ieee spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
SPECTRUM.IEEE.ORG
5.11

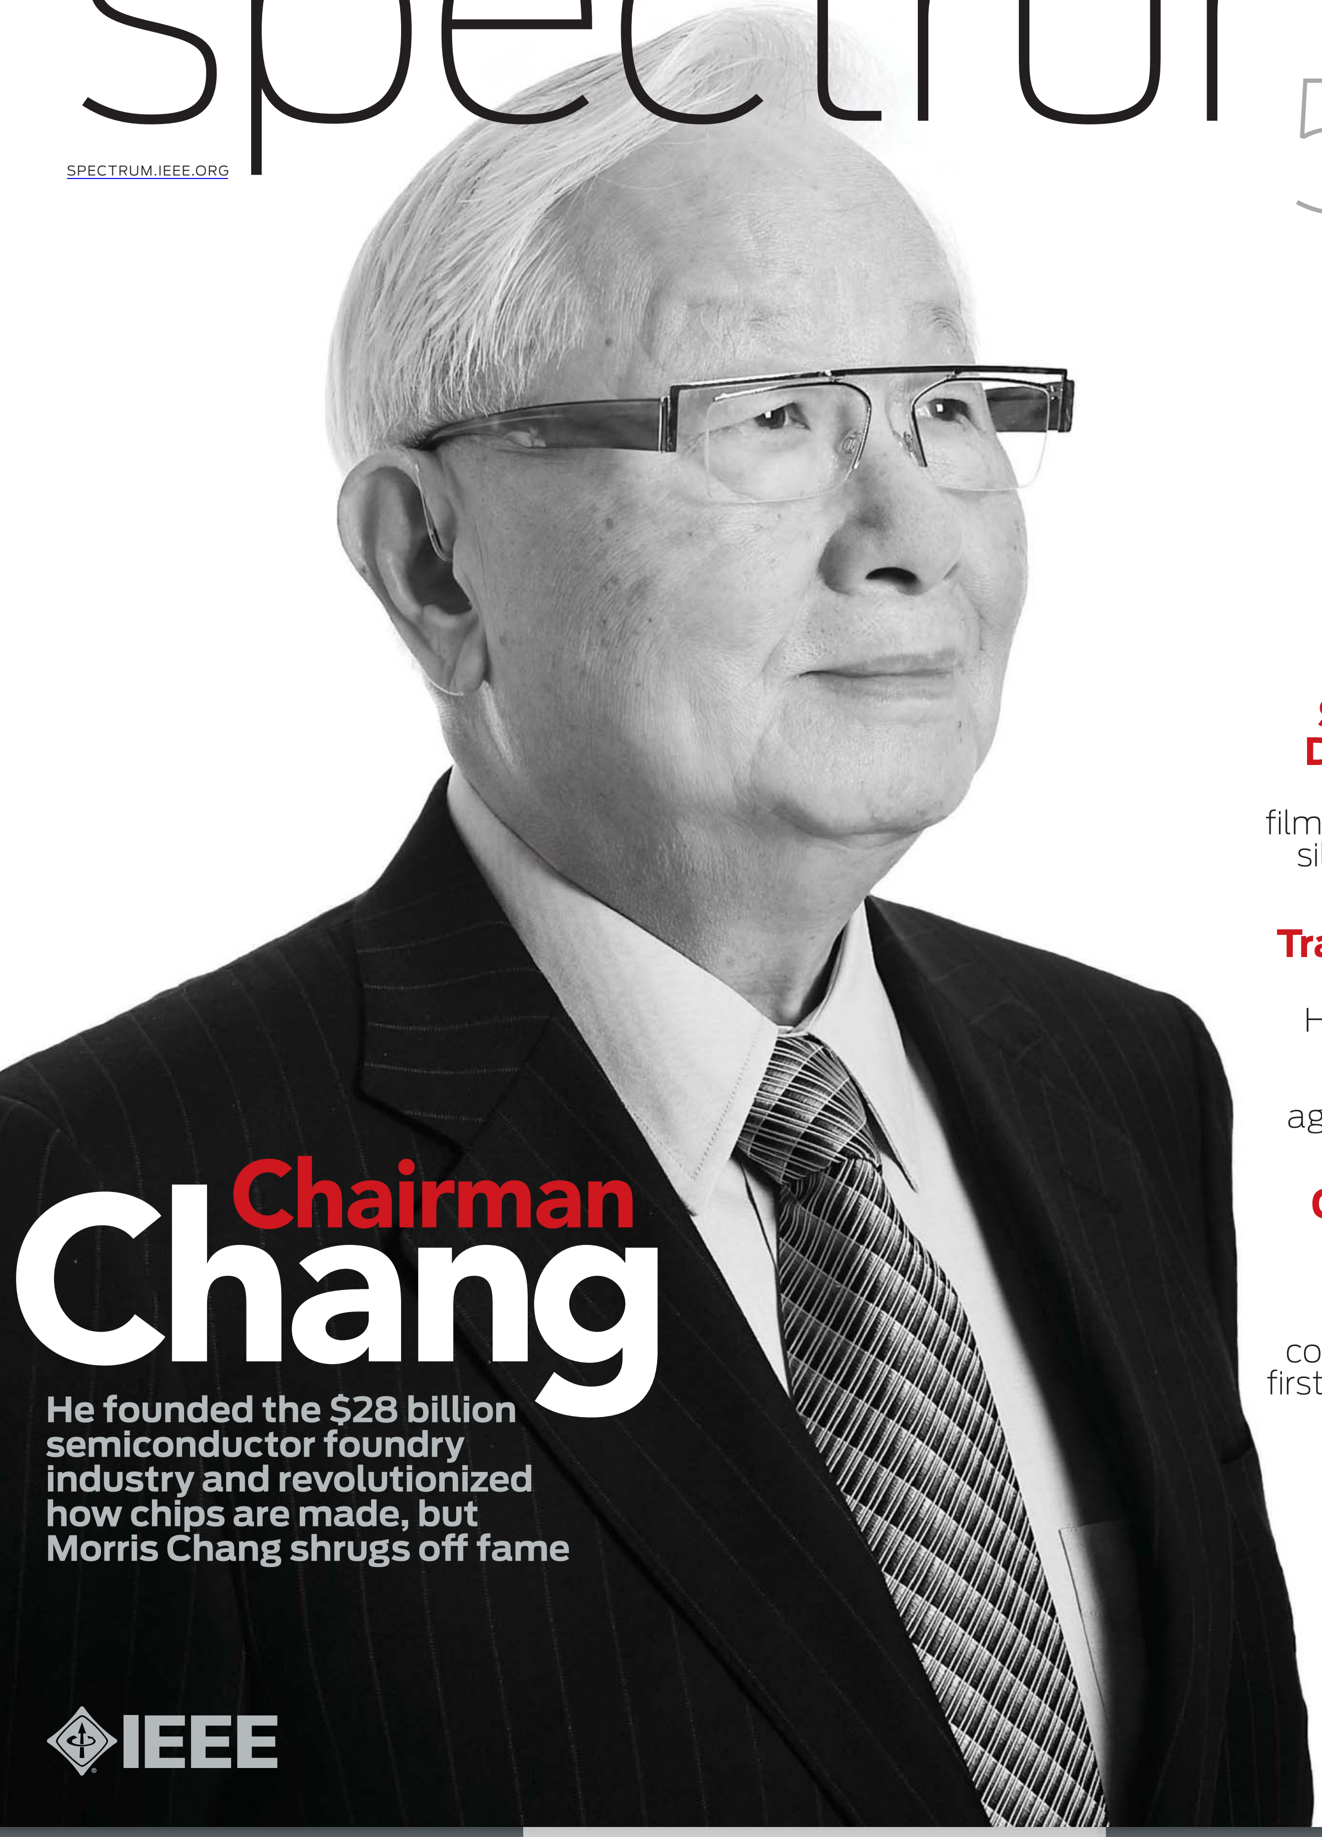

Chairman Chang

He founded the $28 billion semiconductor foundry industry and revolutionized how chips are made, but Morris Chang shrugs off fame

Faster, Sharper Displays

New thin films do what silicon can't

Transistor Decay

How much life is left in your aging chips?

Circuits of the Past

Relics of computing's first centuries

IEEE

ieee spectrum
volume 48 number 5 international
5.11

46

COVER: GARRET M. CLARKE

THIS PAGE, CLOCKWISE FROM LEFT: GARRET M. CLARKE; MARK RICHARDS; ERIC SUNDHOLM

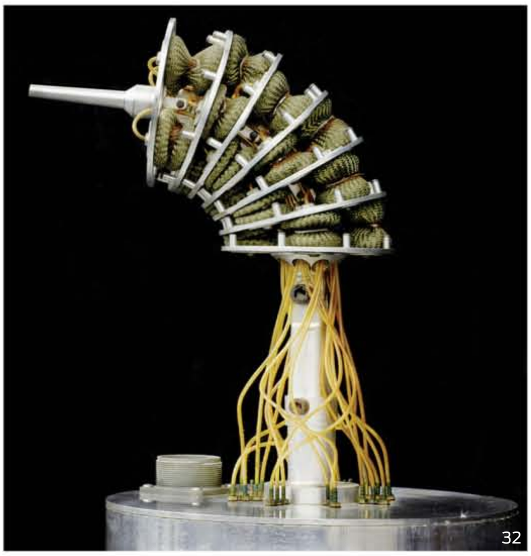

32

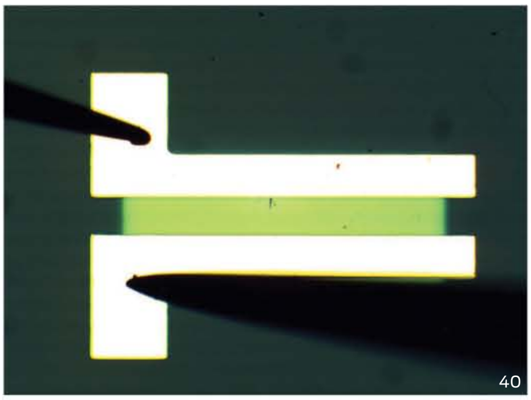

40

Thin Film Technology
Subsidiary of Williams Advanced Materials
incorporated
ENGINEERED MATERIALS
WILLIAMS ADVANCED MATERIALS
ENGINEERED MATERIAL SYSTEMS

ieee spectrum

volume 48 number 5 international

5.11

11

UPDATE

OPINION

DEPARTMENTS

CLOCKWISE FROM LEFT: JAGADEESH NV/EPA/LANDOV; AIR PHOTO SERVICE/REUTERS; HENRY GUTTMANN/GETTY IMAGES

SPECTRUM.IEEE.ORG

AVAILABLE 1 MAY

CONTINUING COVERAGE OF JAPAN'S NUCLEAR EMERGENCY

IEEE Spectrum editors and correspondents continue to follow developments at the stricken Fukushima Dai-ichi nuclear power plant. Frequent updates analyze containment and clean-up efforts. A detailed timeline helps you make sense of events, handy explainers contextualize the engineering challenges, and expert-opinion pieces parse the broader impact.

IEEE.ORG/THEINSTITUTE

AVAILABLE 6 MAY

MILESTONES GO TO SUPERCONDUCTIVITY AND WIRELESS EXPERIMENTS
The discovery of superconductivity and Marconi's early wireless experiments were honored in April with IEEE Milestones in Electrical Engineering and Computing.

IEEE SCIENCE EXHIBIT OPENS IN INDIA
IEEE recently unveiled a science exhibit at the B.M. Birla Science Centre in Hyderabad, India. It features hands-on activities demonstrating applications of electrical and computer engineering.

SYMPOSIUM EXPLORES FUTURE TECHNOLOGIES
The IEEE Technology Time Machine Symposium, being held 1 to 3 June, in Hong Kong, will examine emerging technologies, their current state of maturity, and scenarios for the future.

ONLINE WEBINARS & RESOURCES

AVAILABLE AT http://spectrum.ieee.org/webinar

June 2011: Smart-Grid Control Systems: Moving Towards a Self-Healing Grid

Digitizing the Power Grid: Fundamentals of Smart Grid Sensing & Analytics

Using MATLAB to Simulate and Optimize Multiphysics Designs

Parameter Estimation and Model Validation of Generator/ Excitation Systems

Kinetis ARM Cortex-M4 MCUs: The Inside Story

Effectively Managing Autostar Development

A Smart Approach to Maximizing ROI in Consumer Electronics Innovation

Best Practices for Delivering Value Through Quality Management

RF and Microwave Heating

Changing the Direction of Embedded Design

Free white papers by Intel:
http://spectrum.ieee.org/intel-resource-library

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2011 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/ institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of American Business Media, the Magazine Publishers of America, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

back story

TIMOTHY ARCHIBALD

Admiring the Obsolete

No tale of bits, hertz, or flops drew photographer Mark Richards to the technological relics displayed at the Computer History Museum, in Mountain View, Calif. "I was struck by the human aspects of these artifacts," he says, comparing a clustering of wires to a spinal column and core memory to a handwoven tapestry. "I like the marks. I like the dust. I like all the inconsistencies, the weirdness," he says.

When Richards was asked to name a few favorites, the anthropomorphic Orm—a robotic arm—was the first to come to mind. The wire-laden limb, which uses tiny air sacs to stretch and reach, appears in a parade of historic technologies in this month's issue [see "Bits of History"]. Richards logged more than two months of work—spread out over two years—snapping photos of 1200 machines from the museum's recently unveiled exhibit, *Revolution: The First 2000 Years of Computing.*

"I spent way too much time under fluorescent lighting," he says, "and my feet won't let me stand that long on concrete again." But it was worth it: "How many times are you going to be this close to the stuff that helped to get us to the moon or started the whole video-game industry?"

This is Richards's second collaboration with the museum. His first was a 2007 book, *Core Memory: A Visual Survey of Vintage Computers*. "Geeks are really different," he says of the experts on the museum's staff. "I felt like part of their family, which for a freelance photographer is pretty unusual."

A photojournalist, Richards has most often captured images of singular people and moments in history. In 1975, while in the U.S. Navy, he witnessed the fall of Saigon at the end of the Vietnam War, and in 1983 he interviewed and photographed Ahmad Shah Massoud, who was leading the troops that eventually drove the USSR out of Afghanistan. Photographing the first magnetic memory storage and a prototype *Pong* video-game machine, he says, isn't really all that incongruous. "It's all history," he says, "and it's kind of nice to be there." ❑

CITING ARTICLES IN IEEE SPECTRUM

IEEE Spectrum publishes two editions. In the international edition, the abbreviation INT appears at the foot of each page. The North American edition is identified with the letters NA. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, the The Data is in *IEEE Spectrum,* Vol. 48, no. 5 (INT), May 2011, p. 56, or in *IEEE Spectrum*, Vol. 48, no. 5 (NA), May 2011, p. 60.


Name
Dr. Christian Altenbach, Certified LabVIEW Associate Developer
Job Title
Research Ophthalmologist
Area of Expertise
Biophysics
LabVIEW Helped Me
Analyze and visualize data interactively using custom algorithms
Latest Project
Mapping molecular structure changes during activation of the light receptor, rhodopsin
NI LabVIEW
LabVIEW makes me better because I can
PROGRAM
the way I think

contributors

JOSEPH CALAMIA is a freelance writer in New York City. The chats he had with the staff of the Computer History Museum while writing "Bits of History" [p. 32] inspired him to dust off his own obsolete gadgets: the 1980s Apple IIGS and the original Nintendo Entertainment System. Neither worked. Calamia holds a master's in science writing from MIT and has also written for *Discover*.

GARRET M. CLARKE already knew quite a bit about the semiconductor business when he photographed Morris Chang, IEEE Medal of Honor winner and founder of Taiwan Semiconductor Manufacturing Co. [p. 46]. While earning an MBA at the National Cheng Kung University, Clarke studied plant design and efficiency in the chip industry. But as a photographer he often roams far beyond the high-tech world: He recently photographed fishermen and ice factory workers in southern Taiwan, striving to capture "romantic images of men who are often seen as the rougher class of society."

JOHN KEANE and **CHRIS H. KIM** met in 2005, when Keane, then a graduate student (and now a component-design engineer for Intel in Hillsboro, Ore.), joined Kim's research group at the University of Minnesota. To study the wear and tear transistors experience, Keane and Kim designed special-purpose chips, which they describe in "An Odometer for CPUs" [p. 26]. Making your own integrated circuits is tricky and expensive, but it's an addictively productive way to do this kind of research, Kim says: "Once you start building chips, it's hard to go back."

BRIAN PROFFITT is the author of "Open Android—for Better and for Worse" [p. 22], which compares the Linux-based Android and Apple's iOS operating systems. Proffitt is an adjunct business instructor at the University of Notre Dame and has written 19 books on consumer technology and one on the philosopher Plato. His latest, *Take Your iPad to Work*, was published in October.

PRACHI PATEL wrote "The Aging Nuclear Workforce" [p. 24], which may bring good news for those considering careers in nuclear engineering. As an *IEEE Spectrum* contributing editor, Patel has written on a variety of topics, including the sequencing of the Neanderthal genome [see "Computing the Caveman," July 2010]. "I'm fascinated by archaeology, so that story was especially fun to write," she says.

JOHN F. WAGER and **RANDY HOFFMAN**, authors of "Thin, Fast, and Flexible" [p. 40], began working together at Oregon State University. Wager is a professor of electrical engineering and computer science there and the lead author of the 2008 book *Transparent Electronics*. His student Hoffman—while working toward his master's in electrical engineering—researched oxide semiconductor–based thin-film transistors, producing the first in a series of key patents. He is now a senior engineer at HP, where he develops oxide-based, printed, and flexible electronics and active-matrix displays.

CLOCKWISE FROM TOP LEFT: CHRIS CARLEY; BARB ZAPPIA; JOEL PREDD; GARRET M. CLARKE


Name
Hector Guajardo Betancourt, Certified LabVIEW Architect
Job Title
Automated Test and Control Engineer
Area of Expertise
Manufacturing Test
LabVIEW Helped Me
Reduce test time by more than 10X
Latest Project
Building a vision-based inspection system for washing machine drums
NI LabVIEW
LabVIEW makes me better because I can deliver multiple
PROJECTS
on time, on spec, and on budget

spectral lines

Fukushima's Positive Impact

AMID THE stubbornly disappointing string of news reports emanating from Japan's Fukushima Dai-1 nuclear complex, there are signs that its melting nuclear fuel rods are inspiring some important and long-overdue developments in global power systems. And there's good news for both nuclear energy supporters and critics.

Hopeful outcome No. 1: Berlin is getting serious about upgrading the balkanized and inadequate transmission grid that represents a serious liability for Germany's renewable energy ambitions.

Chancellor Angela Merkel's decision in March to shut down the country's oldest nuclear reactors and temporarily scrub life extensions for the rest was widely seen as a sop to voters in the state of Baden-Württemberg. But it led to real momentum for grid reform. A document leaked from Germany's Federal Ministry of Economics and Technology revealed plans to revamp the power grid—a precondition to replacing nuclear energy with solar, wind, and other renewable power sources.

Back in 2005, the German Energy Agency estimated that the country needed 850 kilometers of new high-voltage lines by 2015 to absorb growing levels of wind power on its fragmented grid. So far, only 90 km have been built, and the challenge keeps growing: An updated study released in November estimated that 3600 km of new lines must be built by 2025.

Hopeful outcome No. 2: China slows its nuclear building spree somewhat. Beijing reaffirmed its nuclear ambitions immediately after Japan's earthquake and tsunami. But then it began to backpedal, suspending new plant approvals and stepping up safety inspections at existing plants. Whether you're a nuclear energy supporter or not, this is very good news.

That's because while China's nuclear future is assured, its nuclear safety is not.

The problem is that China's nuclear exuberance—it has 27 reactors under construction—has outstripped the process of training sufficient numbers of new nuclear operators and inspectors. This serious disconnect inspired surprisingly frank criticism from China's national nuclear safety administration director, Li Ganjie, two years ago—at the time to little apparent effect.

Confronting proposals to more than double the pace of China's nuclear construction schedule, Li courageously warned at an International Atomic Energy Agency meeting in Beijing that "over-rapid expansions" could diminish reactor quality and safety. Li lost that round, and China's nuclear capacity goal for 2020, already set to jump more than fourfold to 40 gigawatts, shot up to 86 GW.

Now a revised scheme may be in the offing. Chinese state media have been cited reporting that "China is likely to scale back its ambitious plans...under a new policy that stresses safety instead of rapid development." The Associated Press quoted deputy director of the China Electricity Council, Wei Zhaofeng, predicting that the policy change would trim growth by 10 GW.

That 10 GW may look small, but only relative to China's aggressive agenda. U.S. nuclear power proponents would count themselves lucky if the United States added that much reactor capacity by 2020.

Hopeful outcome No. 3: The Chinese are now also talking about doubling their goal for solar power for 2015, from 5 GW to 10 GW, in response to Fukushima. Solar power is finally getting credit for being fast to install and reliable, a big change for an energy source that had often been characterized as "intermittent" and "insecure."

So if it can be said that any good can come out of a disaster like Fukushima, perhaps it will be a renewed global commitment to rigorous nuclear safety assessment and monitoring and a renewed global interest in other carbon-free energy solutions. —PETER FAIRLEY

SCALING BACK: Post-Fukushima, China may adopt a new policy that "stresses safety instead of rapid development." But construction will continue on a number of existing projects, such as this pressurized water reactor, one of six planned in Taishan, Guangdong province.

QILAI SHEN/BLOOMBERG/GETTY IMAGES

Peter Fairley is a freelance journalist specializing in energy and the environment and a contributing editor to IEEE Spectrum.

This article is based on a post that appeared on Spectrum's *Energywise blog on 4 April 2011.*

Correction

In "Good Things in Small Packages" [March], we incorrectly stated, "When a device shrinks by 66 percent, for example...the electronics must shrink to a third or less of their original size." We should have written, "When a device shrinks *to* 66 percent...."

>> Find out how LabVIEW can make you better at **ni.com/labview/better** **800 453 6202**

NATIONAL INSTRUMENTS™

COMSOL MULTIPHYSICS®

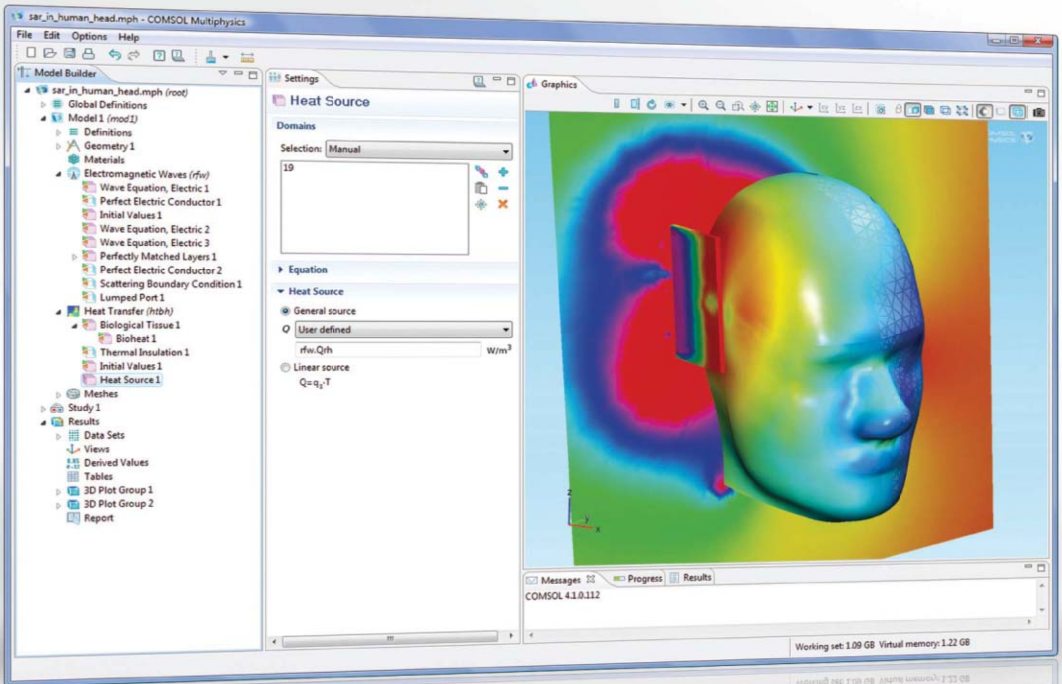

Absorbed Radiation (SAR value) in a Human Head.

Capture the Concept.

With COMSOL Multiphysics® you are empowered to build the simulations that accurately replicate the important characteristics of your designs. The key is the ability to include all physical effects that exist in the real world. This multiphysics approach delivers results—tangible results that save precious development time and spark innovation.

