IEEE SPECTRUM

THE MAGAZINE OF TECHNOLOGY INSIDERS

OCTOBER 2007
www.spectrum.ieee.org

Google's
Robyn Beavers

BRIGHT MOVE

WHY GOOGLE, APPLIED MATERIALS, AND WAL-MART ARE LEADING THE CORPORATE CHARGE INTO SOLAR ENERGY

ALSO

INSIDE A CRASH-TEST DUMMY

WHY COMPUTERS CAN'T BEAT PEOPLE AT GO

IEEE

DISPLAY UNTIL: 3 NOVEMBER 2007
www.ieee.org US $2.99 CAN $3.99
10>
0 73361 64849 5

Contents | Zoom in | Zoom out | For navigation instructions please click here | Search Issue | Next Page

FIRST-PASS SYSTEM SUCCESS

APPLICATION WORKSHOPS FOR HIGH-PERFORMANCE ELECTRONIC DESIGN

ASIA • NORTH AMERICA • EUROPE
SEPTEMBER 12 - NOVEMBER 15, 2007

Join us for First-Pass System Success, Ansoft's technical application workshops for high-performance electronic design.

Explore design techniques and simulation technologies for IC, packaging, RF & microwave, system signal integrity, system-level EMI, low power and much more.

Discover how to design systems that integrate RF/Analog/Digital SoC solutions with memory, graphics, storage, antennas, displays, cameras, and MP3 players.

Learn how to integrate electromagnetic components, power electronic circuits, and controllers into high-performance electromechanical systems.

Discuss with industry leaders how to achieve first-pass system success!

Get more information and register online: **ansoft.com/firstpass**

IEEE SPECTRUM

OCTOBER 2007

VOLUME 44 NUMBER 10 INTERNATIONAL EDITION

36
This electromechanical wonder could help save your life.

COVER: JOSON
THIS PAGE: CHRIS MUELLER

This month on
SPECTRUM ONLINE
www.spectrum.ieee.org

Sputnik's 50th Anniversary
On 4 October 1957, the Soviet Union put the 83.6-kilogram *Sputnik I* into orbit and turned the space race on its ear. We look back on this seminal event with aerospace experts and ask spaceflight visionaries where we might go in the next 50 years.

PARKING LOTS GO HIGH TECH
Combinations of cameras, sensors, and software called smart parking systems are making parking lots more secure and efficient. *IEEE Spectrum*'s Robert Charette investigates.

ALSO ONLINE:
- **Webcasts • Radio • News**
- **Audio Downloads • Podcasts**
- **Jobs • Career Accelerator**
- **IEEE Xplore® digital library**
- **Interviews • Opinions**

