.

CHANGE IS NOTHING NEW to CMOS transistors, but the pace has been accelerating. When the first CMOS devices entered mass production in the 1980s, the path to further miniaturization seemed straightforward. Back in 1974, engineers at the IBM T.J. Watson Research Center in Yorktown Heights, N.Y., led by Robert Dennard, had already sketched out the ideal progression. The team described how steadily reducing gate length, gate insulator thickness, and other feature dimensions could simultaneously improve switching speed, power consumption, and transistor density.

But this set of rules, known as Dennard's scaling law, hasn't been followed for some time. During the 1990s boom

in personal computing, the demand for faster microprocessors drove down transistor gate length faster than Dennard's law called for. Shrinking transistors boosted speeds, but engineers found that as they did so, they couldn't reduce the voltage across the devices to improve power consumption. So much current was being lost when the transistor was off that a strong voltage—applied on the drain to pull charge carriers through the channel—was needed to make sure the device switched as quickly as possible to avoid losing power in the switching process.

By 2001, the leakage power was fast approaching the amount of power needed to switch a transistor out of its "off" state. This was a warning sign for the industry. The trend promised chips that would consume the same amount of energy regardless of whether they were in use or not. Chipmakers needed to find new ways to boost transistor density. In 2003, as the length of transistor channels dropped to 45 nanometers, Intel debuted chips bearing devices made with strain engineering. These transistors boasted silicon channels that had been physically squeezed or pulled to boost speed and reduce the power lost due to resistance. By the next "node"—industry lingo for a transistor density milestone—companies had stopped shrinking transistor dimensions and instead began just squeezing transistors closer together. And in 2007, Intel bought Moore's Law a few more years by introducing the first big materials change, replacing the ever-thinning silicon oxide insulator that sits between a transistor's gate and channel with hafnium oxide.

This better-insulating material helped stanch a main source of leakage current—the tunneling of electrons between the gate and the channel. But leakage from the source to the drain was still a huge problem. As companies faced the prospect of creating even denser chips with features approaching 20 nm, it became increasingly clear that squeezing together traditional planar transistors or shrinking them even further would be impossible with existing technology. Swapping in a new insulator or adding more strain wouldn't cut it. Driving down power consumption and saving Moore's Law would require a fundamental change to transistor structure—a new design that could maximize the gate's control over the channel.

FORTUNATELY, over the course of more than 20 years of research, transistor designers have found two very powerful ways to boost the effectiveness of the transistor gate. As the gate itself can't get much stronger, these schemes focus on making the channel easier to control. One approach replaces the bulk silicon of a normal transistor with a thin layer of silicon built on an insulating layer, creating a device that is often called an ultrathin body silicon-on-insulator, or UTB SOI, also known as a fully depleted SOI.

SMIC Ex. 1017

IPR2020-01003

Page 3 of 8

A second strategy turns the thin silicon channel by 90 degrees, creating a “fin” that juts out of the plane of the device. The transistor gate is then draped over the top of the channel like an upside-down U, bracketing it on three sides and giving the gate almost complete control of the channel. While conventional CMOS devices are largely flat, save for a thin insulating layer and the gate, these FinFETs—or Tri-Gate transistors, as Intel has named its three-sided devices—are decidedly 3-D. All the main components of the transistor—source, drain, channel, and gate—sit on top of the device’s substrate.

Both schemes offer the same basic advantage: By thinning the channel, they bring the gate closer to the drain. When a transistor is off, the drain’s electric field can take one of two paths inside the channel to zero-voltage destinations. It can propagate all the way across the channel to the source, or it can terminate at the transistor’s gate. If the field gets to the source, it can lower the energy barrier that keeps charge carriers in the source from entering the channel. But if the gate is close enough to the drain, it can act as a lightning rod, diverting field lines away from the source. This cuts down on leakage, and it also means that field lines don’t penetrate very far into the channel, dissipating even more energy by tugging on any stray carriers.