Watch Tutorial

Joule Heating of a Circuit Board Fuse
www.comsol.com/showcase

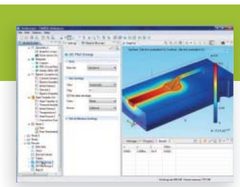

COMSOL

update

more online at spectrum.ieee.org

THUMB DRIVE: A program aiming to provide biometric IDs to all Indians has begun with a few million rural residents. This woman from a village near Bangalore is having her fingerprints and irises scanned and her photo taken. When the project is complete, India will have the largest biometric ID database in the world.

PHOTO: JAGADEESH NV/EPA/LANDOV

Fast Start for World's Biggest Biometrics ID Project

In India, a few million people have been registered so far—only a billion left to go

IDENTITY IS EASY to take for granted. Most of us have multiple legal documents or ID cards that prove we are who we say we are. But for many of the world's poorest citizens, the lack of legal identity is a barrier to participating in commerce and receiving services.

In India, an estimated 500 million people have no form of reliable identification. It's a problem the Indian government has set out to fix through a five-year project with a budget of US $430 million for this year. Starting six months ago with rural populations, the government has begun to create a biometric database that will eventually contain an unprecedented hundreds of millions of records. "We are talking about 10 times more than anything else that has been done before," says Anil Jain, an IEEE Fellow and distinguished professor at Michigan State University, who is an expert in biometrics.

From each volunteer participant, the government collects 10 fingerprints, 2 iris images, and a photo, and if the new data don't match any identity already enrolled, it assigns the person a unique 12-digit number. After that, a single fingerprint or iris scan should be all that's needed to verify the identity of any person. As of the end of March, the Unique Identification Authority of India (UIDAI) has registered more than 4 million people this way. The UIDAI hopes to eventually collect biometrics from a majority of the Indian population.

India has many federal and state programs to help people living in poverty, but today it's nearly impossible to be sure that funds and benefits are actually being delivered to those who need them.

400 megawatts Annual thin-film solar cell production capacity at what will be the largest such factory in the United States, according to the plant's owner, GE.

update

The ID project is an attempt to cut down on fraud and graft by increasing accountability and transparency. It's also meant to provide access to banking and the formal economy that many people lack.

Government biometrics programs have been tried before and failed, in India and elsewhere. The United Kingdom's universal ID program, for instance, got bogged down by both costs and privacy concerns and didn't offer tangible benefits to the average citizen. But the UIDAI's universal ID program, or Aadhaar, as it's called, seems to be off to a fast start. As soon as he was appointed in July 2009, chairman Nandan M. Nilekani set the ambitious goal of issuing the first million IDs within 12 to 18 months, and the UIDAI hit that mark by January 2011. Efficiency is not a strength of most government bureaucracies, so Nilekani looked to Silicon Valley for help.

A core group of Indian expats with Silicon Valley start-up experience began working on the problem, as unpaid volunteers. A three-bedroom flat in Bangalore served as the group's living and work space for six months. "We converted the living room into an office and started designing the core of the biometric system," says Salil Prabhakar, a biometrics expert who was recruited for the project. In the meantime, Nilekani's office prepared the paperwork and logistics to give the team official authority.

Prabhakar and his colleagues started by looking at best practices developed by other biometrics researchers. The U.S. National Institute of Standards and Technology, the International Standards Organization, and the FBI "have done a lot of work to come up with good specifications," says Prabhakar. "If UID happened 10 years ago, life would have been much harder for us." By using standard specifications, the Indian project has been able to use many existing devices and data interchange formats and avoid having to rely on a single private biometrics equipment vendor or proprietary format.

SMILE! A villager in southern India is one of the millions whose biometric data have been registered. *PHOTO: JAGADEESH NV/EPA/LANDOV*

Still, the Indian project has presented unique obstacles. Take fingerprints, the most mature and best-understood biomarker. Because almost all scholarly publications had focused on Western subjects for law-enforcement purposes, Prabhakar and his colleagues were in the dark about how well such systems work when data are collected in a variety of locations and conditions from a rural population of Indians, many of whose fingerprints have been obscured or erased by manual labor.

So UIDAI started with smaller proof-of-concept programs to test the accuracy of the ID generation process with tens of thousands of subjects. There are two factors that determine accuracy: the false positive rate, which is how often a newly registered person's record is judged to be a duplicate of someone else's record, and the false negative rate, which is the frequency that true duplicate records—for instance, if someone registered twice under two names—are not recognized as such. In the pilot study, the researchers found that by adding the iris scans to the 10 fingerprints, they could decrease the false negative rate by a factor of 50 over the use of irises alone and by a factor of 25 over the use of fingerprints alone. "It's hard to make a system error free," says Jain, "but the important thing is to minimize the error."

When a duplicate record is detected, it is flagged for manual verification. In the pilot study, the system generated false positive errors in 0.0025 percent of cases. That rate would generate 25 erroneous records daily, if the project makes the goal of 1 million IDs per day by October 2011 set by Finance Minister Pranab Mukherjee.

In order to scale up so quickly, UIDAI has piggybacked on existing government infrastructure: Biometrics data are collected at government offices by government employees or private enrollment agencies. These workers use officially approved sensors costing $1000 to $2000. A single sensor can collect data from 50 people per day.

So far, the project has done better with speed than accuracy. Raj Mashruwala, part of the core volunteer team, says that the biggest challenge has been consistently getting high-quality data from the enrollment operators. "I don't think they are yet at a stage where it's working flawlessly," he says, but UIDAI will be updating training and procedures to make improvements.

All of this is still just the first phase of the program. Additional pilot tests are ongoing to figure out the best way to run verifications against the database. Despite the rapid early pace, it will be years before the project can succeed in all its goals.

—Joshua J. Romero

DOWNED LINES: The 11 March earthquake and tsunami did lasting damage to Japan's grid.
PHOTO: PHILIPPE LOPEZ/AFP/GETTY IMAGES

Fixing Japan's Fragmented Grid

An east-west frequency divide prevents power flows

The earthquake and tsunami that struck Japan in March blew a large hole in the country's power supply. Despite strenuous efforts, by this summer Tokyo Electric Power Co. (TEPCO) could be dealing with an 8- to 9-gigawatt shortfall under a peak load of up to 60 GW.

TEPCO's supply situation would look less grim were it not for a quirky split that divides Japan's power system in half: While Tokyo and the rest of eastern Japan run on 50-hertz electricity, the big cities southwest of Tokyo and the rest of the country run at 60 Hz. It's this historical accident from the 19th century that's constraining TEPCO's ability to seek help from the 56 percent of Japan's generating capacity in its west.

"It's a shame. The western grids can supply a lot. I think they could cover [TEPCO's] peak demand," says Kent Hora, executive vice president for Mitsubishi Electric Power Products.

Japan's grid managers are now hatching plans to beef up the country's west-to-east power flow capabilities. As it stands, just three small installations are able to squeeze power across Japan's AC frequency frontier. These pull alternating current off one grid, convert the power to high-voltage direct current (HVDC), and then synthesize a novel AC wave to add the power to the other grid. Together these three facilities can push up to 1.2-GW of power east or west.

"Under normal conditions, these kinds of systems would take 18 to 24 months" to install, says Greg Reed, a power engineering professor at the University of Pittsburgh. "Could we get one installed and in service in less than 12 in an emergency situation like this? Absolutely."

The fastest way to add conversion capacity, according to Reed, is voltage-source converter (VSC) technology. VSC-based HVDC uses relatively advanced switches, such as insulated-gate bipolar transistors, to simultaneously transmit DC power and regulate the voltage of neighboring AC lines. VSC costs about 25 percent more than traditional HVDC, but it's quicker to install.

One reason VSC may be faster is that it takes up about 35 percent less space than HVDC systems. That's because VSC produces a cleaner synthetic AC wave and therefore requires less filtering equipment. So the new converters could be set beside the three existing stations, simplifying any required expansion of the AC lines feeding the stations.

Jan Johansson, an HVDC expert with the European engineering firm ABB, stresses another VSC advantage—the fact that VSC-based converters do not require time-consuming grid-modeling studies. Japan's present HVDC requires that the AC grid be stabilized in order to avoid voltage fluctuations, but VSC systems actively stabilize the grid on their own.

Even if VSC wins the day, experts disagree on how much capacity should be added. Mitsubishi's Hora advocates a major expansion to help Japan's grids respond to future emergencies. But Akihiko Yokoyama, a power systems expert at the University of Tokyo, endorses a plan to add a modest 300 megawatts of capacity. More ambitious proposals, he says, are both excessively expensive and impractical.

It might take two years before any extra capacity kicks in. But with so much damage to the country's infrastructure, eastern Japan may need the power even then. —Peter Fairley

A version of this article appeared on the Web in April as "Why Japan's Fragmented Grid Can't Cope."

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

$900 million Amount that Google bid for Nortel Networks' patent portfolio. The telecom-equipment maker went bankrupt in 2009 and has been selling pieces of itself to satisfy creditors.

update

China's Godson Gamble

Mesh structure sets supercomputer processor apart from the competition

The Dawning 6000 supercomputer, which Chinese researchers expect to unveil in the third quarter of 2011, will have something quite different under its hood. Unlike its forerunners, which employed American-born chips, this machine will harness the country's homegrown high-end processor, the Godson-3B. With a peak frequency of 1.05 gigahertz, the Godson is slower than its competitors' wares, at least one of which operates at more than 5 GHz, but the chip still turns heads with its record-breaking energy efficiency. It can execute 128 billion floating-point operations per second using just 40 watts—double or more the performance per watt of competitors.

The Godson has an eccentric interconnect structure—for relaying messages among multiple processor cores—that also garners attention. While Intel and IBM are commercializing chips that will shuttle communications between cores merry-go-round style on a "ring interconnect," the Godson connects cores using a modified version of the gridlike interconnect system called a mesh network. The processor's designers, led by Weiwu Hu at the Chinese Academy of Sciences, in Beijing, seem to be placing their bets on a new kind of layout for future high-end computer processors.

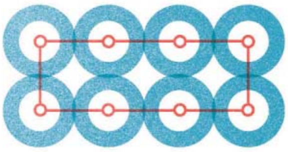

8-CORE RING: In most microprocessors, information circles around a ring-shaped interconnect to reach processor cores.

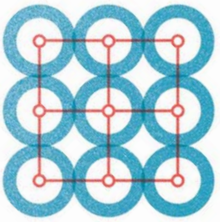

9-CORE MESH: Some new high-end processors use a mesh network. These can be more complex but also more energy efficient.

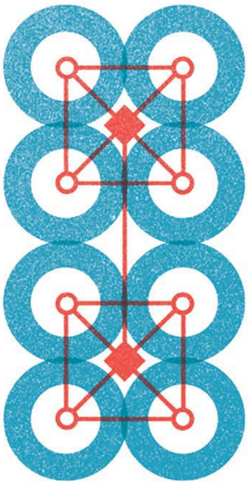

8-CORE GODSON: The Chinese processor relies on a modified mesh network that contains extra direct connections to move data efficiently.

ILLUSTRATIONS: GAVIN POTENZA

A mesh design goes hand in hand with saving energy, says Matthew Mattina, chief architect at the San Jose, Calif.–based Tilera Corp., a chipmaker now shipping 36- and 64-core processors using on-chip mesh interconnects.

Imagine a ring interconnect as a traffic roundabout. Getting to some exits requires you to drive nearly around the entire circle. Traveling away from your destination before getting there, says Mattina, requires more transistor switching and therefore consumes more energy. A mesh network is more like a city's crisscrossed streets. "In a mesh, you always traverse the minimum amount of wire—you're never going the wrong way," he says.

On the 8-core Godson chip, 4 cores form a tightly bound unit—each core sits on a corner of a square of interconnects, as in a usual mesh. Godson researchers have also connected each corner to its opposite, using a pair of diagonal interconnects to form an *X* through the square's center. A "crossbar" interconnect then serves as an overpass, linking this 4-core neighborhood to a similar 4-core setup nearby.

Godson developers believe that their modified mesh's scalability will prove a key advantage, as chip designers cram more cores onto future chips. Yunji Chen, a Godson architect, says that competitors' ring interconnects may have trouble squeezing in more than 32 cores.

Indeed, one of the ring's benefits could prove its future liability. Linking new cores to a ring is fairly easy, says K.C. Smith, an emeritus professor of electrical and computer engineering at the University of Toronto. After all, there's only one path to send information—or two in a bidirectional ring. But sharing a common communication path also means that each additional core adds to the length of wire that messages must travel and increases the demand for that path. With a large number of cores, "the timing around this ring just gets out of hand," Smith says. "You can't get service when you need it."

Of course, adding more cores in a mesh also stresses the system. Even if you have a grid of paths providing multiple communication channels, more cores increase the demand for the network, and more demand makes traveling long distances difficult: Try driving across New York City at rush hour. Still, the bandwidth scaling of a mesh interconnect is superior to that of a ring, Tilera's Mattina says. He notes that the total bandwidth available with a mesh interconnect increases as

$625 million Amount that Apple will *not* have to pay to Mirror Worlds in a patent infringement case. Had Apple lost, it would have resulted in a new record for a patent infringement trial.

update

you add cores, but with a ring interconnect, the total bandwidth remains constant even as the core count increases. Latency—the time it takes to get a message from one core to another—is also more favorable in a mesh design, Chen says. In a ring interconnect, latency increases linearly with the core count, he says, while in a mesh design it increases with the square root of the number of cores.

Reid Riedlinger, a principal engineer at Intel, points out that a ring interconnect has its own scalability benefits. Intel's recently unveiled 8-core Poulson design employs a ring not only to add more cores but also to add easy-to-access on-chip memory, or cache. As long as the chip has the power and the space, Riedlinger says, a ring makes it easy to add each core and cache as a module—a move that would require more complicated validity studies and logic modification in a mesh. "Adding the additional ring stop has a very small impact on latency, and the additional cache capacity will provide performance benefits for many applications," he says.

For those who are not building a national supercomputer, Riedlinger also points out that a ring setup is more easily scalable in a different direction. "You might start with an 8-core design," he says, "and then, to suit a different market segment, you might chop 4 cores out of the middle and sell it as a different product."

—Joseph Calamia

Sony vs. the Hackers

The game titan has taken the gloves off in its fight to protect the PlayStation 3

OVER THE PAST FEW MONTHS, Sony's fight against PlayStation 3 hackers has become a meta-game of its own. The battle pits the game industry giant against geeks who claim to be more interested in exploring the console than in running pirated wares. At stake is Sony's need to prevent the piracy of its software partners' games and PS3 owners' right to do what they will with a product they purchased. —*David Kushner*

GAVIN POTENZA

11 NOV. '06 Sony releases PlayStation 3. Firmware includes **Other OS**, a program that allows players to run operating systems such as Linux on PS3.

22 OCT. '08 Hacker StreetskaterFU unpacks and decrypts the beta client for the PS3's **Home** social gaming service.

29 DEC. '09 Hacker George "GeoHot" Hotz challenges other hackers: "The PS3 has... **yet to be hacked**. It's time for that to change."

3 JAN. '10 Hacker MysticHades releases a method allowing **copied Blu-ray games** to be played on the PS3.

25 JAN. '10 Hotz tells the BBC he has hacked the PS3. He accomplishes this by gaining the ability to read from and write to the PS3 memory and then securing high-level access to the processor. **"I can now do whatever I want with the system,"** he says. "It's like I've got an awesome new power—I'm just not sure how to wield it."

26 JAN. '10 Hotz **releases his PS3 hack** to the public, allowing users to run their own operating systems and other software on the machine.

28 MAR. '10 Sony announces the release of the firmware 3.21 update to **disable Other OS**—the firmware that allows for alternate operating systems—"due to security concerns."

7 APRIL '10 Hotz comes up with a way to **restore Other OS** to PS3 and posts a video of his work.

19 AUG. '10 Australian hackers OzModChips announce release of **PS Jailbreak**, a USB dongle that allows players to boot up both backup and bootleg versions of games.

29 DEC. '10 The German hacker team failOverflow reports that it was able to discover the **encryption key** that allows users to run their own software on the PS3. The team does not, however, release the key.

6 JAN. '11 Hotz publishes the encryption key and uploads video showing him running **home-brew software** on PS3's latest firmware.

27 JAN. '11 Sony sues Hotz and failOverflow for **copyright infringement**.

3 MAR. '11 The Northern California U.S. District Court approves Sony's subpoena request requiring **Twitter**, **Google**, **YouTube**, and **Bluehost** to supply information on Hotz's hacking efforts. The subpoena includes a demand for the IP addresses of anyone who visited his blog.

31 MAR. '11 Sony and Hotz **quietly settle** the suit. News of the settlement doesn't leak out until more than a week later.

6 APRIL '11 The hacker group **Anonymous** attacks Sony Web sites as revenge for Sony's suit against Hotz.

update

The Scientist as Space Tourist

Private rockets like *SpaceShipTwo* will offer space-based science on the cheap

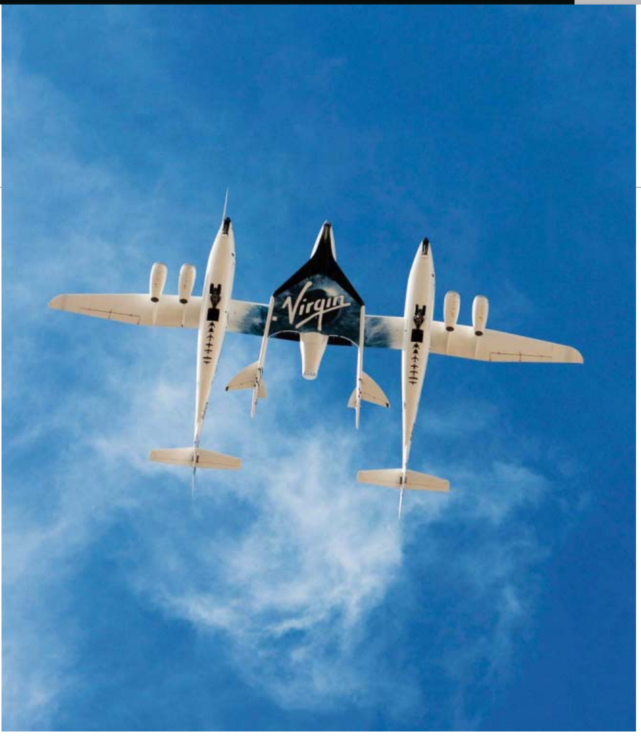

UP, UP, AND AWAY: Virgin Galactic's *SpaceShipTwo*, here carried beneath its launch plane, may make a good flying laboratory. PHOTO: BLOOMBERG/GETTY IMAGES

POWERED FLIGHT tests are slated later this year for *SpaceShipTwo*, a craft built to carry humans on brief excursions more than 100 kilometers into space. And although publicity indicates that these travelers will be rich thrill-seeking "space tourists," some of them are, in fact, going to be scientists, engineers, and probably even graduate students on funded research programs, according to S. Alan Stern, a scientist now spearheading the development of this new concept.

SpaceShipTwo is the descendant of *SpaceShipOne*, the X-Prize–winning craft designed by Burt Rutan and Scaled Composites. Now funded by Virgin Galactic, the space company owned by British aviation mogul Sir Richard Branson, the new vehicle was geared to wealthy adventurers willing to pay US $200 000 for a ride into space.

But an entirely new category of passengers has recently emerged. These space commuters would use the flights to expand on research that has been performed since the 1950s aboard unmanned sounding rockets. The combination of the advantages a human experimenter has over a robotic one and the competitive cost of a ticket on *SpaceShipTwo* make such research flights an attractive idea.

Stern intends to be one of those experimenters and to fly on multiple missions in support of research customers. A veteran planetary scientist and once the director of NASA's sounding rocket program, Stern is now an associate vice president with the Southwest Research Institute (SWRI), headquartered in Austin, Texas. He recently helped organize a conference in Orlando, Fla., to promote the concept of expanding standard lab work to the space environment. Stern sees this new era of human space research as revolutionary in its ease and economy. "The infrastructure is like flying on the K-bird" [NASA's KC-135 zero-*g* aircraft], he says. "It's never complicated or burdensome," as space shuttle–based science has proved to be. Repeated flights with fast turnaround will revolutionize the field, he believes.

As with all space science, robots and other automation are options, but Stern thinks they are not good ones. Automation takes time, costs money, and restricts what is likely to be learned, he asserts. "NASA spends $50 million a year on sounding rockets and performs 20 to 25 missions," he told *IEEE Spectrum*. "That's an average cost of $2.5 million per flight." A single seat on a space tourist vehicle provides the equivalent duration and altitude for one-tenth as much.

Like scientific sounding rockets, tourist flights on *SpaceShipTwo* could perform observations of the upper atmosphere and of the sun and other celestial objects. Onboard apparatus could also exploit the approximately 4 minutes of pure microgravity to explore physical features ranging from fluid and flame dynamics to simulations of the dust clumping that led to the formation of planets.

Beyond that, such flights could help troubleshoot and validate crew equipment—including scientific instruments, tools, and other mechanisms—intended for later long-term use on the International Space Station. This process alone, says Stern, could speed up improvements to space gear while significantly reducing costs.

SWRI has signed a contract with Virgin Galactic for a number of flights by Stern and two other researchers, beginning as early as next year. Already, SWRI has developed three flight-ready experiments to help characterize the capabilities of on-the-scene researchers as well as produce genuine scientific results. One is a self-contained ultraviolet telescope that's aimed and activated manually.

In addition to Virgin Galactic, other carriers, such as XCOR Aerospace, in Mojave, Calif., have begun designing payload accommodation features on their vehicles. These include in-cabin experimental setups and racks for mounting instruments outside the pressurized cabin. There could also be dispersal mechanisms for chemical tracer releases that allow measurements of the upper ionosphere—the most poorly explored region of Earth's atmosphere.

"There is a market for commercial payload specialists," Stern says. In addition to performing experiments for clients, his team will provide flight escorts for those who want to fly themselves, with assistance. —JAMES OBERG

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

hands on

Beehackers

Cheap widgets are like honey to hive keepers

PITY THE poor beekeeper. While bee researchers play with high-frequency dancing robotic bees, DNA scanners, and forensic pollen analysis, beekeepers must scavenge 19th-century feed scales off eBay.

The problem is money. Even though bees play a crucial role in the pollination of agricultural products worth billions of dollars, a hive typically produces honey that's worth no more than US $1000 a year at retail. A few lucky beekeepers get hired by farmers to pollinate their crops, but the overall margin is still far too slim for fancy modern equipment. So beekeepers typically are able to track the health and honey-making performance of their charges in only the crudest of ways.

And although that may have been all right in previous centuries, in this one honeybees have come under increasing pressure from disease, pesticides, fragmentation of foraging space, and even a mysterious ailment known simply as colony collapse disorder. (Last fall, university researchers collaborating with U.S. Army experts in germ warfare announced that they had discovered a new virus and a common fungus that were present in all hives that suffered from collapse, but that conclusion has been questioned and does not necessarily point to a cure.) Ordinary beekeepers may need high-tech help, but it's not clear how they can afford it.

Tom Rearick, an electrical engineer, and some fellow "beehackers" are trying to change all that. He wants his site, BeeHacker.com, to become a hub of on-the-cheap development of appropriate technology for beekeepers, with projects ranging from simple hive scales to laser-based bee tracking. For example, a $20 luggage scale augmented with $5 to $10 of scrap hardware can check the weight of dozens of hives a day. That would give a rough idea of how much honey the bees are producing and of the general health of the hive. With Rearick's hack, you just lift one side of the hive gently with a pry bar connected to the scale by a cable. Assuming that honey and bees are evenly distributed inside the hive, the scale will stabilize at half the hive's weight.

Even better would be real-time monitoring of hive weight, which would allow not only the measuring of honey production but also tracking of the aggregate departure and arrival of the thousands of bees that forage from a hive every day. You can disassemble a $20 bathroom scale to yield four perfectly good strain gauges, but strain-gauge output tends to drift under constant load. Rearick thinks that some clever beehacking could get around the drift problem by focusing (with somewhat lower accuracy) on daily ups and downs, but the idea awaits more spare time and warm-weather testing.