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

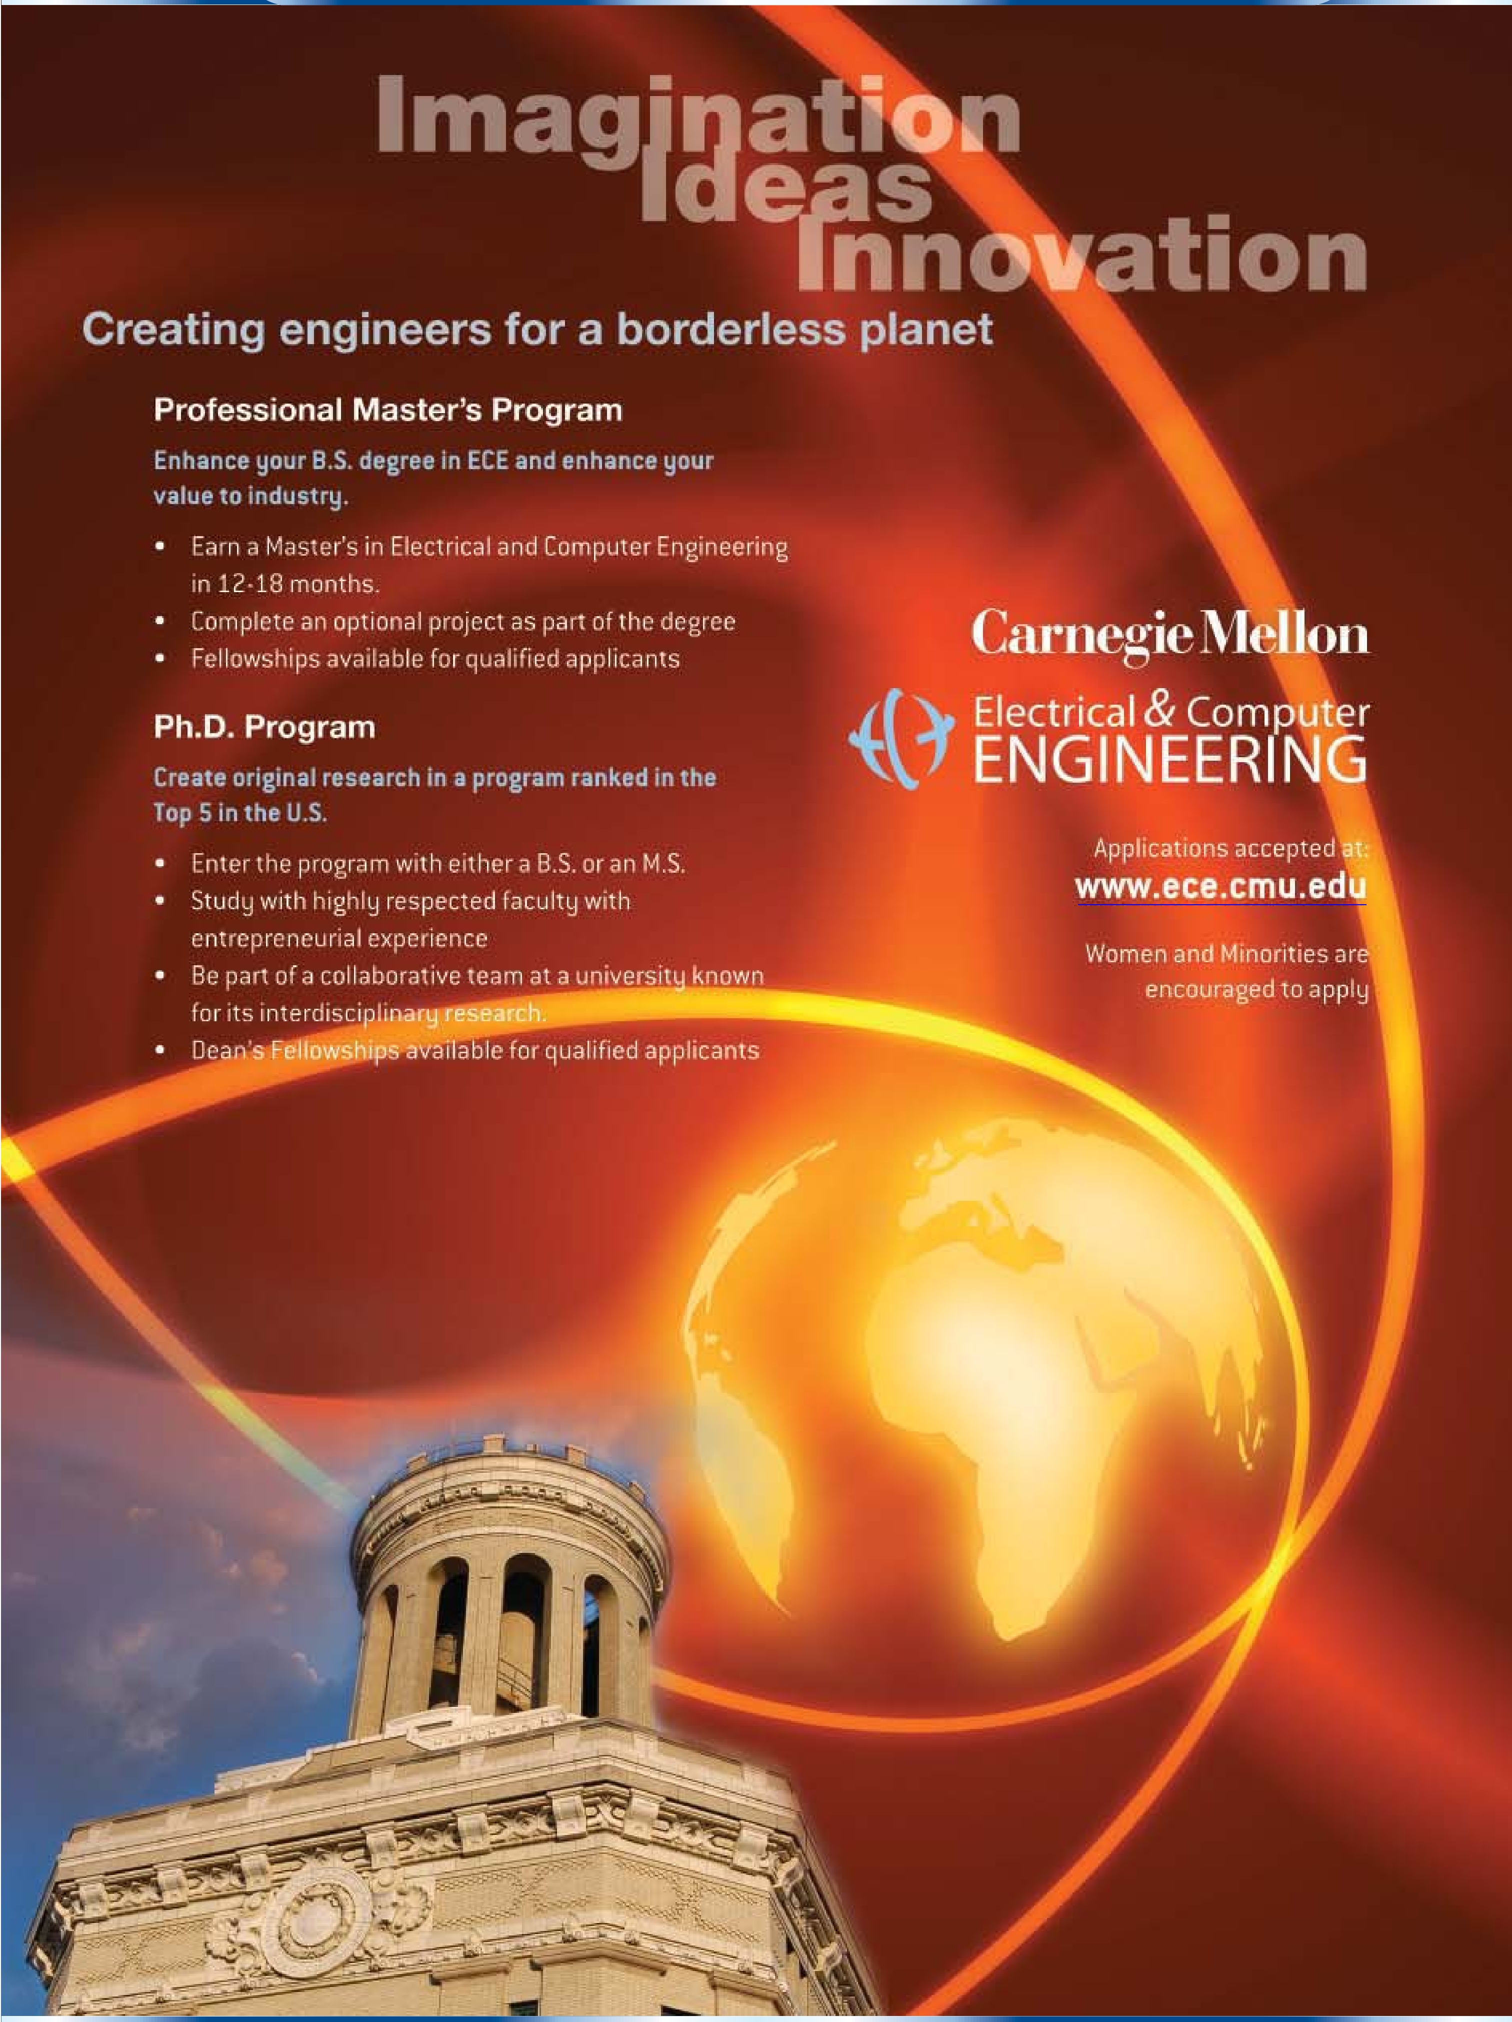
Imagination
Ideas
Innovation
Creating engineers for a borderless planet
Professional Master's Program
Enhance your B.S. degree in ECE and enhance your value to industry.
• Earn a Master's in Electrical and Computer Engineering in 12-18 months.
• Complete an optional project as part of the degree
• Fellowships available for qualified applicants
Ph.D. Program
Create original research in a program ranked in the Top 5 in the U.S.
• Enter the program with either a B.S. or an M.S.
• Study with highly respected faculty with entrepreneurial experience
• Be part of a collaborative team at a university known for its interdisciplinary research.
• Dean's Fellowships available for qualified applicants
Carnegie Mellon
Electrical & Computer ENGINEERING
Applications accepted at:
www.ece.cmu.edu
Women and Minorities are encouraged to apply