The first 3-D transistor was sketched out by Digh Hisamoto and others at Hitachi, who presented the design for a device dubbed a Delta at a conference in 1989. The

UTB SOI’s roots extend even further back; they are a natural extension of early SOI channel research, which began in the 1980s when researchers started experimenting with transistors built with 200-nm thick, undoped silicon channels on insulating material.

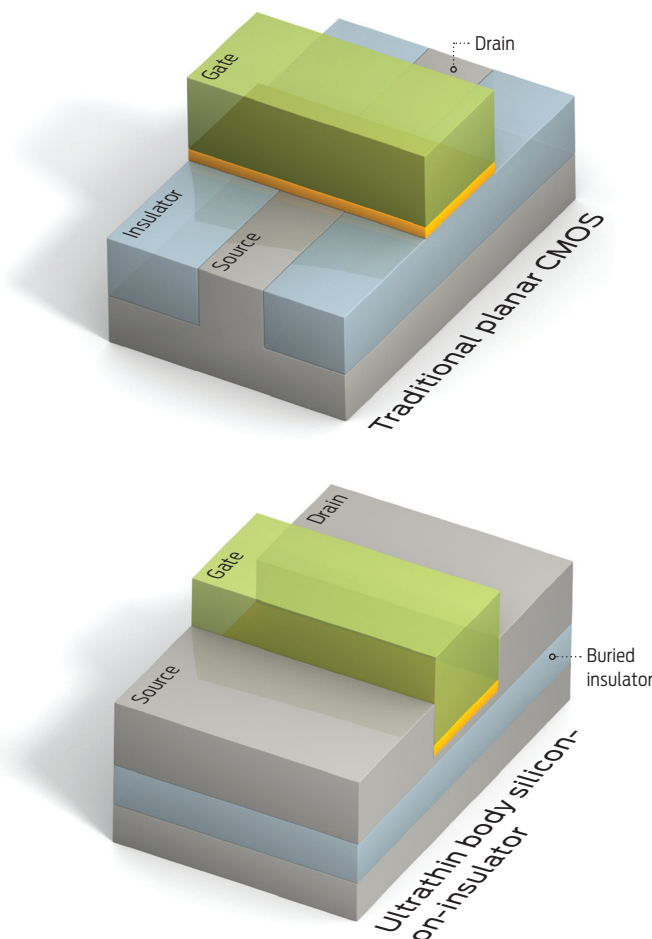
But the promise of both of these thin-channel approaches wasn’t fully appreciated until 1996, when Chenming Hu and his colleagues at the University of California, Berkeley, began an ambitious study, funded by the U.S. Defense Advanced Research Projects Agency, to see how far these designs could go. At the time, the industry was producing 250-nm transistors, and no one knew whether the devices could be scaled below 100 nm. Hu’s team showed that the two alternate architectures could solve the power consumption problems of planar CMOS transistors and that they could operate with gate lengths of 20 nm—and later, even less.

The FinFET and the UTB SOI both offer big gains in power consumption. Logic chip designs typically require that a transistor in its on state draw at least 10 000 times as much current as the device leaks in its off state. For 30-nm transistors—about the size that most chipmakers are currently aiming for—this design spec means devices should leak no more than a few nanoamperes of current when they’re off. While 30-nm planar CMOS devices leak about 50 times that amount, both thin-channel designs hit the target quite easily.

But the two architectures aren’t entirely equal. To get the best performance, the channel of a UTB SOI should be no more than about one-fourth as thick as the length of the gate. Because a FinFET’s gate brackets the channel on three sides, the 3-D transistors can achieve the same level of control with a channel—or fin—that’s as much as half as thick as the length of the transistor gate.

This bigger channel volume gives FinFETs a distinct advantage when it comes to current-carrying capacity. The best R&D results suggest that a 25-nm FinFET can

Continued on page 63



ELIMINATING EXCESS: In the next few years, traditional planar CMOS field-effect transistors [left] will be replaced by alternate architectures that boost the gate’s control of the channel. The UTB SOI [bottom left] replaces the bulk silicon channel with a thin layer of silicon mounted on insulator. The FinFET [bottom right] turns the transistor channel on its side and wraps the gate around three sides.