Ultimately, a hacked hive would be able to report the entrance and exit of individual bees, and perhaps external sensors could track where they gather their pollen and what troubles they encounter along the way. Meanwhile, internal sensors would report temperature and humidity; provide the data to diagnose mite, fungus, and other infestations; and keep tabs on honey production—all using scavenged parts supervised by a few cheap microcontrollers. And even during the winter, when a hive is dormant, a microphone could monitor the sound levels of worker bees flexing their wing muscles to generate heat to warm the rest of the colony. (This is another area where beehackers will have to gather much more data before offering analyses

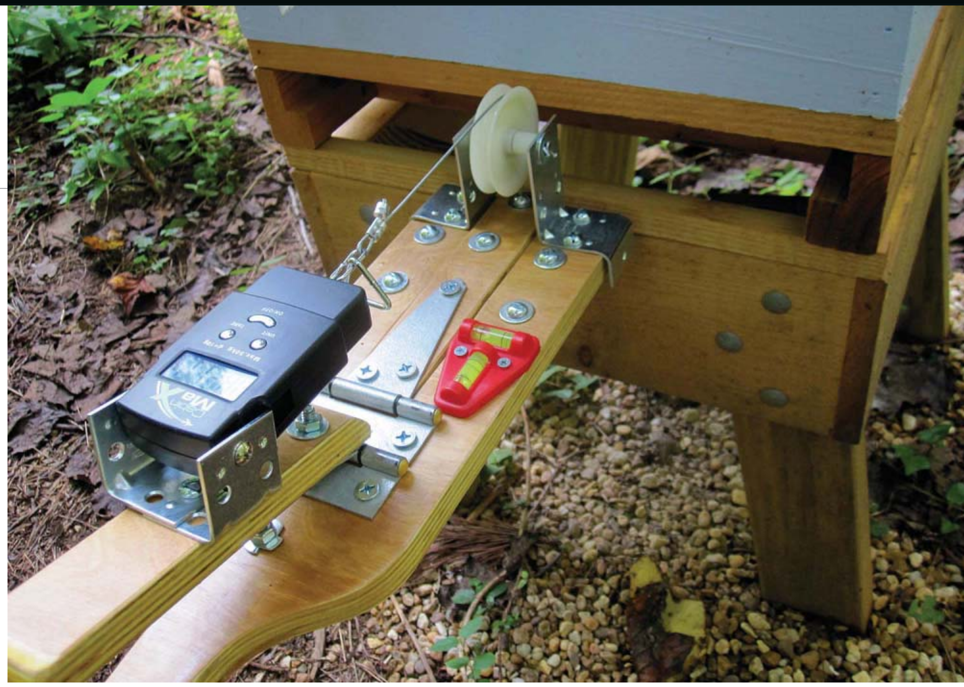

PRY HARDER: Beehacker Tom Rearick [opposite] uses a luggage scale [above] to pull a center tongue upward to weigh a bee hive. Changes in weight help assess the relative health of the hive. *PHOTOS: TOM REARICK*

deeper than "buzzing, good; no buzzing, bad.")

Rearick began his EE career at an aerospace company, where he worked on missile targeting. Next he moved to another aerospace company to run a computer vision group, and then he founded a small company to convert word-processing documents to hypertext. After IBM bought him out, he got an MBA and started another company that built a user-support chatbot for corporate Web sites. He sold that company and started yet another, this time making video management systems for (thus far) police interrogation rooms and car dealerships.

Rearick, who lives in Roswell, Ga., drifted into beehacking almost by chance: "Being a couple of foodies, my wife and I were considering getting chickens—for fresh eggs—or bees for honey. I decided that bees would be a lot less trouble." But as he worked his way into his new hobby, Rearick soon found out that beekeeping was not nearly as well documented a field as he had imagined, so his engineering instincts kicked in. He compares the current state of beekeeping to winemaking in the late 19th century, when phylloxera mites were devastating European vineyards, with no solution in sight. Since then, he points out, enology has been put on a firm scientific basis—those vineyards now use mite-resistant root stock, for instance.

And even as he faces the practical problems of hacking beehives, Rearick has found himself captivated by the puzzle of bee intelligence. Individual bees have just rudimentary brains and live for only about six weeks, yet the superorganism that is a swarm can make rapid decisions about rarely faced problems, such as the location of a new hive. For an engineer whose previous work involved repeated flirtations with artificial intelligence, this conundrum is more than he could have hoped for from a hobby initially intended to provide a few dozen annual jars of sweet goop. —Paul Wallich

http://www.beehacker.com

SO YOU WANT TO KEEP BEES

FIRST, ARE BEES permitted where you live? *Bee Culture* magazine (see links below) lists 95 U.S. cities that don't allow beekeeping.

Assuming you're not in a no-bee zone, you'll need:

- ☐ Hive boxes;
- ☐ A site near enough to flowering plants and trees so that your bees can gather nectar, but far enough from neighbors who might get stung, flee in terror, or complain to local authorities;
- ☐ Bees (there are, of course, different breeds to choose from);
- ☐ Protective gear (at a minimum, gloves, helmet, and veil); and
- ☐ Feeding systems for times when nectar just isn't available.

You also need a tolerance for insects. With as many as tens of thousands of bees in a healthy hive, some of them are going to get into your clothes, car, house, and barbecue. But if you can keep your hives alive (even professional beekeepers report losing about one in six hives in a normal year), you will end up with both a surfeit of honey and the warm glow of contributing a vital service to your local flora.

Beekeeping sites are rife with advice for the tyro, but as Rearick found, there are nearly as many opinions as there are beekeepers. And some of the vital bits, he notes, aren't really written down at all but absorbed by apprenticeship. So find a beekeeper and start asking questions. Or combine the strategies—*Bee Culture*'s site also has a "Find a Beekeeper Near You" link. *—Paul Wallich*

BEE CULTURE
http://www.beeculture.com/index.cfm

BEESOURCE
http://www.beesource.com

THE BEEKEEPING RESOURCE GUIDE
http://www.gardenbuildingsdirect.co.uk/Article/the-beekeeping-resource-guide

HOMESTEAD.ORG
http://www.homestead.org/KimFlottum/BeginningBeekeeping/Bees.htm

SCIENTIFICBEEKEEPING.COM
http://www.scientificbeekeeping.com//index.php

tools & toys

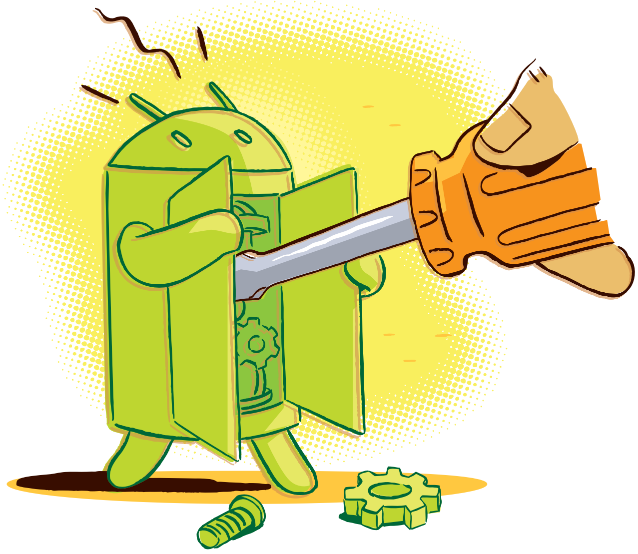

OPEN ANDROID—FOR BETTER AND FOR WORSE

Customization is key to Android's popularity—and to some inconsistent user experiences

BESIDES ITS cute little green robot mascot, what is it about Android that makes it so popular among device manufacturers and developers alike?

In a word, openness.

For developers, Android's license lets them get right into the heart of the platform's code, without signing their creative lives away. Android is based on open source software, allowing both collaboration (among disparate developers) and customization (of one's own offering). Neither is easily done on iOS, Apple's operating system for its iPads, iPhones, and iPods. Hardware manufacturers appreciate the open nature of Android, too. Android falls under the stewardship of the Open Handset Alliance, a consortium of device makers, mobile operators, and semiconductor, software, and other suppliers.

This openness, though, leads to some problems for both customers and developers. Because designers are allowed to modify Android to suit their needs—which they often do, in the hopes of differentiating their devices from those of other Android makers—different Android tablets can do the same thing in different ways. For example, configuring network settings is a rather different procedure on Samsung's and Archos's tablets, even if you're setting up the same network. Worse, the programming environments that manufacturers come up with for developers can also vary. HTC Corp.'s menus are different from those of their Motorola counterparts, for example, and the ways windows are managed can differ as well. That makes it harder to move from one tablet to another, which is presumably one of the selling points of a single OS. Imagine, for example, the Windows File menu being different on a Dell computer and a Lenovo one.

Yet so far, this fragmentation doesn't seem to be hurting the platform. Android now outsells Apple in smartphones and will do so in tablets, and its primary commercial sponsor, Google, often touts the more than 50 000 apps available for Android. (To be sure, nearly all were designed with smartphones in mind, but that will soon change.) Unlike Apple, which uses its App Store to control which apps are permitted on iOS devices, Android's alliance lets developers create and deploy whatever applications they like, as fast as they'd like. Of course, it's not all a bed of roses: That same free environment has led to applications that are sketchy at best and dangerous at worst. Luckily, the Android community is self-correcting, and such apps don't last long in the light of day.

Unlike iOS, which runs only on Apple's proprietary hardware, Android gives tablet makers hardware flexibility—different processors, for example—and interface options, such as how to sync with a computer.

Of course, more choices mean more caveats for consumers to emptor. So here are some things to look for when considering an Android tablet.

The version. Android's latest release is Android 3.0—its dessert-oriented code name is Honeycomb—and it will be in all the new tablets from the Motorola Xoom onward. Look for Ice Cream Sandwich Android devices in the latter half of 2011. Android 3.0 is the first tablet-specific version, which is why we're only now seeing applications that take advantage of a tablet's size and interface.

Flash or no Flash. Unlike iOS devices, Android devices will run Adobe Flash. But it's a mixed blessing, as running Flash movies depletes power rapidly. Apple is relying on HTML5 as its video standard, and you may want to do so as well.

Application store. Google, the commercial sponsor of Android and its strongest developer, has put together a very good app store, known as Android Market. Hardware vendors usually use some form of Market to enable users to download and buy Android apps. Do some digging and see how many apps are really available for that device. Check especially how many tablet-friendly apps are available—the overall number of Android apps specifically made for tablets is very small. Fewer than 50 were available at press time.

—BRIAN PROFFITT

CARL WIENS

NAME	Apple iPad	Samsung Galaxy Tab 10.1	Motorola Xoom	BlackBerry PlayBook	HP TouchPad	Apple iPad 2
PRICE	US $399/$599/ $729*	$499/$599	$600 (Wi-Fi only)/$799 (3G)	$499/$599/ $699	Info not available	$499/$599/ $699*
OPERATING SYSTEM	iOS 4.3	Android 3.0	Android 3.0	QNX	webOS	iOS 4.3
DISPLAY SIZE (diagonal, inches)	9.7	10.1	10.1	7.0	9.7	9.7
STORAGE (gigabytes)	16/32/64	16/32	32	16/32/64	16/32	16/32/64
NETWORK	Wi-Fi, 3G	Wi-Fi, 3G	Wi-Fi, 3G/4G	Wi-Fi	Wi-Fi, 3G/4G	Wi-Fi, 3G/4G
BATTERY LIFE (hours)	9–10	7	10 (over Wi-Fi)	Info not available	8–10	10 (for video)
WEIGHT (kilograms)	0.68	0.60	0.73	0.43	0.74	0.60

** Add $130 for 3G models.*

Tablet Table

BECAUSE OF Apple's and Samsung's early success with the iPad and Galaxy Tab, respectively, many new tablets are staying fairly close to these pioneering products' price, capacity, and capability. Based on all the leaked information as of early April, expect to see these similarities between devices:

- An average price of around US $613
- A screen size of about 9.7 inches, diagonal (25 centimeters)
- 32 gigabytes of storage capacity
- At least Wi-Fi and optional or included cellular connectivity
- About 10 hours of battery life (for video playback)
- An average weight of 0.6 kilograms
- At least one camera

The most dramatic differences are in the operating systems. While Apple devices will remain forever true to iOS, and all the other existing tablets use a version of Android, HP's upcoming TouchPad will support webOS (which, like Android, is derived from Linux via Palm smartphones), while the Research in Motion (BlackBerry) PlayBook will sport QNX, a 20-year-old Unix operating system newly transported to mobile devices.

The different operating systems will manifest themselves in the hardware, the look and feel of the interface, and the application ecosystem available for the devices. Both iOS and Android already have well-populated app stores, in part benefitting from their earlier presence on mobile phones. QNX and webOS also have a history from BlackBerry and Palm devices, respectively, but it will take a while to see how strong developer interest is in rewriting their smartphone apps for tablets or coming up with new ones. At press time, both the PlayBook and TouchPad are still in rumor mode, and there's always the chance RIM or HP will push out something truly unexpected for the devices' expected releases this spring.

Then there's the iPad 2, released back in March. Apple held steady on prices, but the memory capacity and number of cores doubled, while the number of cameras went from zero to two and cellular networking went from three to four—3G to 4G, that is.

—BRIAN PROFFITT

careers

THE AGING NUCLEAR WORKFORCE

A third of all workers in U.S. plants will be eligible to retire in the next five years

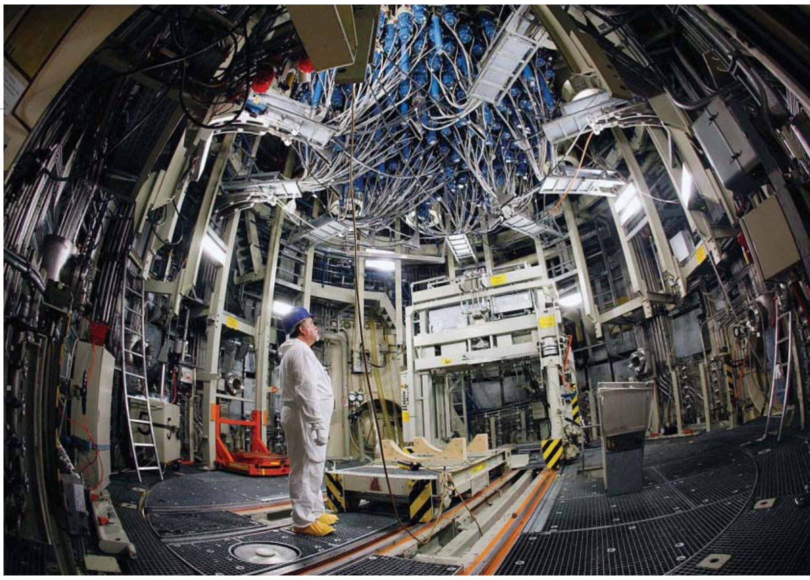

AGING PLANTS: Nuclear plants require constant maintenance and upkeep. This 24-year-old reactor in Kruemmel, about 30 kilometers southeast of Hamburg, Germany, was shut down for more than two months after a simple fire in the plant's transformer station in 2007.
PHOTO: KRAFFT ANGERER/GETTY IMAGES

THE PARTIAL core meltdowns at the Fukushima nuclear power plants in Japan came just when nuclear engineering had been on the rebound. The U.S. Nuclear Regulatory Commission has 19 license applications under active review for 26 new nuclear power plants. The Nuclear Energy Institute, a trade group in Washington, D.C., had expected four new plants to come on line by 2019, and construction is rampant all around the world [see "Will Nuclear Energy Charge Ahead?" The Data, in this issue].

But there's a problem—a problem that's also an opportunity for engineers. A third of all workers at the 104 currently operating U.S. plants could retire in the next five years, says Elizabeth McAndrew-Benavides, manager of industry infrastructure at the Nuclear Energy Institute.

Many of the tens of thousands of jobs created will be for engineers—not just nuclear but also civil, mechanical, and electrical. In fact, only about 10 percent of engineers at operating power plants are nuclear engineers by training. But nuclear engineers will also be needed in the United States at the Department of Energy's research labs, the Nuclear Regulatory Commission, defense agencies, and in the medical technology field.

Fortunately, students are "looking at nuclear engineering with a fresh set of eyes," McAndrew-Benavides says. "There's now a lot of energy and excitement surrounding nuclear careers, but more important, there is an understanding of the safety nuclear engineers provide to the world." Close to 400 bachelor's degrees in nuclear engineering were awarded in the United States in 2009, roughly twice as many as at the start of the decade, according to the Oak Ridge Institute for Science and Education. And undergrad enrollment in 2009 was the highest it has been since the mid-1980s.

The Department of Energy has awarded more than US $80 million to universities for education, new equipment, and upgrades to research reactors since 2008. And nuclear power companies are partnering with universities and community colleges, giving millions of scholarship dollars, providing internships, and hiring graduates.

University programs are adapting. Richard Lester, head of MIT's nuclear science and engineering department, says that MIT, based on feedback from industry leaders, is expanding its core curriculum to include nuclear reactor simulation and modeling and materials science in extreme environments. Students are also learning more about risk assessment, licensing and regulation issues, and newer reactor technologies, says Per Peterson, chair of the University of California, Berkeley's nuclear engineering department.

Before the 11 March earthquake, nuclear engineering's upswing had been apparent in Europe as well. Sweden and Italy had ended their nuclear power bans and planned to build new reactors. Finland, Spain, and the United Kingdom were also ready to expand their nuclear energy programs. Following the Japanese nuclear emergency, however, Italy said it would be toning down its expansion plans. Presumably other European countries will be reviewing their plans, too.

East Asia is the leader in actual construction, with China alone accounting for 27 of the 65 plants under construction worldwide. India and South Korea also have ambitious expansion plans. Meanwhile, countries such as Indonesia, Jordan, and the United Arab Emirates (UAE) plan to build their first nuclear power plants.

U.S. nuclear engineering expertise will play a role in these burgeoning national nuclear programs, according to McAndrew-Benavides. Many nuclear engineering schools in the United States have research ties in place with those in other countries: The UAE is consulting with U.S. educators and industry executives as it establishes nuclear engineering education programs, for example. Westinghouse Electric Co., which has a contract to design four nuclear power plants in China, has representatives there to help during construction as well as training. —PRACHI PATEL

BY ROBERT W. LUCKY

MICK WIGGINS

Are Social Technologies Really Invented?

IN A memorable scene from the movie *The Social Network*, Mark Zuckerberg is deposed by a lawyer for the Winklevoss twins, who are suing him, claiming that Zuckerberg had stolen their idea for Facebook. Zuckerberg's attention wanders. When he is accused of being inattentive, he comes out of his fog and points his finger across the table at the Winklevoss twins and their lawyer, saying, "You have part of my attention. You have the minimum amount. The rest of my attention is back at the offices of Facebook, where my colleagues and I are doing things that no one in this room, including and especially your clients, is intellectually or creatively capable of doing."

My eyes were on the screen, but my mind was reflecting on the question of who should get the credit. Even if I had been physically present for all the events portrayed, I still wouldn't know for sure to whom all those billions of dollars should be assigned.

Credit can be a complex and murky thing. Once it is decided, however, it becomes simple and robust. The 1870s race to invent the telephone between Elisha Gray and Alexander Graham Bell, among others, was as murky as it gets, but the outcome was as clear as day—Bell won his patent. Had it been the other way around, none of us would have heard of Alexander Graham Bell.

Facebook is an example of something that happens periodically in cyberspace—an invention built on a technology platform but not necessarily involving technological innovation. Its success relies on a "tipping point" phenomenon, in which the winner takes all. In the regular world, we benefit from competition and even have antitrust laws to ensure it, but in the networking world, we all benefit from being in the same place. When all your friends are on Facebook, why go anywhere else?

In my experience, while great ideas are rare, good ideas are rather plentiful, and what really matters is what you do with them. Zuckerberg took a good idea and made it better. It wasn't the first social network; sites like Myspace and Friendster already existed. But Zuckerberg captured the world with good features, good business strategy, and—one of the most important ingredients—good timing. In such cases, moreover, we shouldn't forget the roles of happenstance and luck.

However, if Zuckerberg hadn't created Facebook, some similar site would be offering the same features. Maybe we'd all be on Myspace, which might have become even better than what we have today. I often reflect on the transience of even great ideas and inventions. If Beethoven hadn't written his symphonies, they wouldn't exist. Great art, music, and literature are like that, but science and technology uncover what nature has intrinsically made possible. It's as if our inventions have been lying under rocks and the first person to turn over the rock gets the credit. The second person who comes along...well, it's too late. Does anyone think that if Shockley, Bardeen, and Brattain hadn't invented the transistor in 1947 it wouldn't exist today?

Perhaps the measure of any invention or discovery is the amount of time that passes before someone else ends up doing much the same thing. Paul Baran, who died in March, had as much to do with the founding of the Internet as anyone. Yet Donald Davies in the United Kingdom and Len Kleinrock at the University of California, Los Angeles, were working independently on Baran's key idea, packet switching. Even Einstein's general theory of relativity was a law of nature waiting to be discovered, and we undoubtedly would know it today. Perhaps the state of knowledge now would be no different, in spite of the fact that this was surely one of the most transformative ideas in the history of science.

It's nice to be the first person to turn over the rock and see something novel underneath, but unless you do something with that find, maybe you don't deserve the credit. In the case of Facebook, it appears that Zuckerberg does deserve the credit—or at least the money. ❑