IEEE SPECTRUM

OCTOBER 2007
VOLUME 44 NUMBER 10 INTERNATIONAL EDITION

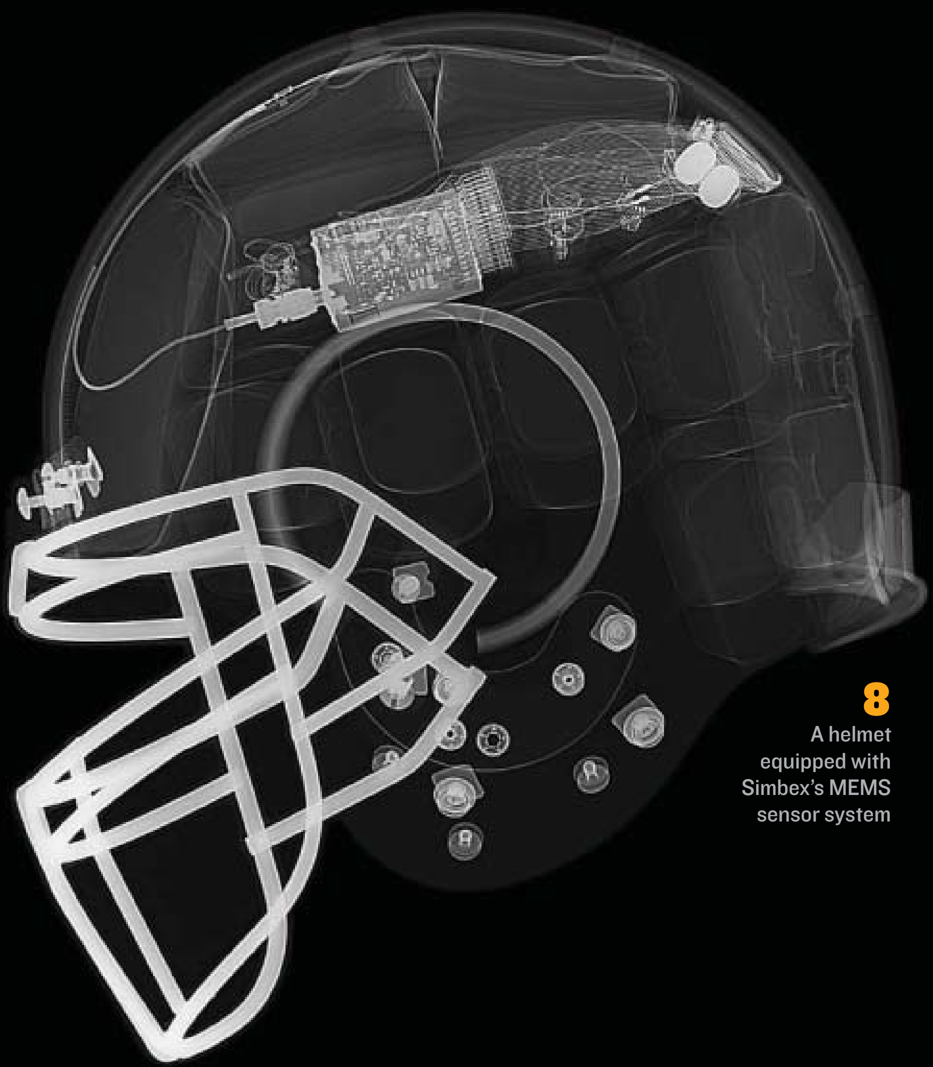

8 A helmet equipped with Simbex's MEMS sensor system

Available 5 October on
THE INSTITUTE ONLINE
www.ieee.org/theinstitute

Innovative Members In the News
IEEE members doing cutting-edge work are often featured in the media. Read about a few who are advancing technologies that can improve the lives of people with health problems.

IEEE.TV UPGRADE
Get the details on a new viewer just launched by IEEE.tv.

GLOBECOM CONFERENCE
Celebrating its 50th year, the conference continues to cover the latest on sensor networking, multimedia communications, optical networks and systems, and signal processing. Check it out from 26 to 30 November in Washington, D.C.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2007 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: $19.50 included in dues. Libraries/institutions: $205. 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 W224-N3322 Duplainville Rd., Pewaukee, WI 53072-4195, 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 the Society of National Association Publications.

SIMBEX

IEEE Spectrum Editorial Offices
3 Park Ave., 17th Floor, New York, NY 10016-5997 U.S.A.
Tel: +1 212 419 7555 Fax: +1 212 419 7570
Bureau: Palo Alto, Calif.; Tekla S. Perry +1 650 328 7570
http://www.spectrum.ieee.org

IEEE Operations Center 445 Hoes Lane, Box 1331, Piscataway, NJ 08855-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

THE BACK STORY

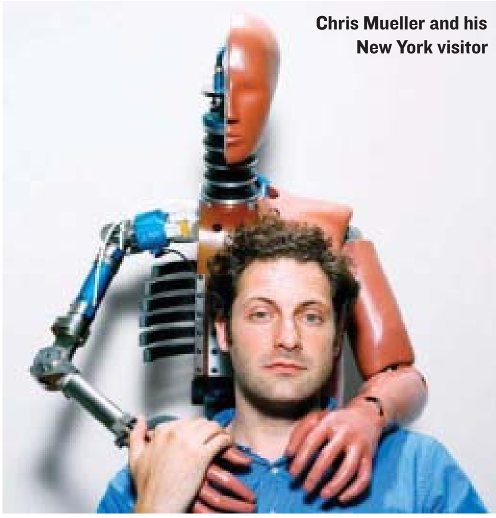

DAWN BLACKMAN

Chris Mueller and his New York visitor

When Chris Met Fred

Photographer Chris Mueller has had a few characters show up at his door in New York City. But none quite like Fred.

Fred is a crash-test dummy—and not just any crash-test dummy. He's a Hybrid-III adult male model, with half of his vinyl flesh removed to show off his sensor-packed steel skeleton. He was in town to be photographed for this issue's "Anatomy of a Crash-Test Dummy."

Fred, whose home is at the headquarters of dummy maker Denton ATD, in Rochester Hills, Mich., has been to trade shows all over the world. He normally travels by truck, but he's also traveled by car, sitting next to the driver, and once Denton bought him a seat on a flight to Europe.

This was Fred's first trip to New York, and he was delivered to Mueller in a refrigerator-size crate. "It was giant," says Mueller, who has photographed airliners, racing cars, human brains, and NASA engineers, but never dummies. How to get the thing to his studio, 10 blocks north, where the photo shoot was to take place?

"I opened the crate, put Fred on a hand truck, and rolled down the street with him," he says. "I covered him with a black canvas, and people kept looking over, catching glimpses of Fred."

At the photo session Fred proved very easy to work with. "He just listened and did what he was told," says Mueller, who after the shoot was off to an assignment in California. As for Fred, he was ready to go home, back to his Hybrid-III female companion, Frida. ■

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 first News page is in *IEEE Spectrum*, Vol. 44, no. 10 (INT), October 2007, p. 8, or in *IEEE Spectrum*, Vol. 44, no. 10 (NA), October 2007, p. 10.

FORUM

PRAKASH SINGH/AFP/GETTY IMAGES

"India's problems stem from three causes: overpopulation, rampant corruption, and nepotism"
—Vittal P. Pyati

BUREAUCRACY BLUES

I couldn't disagree more with Nirode Mohanty's assertion that India's colonial past is to blame for its present predicament [Forum, August]. India gained freedom in 1947. Sixty years strikes me as a sufficient period in which to shed a colonial past and strike out on an independent path. Blaming the British for India's difficulties is simply grossly unjust. India's problems stem from three causes: overpopulation, rampant corruption, and nepotism. The most basic of these is that India is one-third the size of the United States and has three times its population.

In fact, there were some unintended and significant benefits when the British took over. India's status as the largest democracy in the world is in no small measure due to the infrastructure established by the British for their own benefit. The British also left a railroad covering the entire country, a legal framework and secular courts, an education system second to none, a parliamentary government, and more. Before the Raj, India was a loose federation of kingdoms in perpetual war with one another based on differences in religion, language, and so on. The Moguls were the predominant power, and India would probably be an Islamic country today had not the British stepped in, uninvited as they were.

Vittal P. Pyati
Beavercreek, Ohio

FOILED AGAIN

The letter "Newton, Not Bernoulli" [Forum, August] rails against misinformation but seems to be guilty of that very transgression. Using the examples of an inverted airfoil and a fully symmetrical airfoil, the writer argues that the Bernoulli principle cannot explain how an inverted airfoil could generate lift. But in fact, the Bernoulli principle does so handily, according to NASA. A search of the NASA Web site on the topic of lift generation returns results that leave no doubt that NASA still endorses the Bernoulli effect as the primary source of lift by an airfoil.

NASA does not discount the lift generated by the Newtonian equal-and-opposite downward force due to angle of attack; rather, it acknowledges this as a component of the total lift. The NASA Web site includes material for schoolteachers, along with documentation for its various airfoil software packages.

Kirt Blattenberger
IEEE Member
Mt. Airy, N.C.

DOWN TO THE CORE

The Big Picture in the July issue shows the delay-line memory in UNIVAC. But the accompanying title, "Core Memories," might confuse some readers. The term "core memory" came about from the development of a different machine, also in the late 1940s and early 1950s. The Memory Test Computer (MTC) was built at the Massachusetts Institute of Technology's Lincoln Laboratory for the express purpose of testing the Multiplanar Coincident-Current Magnetic-Core Memory, and also as a general-purpose computer under a U.S. Air Force contract. It was there that I worked with the MTC between 1952 and 1956. And it was there, in 1954, that I completed my master's thesis, based on a compact magnetic-matrix switch to drive the core memory, instead of an assortment of big and hot 5998 vacuum tubes.