Transistor Wars

Continued from page 53

carry about 25 percent more current than a UTB SOI. This current boost doesn't matter much if you have only a single transistor, but in an IC, it means you can charge capacitors 25 percent faster, making for much speedier chips. Faster chips obviously mean a lot to a microprocessor manufacturer like Intel. The question is whether other chipmakers will find the faster speeds meaningful enough to switch to FinFETs, a prospect that requires a big up-front investment and an entirely new set of manufacturing challenges.

THE SINGLE BIGGEST HURDLE in making FinFETs is manufacturing the fins so that they're both narrow and uniform. For a 20-nm transistor—roughly the same size as the one that Intel is putting into production in 2012—the fin must be about 10 nm wide and 25 nm high; it must also deviate by no more than half a nanometer—just a few atomic layers—in any given direction. Over the course of production, manufacturers must control all sources of variation, limiting it to no more than 1 nm in a 300-millimeter-wide wafer.

This precision is needed not only to manufacture the fin; it must also be maintained for the rest of the manufacturing process, including thermal treatment, doping, and the multiple film deposition and removal steps needed to build the transistor's gate insulator and gate. As an added complication, the gate oxide and the gate must be deposited so that they follow the contours of the fin. Any process that damages the fin could affect how the device performs. The resultant variation in device quality would force engineers to operate circuits at a higher power than they're designed for, eliminating any gains in power efficiency.


The unusual geometry of the FinFET also poses challenges for doping, which isn't required but can help cut down on leakage current. FinFET channels need two kinds of dopants: One is deposited underneath the gate and the other into the parts of the channel that extend on either side of the gate, helping mate the channel to the source and drain. Manufacturers currently dope channels by shooting ions straight down into the material. But that approach won't work for FinFETs. The devices need dopants to be distributed evenly through the top of the fin and the side walls; any unevenness in concentration will cause a pileup of charges, boosting the device's resistance and wasting power.

Doping will get only more difficult in the future. As FinFETs shrink, they'll get so close together that they will cast "shadows" on one another, preventing dopants from permeating every part of every fin. At Applied Materials' Silicon Systems Group, we've been working on one possible fix: immersing fins in plasma so that dopants can migrate directly into the material, no matter what its shape is.

Because UTB SOI devices are quite similar to conventional planar CMOS transistors, they are easier to manufacture than FinFETs. Most existing designs and manufacturing techniques will work just as well with the new thin-silicon transistors as they do with the traditional variety. And in some ways, UTB SOIs are easier to produce than present-day transistors. The devices don't need doped channels, a simplification that can save planar CMOS manufacturers some 20 to 30 steps out of roughly 400 in the wafer production process.

THE INDUSTRY STANDARD SOFTWARE PACKAGES IN GROUNDING AND ELECTROMAGNETIC INTERFERENCE

Are now more powerful with the addition of new standard components such as arbitrary transformers and cables.



Classical Equipotential Grounding



Multiple grounding systems having any shape, in simple or complex soils, including any number of layers or heterogeneous soil volumes.

- AutoGroundDesign
- AutoGrid Pro
- AutoGround & MultiGround

Complex Frequency Domain Grounding, Interference Analysis & Environmental Impact Assessment

Perform fast, yet complex and accurate, Interference analyses on pipelines, railways, etc. Examine electromagnetic impacts to the environment.

- SESTLC
- Right-Of-Way
- MultiLines & SES-Enviro
- MultiGroundZ



Frequency & Time Domain Arbitrary Network Analysis

Upgrade to the full power of MultiFields or CDEGS to tackle any electromagnetic problem involving aboveground and buried conductor networks including lightning or switching surge analysis.