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

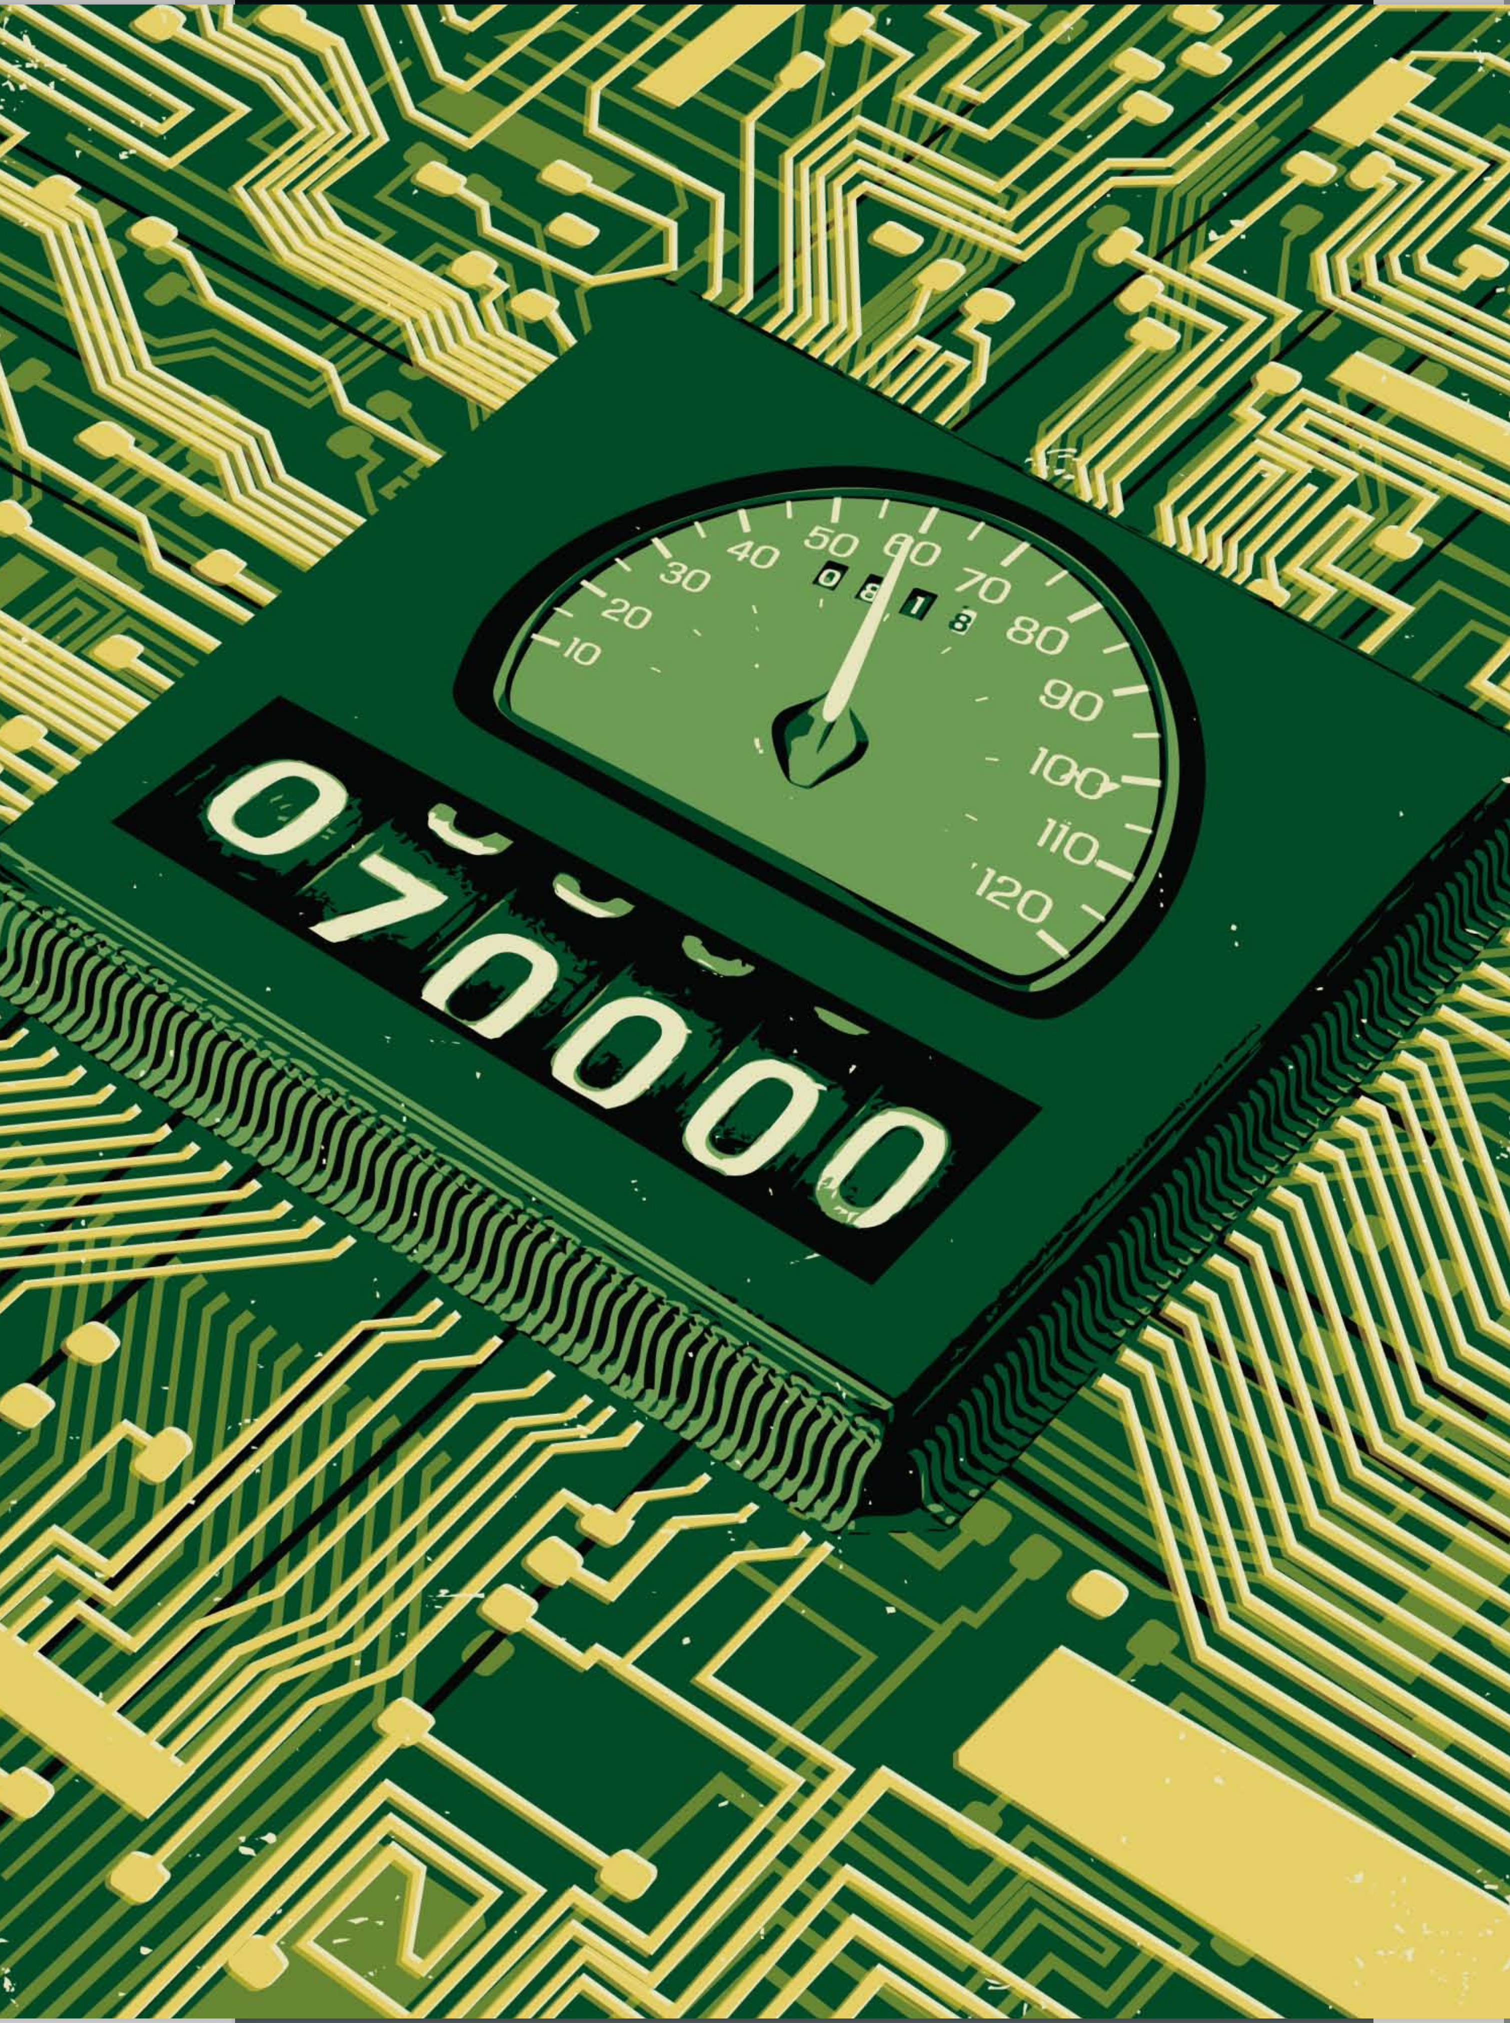
10
20
30
40
50
60
70
80
90
100
110
120
070000

AN ODOMETER FOR CPUS

Measuring the degradation of microprocessors is tricky. Doing it better would unleash more processing power

BY JOHN KEANE & CHRIS H. KIM

YOU KNOW WHEN it's time to get a new car. Your odometer is far into six digits, perhaps the engine is burning lots of oil, or the transmission is growling. Fixing all that might well cost quite a bit more than your ancient vehicle is worth.

But what about your microprocessor? Unlike automobiles, microprocessors don't have convenient little gauges that reflect how much wear and tear they've endured. And wear they do—though you'll probably never notice it. The degradation of their transistors over time leads slowly but surely to decreased switching speeds, and it can even result in outright circuit failures.

DANIEL HERTZBERG

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

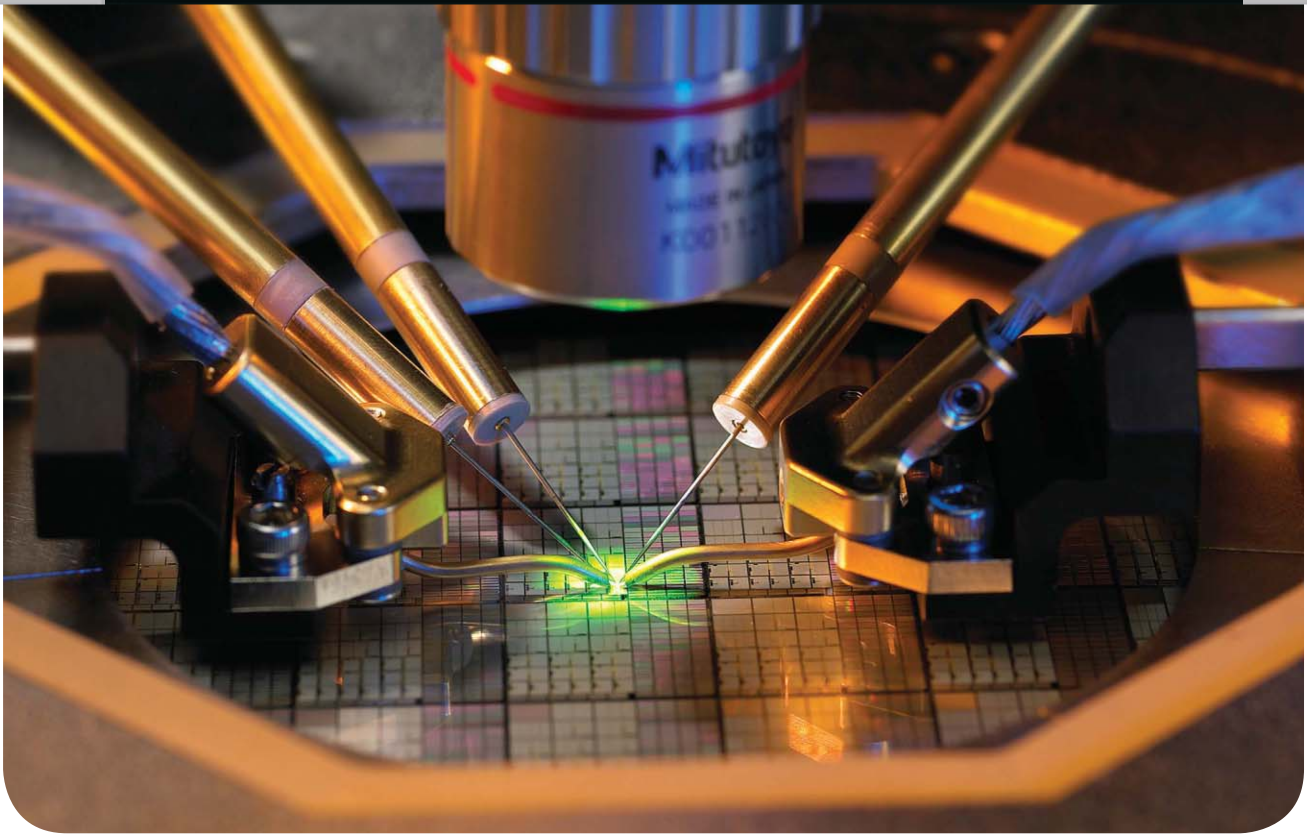

SQUINT HARD: The traditional method of monitoring transistor-aging effects in chips requires the careful placement of tiny probes, which are manipulated while being viewed under a microscope.
PHOTO: CASCADE MICROTECH

You generally don't perceive this deterioration because semiconductor companies always play it very safe—they set the clock-speed rating of their microprocessors so conservatively that almost every one of their products will continue to operate flawlessly throughout its intended lifetime. That strategy works. But it's kind of like never taking your Ferrari out of the slow lane because you're concerned that its engine might throw a rod 10 years down the road.

Several different phenomena can degrade the transistors on chips. How those phenomena combine to diminish a chip's functioning depends on such factors as the circuit arrangement of the aging transistors as well as the voltages and temperatures they're exposed to. With all these variables, it's difficult to predict how the peak performance of a given microprocessor will decline over time.

We and other researchers are trying to improve that situation. One critical aspect of the work we did at the University of Minnesota was to develop better ways to study the different physical mechanisms of transistor aging. Today semiconductor engineers measure those aging effects primarily by examining transistors one at a time, using microscopic electrodes to probe a silicon wafer. The necessary equipment can cost tens of millions of dollars, and probing transistors individually is arduous when you're trying to gather many observations. Sometimes you can't do those measurements well, no matter how much time you spend.

We need better techniques. And we need them sooner rather than later. Microprocessors now contain billions of transistors, sometimes operating at clock speeds in excess of 3 billion cycles per second. The blazingly fast clocks mean that the transistors are exposed to lots of heat, which accelerates their decline. Another worry is that there are precariously small voltage differences between supply levels and the threshold at which the transistors turn on. Also, various improvements in the way silicon logic is fabricated have introduced new concerns about degradation. And transistors scaled down to today's tiny dimensions experience more variation than ever in their operating conditions, which in turn leads to great differences from one transistor to another in how fast they wear out.

With better ways to measure transistor aging, chipmakers could let their microprocessors run faster—appreciably faster—than they do now. In the future it might even be possible to use these techniques to build circuits into microprocessors that continuously measure the subtle effects of aging and adjust clock frequency or operating voltages so that the transistors, old or new, could always run at peak speeds.

WHY SHOULD TRANSISTORS age at all? The kinds of transistors we're talking about here—the good old metal-oxide semiconductor field-effect transistors that are the basis of ordinary CMOS chips—function as electrical switches. A MOSFET has four terminals, called the body, the gate, the source, and the drain, although the source and body are often connected. The voltage that's applied to the gate determines whether current can flow between the

spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

source and drain. Although a thin layer of dielectric material electrically insulates the gate, the electric field applied across it alters the conductivity of the underlying semiconductor channel connecting the source and drain.

And that brings us to our first degenerative mechanism: Over time, charge carriers (electrons for negative, or *n*-channel, MOSFETs; holes for positive, or *p*-channel, MOSFETs) with a little more energy than the average will stray out of the conductive channel between the source and drain and get trapped in the insulating dielectric. This process, called hot-carrier injection, eventually builds up electric charge within the dielectric layer, increasing the voltage needed to turn the transistor on. As this threshold voltage increases, the transistor switches more and more slowly.

There's a second mechanism that can also trap charge in the dielectric, and it doesn't require any current to flow between the source and drain. Whenever you apply voltage to the gate, a phenomenon called bias temperature instability can cause a buildup of charge in the dielectric, along with other subtle problems. After that gate voltage is removed, though, some of this effect spontaneously disappears. This recovery occurs within a few tens of microseconds, making it difficult to observe during routine experiments, where you stress the transistor but measure the resulting effects only after the stress is removed.

Yet another aging mechanism comes into play when a voltage applied to the gate creates electrically active defects, known as traps, within the dielectric. If they become too numerous, these charge traps can join and form an outright short circuit between the gate and the current channel. This kind of failure is called oxide breakdown, or more verbosely, time-dependent dielectric breakdown. Unlike the other aging mechanisms, which cause a gradual decline in performance, the breakdown of the dielectric can lead to the catastrophic failure of the transistor, causing the circuit it's in to malfunction.

As if the aging of transistors wasn't enough to worry about, semiconductor engineers also have to grapple with the metal connections between transistors wearing out over time. The concern here is a phenomenon called electromigration, which damages the copper or aluminum connections that tie transistors together or link them to the outside world.

Electromigration occurs when a surge of current knocks metal atoms loose and causes them to drift along with the flow of electrons. This depletes the metal of some of its atoms upstream, while causing a buildup of metal downstream. The upstream thinning of the metal increases the resistance of the connection, sometimes to the point that it can become an open circuit. While downstream deposition isn't similarly catastrophic, it can cause the metal to bulge out of its designated track.

SILICON'S DECLINE

Three distinct physical phenomena cause MOSFET transistors to age. Usually, the aging process merely slows the speed at which transistors turn on and off. But it can also break them outright.

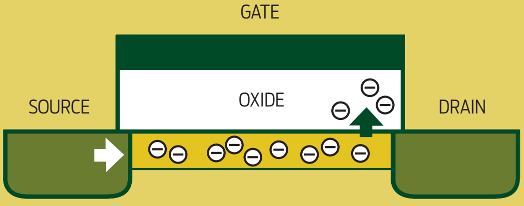

HOT-CARRIER INJECTION

Some of the more energetic electrical charges flowing between a MOSFET's source and drain terminals can end up in the insulating gate oxide, which increases the voltage needed to turn the transistor on. This process, called hot-carrier injection, slows the speed at which the transistor can operate.

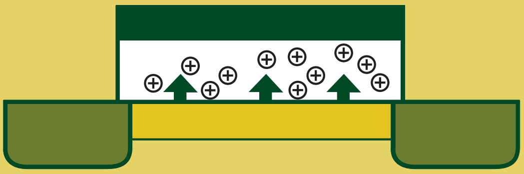

BIAS TEMPERATURE INSTABILITY

Even when no current is flowing from source to drain, a voltage applied to the gate can cause charges to migrate into the insulating gate oxide. This phenomenon, called bias temperature instability, is partly reversible: Many of those errant charges quickly leave the oxide after the gate voltage is removed, which makes this effect difficult to measure.

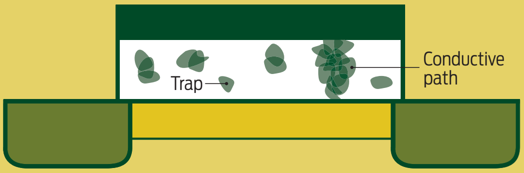

OXIDE BREAKDOWN

The application of voltage across the gate can also cause electrically active defects within the oxide layer, which trap charges. If enough of those charges accumulate in the gate oxide, they can create a short circuit, causing the catastrophic failure of the transistor.

DANIEL HERTZBERG

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

BEAT GENERATION

Measuring slight differences in the frequencies of two signals can be a snap if you beat them together

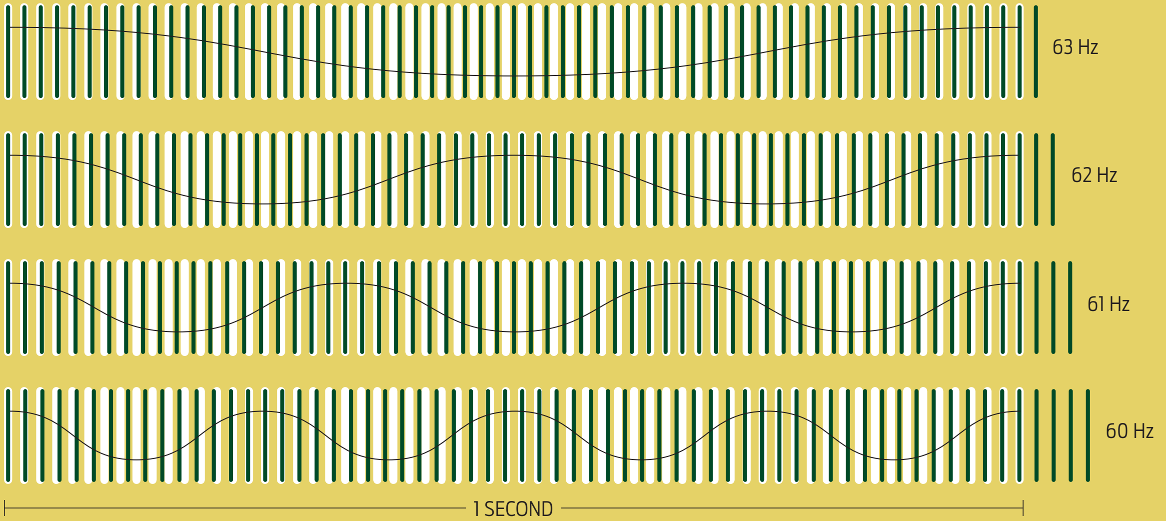

THE AUTHORS' silicon odometer senses small changes in the switching frequency of two digital oscillator circuits when one is put under stress. It does that by measuring their beat frequency. The beat phenomenon can readily be seen in this schematic illustration of low-frequency signals. When a 63-hertz signal [green bars, top] is superimposed on a 64-Hz signal [white bars] they first add together, then fall out of sync, then add together again—going through one such beat cycle during a 1-second period. A 62-Hz signal and a 64-Hz signal give two beat cycles, and so forth. Measurements of beat frequency can thus show subtle differences in frequency of much faster oscillations.

DESPITE THE MANY efforts of process engineers to create long-lasting transistors and metallic connections, a certain amount of wear is unavoidable. So it must be reckoned with. Probing chips to study how their transistors degrade provides only limited information on how bad the problems are inside a real microprocessor. A better experimental approach, just now taking hold in industry, is to fabricate special chips for the sole purpose of testing how the transistors on them function over time. These chips contain what are known as ring oscillators, consisting of many inverter circuits chained together in a loop. Each inverter outputs the opposite of whatever signal is applied to its input. So when an odd number of them are wired into a ring, the circuit oscillates, at a frequency that depends on how fast the constituent transistors can switch states.

To measure the slowdown that arises as transistors age, engineers subject chips containing such ring oscillators to extreme conditions, raising the supply voltage or operating temperature so that they wear out in a matter of hours or days. Engineers can measure how transistor-switching times increase, or how the average time to failure decreases, during several different accelerated-aging experiments conducted using a variety of stress levels. That allows them to extrapolate their results to the real world, where transistors age much more slowly than they do during these tests.

Testing ring oscillators provides more insight into how circuits age than does probing transistors individually, but it still has flaws. One is that it takes a relatively long time to make a measurement—a large fraction of a second. That's because you're looking for very subtle changes in the frequency of the oscillations, so you need to count a lot of cycles to measure the shift. Taking a second might seem quick enough, but remember, most of one particular aging effect (bias temperature instability) lasts for only a few tens of *microseconds* after the stress is removed.

Most often, you would stress the circuit by raising its supply voltage; elevating the temperature works less well. But you then need to return the voltage to its normal level to judge how much the transistors have changed. So most of the effects of bias temperature instability will disappear before you'll have a chance to observe them. We've worked hard to remedy this problem and also to improve the way transistors are examined in general.

To measure the shifting frequency of ring oscillators much more quickly, we developed what we call a silicon odometer. It is based on a pair of ring oscillators and works by measuring their beat frequency. If you're at all

spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

musical, you're probably familiar with the beat phenomenon: When two very similar notes are played simultaneously, you hear just one note whose amplitude changes rhythmically—it beats. The number of beats per second is equal to the difference in frequency between the two original notes.

In our silicon odometer, we measure the difference between the frequency of an unstressed ring oscillator and the one we have stressed by increasing its operating voltage. To measure actual beats, we'd have to sum the two outputs and run the results through an analog-to-digital converter. Rather than doing that, we do something similar that is easier to accomplish: We sample the output of one oscillator at intervals set by the output of the other. That gives us a digital signal that oscillates at the beat frequency, which is easy to measure with simple circuitry. This approach allows us to measure changes in transistor switching times as small as one part in 10 000 in less than a microsecond, which is short enough to capture even the most fleeting effects of aging.

In addition to the short measurement times, our silicon odometer has another nice attribute: It's essentially immune to the gradual changes in operating voltage or temperature that often take place during long stress experiments. Any stray variations in either will shift the frequencies of the two oscillators by roughly the same amount. So their difference, which is all we measure, will continue to reflect just the effects of stress. We've gathered a lot of valuable aging data with this odometer. And with a somewhat different design, which we've dubbed the all-in-one odometer, we were able to separate the simultaneous effects of hot-carrier injection from those of bias temperature instability.

To get a better idea of how aging varies from place to place on a chip, we developed a third design, a statistical silicon odometer, which uses arrays of stressed ring oscillators. The circuitry we've devised allows us to connect any of those ring oscillators to a reference oscillator on the chip that runs at about the same frequency. We can thus measure subtle aging effects at many different locations on the silicon die.

Beat-frequency detection systems are not the only on-chip monitors we have implemented, though. We've put together other circuits to study the statistics of oxide breakdown, which is rarely extensive enough to disable a transistor. But rare events count when you're dealing with the billions of transistors in a modern microprocessor, which make for billions of opportunities for failure. To study that, we have designed several different arrays of transistors that are particularly prone to having their gate dielectrics wear out and finally break down. By stressing all the transistors in these arrays at the same time, we can easily take measurements of their time to breakdown. This information is useful to chip designers, whose usual recourse is to deal with this problem by adjusting fabrication procedures.

OUR ON-CHIP CIRCUIT-AGING sensors are providing valuable information about how real circuits age. Not long from now, we expect it will be possible to design similar circuits into microprocessors for use as real-time aging monitors, which could trigger adjustments in clock speeds or operating voltages to ensure that these chips work at peak levels, even as they grow older. That would eliminate the need for slowing the chips down from the start.

The semiconductor industry is just beginning to dabble with such ideas. Intel, for example, has engineered a technology it code-named Foxton, which measures the amount of propagation delay along certain critical signal paths within a microprocessor. Specialized circuitry on the microprocessor then uses that information to adjust supply voltage and clock frequency. Intel had planned to include Foxton technology in its dual-core Itanium 2 processor, released in 2006, but undisclosed complications delayed the introduction. We suspect that using something like our silicon odometer might help overcome the difficulties Intel faced.

In that vein, we are collaborating with groups at Intel and also at such companies as Advanced Micro Devices, Broadcom, Freescale Semiconductor, Texas Instruments, and (with help from the Defense Advanced Research Projects Agency) IBM, which are all striving to eliminate the compromises they routinely make to accommodate transistor aging. Silicon odometers, or something similar, might well appear in some of their new CMOS chips.