In the core memory, each cell was a toroidal ceramic magnet about 2 to 3 millimeters in diameter whose polarity represented a 0 or a 1, and could be switched by current in wires going through the cores. The computer word length was 16 bits, for a total of 64 x 64 x 16 = 65 536 bits of random access memory using 64 x 64 x 17 = 69 632 cores. The 17th bit checked for even or odd parity. I don't think any of us foresaw at the time the gigabyte memories of today.

The MTC project grew out of the invention and development work led by Jay W. Forrester at the MIT Servomechanisms Laboratory under Project Whirlwind. The Whirlwind I computer was retrofitted with a magnetic-core memory about the same time I was there.

Both the MTC and Whirlwind I computers, like UNIVAC I, were built with vacuum tubes. Until computers were constructed using transistors, their reliability was a constant problem. Each of these early computers occupied a large room full of electronics, with another large room full of air-conditioning equipment. Because of their size and appearance, we called the earliest magnetic-core memories "shower stalls." One of those shower stalls was exhibited at the Smithsonian Institution in Washington, D.C., for many years and at the MIT Museum more recently. By then, some of the vacuum tubes had been either broken or stolen as souvenirs.

Arthur D. Hughes
IEEE Life Member
Gladwyne, Pa.

CORRECTION

In "China Reaches for the Red Planet" [News, August], the orbit of Phobos should have been stated as 5989 kilometers above the surface of Mars, not 5989 meters. *—Ed.*

Letters do not represent opinions of the IEEE. Short, concise letters are preferred. They may be edited for space and clarity. Additional letters are available online in "And more Forum" at http://www.spectrum.ieee.org. Write to: Forum, *IEEE Spectrum*, 3 Park Ave., 17th floor, New York, NY 10016-5997, U.S.A.; fax, +1 212 419 7570; e-mail, n.hantman@ieee.org.

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

SPECTRAL LINES

Solar Gets Googled

Corporate America has discovered solar energy. Is it a fad or the start of a fundamental shift toward renewable energy? Associate Editor Sandra Upson explores the intricacies of big companies going solar in "The Greening of Google," in this issue.

Google garnered frenzied media attention in June when it fired up 9000-plus polysilicon panels mounted on rooftops at its complex in Mountain View, Calif. (check out the Google Solar Panel Project monitor at http://www.google.com/corporate/solarpanels/home to see a real-time tally of their output). When fully operational, the panels will be able to generate 1.6 megawatts of electricity, about 30 percent of the total needed to run the buildings they sit on at the headquarters, known as the Googleplex. The company is also using solar power to charge a fleet of hybrid electric vehicles to promote the development and use of plug-in hybrids. And it plans to generate an additional 50 MW of renewable energy by 2012. Google can take such steps more easily than many other companies because it is flush with cash and doesn't need an immediate return on its investment, and because California subsidizes solar investments more than any other U.S. state.


Google's Sergey Brin and Larry Page plugging in a RechargeIT hybrid electric car.

What makes the Google solar project more than a green publicity stunt is that it is part of a larger energy strategy that encompasses conservation efforts. The company plans to continue improving the energy efficiency of its densely packed, power-intensive data centers, home to hundreds of thousands of servers worldwide. It has been installing more efficient lighting and building control systems in all its corporate locations. And it runs a biodiesel-fueled shuttle service for its employees at the Mountain View location. Google also has situated its new Oregon data center precisely where hydropower is cheap and abundant.

Reaching out beyond the problem of its own energy issues, Google this year joined forces with Intel and a number of other technology heavy hitters, including Dell, Lenovo, Microsoft, and Pacific Gas & Electric, to launch the Climate Savers Computing Initiative, a consortium that plans to set new efficiency targets for computers, which are notoriously wasteful of energy. Another industry consortium called Green Grid is pursuing similar goals.

Google says it wants to be carbon-neutral by 2008, a huge challenge given the power-hungry nature of its businesses. But if Google and the Microsofts and Wal-Marts of the world continue their concerted efforts to save energy and invest in renewable sources, and government policies continue to encourage them to do so, this could be the start of something big.

So whatever you think about the economic logic of solar energy or the wisdom of subsidizing renewable energy—or about the fact or fiction of climate change, for that matter—you've got to give Google and the rest of the participating corporate giants credit for taking on this important work. Maximizing energy efficiency and pursuing workable alternatives to fossil fuel–based energy isn't just green. It's good engineering sense. ■

MacCready's Last Flight

Paul MacCready, prolific inventor of human- and sun-powered machines, died in August. He built the *Gossamer Condor*, the first practical human-powered flight machine; the *Gossamer Albatross*; the *Gossamer Penguin*, the world's first successful solar-powered airplane; and the *Solar Challenger*, which awakened the world to the possibilities of solar energy. In 1981 the *Challenger* flew 262 kilometers, from France to England. In 1987, MacCready's team also designed the solar-powered Sunraycer car for General Motors, to compete in the Solar Challenge, the first competition to cross Australia from Darwin south to Adelaide. The Sunraycer won, and its success led MacCready to work on the EV-1 line of electric-powered cars for GM.

Contributors to a special section for condolences on the Web site of his company, AeroVironment, refer to him as a cherished inspiration. "More than anyone I know, he was aware of the dangers we all face due to environmental abuse, and he was aware of the possibilities for solving these problems," says Wally E. Rippel, principal power electronics engineer at Tesla Motors in San Carlos, Calif., who worked with MacCready on the EV-1 project. "It is my desire that people will remember him not just for his aeronautical accomplishments but also for his environmental vision and achievements. May others follow passionately in his footsteps." We couldn't agree more. ■

GOOGLE

NEWS


JED JACOBSOHN/GETTY IMAGES

Helmets Sense The Hard Knocks

Wireless device will let coaches pull football players before they suffer brain damage

In the United States' National Football League, SUV-sized men are paid astronomical sums to delight stadium crowds with their ability to run down and demolish the opposing teams' ball carriers. The automotive analogy is quite apt. When they collide, the forces that thickly muscled behemoths such as the San Diego Chargers' Shawne "Lights Out" Merriman exert on each other regularly exceed 100 times the force of gravity—the kind of jarring that passengers experience in a car crash. The result is roughly 230 000 concussions among professional, college, and youth football players each year.

Concern is growing over the long-term effects of skull-rattling tackles where a brain injury occurs, but the signs—including headache, nausea, and short-term memory loss—are difficult for coaches and trainers to spot; the injuries are unlikely to be reported by players because of the gladiator mentality that makes them keen to shake off any injury and get back into the game.

In the absence of hard medical data for assessing the severity of a player's head injury, coaches and trainers have to wrestle with tough-to-answer questions: When should a player sit out the remainder of a game? The remainder of the season? It's still a judgment call. But now a device installed in a player's helmet, which measures each blow to the head and reports the force of the impact, could make it a simple question

CRUNCH: Accelerometers in helmets [inset] measure the potential of hits like this to cause brain injury.

of physics. The device would provide, for the first time, a data set large enough for researchers to understand when a concussion is likely to have occurred, according to how hard a hit was and to which part of the helmet. Armed with that knowledge and real-time data from the helmets, sideline staff will know when a player should be brought off the field for neurological tests.