World Leader in Grounding & EMI
www.sestech.com

800-668-3737

450-622-5000

CALL FOR ABSTRACTS

DEADLINE: Nov 30, 2011

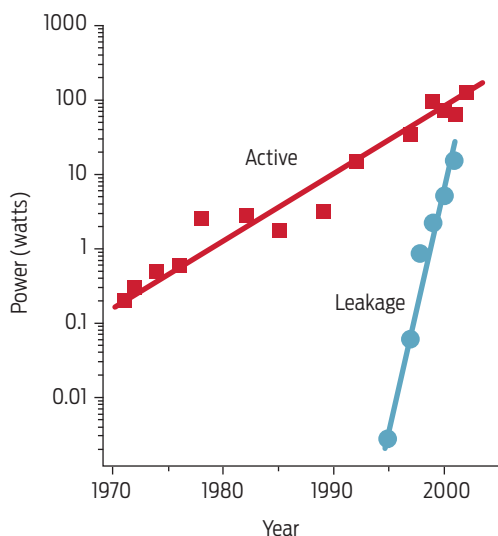
The National Trauma Institute invites scientists, engineers and providers to submit abstracts for presentation at its "Advances in Trauma" conference. Topics related to the use of computers, systems and IT in the critical care environment may include:

- Automation in the ICU and OR
- Fluid delivery, drug titration and hemodynamic control
- Patient dashboard technology for the ICU
- Advanced monitoring and computer derived vital signs and indices
- See NTI website for additional categories

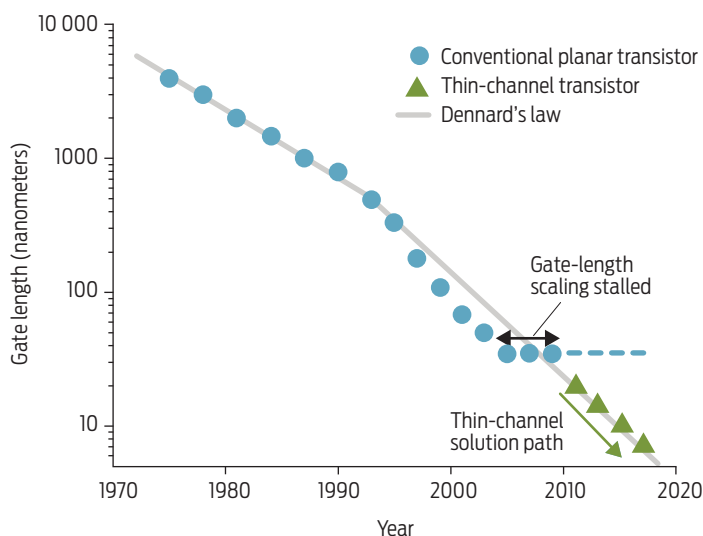
NTI presents
"Advances in Trauma"
May 9-10, 2012 -- San Antonio, Texas

NationalTraumaInstitute.Org





Source: Gordon Moore, Intel; IEEE



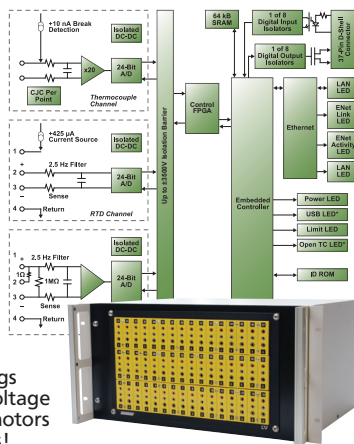
Source: Intel; Khaled Ahmed, Applied Materials

SHRINKING RETURNS: As transistors got smaller, their power demands grew. By 2001, the power that leaked through a transistor when it was off was fast approaching the amount of power needed to turn the transistor on [left], a warning sign for the chip industry. As these Intel data show, the leakage problem eventually put a halt to the transistor scaling [right], a progression called Dennard's law. Switching to alternate architectures will allow chipmakers to shrink transistors again, boosting transistor density and performance.

But the UTB SOI comes with its own challenges, chiefly the thin channel. The requirement that UTB SOI channels be half as thick as comparable FinFET fins makes any variations in thickness even more critical for these devices. A firm called Soitec, headquartered in Bernin, France, which has

been leading the charge in manufacturing ultrathin silicon-on-insulator wafers, is currently demonstrating 10-nm-thick silicon layers that vary by just 0.5 nm in thickness. That's an impressive achievement for wafers that measure 300 mm across, but it will need to be improved as transistors shrink.