With such technology, microprocessors could be made to always perform at their peak levels, even though those levels would still decline slowly over time—as is the case with so many of our other machines. This will take some getting used to, of course, particularly if you ever attempt to sell your used computer. Maybe sometime in the not-too-distant future we'll be reading classified ads that say things like, "For sale: Five-year-old laptop, but with like-new processor. Used only to compute at church on Sundays." ❑

DANIEL HERTZBERG

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

BITS *of* HISTORY | *Artifacts from the first 2000 years of computing*

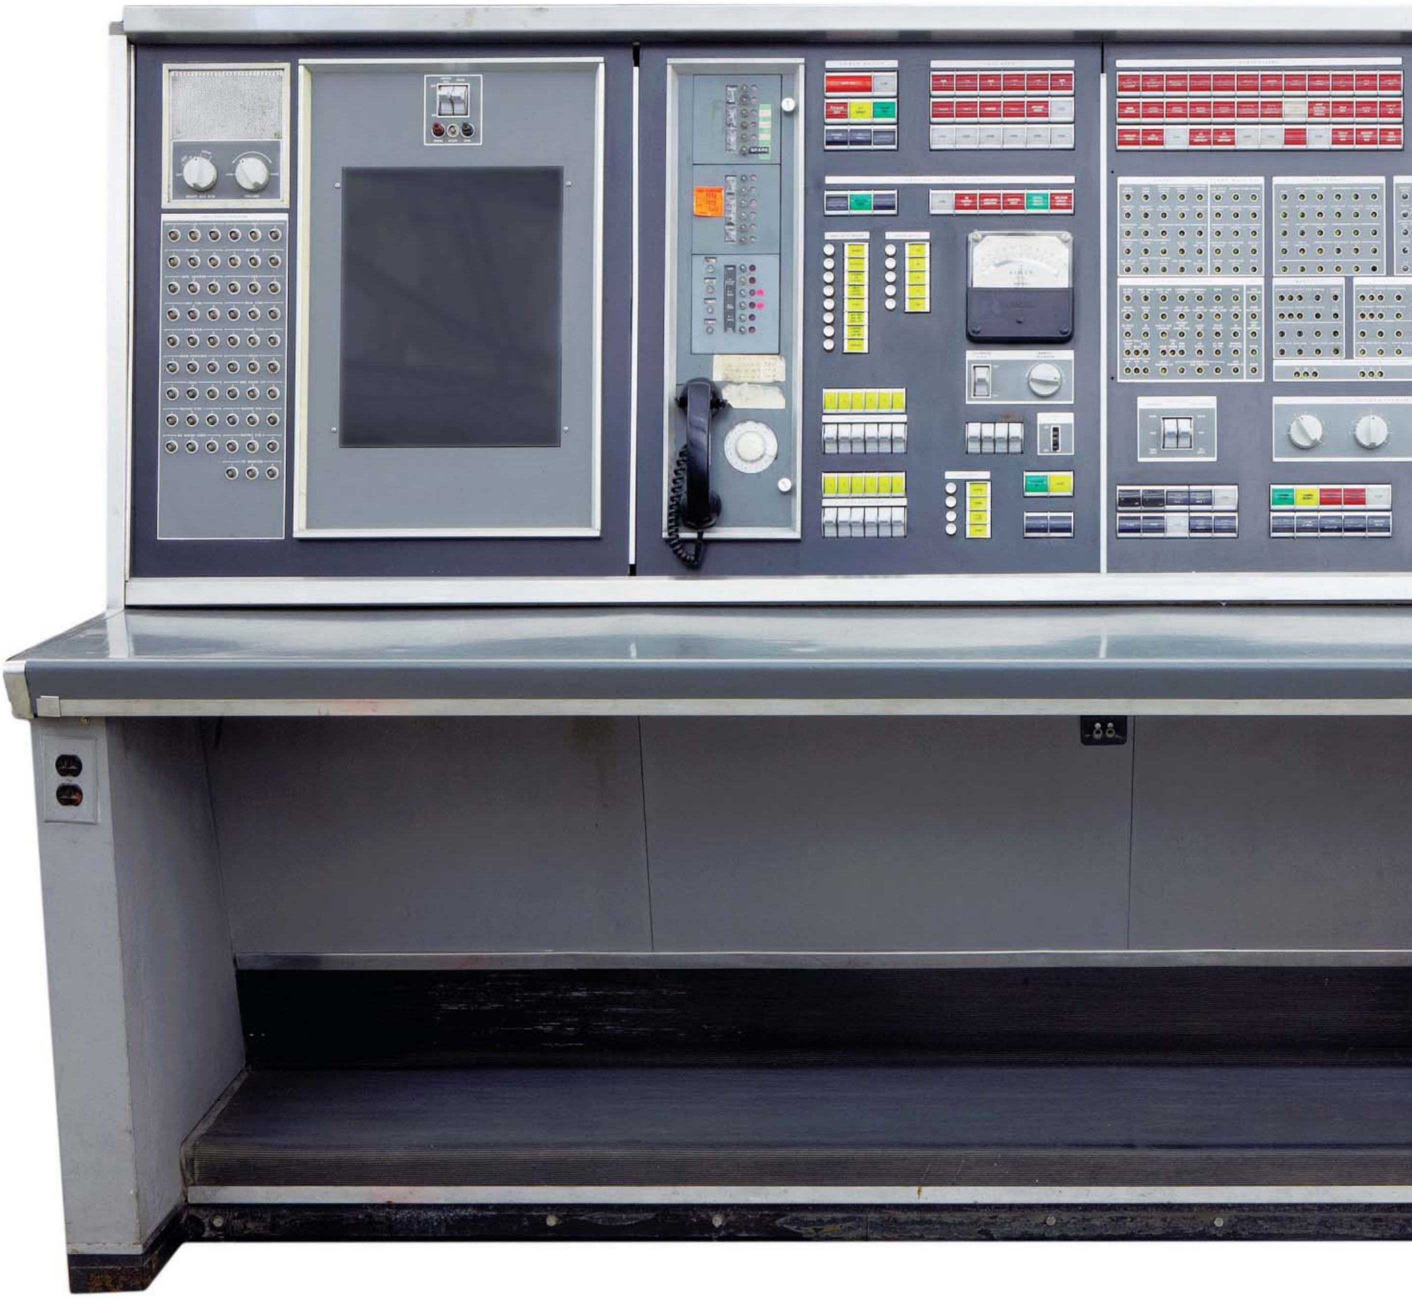

1950s | *SAGE*

COLD WAR COMPUTER The biggest pieces in the Computer History Museum's collection belong to the Semi-Automatic Ground Environment (SAGE) system, developed for the U.S. Air Force. SAGE, explains senior curator Dag Spicer, was really a network of 23 Costco-size warehouses located throughout the United States and Canada. The system stored flight information for all authorized

ieee spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

by JOSEPH CALAMIA | *Photography by* MARK RICHARDS

IN SOME WAYS, COLLECTING old computers isn't much different from collecting anything old: You have to take care of the stuff. "Is it decaying?" asks Dag Spicer, senior curator at the Computer History Museum, in Mountain View, Calif. He describes the remains of sound-dampening foam that once hushed the whir of cooling fans in 1960s and '70s mainframes. "It turns into a tarry mess—really gross, black sludge," he says. That's relatively easy to clean out, but some troubles require more technical expertise. Reading the information on a 1950s disk stack might be hard, says Spicer, a circuit designer turned historian, but harder still is making sense of it. "Do you recognize what these bits are?" he asks, explaining the need for both obsolete hardware and outdated operating systems.

Despite such challenges, Spicer has helped the Computer History Museum acquire more than 100 000 technological artifacts, building the largest collection of its kind. Of these the museum has selected 1200 for a new exhibit called *Revolution: The First 2000 Years of Computing.*

The curator estimates that 3000 people come each week to see *Revolution*, which opened in January 2011. The stories behind the artifacts attract all those visitors, says the museum's president, John Hollar, who particularly enjoys exhibits that highlight tales of engineering triumph. A favorite artifact is a piece of the Apollo Guidance Computer, which helped put men on the moon despite having only 36 kilobytes of memory. But Spicer admits that the computer relics draw visitors for simpler reasons too. Sans sludge, "they're very beautiful objects," he says. ❑

commercial and military flights and flagged any "unknowns" spotted by radar, in an attempt to protect the United States from Soviet bombers. "It probably cost more than the Manhattan Project, and yet very few people know about it," Spicer says. "At one point, something like three-quarters of all programmers in the country were working on SAGE, so it trained a whole generation."

Arm Waving

From the Norwegian for "snake," the Orm was an early attempt at a computer-controlled robotic arm. Created in 1969 by Stanford engineers Victor Scheinman and Larry Leifer, the arm could extend by inflating several of 28 air sacs sandwiched between seven metal disks. Despite its elegant design, accuracy was never its forte. Scheinman would go on to produce the Stanford Electric Arm, which was capable of building a Ford Model T water pump. Manufacturers also consulted him on many of the key industrial robot designs of the 1970s and '80s, Spicer says.

1940s | *Enigma*

CODE MAKER

One of Spicer's favorite artifacts is this World War II German enciphering machine, called the Enigma. After the operator arranged wires on the device's front and typed a letter on the attached keyboard, say a *T*, a light would blink above the corresponding letter, say an *E*, of the encrypted message. A recipient with another Enigma could then decipher the message by using an identical arrangement of the wires on his machine. Typing a letter, such as an *E*, in a seemingly nonsensical coded sentence would then reveal the sender's true intent, a *T*.

"It doesn't even count as a computer—it's basically just a flashlight circuit," Spicer says, "but this machine had an enormous significance for humanity." Some historians believe that by cracking Enigma's codes, the Allies may have shortened the war by as much as two years. "Gut-wrenching decisions," Spicer says, "came out of this very simple, tiny little box."

1954 | *TR-1*

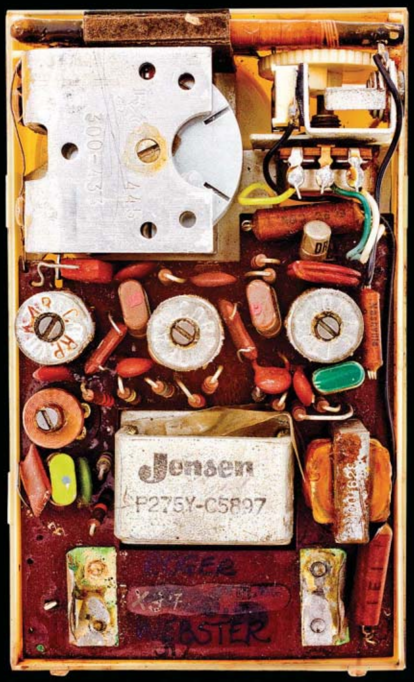

TRANSISTOR TRANSMISSIONS

In October 1954, this radio helped to bring the word *transistor* into popular parlance. Texas Instruments' Regency Model TR-1 radio was the first to forgo vacuum tubes, replacing them with germanium transistors.

Previously, manufacturers had made transistors by tediously assembling tiny semiconducting bars, but TI developed a method for mass-producing them in a furnace. The company sold around 100 000 of the pocket radios, but Spicer says they're now hard to come by. "They didn't cost very much," he says, "and they were kind of throwaway things." He points out the slight rusting inside the museum's radio, noting the device's popularity among beachgoers.

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

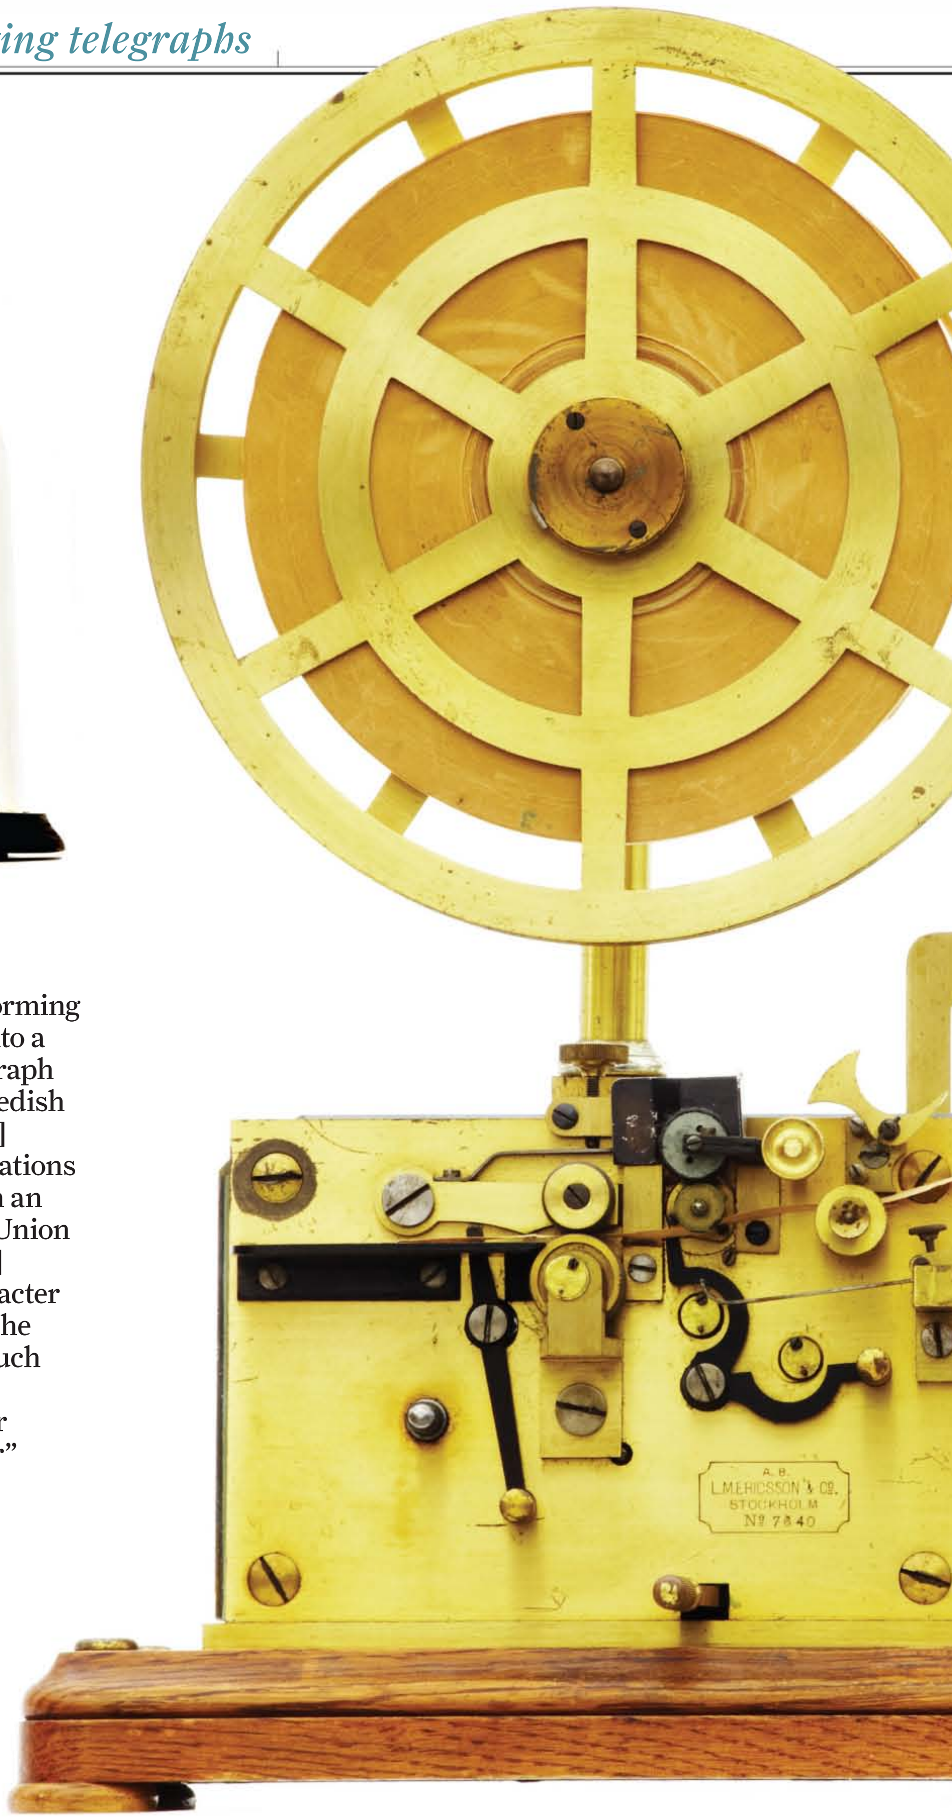

A Parade of Tickers

Many machines existed for transforming the telegraph's dots and dashes into a readable printout. An 1890s telegraph receiver manufactured by the Swedish firm L.M. Ericsson and Co. [right] converted Morse code into indentations on paper tape as it unwound from an ornate wheel. An 1870s Western Union Telegraph Co. stock ticker [above] printed selling prices, at one character per second, live from the floor of the stock exchange. The sound that such devices produced while printing inspired a descriptive moniker for both the machines and the "ticker" tape they unfurled.

ONLINE SLIDESHOW
For more computer relics, check out *http://spectrum.ieee.org/computerhistory0511*.

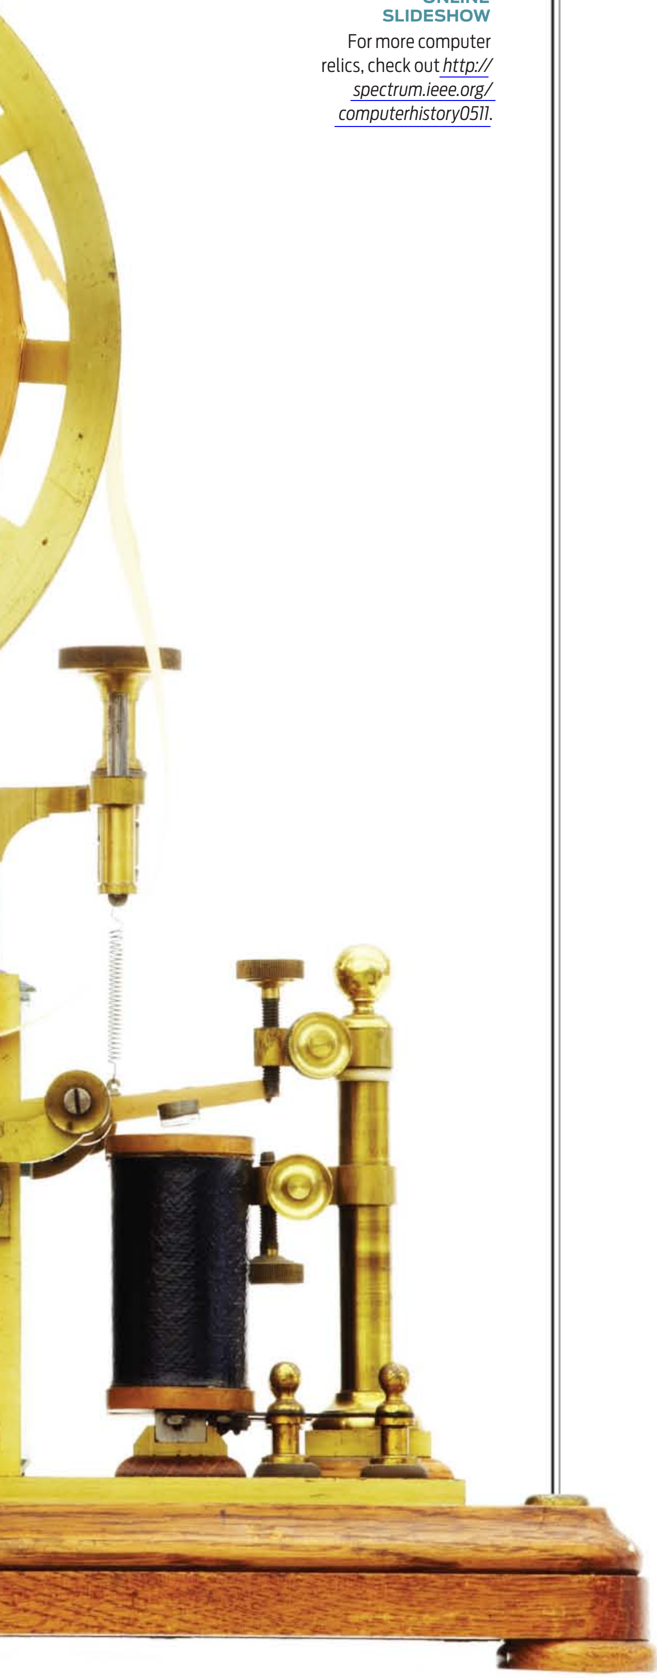

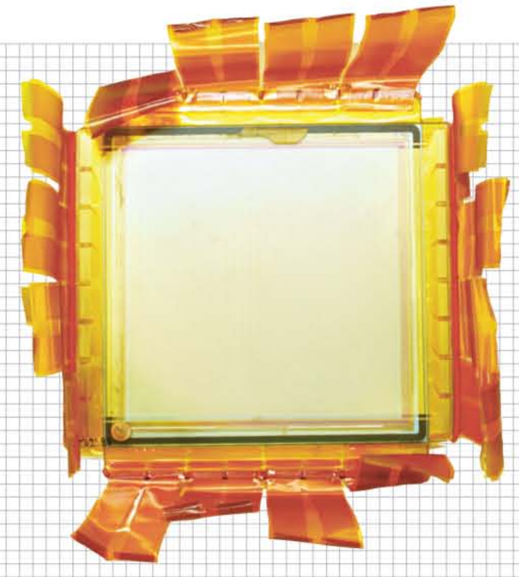

1970s *Plasma display panel*

A NEW FACE FOR EDUCATION

The orange glow from these 1970s plasma displays welcomed students who used the University of Illinois's networked teaching terminals. For more than a decade, the Programmed Logic for Automated Teaching Operations (PLATO) system provided users with basic instruction in a range of subjects—including genetics, the operation of nuclear power plants, and aircraft piloting—before microcomputers finally replaced the system. The displays' ability to store images without external memory made them especially attractive.

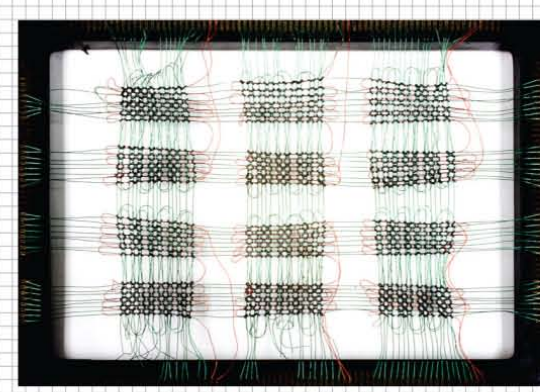

1965 *Magnetic core panel*

RING RECORDINGS

This 1965 IBM artifact is core memory, which recorded bits as the orientation of the magnetization in tiny ferromagnetic rings. The computer read and wrote bits using the wires attached to each ring in the grid. The example shown here contains 1536 magnetic memory cores. Even today, in the age of semiconductor memory, some people refer to a file that holds the working memory's contents after a computer crash as a "core dump."

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page
Qmags THE WORLD'S NEWSSTAND®

1953 *Johnniac Selectron tube*

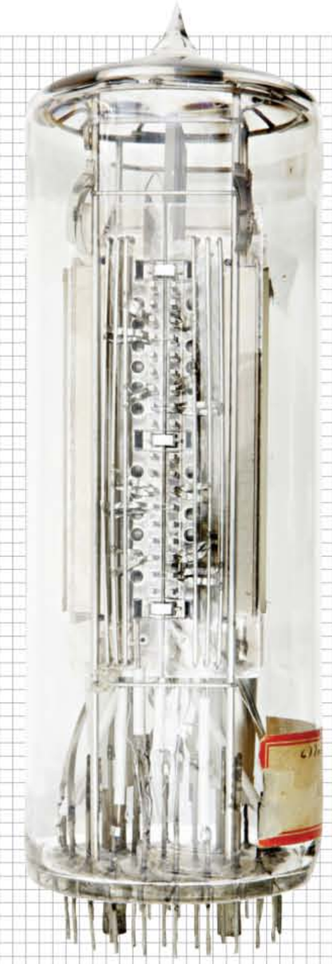

Remembering Vacuum Tubes

Isolated gates, called eyelets, stored charge in this Selectron vacuum-tube memory, which Jan Rajchman started developing in 1946 while at RCA Laboratories. The 1953 version, shown here, held 256 bits—a far cry from the device's original design goal of 4096 bits.

The only machine ever to use the expensive Selectron memory was the Johnniac, or John von Neumann Numerical Integrator and Automatic Computer, which the Rand Corp. built in 1953. Spicer says that the custom-made computer, which operated continuously for 13 years, only narrowly avoided the scrap heap before finding a home at the museum. "It was sitting in a parking lot...and someone who had worked on it recognized it and rescued it. It was pretty dramatic," Spicer says. "This is an irreplaceable, priceless computer—it was like finding a Fabergé egg in your recycle bin."