The Head Impact Telemetry (HIT) System developed by Simbex, of Lebanon, N.H., combines six strategically positioned MEMS accelerometers, a temperature sensor, a wireless transceiver, nonvolatile onboard memory, and a nickel-metal-hydride battery pack in a halo that fits inside a helmet. The package allows the spring-mounted accelerometers to sit right up against a player's head, so the movement of the skull is measured instead of the movement of the helmet itself. It adds just 170 grams to a 1- to 2-kilogram helmet's weight and does not significantly alter its fit.

The system automatically generates a data report when any single sensor detects an acceleration that exceeds 10 gravities. The report includes 12 milliseconds of data from before the system is triggered and 28 ms following the instant of impact. A controller connected to a laptop on the sidelines receives this information wirelessly.

An early version of the HIT system was introduced during the 2003 football season: four football players from Virginia Polytechnic Institute and State University, in Blacksburg, Va., were monitored through 35 practices and 10 games. Researchers recorded roughly 3300 head hits and found that, on average, players endured 50 impacts strong enough to trigger the system during the course of a single game. The average acceleration caused by those hits was 40 *g*'s per blow, the same level of impact delivered by the gloved fist of a professional boxer. At least twice a game, the players took shots to the head with forces on the level of a car crash. The data reports include parameters such as the Gadd Severity Index (GSI), a method developed by automobile crash researchers for describing just how jarring a blow someone has received. A human head can withstand GSI values as high as 1000 without serious injury; the blows endured by the Virginia Tech players ranged from 1 to 1599.

Simbex, working with researchers at engineering and medical schools at Virginia Tech and Brown University, in Providence, R.I., improved the communication system, allowing as many as 64 players to be monitored simultaneously with a single controller on the sideline. In 2004, nearly 500 players at five colleges—including Virginia Tech; the University of North Carolina, Chapel Hill; and the University of Oklahoma, Norman—wore the device in their helmets throughout the season. Simbex says 15 schools will take to the field this fall with the HIT System, capturing a snapshot of the forces at play in each tackle.

With a large data set from on-field collisions, "we may be able to develop predictive algorithms, using a player's impact history, to remove players before they get seriously injured," says Simbex's director of engineering, Jeffrey J. Chu. Duke University neurologist Joel C. Morgenlander, who recently joined the NFL's brain injury committee, thinks this is a good idea. "The medical decision has to be separated from the heat of the moment," he says.

Stefan Duma, a professor of mechanical engineering and director of Virginia Tech's Center for Injury Biomechanics, says getting even greater numbers of players outfitted with the device is important because there is "great potential for prevention of brain injuries, but the lack of scientifically sound, evidence-based studies is a barrier to improved prevention and treatment." The data already collected have been enough to demonstrate that, contrary to conventional wisdom, most impacts that result in concussion occur on the side of the head rather than the front or the top. Riddell Sports Group, in Rosemont, Ill., which makes helmets for the NFL and dozens of college teams, has already responded with a new helmet design, called the RevolutionT, which extends farther down the jawline and includes the same kind of shock-absorbing padding near the jaw that is found in the helmet's crown. Previously, the only purpose for the padding on the side of its helmets was to improve the fit.

Asked if the NFL had any plans to use the HIT System or to make Riddell's new helmet a required part of every player's uniform, league spokesman Greg Aiello noted that each player gets to choose the type of helmet he wears. As for the monitoring system, Aiello says the league's Committee on Mild Traumatic Brain Injury (MTBI) is currently reviewing it but would offer no details regarding when the review would be completed or whether the group had reached any preliminary conclusions.

For its part, the NFL has been sponsoring studies aimed at quantifying on-field collisions using methods other than the HIT System. Chris Withnall, a senior engineer at Biokinetics and Associates, in Ottawa, conducted studies on behalf of the NFL that used video footage from games to recon-

SHOCK ABSORBER: The inside of this helmet is outfitted with a wireless sensor system by Simbex.

SIMBEX (2)

NEWS

struct tackles using sensor-laden crash-test dummies [see "Anatomy of a Crash-Test Dummy," elsewhere in this issue]. Withnall and his colleagues concluded that, on average, players diagnosed with concussions had their heads suddenly whipped in one direction, with acceleration greater than 80 *g*'s. They found that blows below that threshold were much less likely to result in a brain injury.

Epidemiological studies suggest that there is a link between football-related concussions and subsequent memory problems and other brain dysfunction. But based on studies produced by its own MBTI committee, the NFL remains adamant that there is no such link and no proof that a player who has suffered a concussion is at much greater risk of subsequent brain injuries.

Informal evidence is cropping up, however, that contradicts the NFL's stance. Former NFL players who suffered multiple concussions during their playing days have begun reporting signs of memory loss, slurred speech, depression, and other signs of early-onset dementia. In November 2006, Andre Waters, a 44-year-old former Philadelphia Eagles player who suffered from depression, committed suicide. An autopsy revealed that his brain had deteriorated to an extent comparable to that of an 85-year-old with Alzheimer's disease. The pathologist who conducted the autopsy attributed the damage to repeated blows to the head over the course of the player's career. **—WILLIE D. JONES**

Fairchild Turns 50

This month Fairchild Semiconductor celebrates 50 years in the business. Executives and engineers from Fairchild founded many of the most influential technology firms in Silicon Valley, including microprocessor rivals Intel and AMD, reconfigurable chip leader Xilinx, and one of the best-known venture capital firms, Kleiner Perkins Caufield & Byers. The company was founded by the "Traitorous Eight"—a group of engineers who abandoned William Shockley's semiconductor firm en masse. Among the eight were Gordon Moore, whose eponymous law has been a guiding force in the chip industry; Robert Noyce, the co-inventor of the integrated circuit; and Jean Hoerni, the inventor of the process that made silicon the dominant semiconductor [look for a profile of Hoerni in our December issue].

Many of the companies these eight and others from Fairchild founded are still going strong, while some have been acquired by larger firms, and some have simply faded away. Fairchild itself was purchased by National Semiconductor, of Santa Clara, Calif., in 1987. Ten years later it was spun out as an independent company, focused on power-related chips and headquartered at one of Fairchild's original manufacturing sites in South Portland, Maine.