Poof!
Unknown high voltage transients cause grief and lost money.
...but there is a solution



Measurement of sensor signals in industrial settings can provide **big surprises**. High common mode voltage from transients caused by nearby generators or motors can **wipe-out instrumentation**. We have seen this!

Using **galvanic isolation** to eliminate this high voltage makes sensor measurement straightforward.

We make energy solutions a reality.

MEASURpoint™ provides isolation up to $\pm 3500V$ continuously or $5000V$ for transients, protecting your analog sensor signals. Ethernet and application software included.

800-525-8528
 WWW.DATATRANSLATION.COM

DATA TRANSLATION®

IEEE WAS HERE



Innovation doesn't just happen.
 Read first-person accounts of IEEE members who were there.

IEEE Global History Network
www.ieeeahn.org



SMIC Ex. 1017
IPR2020-01003

SPECTRUM.IEEE.ORG



DOWN AND UP: A cross section of UTB SOI transistors [top] and a micrograph of an array of FinFET transistors [bottom].

PHOTOS: TOP, STMICROELECTRONICS; BOTTOM, INTEL

And it's not clear how precise Soitec's technique—which involves splitting a wafer to create an ultrathin silicon layer—can ultimately be made.

Another key stumbling block for UTB SOI adoption is the supply chain. At the moment, there are few potential providers of ultrathin SOI wafers, which could ultimately make manufacturers of UTB SOI chips dependent on a handful of sources. Intel's Mark Bohr says the hard-to-find wafers could add 10 percent to the cost of a finished wafer, compared to 2 to 3 percent for wafers bearing 3-D transistors (an estimate from the SOI Industry Consortium suggests that finished UTB SOI wafers will actually be less expensive).

GOING FORWARD, we expect that chipmakers will split into two camps. Those interested in the speediest transistors will move toward FinFETs. Others who don't want to invest as much in a switch will find UTB SOIs more attractive.

UTB SOI transistors have an additional feature that makes them particularly appealing for low-power applications:

25th Annual
**IEEE Canadian Conference on
 Electrical & Computer Engineering**

April 29–May 2, 2012
Montréal, Québec

Symposia	Featured Speakers
▶ Circuits, Devices & Systems	Alessandro Astolfi Imperial College London
▶ Control & Robotics	Gregor Bochmann University of Ottawa
▶ Power Electronics & Energy Systems	Peter Caines McGill University
▶ Signal & Multimedia Processing	Vahid Tarokh Harvard University
▶ Biomedical & Health Informatics	Ke Wu École Polytechnique de Montréal
▶ Computers, Software & Applications	
▶ Communications & Networking	

Important Dates	Contact
Tutorial/Workshop Proposals ▶ due Dec. 2, 2011	Amir Aghdam, Conf. Chair (514) 848-2424 Ext. 4137
Full-length Papers ▶ due Jan. 7, 2012	aghdam@ece.concordia.ca

<http://www.ccece2012.org>

Download **free**
white papers on

ieee
spectrum
For Tech Insiders **online**

Expert Information from Experts.

**Download a free
 white paper today!**

www.spectrum.ieee.org/whitepapers

IEEE



Where today's technology gurus converge.

The brightest minds discussing the biggest topics.

Earn CEUs and PDHs by attending a Tech Insider Webinar!

UPCOMING NOVEMBER WEBINARS

Big Demand Drives Small Solutions – Scalable Baseband Solutions • 16 November

In this webinar, Freescale will review the challenges engineers face while trying to create the next generation of wireless networks and how smart silicon approaches can help to simplify the task.

spectrum.ieee.org/webinar/1942325

Electrostatic Actuation with COMSOL Multiphysics • 17 November

For microscale devices, an applied voltage difference can be used to actuate a structure. This webinar shows how to model such electrostatic actuation, as well as adding subsequent analyses such as computing the AC electric response, the biased eigenvalues, and biased frequency response, as well as damping effects. spectrum.ieee.org/webinar/1927822