1975 *Altair*

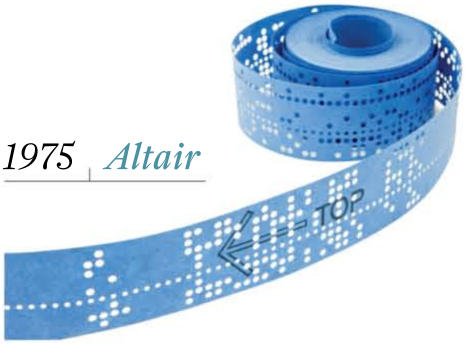

BEFORE WINDOWS

Inspired by a *Popular Electronics* cover story on the Micro Instrumentation and Telemetry Systems (MITS) Altair 8800—a do-it-yourself personal microcomputer—Honeywell programmer Paul Allen and Harvard sophomore Bill Gates set out to write software for the machine. After the pair endured several all-day programming sessions, Allen completed the BASIC software while flying to MITS headquarters in Albuquerque. They dubbed the partnership they had formed "Micro-Soft." The 1975 Altair BASIC Interpreter, on the source tape shown above, would prove a big seller into the early 1980s and a popular piracy target of hobbyists.

1956 *IBM RAMAC*

STORAGE STACK

Introduced in 1956, IBM's 305 Random Access Memory Accounting System, or RAMAC, included the world's first hard disk drive. Fifty 60-centimeter disks spun at 1200 rotations per minute and could hold 5 million characters of information—the equivalent of 62 500 punch cards. Reminiscent of a jukebox, the drive had a pneumatic "access mechanism" to select a particular disk for recording at up to 40 bits per centimeter. By 1958, IBM offered a second disk stack for additional storage in each system, complete with a second access arm.

The 305 RAMAC was one of IBM's last vacuum-tube computers. The company produced more than 1000 of the machines.

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®


© 2011 Marriott International, Inc.
INTRODUCING THE REIMAGINED LOBBY
WORK, CONNECT, OR RELAX IN AN AMAZINGLY FLEXIBLE SPACE.
MASTER THE LONG TRIP.®
RESIDENCEINN.COM
Residence Inn
Marriott

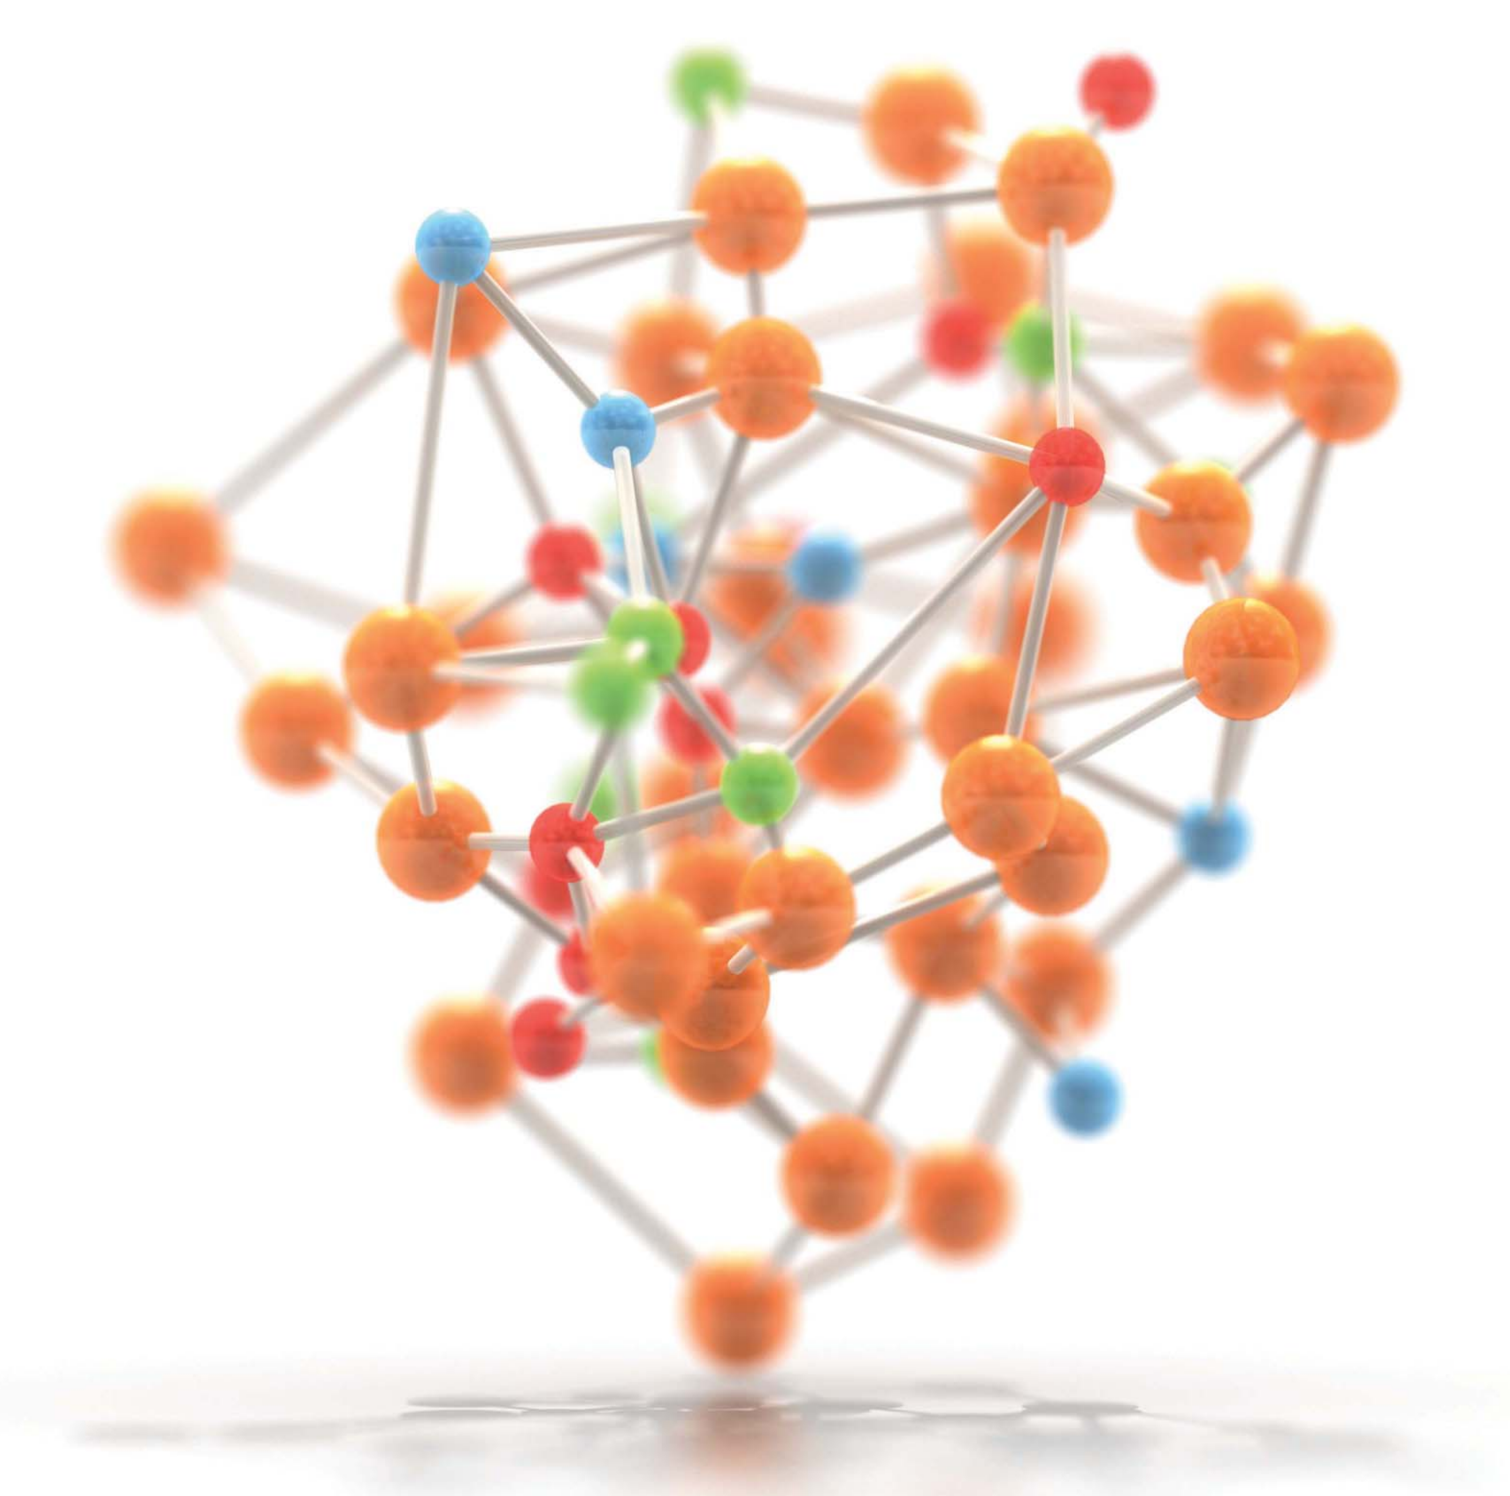

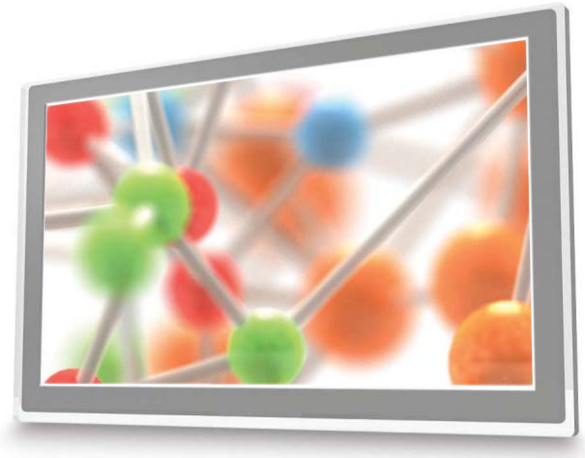

Thin, Fast, and Flexible

Amorphous oxide semiconductors promise to make flat-panel displays faster and sharper than today's silicon standby

BY JOHN F. WAGER & RANDY HOFFMAN

ALL ILLUSTRATIONS: BRYAN CHRISTIE DESIGN; MONITOR: ISTOCKPHOTO

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

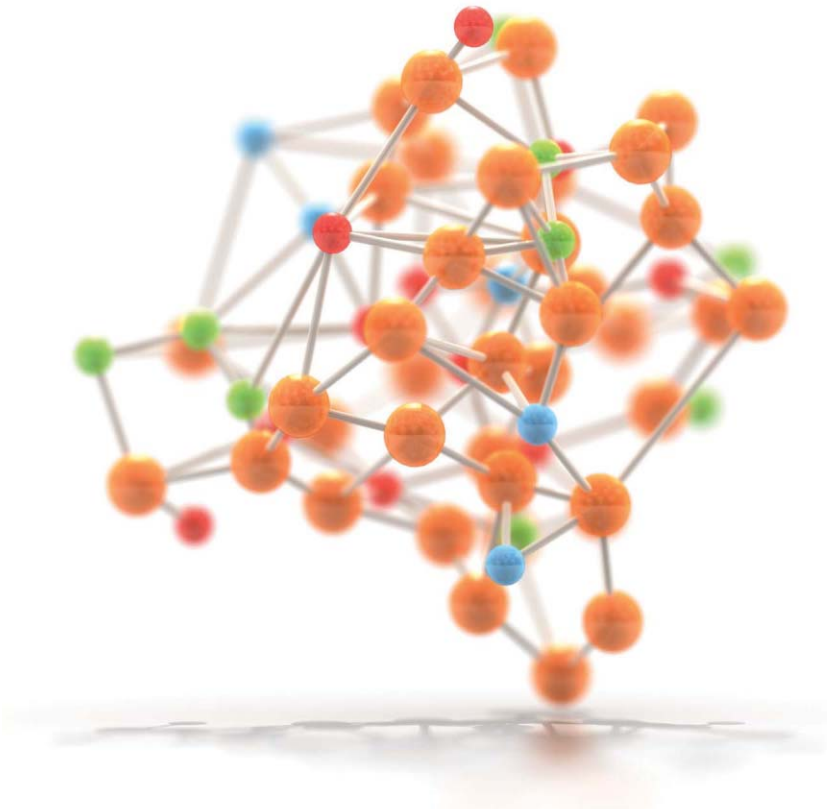

AMORPHOUS SILICON has long been the king of flat-panel displays. It began its reign in PC monitors and high-definition TV, then conquered netbooks, e-readers, and smartphones. No other substance was as suitable for the thin-film transistors that sit behind a display's hundreds of thousands of pixels, turning each one on or off.

But soon the dominion of amorphous silicon will pass, because it can't provide what coming generations of electronic products will require. For one thing, it isn't fast enough. Next-generation LCD TVs will be refreshed at least 240 times a second, which is two to four times as quick as today's versions; that way, they'll provide sharper fast-action sports and movies. Three-dimensional displays will need refresh rates twice again as high, to provide all that fast-motion goodness to each eye.

Nor are today's thin-film transistors stable enough for displays that use the organic light-emitting diode (OLED), a thin, efficient, high-contrast technology. Stability matters because the "threshold voltage" that an amorphous silicon transistor needs to turn on tends to drift as the transistor works. And both the problems of slow switching and drift get worse when amorphous silicon devices are made on flexible plastic, which is a critical design requirement for tomorrow's roll-up displays. Such displays will enable a laptop-size screen to fold away for storage inside a smartphone.

For these reasons, researchers and display manufacturers need a replacement in hand when the day comes for amorphous silicon to step down from its throne. And they already have their eyes set on a promising heir—in fact, a whole family of materials, known as amorphous oxide semiconductors. They're amorphous because like today's silicon standby, they lack a regular crystalline structure, and they're oxides because they're made of oxygen compounded with two or three metals, most commonly selected from zinc, indium, gallium, and tin.

Amorphous oxides can form thin films that are transparent and electrically conductive, which is why they already serve as the see-through electrode layer in displays and solar cells. It was this quality that led to the surge in research that began in 1996, when Hideo Hosono and his colleagues at the Tokyo Institute of Technology first noted the merits of amorphous transparent conducting oxides.

We believe that amorphous oxides could do more than simply serve as a passive electrode. They could also replace amorphous silicon as the active semiconducting material that does the heavy lifting as the channel in thin-film transistors.

HERE'S WHAT'S GREAT about amorphous oxides: First and foremost, charge can zip through them 20 to 40 times as fast as in amorphous silicon. This speed is defined by a material's charge-carrier mobility, and the higher it gets, the faster transistors can switch.

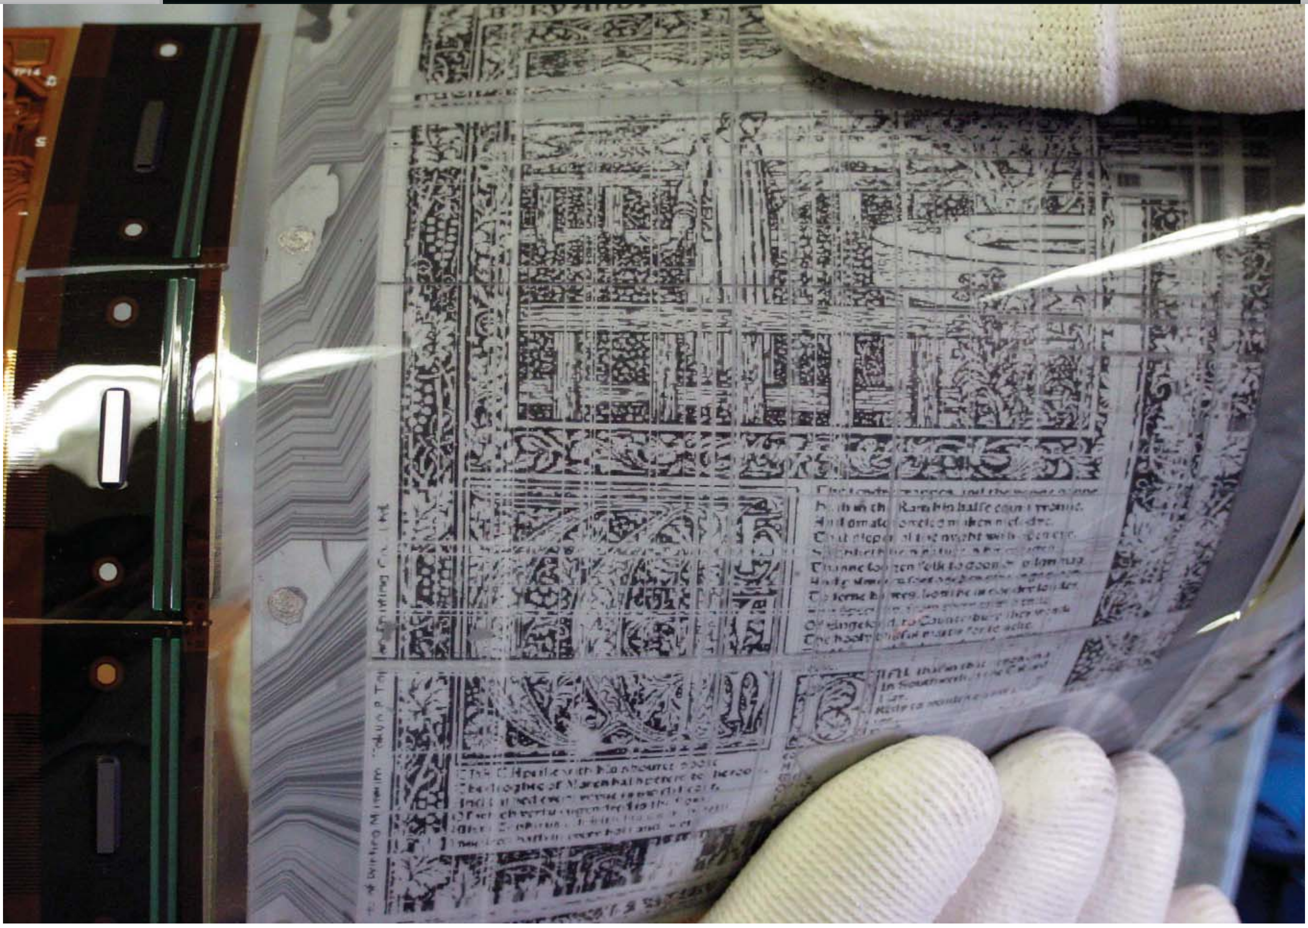

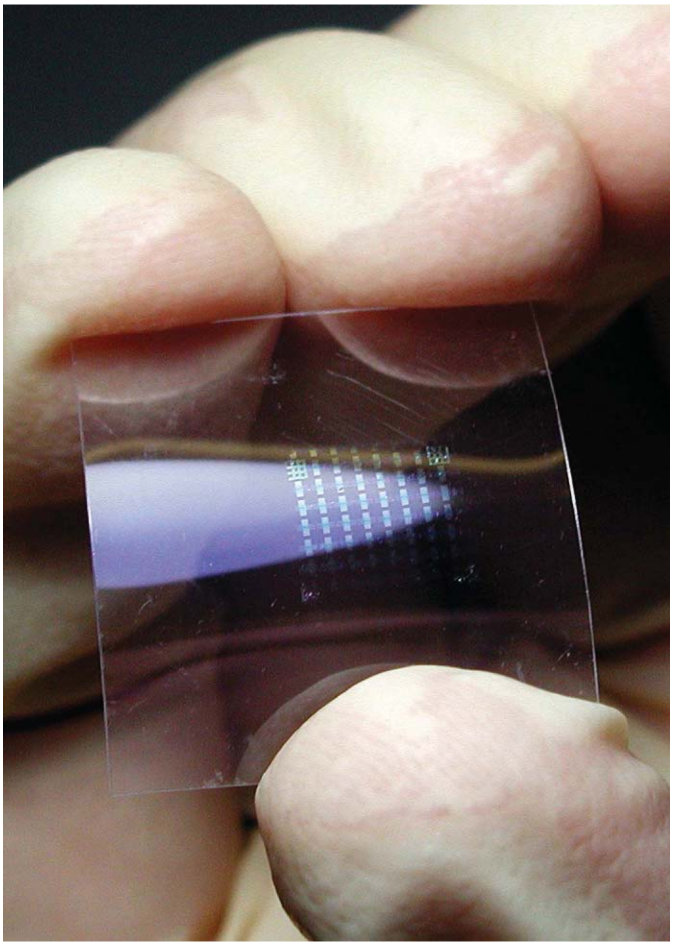

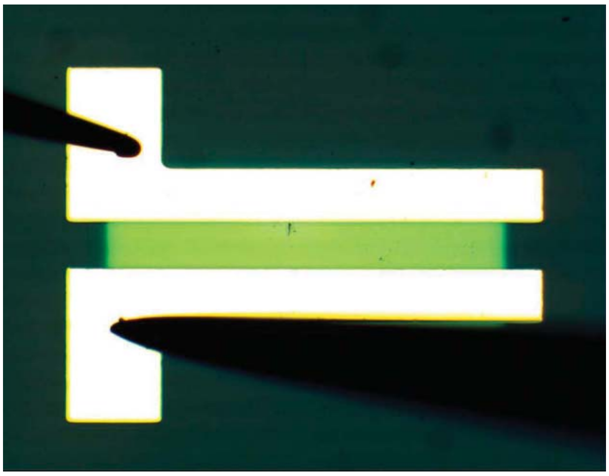

Flexible Displays

such as this 150-pixel-per-inch resolution black-and-white electronic paper [top] by Japan's Toppan Printing Co. use amorphous oxide semiconductor thin-film transistors [bottom left], devices that were first reported by researchers at the Tokyo Institute of Technology; at an Oregon State University lab in Corvallis, tests are performed on a single transistor [bottom right].

PHOTOS, CLOCKWISE FROM TOP: TOPPAN PRINTING CO.; ERIC SUNDHOLM; TOKYO INSTITUTE OF TECHNOLOGY

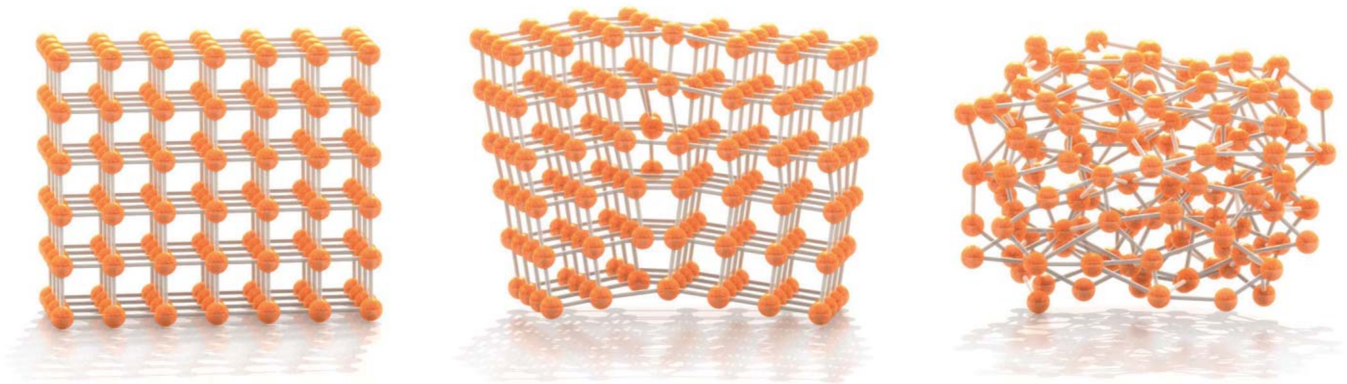

Crystalline *Polycrystalline* *Amorphous*

Solid Insight

Atoms line up in an orderly 3-D array in a crystalline solid [left], while a polycrystalline solid [middle] is made of smaller crystalline chunks. In an amorphous solid, perfect order goes out the door: The irregular arrangement lends itself to the formation of smooth, uniform films but also slows down charged particles—a mixed blessing for thin-film transistors.

Second, unlike amorphous silicon, amorphous oxides can be deposited at low temperatures without compromising their electronic properties. That means they can be laid down or even printed on pliable plastic sheets to make paperlike displays. Such printing could make flexible electronics cheap and ubiquitous. Finally, because the materials are transparent, they could be used in electronic-ink screens that could be laminated on windshields, windows, and eyeglasses.

Even some of the first thin-film transistors made from amorphous oxide semiconductors showed amazing promise. These devices, first reported in November 2004, also by Hosono's group in Tokyo, consisted of indium gallium zinc oxide (IGZO) built at low temperatures on flexible plastic substrates. They had more than 10 times the charge-carrier mobility of amorphous silicon transistors. Two months later, our research group at Oregon State University reported zinc tin oxide transistors with mobility up to *four times* as high as that.