—SAMUEL K. MOORE

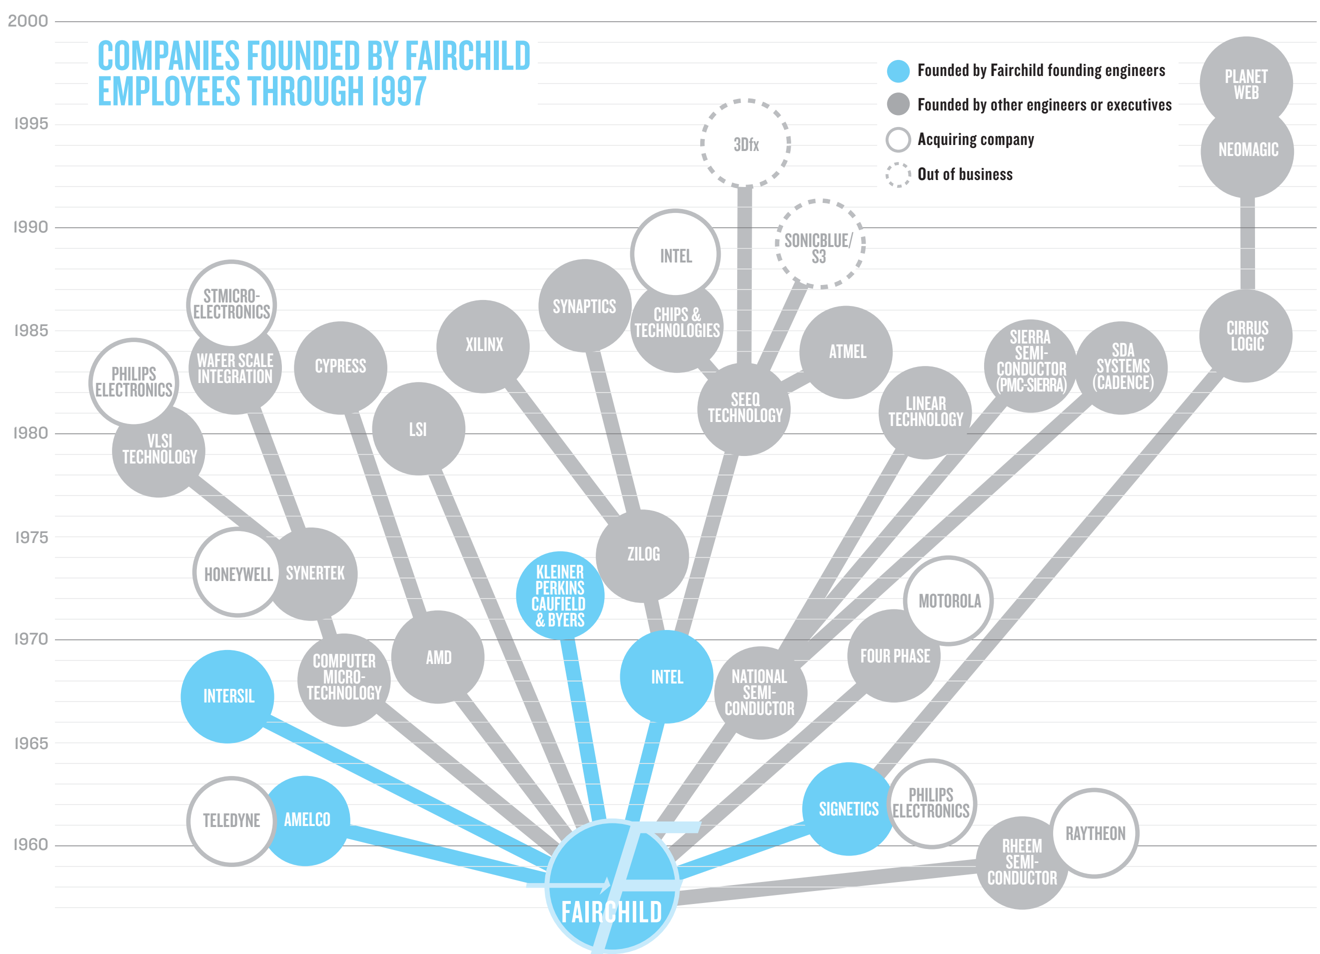

SOURCES: FAIRCHILD SEMICONDUCTOR, COMPUTER HISTORY MUSEUM

BRYAN CHRISTIE DESIGN

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

COMMENTARY

The iPhone's Missing Buttons

Valuable real estate on the main screen won't stay vacant for long

BY STEVEN CHERRY

When I look at the main screen of the iPhone, I see something missing: four buttons. In the main area, there's a 4-by-3 grid of 12 buttons: SMS, Calendar, Photos, and so on. Along the bottom, there are four buttons that give pride of place to the most essential services: Phone, Mail, Safari (Web), and iPod. In between, there's room for four more buttons, turning the top portion of the screen into a 4-by-4 grid.

Those nonexistent buttons are surely worth hundreds of millions. If you're AOL or Microsoft, can you really afford not to be on the iPhone's main screen? After all, the iPhone had one of the most successful product launches in history: Apple sold 270 000 in just the first two days (29 and 30 June).

Early last month Apple cut the price of an 8-gigabyte iPhone from $599 to $399. (At the same time it introduced a WiFi-enabled touch-screen version of the iPod). But the iPhone is still thought to be highly profitable. According to a teardown analysis by market research firm iSuppli Corp., in El Segundo, Calif., the hardware costs for the 8-GB version are only US $266. There may be other expenses, such as royalties or shipping, but there may also be other revenue—big revenue. An analyst for the Minneapolis-based securities firm Piper Jaffray has opined that Apple gets a commission of $11 per month from AT&T for each iPhone contract. Apple expects to sell 10 million phones by the end of 2008, so that would come to $1.32 billion in 2009 alone.

Though the company won't say, I think that there's even more revenue coming to Apple, and it's contained in the value of those home-screen buttons. Right now, Google Maps is one of the honored 12. But the label doesn't say "Google Maps," it says "Maps." Imagine the CEO of AOL telling Apple's Steve Jobs, "We're spending $25 million to make our MapQuest service better than Google Maps, and we're willing to give you another $25 million to put us there instead of them." After all, mapping is a major gateway to shopping and spending. Where's the nearest department store or gas station or fast food outlet? Ask your phone. The age of mobility is still in its infancy, and some of these applications haven't been written yet. But as cellphone users come to rely on always-on data connections as they move around, such services will become increasingly important—and lucrative.

Ian Lao, a senior analyst at Scottsdale, Ariz.–based In-Stat, a market research firm that specializes in telecommunications, agrees. "There's a lot of value locked up in each of those buttons," he says. "But it's hard to put a value on them. It's easily over $10 million."

We saw a hint of the iPhone's potential as a gateway for e-commerce with September's announcement of a new feature. When the iPhone, or its new sibling, the iPod Touch, comes within Wi-Fi contact of a Starbucks, an additional button, labeled iTunes, will appear on the main screen. Pressing it instantly downloads the song currently playing in the coffeehouse.

Apple's control over the iPhone's 12—or 13 or 16—buttons is controversial in two ways. The first is reminiscent of the Microsoft antitrust case. A key point of dispute was that Microsoft gave center stage to its own Web browser, Internet Explorer. That positioning was generally thought to be the single biggest reason IE's market share quickly overtook all its competitors, including one-time browser giant Netscape Communications Corp. Regulators in the United States and especially in Europe fought to force Microsoft to divorce the browser from the operating system and keep it from being the default Windows browser.

APPLE

ROOM FOR MORE? The iPhone looks as if it could fit another row of buttons. They'd surely be worth tens of millions of dollars each to companies providing online services, like Yahoo and Google.

The iPhone will never have anything like the market share Windows has. In fact, Apple's 10-million-unit goal amounts to a mere 1 percent of the 1-billion-unit global cellphone market. All the same, wouldn't it be better for consumers—to say nothing of AOL—if the Maps icon were neutral with respect to mapping services, and similarly for other buttons? When you first get the phone, you would go through a series of questions. Do you want Google Maps, MapQuest, or an application of your choosing that you can download? Should your Weather button lead to The Weather Channel instead of AccuWeather? How about a user-defined button that you could assign to AOL's Moviefone on your iPhone but that your daughter could assign to Facebook on hers?

The second controversy surrounds what's called network neutrality. The issue is whether a telecom carrier like Verizon Communications, in New York City, should be allowed to discriminate among the different data packets that cross its network. For example, given that Verizon offers cablelike television services to its customers, it might want to cut off or impede the delivery of data from companies that compete with it in providing movies on demand, such as Netflix or Blockbuster. Or Verizon might, in exchange for a hefty sum from Google, speed the results of Google searches but not Yahoo searches. Yet Apple is already favoring Google for its maps service as well as Starbucks and Apple's own iTunes music service.