AVAILABLE ON DEMAND WEBINARS

Getting More Benefits From Modeling Vacuum Electronic Devices

Learn about the modeling of vacuum electronics devices such as Magnetrons, Klystrons, and Traveling Wave Tubes, using the VORPAL conformal-boundary particle-in-cell software. spectrum.ieee.org/webinar/1937415

Power System Contingency Analysis and Security Constrained Optimal Power Flow in MATLAB

This webinar will demonstrate how Optimization Toolbox and Parallel Computing Toolbox can be used for performing automated N-1-1 contingency analysis and SCOPF. spectrum.ieee.org/webinar/1911819

Plasma Modeling with COMSOL Multiphysics

In this webinar learn about the process of setting up a model of an inductively coupled plasma with help of application specific examples. spectrum.ieee.org/webinar/1921494

Requirements Management's key role in the Nuclear Industry

Learn how formalizing Requirements Management (RM) makes it possible to add a new dimension of safety, reliability and predictable project completion, to the nuclear power generation process. spectrum.ieee.org/webinar/1924206

SPONSORS:



Sign up today!
www.spectrum.ieee.org/webinar

A small voltage can easily be applied to the very bottom of a chip full of UTB SOI devices. This small bias voltage alters the channel properties, reducing the electrical barrier that stops current flowing from the source to the drain. As a result, less voltage needs to be applied to the transistor gates to turn the devices on. When the transistors aren't needed, this bias voltage can be removed, which restores the electrical barrier, reducing the amount of current that leaks through the device when it's off. As Thomas Skotnicki of STMicroelectronics has long argued, this sort of dynamic switching saves power, making the devices particularly attractive for chips in smartphones and other mobile gadgets. Skotnicki says the company expects to release its first UTB SOI chip, which will use 28-nm transistors to power a mobile multimedia processor, by the end of 2012.

That said, few companies have committed to one technology or the other. STMicroelectronics—as well as firms such as GlobalFoundries and Samsung—is part of the International Semiconductor Development Alliance, which supports and benefits from device research at IBM and is investing in both FinFETs and UTB SOIs. Exactly how the industry will split up and which design will come to dominate will depend on decisions made by the biggest foundries and how quickly standards are developed. Reports suggest that Taiwan Semiconductor Manufacturing Co., which dominates bespoke manufacturing in the chip industry, will begin making 14-nm FinFETs in 2015, but it's not clear whether the company will also support UTB SOI production. Switching to FinFET production requires a substantial investment, and whichever way TSMC swings, it will put pressure on other manufacturers, such as GlobalFoundries, United Microelectronics Corp., and newcomers to the foundry business such as Samsung, to choose a direction.

Also still unclear is how far each technology can be extended. Right now it looks like both FinFETs and UTB SOIs should be able to cover the next three generations of transistors. But UTB SOI transistors may not evolve much below 7 nm, because at that point, their gate oxide would need an effective thickness of 0.7 nm, which would require significant materials innovation. FinFETs may have a similar limit. In 2006, a team at the Korea Advanced Institute of Science and Technology used electron-beam lithography to build 3-nm FinFETs. But crafting a single device isn't quite the same as packing millions together to make a microprocessor; when transistors are that close to each other, parasitic capacitances and resistances will draw current away from each switch. Some projections suggest that when FinFETs are scaled down to 7 nm or so, they will perform no better than planar devices.

Meanwhile, researchers are already trying to figure out what devices might succeed FinFETs and UTB SOIs, to continue Moore's Law scaling. One possibility is to extrapolate the FinFET concept by using a nanowire device that is completely surrounded by a cylindrical gate. Another idea is to exploit quantum tunneling to create switches that can't leak current when they're not switched on. We don't know what will come next. The emergence of FinFETs and UTB SOIs clearly shows that the days of simple transistor scaling are long behind us. But the switch to these new designs also offers a clear demonstration of how creative thinking and a good amount of competition can help us push Moore's Law to its ultimate limit—whatever that might be.

POST YOUR COMMENTS
online at <http://spectrum.ieee.org/finfet111>