Reports on transistors made from zinc indium oxide and many other combinations followed in quick succession. These devices have consistently performed better than their amorphous silicon counterparts.

The advantages of oxide semiconductors over amorphous silicon are causing a stir in the display industry. Just two years after the first oxide-based transistors were reported, Korea's LG Electronics Co. revealed a prototype OLED display that used IGZO transistors to drive its pixels. Other companies followed quickly with oxide-based displays of their own. Samsung presented a 19-inch OLED and a 17-inch LCD in 2009. AU Optronics Corp., in Taiwan, introduced a 37-inch LCD in early 2010. Samsung followed later that year with a 70-inch ultrahigh-definition (4000 by 2000 pixels) 3-D LCD panel with a 240-hertz refresh rate.

Several companies have also reported flexible displays built on oxide transistor backplane technology. Toppan Printing Co., in Tokyo, made the first black-and-white plastic e-paper displays using oxide transistors in 2005. It has since made an e-paper display measuring 5 centimeters diagonally with a resolution of 160 pixels per centimeter. LG and Dai Nippon Printing Co., also in Tokyo, have reported small, flexible OLED displays fabricated on stainless steel.

Researchers still do not fully understand the materials science or the device physics, but it's apparent that today's best IGZO devices would do well in LCD backplanes and are nearly up to the needs of OLED displays. For that latter application, the key question is device stability—that is, the problem of drift in the transistors' threshold voltage. Recent progress suggests that stability will not be a showstopper. *Continued on page 51*

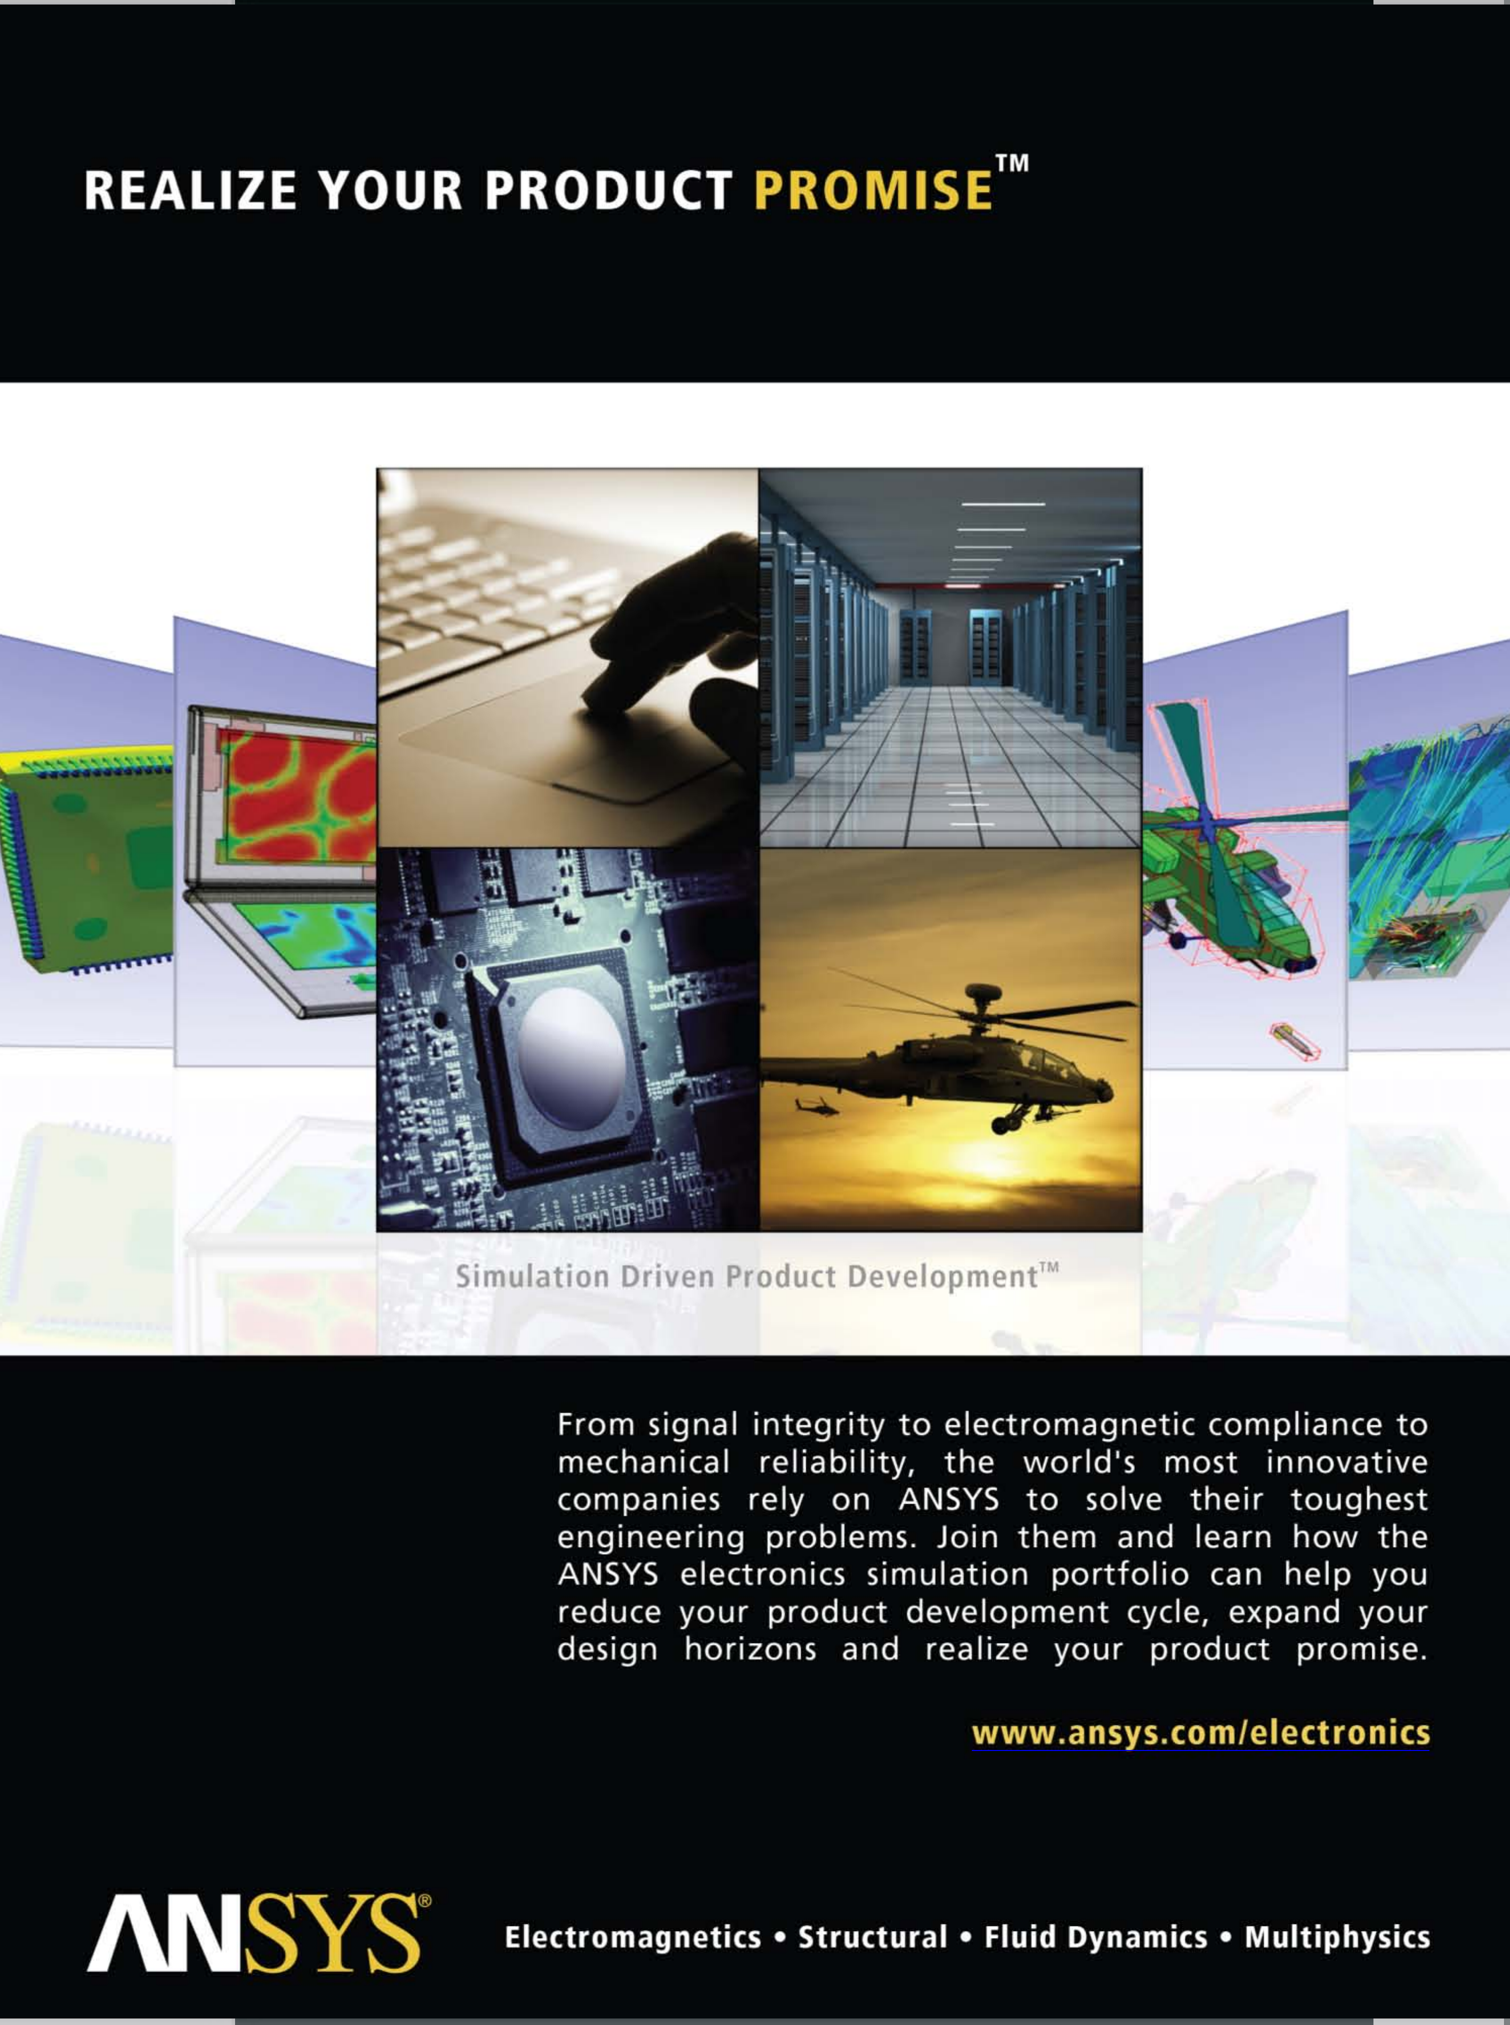
REALIZE YOUR PRODUCT PROMISE™
Simulation Driven Product Development™
From signal integrity to electromagnetic compliance to mechanical reliability, the world's most innovative companies rely on ANSYS to solve their toughest engineering problems. Join them and learn how the ANSYS electronics simulation portfolio can help you reduce your product development cycle, expand your design horizons and realize your product promise.
www.ansys.com/electronics
ANSYS®
Electromagnetics • Structural • Fluid Dynamics • Multiphysics

IEEE MEDAL *of* HONOR

Foundry Father

Morris Chang pioneered the $28 billion semiconductor foundry industry, kicking off a revolution in the business of chips / BY TEKLA S. PERRY

Even in the very anomalous category of tech entrepreneurs who've become pop stars, Morris Chang is an anomaly. He's 79 years old, and unlike Steve Jobs, who makes stylish consumer gear, or Mark Zuckerberg, who runs the world's largest social network and is the subject of a major motion picture, Morris Chang runs a semiconductor foundry. It's a big one, to be sure: Taiwan Semiconductor Manufacturing Co. (TSMC), in Hsinchu, is by far the world's largest foundry company, with annual revenues last year of about US $13.3 billion. And yet, the foundry business isn't what most people would regard as the glamour sector of the semiconductor industry—a foundry is basically a fab-for-hire. Companies create their own designs and pay the foundry to manufacture the chips.

Chang is credited with pioneering the foundry concept, in 1987, when he founded TSMC with backing from the Taiwanese government and Netherlands-based Philips Electronics. For that achievement, Chang will receive the 2011 IEEE Medal of Honor.

Like many radical ideas, the foundry concept did not seem so brilliant at first. But the business took off, with TSMC dominating it in a way that few companies have ever dominated any industry. Today TSMC's revenues account for almost half those of the entire foundry industry. And its operating profits—about $5.3 billion last year—were an astounding 90 percent of the whole global foundry industry's.

As TSMC's revenues soared, so too did Morris Chang's stature in Taiwan, where he is a national hero unlike any other. His smiling face appears on billboards endorsing consumer gadgets, real estate, and other goods. Even in his own company, work stops and people stare when he walks down a hall.

Checking out of my hotel in Hsinchu, I decide to test Chang's fame. "Is this your first visit to Taiwan?" the desk clerk asks politely. She looks to be about 25 years old. "Yes," I say. "I was here to meet with Morris Chang." Her voice jumps nearly an octave. "Oh, really? How exciting!" she yelps. "That must have been so, so...." Her English fails her. I feel sorry that

PHOTOGRAPHY BY GARRET M. CLARKE

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

I hadn't mentioned it sooner; I could have gotten his autograph for her and really made her day.

Necessity is the mother of invention, as the old saying goes. Chang wasn't trying to reinvent the semiconductor industry when he started TSMC. He did it because government officials in Taiwan wanted him to start a semiconductor business, and with Taiwan weak in both design and marketing, he didn't see any other way.

Chang had no idea that he was launching a technological revolution, the likes of which the semiconductor industry had not seen since, arguably, the rollout of user-programmable logic devices during the previous decade. Suddenly, entrepreneurs could create semiconductor businesses around chips without the huge expenditure of cash and effort it takes to open a semiconductor fabrication facility; all they had to do was design a chip and then market it. Thus was born the "fabless" semiconductor industry, which racked up $73.6 billion in revenues last year. A surprising number of today's hottest high-tech companies got their start in TSMC's factories—telecommunications pioneers Broadcom and Qualcomm, graphics powerhouses Nvidia and ATI, mobile device innovator Marvell, programmable logic creator Altera—and their products continue to roll off TSMC's manufacturing lines today.

"Morris Chang completely changed the landscape of the semiconductor industry," says James Plummer, dean of the school of engineering at Stanford. "He enabled start-ups to start with a few million dollars rather than a few hundred million. That makes a huge difference."

Chang aspired to be a writer as a child. Born in Ningbo, just outside of Shanghai, Chang moved several times as a youngster: to Nanjing, for his father's work in banking; to Canton (now Guangzhou); to Hong Kong, to escape the Japanese bombing during the second World War; and eventually to Chongqing in southwestern China. He thought he'd be a novelist, maybe, or a journalist, but perhaps it was "just a young man's fancy," he says. Then, in 1949, Chang, with the help of an uncle in Boston, was accepted at Harvard University.

"My reaction entering Harvard was sheer ecstasy, almost disbelief," he recalls. "What a country! The United States was at its peak in its moral leadership, in its political leadership in terms of democracy—and it was the richest country in the world." It was then he gave up his childhood dream of becoming a writer. "In the early '50s in the United States, there were Chinese laundrymen, Chinese restaurateurs, Chinese engineers, and Chinese professors," he says. "Those were the only respectable professions for Chinese—no lawyers, no accountants, no politicians"—and certainly no writers. Chang decided he'd be an engineer, declaring "applied physics and engineering sciences" as his major.

After a year at Harvard, he realized that the school's general technical education wasn't going to land him a top engineering job, so he transferred to MIT. He majored in mechanical engineering because it seemed the most general of the specialties. After failing his Ph.D. qualifying exam—he simply didn't study hard enough, he says—he went out looking for jobs. Ford Motor Co. offered him $479 a month; Sylvania Electric Products' then tiny semiconductor division offered $480. Chang asked Ford to beat the Sylvania offer. It refused. He took the job with Sylvania. It was 1955.

His first assignment was to work on germanium transistors, first on improving manufacturing yields and later on developing devices. Chang started working toward yield improvement by looking at bonding, the process by which electrical contacts are attached to the transistors. At the time, a technician on the assembly line connected those electrodes by soldering on a wire, and the heat often damaged the transistor. Chang devised a way to connect the wire using indirect heat instead.

It was a basic mechanical engineering problem; he didn't have to understand semiconductor theory to figure it out. But he started studying semiconductors anyway, using as his text *Electrons and Holes in Semiconductors with Applications to Transistor Electronics* (1950) by William Shockley.

Parts of the book baffled him. But he soon discovered that a senior engineer at Sylvania spent hours every evening drinking heavily at the bar in the same hotel where Chang was living, in Ipswich, Mass. So Chang's typical evening in those days went like this: He'd spend a couple of hours holed up in his room poring over a section of the book and trying to solve the problems on his own. Then he took his questions down to his friend in the bar. "He didn't solve all my problems, but he solved enough so I could move ahead. He was my main teacher," Chang says.

Chang eventually moved on to designing new transistors. He attended technical conferences, served on a standards committee. By 1958 he'd grown tired of developing what he saw as good transistors that went nowhere because of Sylvania's lack of marketing prowess, so he accepted an invitation to visit Texas Instruments, in Dallas.

"My future boss actually met me at the airport," he recalls. Then 27 years old, Chang was struck by the relative youth and exuberance of TI's corporate culture. The majority of the people in the semiconductor business at Sylvania had moved over from vacuum tubes and were in their 50s. He took a job at TI in production, working on four transistors being manufactured for IBM, for which IBM had provided the design and process. In essence, it was an early foundry deal.

One of the transistors, a particularly difficult one to produce, was manufactured by both TI and IBM, which operated a pilot line elsewhere just for that device. Chang took charge of TI's production line for that troublesome transistor. At IBM, at best 10 percent of the transistors produced on the pilot line worked, while at TI, the yield on Chang's line was pretty much zero, occasionally getting as high as 2 to 3 percent. Months passed while Chang tinkered with the process.

"The supervisor was concerned. The operators were concerned. Everybody was concerned," Chang recalls, but he says he wasn't too worried, because he knew that getting a 10 percent yield was at least possible.

Finally, about four months into the project, Chang began testing yet another combination of process parameters, changing the temperature and the time in the furnace and experimenting with different alloys and diffusion materials. The yield jumped to more than 20 percent.

Suddenly, even TI president Pat Haggerty knew his name. IBM thought Chang had just gotten lucky, but when the company sent engineers to talk to him, he described the theories he'd been testing and explained why his experimental process worked. The achievement propelled him into his first management job, creating a germanium transistor development department with 20-plus engineers.

Chang loved management. He liked being able to do bigger things than he could as an individual, and frankly, he aspired to a higher standard of living—what he calls a "reasonably comfortable life." He was 28 years old and entertaining dreams, and not

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

unrealistic ones, of being vice president of R&D someday. His boss at the time told him that if he aspired to that job, he'd need to get a Ph.D.—not because his superiors cared but because the R&D teams at TI definitely would. So, in 1961, at age 30, Chang went to Stanford, with TI paying his tuition and expenses on top of his full salary while he pursued his degree full time.

By spring of 1964 Chang was back in Texas with his Ph.D. in electrical engineering. He quickly moved up, eventually having all 3000 people in the germanium transistor department under him. He had a lot to learn—product marketing and finance, for example, weren't taught in engineering school.

"Morris's foundry model unleashed an army of engineers and made people's dreams come true."

—ROAWEN CHEN
vice president of manufacturing, Marvell Technology Group

For the next 19 years Chang moved through various management jobs at TI, working his way up. He pioneered the then controversial idea of pricing semiconductors ahead on the cost curve—that is, sacrificing profits early on to grab market share to achieve manufacturing yields that in turn would yield greater long-term profits. By 1972 he was running the company's entire semiconductor business—and aspiring to take over as CEO one day.

But in the mid-'70s, TI decided to move more heavily into the consumer electronics business, with calculators and wristwatches. In 1978 the company picked Chang to run the consumer business. He was proud of Speak & Spell, the breakthrough speech-synthesis toy that came out that year, which used the world's first single-chip speech synthesizer [see "25 Microchips That Shook the World," *IEEE Spectrum*, May 2009]. But overall, he says, the move was a mistake, for the company and for him personally.

Mark Shepherd, then chairman and CEO of TI, agreed with the prevailing wisdom at the time that a good manager could manage anything. In this case, Chang says, "I think he was wrong. I found the consumer business to be very different. The customer set—completely different. The market—completely different. And what you need to get ahead in that business is different, too. In the semiconductor business, it's just technology and costs; in consumer, technology helps, but it's also the appeal to consumers, which is a nebulous thing."

In 1983, with TI's consumer business struggling and Chang sidelined as head of "quality and people effectiveness," he knew his path at TI no longer led to the executive suite. He decided to leave.

He had no real plans, but he was confident he'd be a hot job prospect. And so he was; after his departure hit the news, his phone started ringing, he says, and didn't seem to stop for days. Eventually, he had two offers he seriously considered—one from a venture capital firm and one from General Instrument, an East Coast manufacturer of semiconductors and cable set-top boxes.

He joined GI in 1984 as president and chief operating officer and moved to New York City; he'd once dreamed of living there and walking to work. He moved to an apartment on Fifth Avenue, just a few blocks away from his office in the General Motors building. And he had, he thought, free rein to build up the company's R&D operation.

He didn't. The company, he says, wasn't really serious about R&D, preferring to buy businesses rather than grow them internally. And the pleasure of walking to work paled quickly in the rain and winter slush. Meanwhile, he had separated from his then wife, Christine.

He left GI after just a year. Once again, he quit without a plan. "After these two setbacks, at TI and at GI," he says, "I did not think that my aspiration to be the CEO of a major U.S. company was in the cards." He contemplated going to a venture capital firm, but, he says, "I didn't think I would be very good at it."

Then a government official in Taiwan, K.T. Li, contacted him. Chang had met Li several years previously, when TI was negotiating to build an assembly plant in Taiwan. Li was looking for someone to manage the decade-old Industrial Technology Research Institute (ITRI), turning it into the Bell Labs of Taiwan.

Just about everybody he knew told him not to do it. His soon-to-be ex-wife thought he'd lost his mind for even seriously considering the offer. But Chang, thinking he had little to lose, decided to give it a try. He took over as president of ITRI in 1985 and immediately began shaking things up. For starters, he broke the "iron rice bowl," challenging the assumption that once you got a job in a government organization, you'd be set for life. Instead, Chang instituted a review process that put the lowest 2 percent of performers on probation, with the possibility of dismissal looming if they didn't improve. He also began transitioning ITRI toward partial funding by industrial contracts, not just by government subsidy. Not surprisingly, these changes didn't go over well with the staff. Hate mail began to pour in to Taiwanese legislators, complaining about Chang and his policies. "Back 25 years ago," he says, "they considered me a foreigner who suddenly became their boss. They were scared of me."

"He's known not to tolerate fools," says Chenming Hu, a professor emeritus at the University of California, Berkeley, and former chief technology officer for TSMC. "And that engenders both respect and fear."

The situation at ITRI had only just started to settle down when Li came back with another proposal. The government, Chang told him, wanted Taiwan to have a semiconductor company and thought that if anyone could get it off the ground, Chang could.

Chang knew it wouldn't be easy. TI was still the dominant company in the semiconductor business, with Intel rising fast and National Semiconductor Corp. and Advanced Micro Devices also strong players. Just to survive among these giants would be a challenge.

And he would be starting at a huge technical disadvantage. In 1975, ITRI had licensed from RCA Corp. a 7-micrometer semiconductor processing technology that at the time was already out of date. By 1985, research at ITRI had advanced to a 3-µm process; however, the rest of the industry had moved forward even faster—about two generations ahead, with advanced products being manufactured in 1.5-µm technology.

But Chang said yes anyway. "It's like in the movie *The Godfather*," he says, "an offer you can't refuse," explaining that to say no would have tagged him as someone completely without ambition and would likely have dimmed his future in Taiwan.

Taking on the semiconductor behemoths head-to-head was out of the question. Chang considered Taiwan's strengths and weaknesses as well as his own. After weeks of contemplation he came up with what he calls the "pure play" foundry model.