To be sure, achieving the equivalent of network neutrality on the iPhone runs up against a special complication, because these buttons often lead to services that are customized for the device's unique touch screen, which is larger than the display on any other PDA-like phone. Still, if the iPhone continues to be a big success, a service provider like AOL would probably be happy to customize MapQuest and Moviefone for it. We can expect companies to start clamoring for this kind of neutrality by January, when the MacWorld conference is held. Meanwhile, Apple zealots are already clamoring for iPhone 2.0, without, it seems, any regard for how many buttons it has. ■

NEWS

THE ARMY YOU WANT: Stanford University's Junior will compete in DARPA's Urban Challenge this month.

Cars Get Street Smart

No thanks to DARPA

"Robot Cars Drive Themselves!" Pretty grabby headline, right? Throw in a few million dollars, a lot of publicity, and you've got a great story. The TV footage is compelling: brightly colored vehicles without drivers, bristling with cameras and sensors, driving themselves over dunes, down rutted trails—and, late this month, through simulated suburbs and city centers, in the US $2 million DARPA Urban Challenge.

There's just one problem with the imagery: the technologies likely to win the Challenge—those expensive cameras and sensors—probably won't be the ones that let future passenger vehicles "drive themselves." Instead, automakers expect that cars of the future will pay less attention to the lay of the land and more to each other, informing other vehicles about what they're doing several times each second by transmitting data that cars today already gather, via cheap wireless transponders roughly equivalent to your US $40 Wi-Fi router.

DARPA's interests are not in replacing commuters but in providing new and better technology for waging war. The appeal of an autonomous tank or rocket launcher is obvious: without soldiers inside, the potential casualties are reduced to zero. And the Department of Defense is under a 2015 deadline for making 30 percent of the U.S. military's land vehicles autonomous.

The challenge is substantial. An autonomous military vehicle must negotiate every kind of terrain: sandy desert, muddy forest, and dense urban core. To a tank, everything is a potentially hostile obstacle. Aside from its own location, tracked via the Global Positioning System, it has to figure out where everything in its surroundings is in real time.

Passenger cars, on the other hand, operate in far more limited circumstances: they stay on roads, almost all paved. They have no need to hide themselves, operate stealthily, or attack other objects. (In fact, making themselves known leads to avoidance, and hence safety.) And there are 250 million vehicles in the United States alone, according to the U.S. Bureau of Transportation Statistics, so traveling at high speeds among many adjacent moving objects with constantly changing trajectories is crucial.

Modern cars are stuffed with microprocessors and electronic control units that process data from a huge variety of sensors in the engine, transmission, suspension, and other systems—and then deliver the right blend of performance, fuel economy, and safety. Already, many traction-control systems simply ignore what drivers ask the car to do if the actions would cause the car to skid. Their sensors, though, are limited to the mechanical phenomena the car itself is experiencing.

Several safety systems have now added environmental data to the mix. Adaptive cruise control, from Mercedes-Benz and others, is one. It uses radar to calculate the distance to the car ahead and that car's velocity and adjusts its own speed to maintain a safe distance at all times—braking right down to a standstill and then accelerating back to highway speeds.

Another is the Volvo Blind-spot Information System (BLIS), which scans the area around a car's rear corners with side-mounted cameras and alerts the driver if there's a vehicle present. A third is Infiniti's Lane Departure Warning system: it calculates the edges of the lane from images captured by a video camera behind the windshield and alerts the driver if the vehicle is about to drift too far.

All of these systems still presume a vehicular environment that's mute. And that's one thing that will change over the next 10 years. Several initiatives around the world are considering standards for vehicle-to-vehicle (V2V) communications, in which new cars would be fitted with low-cost, short-range wireless transmitters. They would continuously alert sur-

ALL STOP: Vehicle-to-vehicle communication like General Motor's V2V technology keeps cars from crashing.

TOP: STANFORD RACING TEAM; BOTTOM: GENERAL MOTORS

rounding vehicles (as well as elements of the highway infrastructure) to the vehicle's trajectory, the driver's actions, perhaps even the car's ultimate destination. The infrastructure, in turn, would alert cars to accidents, congestion, speed zones, vehicles nearing crossroads, and other conditions ahead.

But how many cars and signposts must communicate to make a difference? Larry Burns, head of R&D for General Motors, says the company's modeling shows safety benefits even with less than 10 percent of vehicles outfitted with transponders. That number might be reached just by retrofitting all rental cars and other large fleets, while new cars with factory-installed transponders gradually raise the overall ratio. Even if only every 10th car is communicating, a vehicle in the fast lane might "hear" that someone just slammed on the brakes 15 cars ahead and start to slow well before the driver can see or react to the braking car ahead, says Burns.

A group of German automakers and component suppliers, along with Deutsche Telekom and several government ministries, is now writing requirements for a test in Hesse, in Germany's Rhine-Main region, that will equip more than 500 vehicles with transponders. Cars will communicate with each other but also with roadside units linked to central traffic control computers. The goal of the project, called SIM-TD, is to get real-world experience, including data that will help settle questions on what information is most useful. For example: Is a vehicle's relative trajectory adequate, or should it also transmit absolute position data from the navigation system? The first vehicles in this project are expected to hit the roads in 2009.

Five large automakers in North America—Chrysler, Ford, GM, Honda, and Toyota—are now defining a V2V message set; they plan to equip 50 vehicles and 20 intersections with communications technology and start on-road testing late next year. Similar efforts are underway in Japan as well.


LOOK OUT! A GM system warns of cars you can't see.

Roughly a year ago, GM offered journalists a glimpse of the potential of its V2V research in a demonstration held at Camp Pendleton, Calif. Each of us was asked to drive a Cadillac sedan at 40 miles per hour (64 km/h) toward another Cadillac stopped ahead in the same lane. As the distance narrowed, a colored indicator on the dash turned from green to amber to red. A warning tone sounded steadily louder and faster as the cars calculated that a collision was imminent.

Shortly after the indicator turned red, the moving car braked itself, staying in lane and coming to a halt just a few car lengths behind the stopped vehicle. Forward-looking radar can do the same thing, of course. But the cost of radar transponders and image-processing software and circuitry is far greater than that of a short-range wireless transponder incorporated into each car—transmitting data already gathered by existing in-car sensors.

The great promise is that one day, a vehicle might—if the driver chooses—even drive itself autonomously. Sure, everyone loves driving down country roads on sunny Sundays. But suppose your car could handle the heavy parts of that grinding, stop-and-go, 40-km suburban commute while you answered e-mail or concentrated on that conference call. The car would speed up, slow down, and choose its routes to minimize fuel usage and emissions. What's more, it would keep traffic flowing smoothly and enable more cars to occupy the limited road space, adding freeway capacity without the need to lay more concrete.

We're a long, long way from that point, of course. But one thing is clear: even if DARPA gets its robotic vehicles in time to meet the Defense Department's 2015 deadline, the car companies that serve everyday drivers won't be adapting military technology for civilian use. **—JOHN VOELCKER**

Quantum Tunneling Creates Fast Lane For Wireless

Star Wars on your iPod in 2.5 seconds

Until recently, a truly wireless existence was beyond what silicon circuits could offer. The bands of the radio spectrum, such as Wi-Fi, that they could reach were too narrow to connect a high-definition TV to a high-definition DVD player. The chips that could do the job, made with exotic semiconductors, were too expensive for consumer electronics. But in the last two years, silicon circuits finally broke into the 60-gigahertz band, which has been shown to allow data-transfer rates of 5 gigabits per second over a distance of 5 meters.