In the mid-'80s there were approximately 50 companies in the world that were what we now call fabless semiconductor companies. The special-purpose chips they designed were fabricated by big semiconductor companies like Fujitsu, IBM, NEC, TI, or Toshiba. Those big firms drove tough bargains, often insisting that the design be transferred as part of the contract; if a product proved successful, the big company could then come out with competing chips under its own label. And the smaller firms were always second-class citizens; their chips ran only when the dominant companies had excess capacity.

But, Chang thought, what if these small design firms could contract with a manufacturer that didn't make any of its own chips—meaning that it wouldn't compete with smaller firms or bump them to the back of the line? And he realized that this pure-play foundry would mean that Taiwan's weaknesses in design and marketing wouldn't matter, while its traditional strengths in manufacturing would give it an edge.

TSMC opened for business in February 1987 with $220 million in capital—half from the government, half raised from outside investors. Its first customers were big companies like Intel, Motorola, and TI, which were happy to hand over to TSMC the manufacture of products that used out-of-date technology but were still in some demand. That way, the companies wouldn't have to take up their own valuable fab capacity making these chips and would face little harm to their reputation or overall business if TSMC somehow failed to deliver.

Soon, Chang says, start-up companies, which would live or die depending on TSMC's manufacturing runs, signed on. Early customers like Broadcom, Marvell, Nvidia, and Qualcomm, Chang says, "started with us when they were small," and being able to tap into TSMC's manufacturing prowess was a big reason for their success.

"Not having us would have slowed down innovation in the industry," Chang maintains, pointing out that those little companies would likely have had to invest in manufacturing capacity instead of R&D or share their intellectual property with a TI or an IBM.

Jen-Hsun Huang, cofounder, president, and CEO of Nvidia, says that TSMC's birth enabled all sorts of creative ideas—in areas like networking, consumer electronics, computers, and automotive technology—to be turned into successful companies, because "the barriers to getting your chips built, to realizing your imagination, disappeared."

Huang started Nvidia assuming that the company's chips would have to be made in partnership with a traditional semiconductor company; Nvidia initially partnered with SGS Thomson Microelectronics (now STMicroelectronics). As those things go, he says, it was a good enough partnership, but there was no denying that the two businesses had different interests.

"When I learned about Morris and TSMC," Huang says, "I was thrilled. I wrote him letters—the only way I had of contacting him. And one day he called me." The two met soon after that phone call and, for Huang, the two companies were a perfect match.

Morris Chang

2011 IEEE Medal of Honor for outstanding leadership in the semiconductor industry

CURRENT JOB: Founder, chairman, CEO of Taiwan Semiconductor Manufacturing Co.

DATE OF BIRTH: 10 July 1931

BIRTHPLACE: Ningbo, China

FAMILY: Wife, Sophie; one daughter from first marriage; two stepdaughters

FIRST JOB: Typing papers in college

MOST RECENT BOOKS READ: *Fall of Giants* by Ken Follett; *Decision Points* by George W. Bush

FAVORITE MUSIC: Classical through Gustav Mahler

COMPUTER USED MOST OFTEN: BlackBerry

FAVORITE RESTAURANT: The Four Seasons, New York City

FAVORITE MOVIES: *The Third Man*, *Sunset Boulevard*, *All About Eve*

"I loved that TSMC's intentions were pure, that their success only came with our success. Morris and I were both building our companies; Nvidia had to move fast to keep up with the competition, and TSMC kept up with our needs," he says. Today Nvidia has revenues of nearly $4 billion a year, and TSMC still does virtually all of its manufacturing.

TSMC was already well established when fabless semiconductor firm Marvell Technology Group launched in 1995. Says Roawen Chen, vice president of manufacturing at Marvell, which brought in $2.8 billion in 2010 revenues: "It would have been impossible for Marvell to start without foundries—it would have needed $100 million to build factories, the entry barrier was so high."

"Morris's foundry model unleashed an army of engineers," says Chen, "and made people's dreams come true."

Today TSMC has a market capitalization of nearly $70 billion. Chang himself received no stock for founding the company. Though he created one of the biggest companies in Taiwan, he doesn't even make the country's *Forbes* 40-richest list. "I like to measure a person's wealth as a portion of the shareholder wealth he created," Chang says. "On that basis, I made out really, really poorly." Still, he isn't complaining; he is well compensated, he says. And several years after TSMC's start, he invested his own savings in the company, profiting as TSMC boomed.

Life is good, to put it mildly. During the week, he lives with his second wife in an apartment in Hsinchu, about 20 minutes from TSMC's huge headquarters complex of factories and offices. On weekends, they stay at their 460-square-meter duplex penthouse in Taipei.

He is also famous, though he shrugs it off. "I don't give a great deal of stock to this thing of being recognized," Chang says. "After all, Taiwan is a pretty small place."

And at nearly 80 years of age, Morris Chang is still running TSMC full time, as CEO and chairman. He starts his days at 6:30 a.m. to check e-mails at his apartment, arriving in the office around 8:30 and working until about 6:30 every evening. He doesn't have any plans to stop. He does contemplate stepping back as CEO at some point while remaining chairman. He smiles, warming to the idea. In Taiwan, he notes, the chairman's role can be defined just about any way someone wants to.

And if Chang is good at anything, it's defining things. After all, he defined an entire industry. ❑

Thin, Fast, and Flexible

Continued from page 44

Of course, the US $100 billion flat-panel display industry has been built on amorphous silicon, and the new materials will have to contend with its 30-year head start. But amorphous silicon is a mature technology; it is extremely unlikely that anyone will overcome the limitations it faces now, because these arise from the material's fundamental physical and chemical nature.

There's little doubt that amorphous oxide semiconductors can oust amorphous silicon; the only question is *when*. The answer depends mainly on how long it will take to develop large-scale manufacturing capability. The oxide thin-film transistor fabrication process is very similar to that used for amorphous silicon devices, so the display industry can leverage much of the existing infrastructure and know-how.

TO UNDERSTAND the promise of amorphous oxide semiconductors, we first have to deconstruct the term. Let's start with amorphous materials in general and why they are relevant for large-area electronics.

Solids can be crystalline, polycrystalline, or amorphous depending on how their atoms are arranged. In a single crystal the atoms stand in perfect order, forming a repetitive 3-D pattern. In a polycrystalline solid, many microscopic mini crystals, or grains, butt up against one another in a haphazard manner.

In an amorphous solid, each atom and its nearest neighbors stand in a not-quite-perfect order, an arrangement that gets increasingly erratic as you zoom out. The atom's geometry with respect to its second-nearest neighbors will be less predictable, still more random with its third-nearest, and so on.

The random arrangement of atoms tends to give amorphous materials smooth surfaces and uniform structure when they are produced as a film, which can measure anywhere from tens to hundreds of nanometers thick. And these films are more robust than polycrystalline ones, whose grain boundaries act like shunt paths for impurities. Water vapor, for instance, would more readily penetrate polycrystalline films. That's why amorphous silicon has wrested the domain of thin-film transistors from crystalline silicon, its electronically superior cousin.

Single-crystal silicon is the unrivaled ruler for integrated circuits used in mainstream electronics, and it has worn its crown since the 1960s. But single crystals are too costly, brittle, and heavy to cover 2-meter-wide glass panels. While thin-film polycrystalline silicon can work well for transistors in large high-performance displays, it requires costly, complex processing.

Amorphous silicon, on the other hand, has just what this job requires. Large glass panels can be cheaply coated with uniform thin films of the stuff at high volumes, producing hundreds of thousands of electronically consistent transistors.

Transistors have three electrodes. A voltage applied to the gate electrode allows charge to flow through a channel between the source and drain electrodes, turning the device on. Like silicon integrated circuits, thin-film transistors are built layer by layer. Films of various materials are coated one over the other on a substrate to form the electrodes and channels. Amorphous silicon films form the semiconducting channel layer in thin-film transistors.

For all its advantages, amorphous silicon really isn't the best electronic material. That's because the more randomly you arrange the atoms in a solid, the harder it is for current to flow through. Charge carriers travel in amorphous silicon at only one-thousandth the speed they do in single crystals. In LCD pixel electronics, that speed limit hasn't mattered much—until now.

THIS IS WHERE amorphous oxide semiconductors come in. They offer all the advantages of amorphous silicon for large-scale electronics, with much better electronic properties to boot.

A TRANSPARENT TFT, or thin-film transistor, uses an amorphous oxide semiconductor channel [orange] to connect two electrodes—the source and the drain. A voltage applied at a third electrode, the gate, allows charge to flow through the channel, switching on the device.

But first, a bit more about these materials. The metal atoms for the oxides are selected from the 15 elements that sit in rows 4 to 6 and columns 11 to 15 of the periodic table. This block includes rare elements like silver and gold as well as environmentally unfriendly ones, like arsenic and lead. Eliminating toxic and costly materials leaves 8 elements that can give 28 dual-oxide combinations.

The individual metal oxides, say, zinc oxide and tin oxide, tend to become polycrystalline when deposited as thin films. But an interesting thing happens when the two are mixed in similar concentrations and deposited. Because they have different crystal structures, the atoms cannot figure out which structure to adopt, and zinc tin oxide ends up as an amorphous film.

A key advantage amorphous oxides hold over amorphous silicon is their higher charge-carrier mobility. Electrons in the oxides fly about in a spherical region around the metal nucleus. The radii of these spherical orbitals are large. So even when atoms are arranged amorphously, neighboring atoms' orbitals overlap to a large extent. Electrons can therefore rapidly skip between atoms in the solid. In amorphous silicon, by contrast, electrons move around the nucleus in thin dumbbell-shaped regions that don't overlap much.

The electron-rich amorphous oxide films are conductive. But in order to make effective transistor channels, they must be insulating. They attain that ability when they are deposited in the presence of oxygen, which mops up free electrons. That way almost no current flows in the transistor's channel unless the proper voltage is applied at the gate.

But when that voltage is applied, the device switches on and the charge flies. In today's oxide transistors, electrons have a mobility of about 10 to 20 square centimeters per volt-second, as opposed to 0.5 for amorphous silicon transistors. Quicker-moving charges translate to faster, higher-current transistors. That means oxide transistors can carry up to 40 times as much current and switch on and off up to 40 times as fast. Further, if a transistor can carry charge quickly, it can be made smaller, which means more pixels can be packed into a square millimeter as ultrahigh-definition displays require. As research and development progress, we expect to get even higher mobilities.

Another advantage of amorphous oxides presents itself at low process temperatures. Making semiconductor thin films usually requires putting down the vapor form of the material on a substrate, where it solidifies. The higher the process temperature, the better the film's electrical properties.

Amorphous silicon films are typically deposited at around 300 to 350 °C. These temperatures work for glass, but they would distort plastic. At plastic-friendly temperatures, below around 200 °C, amorphous silicon transistors degrade tremendously.

Amorphous oxide semiconductor films, however, are typically deposited at benign temperatures below 100 °C. Their performance usually improves when they are baked at temperatures of around 250 to 300 °C, but even films processed at room temperature can easily be 10 times as mobile as amorphous silicon. This opens up the possibility of making roll-up electronic circuits on plastic.

What's more, these oxides can be prepared in liquid form and sprayed onto a surface. If you tried doing that with a silicon-based system, you'd need a vacuum chamber to keep out oxygen, which would combine with the silicon to form silicon oxide, an insulator. But the new materials aren't degraded by oxygen, because they're oxides to begin with. These materials could print thin-film transistors on plastic using inexpensive roll-to-roll processing techniques similar to those used to print newspapers.

SOME ISSUES must still be solved, beginning with device stability. Oxide-based transistors are more stable than the amorphous silicon kind, but they're still not as stable as the industry would like. This can be a problem even with LCDs, which are switched on and off by voltage, with a threshold that is not perfectly constant. It's an even bigger problem with OLEDs, which are switched by current—an arrangement that makes the circuits even more sensitive to transistor stability. We are working with Applied Materials, in Santa Clara, Calif., on ways to make these devices more stable.

Then there's the problem of printing inexpensive, flexible displays. True, Toppan's process can do so for e-paper, but these black-and-white displays can get by with slow switching speeds and low power. Roll-up color LCDs and OLEDs require oxide transistors of higher quality.

We are now working with Oregon State chemistry professor Douglas Keszler and Inpria Corp., in Corvallis, Ore. They are designing high-quality IGZO inks as well as other semiconductor and insulator inks. When these inks are used to pattern transistors, the devices perform almost as well as those made using conventional vacuum processes. The next step would be to pair this printing approach with a high-volume roll-to-roll process on flexible plastic substrates.

We are also working on zinc tin oxide, which incorporates less expensive metals and can be more easily fabricated as thin films. The problem is that its electronic properties are harder to control. We are now trying to tame them.

Finally, dare we mention an unrealized dream of ours? We would love to discover a *p*-type amorphous oxide semiconductor. Unlike today's amorphous oxides, which juggle electrons and are known as *n*-type, such a semiconductor would shuttle around positively charged holes—virtual particles corresponding to the absence of an electron.

Complementary metal-oxide semiconductor (CMOS) technology, which combines *n*- and *p*-type semiconductors, was a game changer for silicon integrated circuits, making possible low-power computer chips and microprocessors. Having something akin to CMOS for amorphous oxide semiconductors would have the same effect. For now, together with Boeing, we are trying to use *n*-type oxide transistors in simple digital and analog logic circuits.

CHANGE IS AFOOT: Although no one has announced it formally, recent developments suggest that display industry leaders are ramping up plans for large-scale production. Companies are already making oxide transistor-based displays with sixth-generation LCD panel manufacturing equipment, which uses glass substrates 185 centimeters wide and 150 cm high. What's more, the scale of recently demonstrated displays suggests that the industry will soon be able to use the eighth-generation equipment that makes today's large high-performance LCDs. Look for the new materials to begin to appear in commercial displays as early as 2012.

In 2010, Samsung introduced a 14-inch OLED screen that was 40 percent transparent. While the company hasn't revealed its pixel-controlling technology, there is no doubt that amorphous oxide transistors are perfect for the application and could bring see-through displays into consumer products.

Given the astonishing developments of the past two years, it's easy to forget that amorphous oxide semiconductors are still in their early adolescence; practical work started only about five years ago. In a few more years, they'll start appearing in your home, in your hand, and perhaps on the lenses of your glasses. ❑

Dean of the Faculty of Electrical Engineering

CZECH TECHNICAL UNIVERSITY IN PRAGUE

Czech Technical University in Prague (CTU) seeks visionary and entrepreneurial candidates for the position of The Dean of the Faculty of Electrical Engineering.

CTU is the premiere national technical research university and ranks 121st worldwide in the Engineering & IT category of the 2011 THES QS. The Faculty of Electrical Engineering (FEE) is the largest faculty at CTU with 3800 undergraduates, 2000 graduate students, 470 faculty members, and an annual budget of over CZK 700 million. FEE includes several internationally recognized departments and engages in research on the world-class level in several disciplines of Electrical Engineering and Computer Science.

The Dean of FEE is expected to define a management strategy, pursue opportunities for industry collaboration, and enrich the curriculum. The ideal candidate will have demonstrated managerial capabilities in large organizations comparable to FEE, proven leadership and administrative skills, and relevant experience in research or teaching in a scientific discipline. Czech fluency is not a requirement, but the Dean will be expected to communicate effectively in a Czech work environment.

The Dean will be appointed to a four-year term starting in June 2011. The Dean's duties include:

- **managing the budget and long-term investments**
- **chairing the FEE Scientific Council**
- **appointing vice-deans and department heads**
- **defining the organizational structure**
- **serving as the Faculty's official representative and as a liaison to industry partners and other Faculties**

Applicants should submit a CV and a management strategy proposal that articulates the candidate's vision before **June 13, 2011**. Candidates will introduce themselves and participate in a question-and-answer session with faculty members on June 22. The Dean will be chosen by a vote of FEE academic senate members.

For more information please visit **http://www.fel.cvut.cz/senat/dean** or contact Dr. Petr Habala (habala@fel.cvut.cz, +420 22435 3365).

ASSISTANT PROFESSOR
in Fiber Optical Communications

The position is placed at the Chalmers FORCE Centre, an unprecedented endeavor in Swedish research, devoted to merging the fields of fiber optic technologies and digital communication theory, driven by the increasing ability of fast electronics to perform advanced signal manipulation at GHz speeds.

Application deadline is May 30, ref no 2011/25. For further information and how to apply, please visit Vacancies at:
chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY

INC

Institute of Network Coding

Postdoctoral Fellows and Research Associates

Positions of Postdoctoral Fellows and Research Associates are open at the Institute of Network Coding (INC) of The Chinese University of Hong Kong (CUHK). Initial appointments are typically for two years, and the commencing date is flexible.

Applicants should have a strong research record in network coding related areas, including theory, applications, or implementation.

For further information please visit the INC home page at http://www.inc.cuhk.edu.hk or contact Prof. Raymond Yeung at whyeung@ie.cuhk.edu.hk

香港中文大學
The Chinese University of Hong Kong

the data

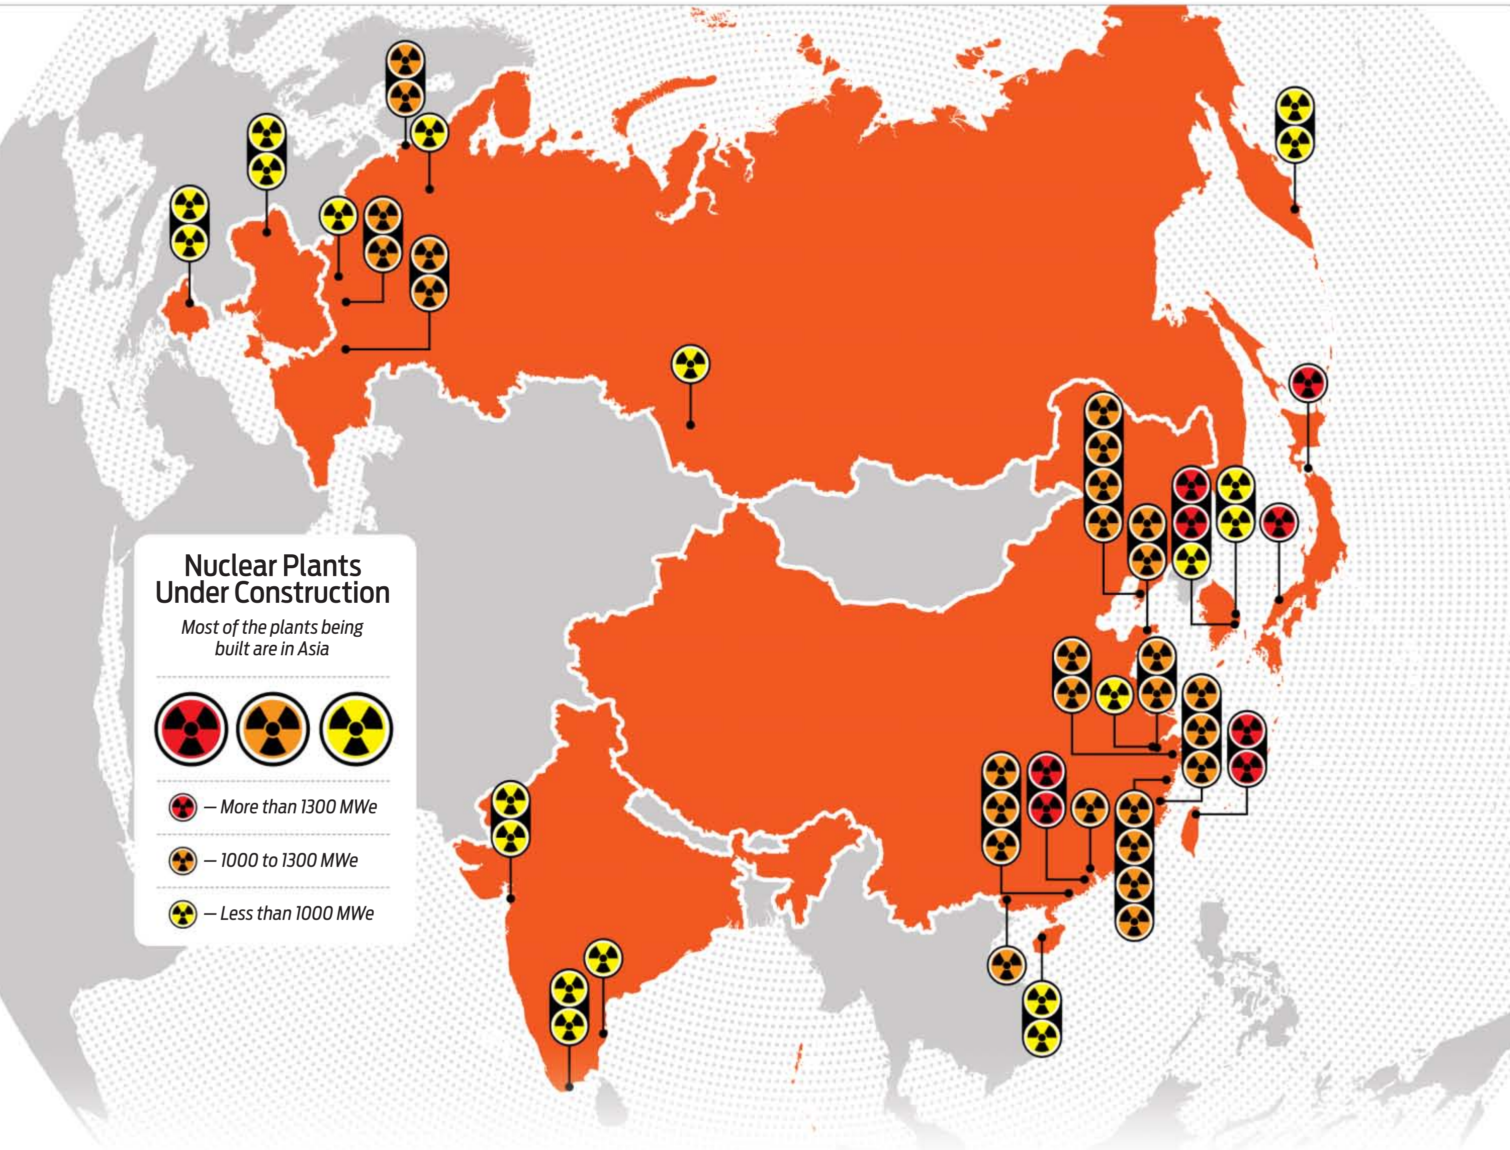

Will Nuclear Energy Charge Ahead?

THE JAPANESE earthquake and its devastating consequences for the Fukushima reactors will surely affect how governments and their citizens feel about nuclear power in the coming months. Nonetheless, as of January, 65 new nuclear power plants, capable of producing over 62 gigawatts, were under construction in 16 countries, according to the Nuclear Energy Institute, a trade group in Washington, D.C. In 2009, construction began on 11 plants, the most in one year since 1987.

Nearly all the plants that have gone on line since 2005—totaling more than 10 gigawatts of new capacity—are in Asia, and more are coming. India and South Korea are building five new plants each, and Japan and Taiwan are each building two.

But the dragon's share of growth is in China, which plans to build 50 reactors by 2020, quadrupling its nuclear capacity to 40 GW. However, energy demand is rising so rapidly that nuclear power will meet less than 10 percent of what China will need by then.

In that respect, China trails many European countries. France gets 75 percent of its energy from nuclear plants, the most of any country, and it is expected to connect one more 1600-megawatt plant in 2012.

Where's the United States in all of this? Its 104 operating nuclear plants represent 20 percent of its electricity capacity, and it has one plant under construction.

—*Prachi Patel*

Sources: Nuclear Energy Institute, World Energy Council

Nuclear Share of Power Generation (2009)

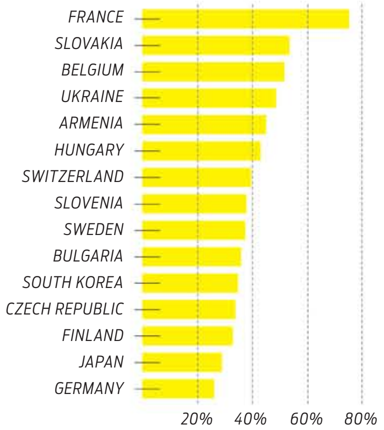

MCKIBILLO


Can semantic technologies
make the Web truly worldwide?
Find the latest telecommunications research in IEEE Xplore
Wherever you find people developing the most advanced telecommunications technology, chances are you'll find them using the IEEE Xplore digital library. That's because IEEE Xplore is filled with the latest research on everything from wireless technology and optical networks—to a semantic Web that can connect people around the world.
When it comes to telecom, the research that matters is in IEEE Xplore.
See for yourself. Read "Toward a New Generation of Semantic Web Applications," only in IEEE Xplore.
Try IEEE Xplore free—
visit www.ieee.org/betterinternet
IEEE Xplore® Digital Library
Information driving innovation
IEEE