Sixty-GHz radios, based on silicon or silicon-germanium chips, are expected to be integrated into TVs, set-top boxes, and other media-linked devices starting in 2009. But a new dark-horse candidate has emerged that claims to be able to make cheap 60-GHz technology without using any semiconductor materials at all—silicon or otherwise.

Boulder, Colo.–based Phiar Corp. (pronounced "fire") uses a proprietary mix of insulators and metals to achieve quantum tunneling, which lets electrons zip through devices in mere trillionths of a second. "We're at a tipping point," says Adam Rentschler, Phiar's director of business development. "Metal insulators are the first viable alternatives to semiconductors since the era of vacuum tubes."

Normally you can imagine an electron as a ball and an insulator as a high hill. Given enough energy, the ball will make it over the hill; this is how an electron punches through insulation. But when the hill—the insulator—is only a couple of atomic layers thick, the rules of classical physics no longer apply: instead of a ball, the electron looks more like a wave. (The wave is actually the function that defines the probability of finding that electron in a specific place.) This wave is wider than the very narrow hill, so it stretches from one side of the insulator to the other. As a result, sometimes the electron simply appears on the other side, having "tunneled" through the insulator. Tunneling happens all too often in the transistors of modern microprocessors, and is a serious problem [see "The High-*k* Solution," else-

GENERAL MOTORS

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

NEWS

where in this issue]. But Phiar took advantage of the phenomenon by finding a way to make electrons tunnel easily in one direction and not at all in the other. The result is a diode made up of two very thin insulators sandwiched between two metal layers.

From the electron's perspective, Phiar's proprietary blend of metals and insulators makes the energy barrier between the metals thinner in one direction than the other. That happens because when the electron travels in the "easier" direction, the interface between the two insulators creates a quantum well, a structure that confines electrons in two dimensions. The quantum well sits at the halfway point between the two metals, and once the electron has tunneled that far, the quantum well boosts the chances that it will make it through the second insulator. But the quantum well appears only when the electron is being pushed in one direction, called the forward bias [see illustration]. When the electron is being pushed in the opposite direction, reverse bias, there is no well, and without it, the insulation is too thick to tunnel through.

The metal-insulator-insulator-metal (MIIM) structure makes electrons "10 billion times more likely to tunnel," says Rentschler. Quantum tunneling lets an electron traverse a device in just 1 femtosecond, thousands of times as fast as an electron traveling through a typical semiconductor transistor. "Semiconductors are called *semiconductors* because an electron doesn't move through them very well," he says. "It spends a lot of time bumping around through a slow, molasses-like atmosphere." Rentschler says the devices are so fast that the 60-GHz band is the lowest frequency the company is interested in and that its circuits have been clocked at 110 GHz.

Phiar contends that its devices can be manufactured directly on top of a CMOS chip, potentially making them a simple addition to an already inexpensive technology. In fact, Rentschler maintains that Phiar devices can be fabricated on almost any substrate. "Tree bark is probably a really bad choice," he says, "but we can pretty much deal with anything."

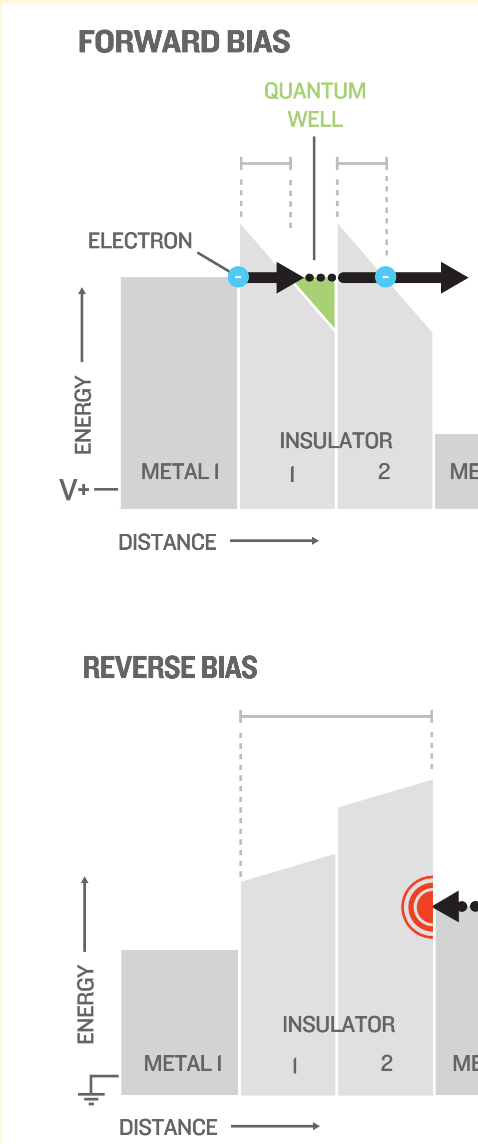

QUANTUM TUNNELING: Phiar Corp.'s gigahertz-frequency diodes rely on the phenomenon of quantum tunneling. In one direction, a quantum well forms between two insulators. An electron that tunnels the short distance to the well can easily tunnel the rest of the way.

In the other direction, no quantum well forms, and the distance is too great for an electron to tunnel through, so no current flows.

BRYAN CHRISTIE DESIGN

Tree bark aside, the company is hoping to build its circuits on a variety of substrates in an effort to get around one of the peculiarities of the 60-GHz band. Radiation at that frequency is absorbed by oxygen in the atmosphere. To get a strong enough connection between, say a video player and an HDTV, the electronics must communicate in a tightly focused, highly directional beam. In some ways this attribute makes 60 GHz perfect for a wireless personal area network—you can beam your home movie from your DVD player to your TV without worrying about your neighbors watching, too. But it also may mean that people walking through the beam can disrupt the link.

Phiar's Rentschler says that instead of being confined to a single chip in a single location in a computer or TV, his firm's antennas and transceivers can be distributed in tiny strips all over a consumer device, allowing it to pick up a signal from any direction.

Of course, other radio developers, such as the leading 60-GHz silicon firm SiBeam, in Sunnyvale, Calif., have solutions for the band's directionality problem, too. SiBeam chief technical officer Jeffrey Gilbert says his company's RF chip looks for the receiver, and if it's not directly in sight, the receiver figures out the best path to ricochet the signal—off a wall or the floor—to get to the target.

NOT EVERYONE is convinced that Phiar's technology will make it into upcoming consumer devices. "I would personally be very skeptical of anyone saying they will put semiconductors out of business anytime soon," says John Cressler, a professor of electrical engineering at Georgia Institute of Technology, in Atlanta, who studies the 60-GHz band. Though Phiar's tunneling-based approach will produce fast circuits, he says, silicon-germanium and CMOS chips have the advantage at 60 GHz if for no other reason than that they are already nearing production.

IBM's Brian Gaucher agrees. His company is developing silicon-germanium-based 60-GHz chips. "I don't doubt the device physics," says the Yorktown Heights, N.Y.–based research staff member, "but I think that traditional silicon, due to its maturity, is the technology that will likely be leveraged to enable the HD-multimedia revolution."

Behzad Razavi, an electrical engineering professor at University of California, Los Angeles, adds that while Phiar has demonstrated an ability to use quantum tunneling to make fast devices, it is not clear whether those fast devices can be integrated into actual applications. "Everything they have in their products is single devices," he says. "Diodes, transistors—these are single components." In an integrated system, countless components running at a 60-GHz frequency must work together flawlessly. But the metal interconnects between them are extremely difficult to manage at such high frequencies. "You have these little components, each good for 60 GHz, but the wires have introduced their own problems," Razavi says. **—SARAH ADEE**

THE BIG PICTURE

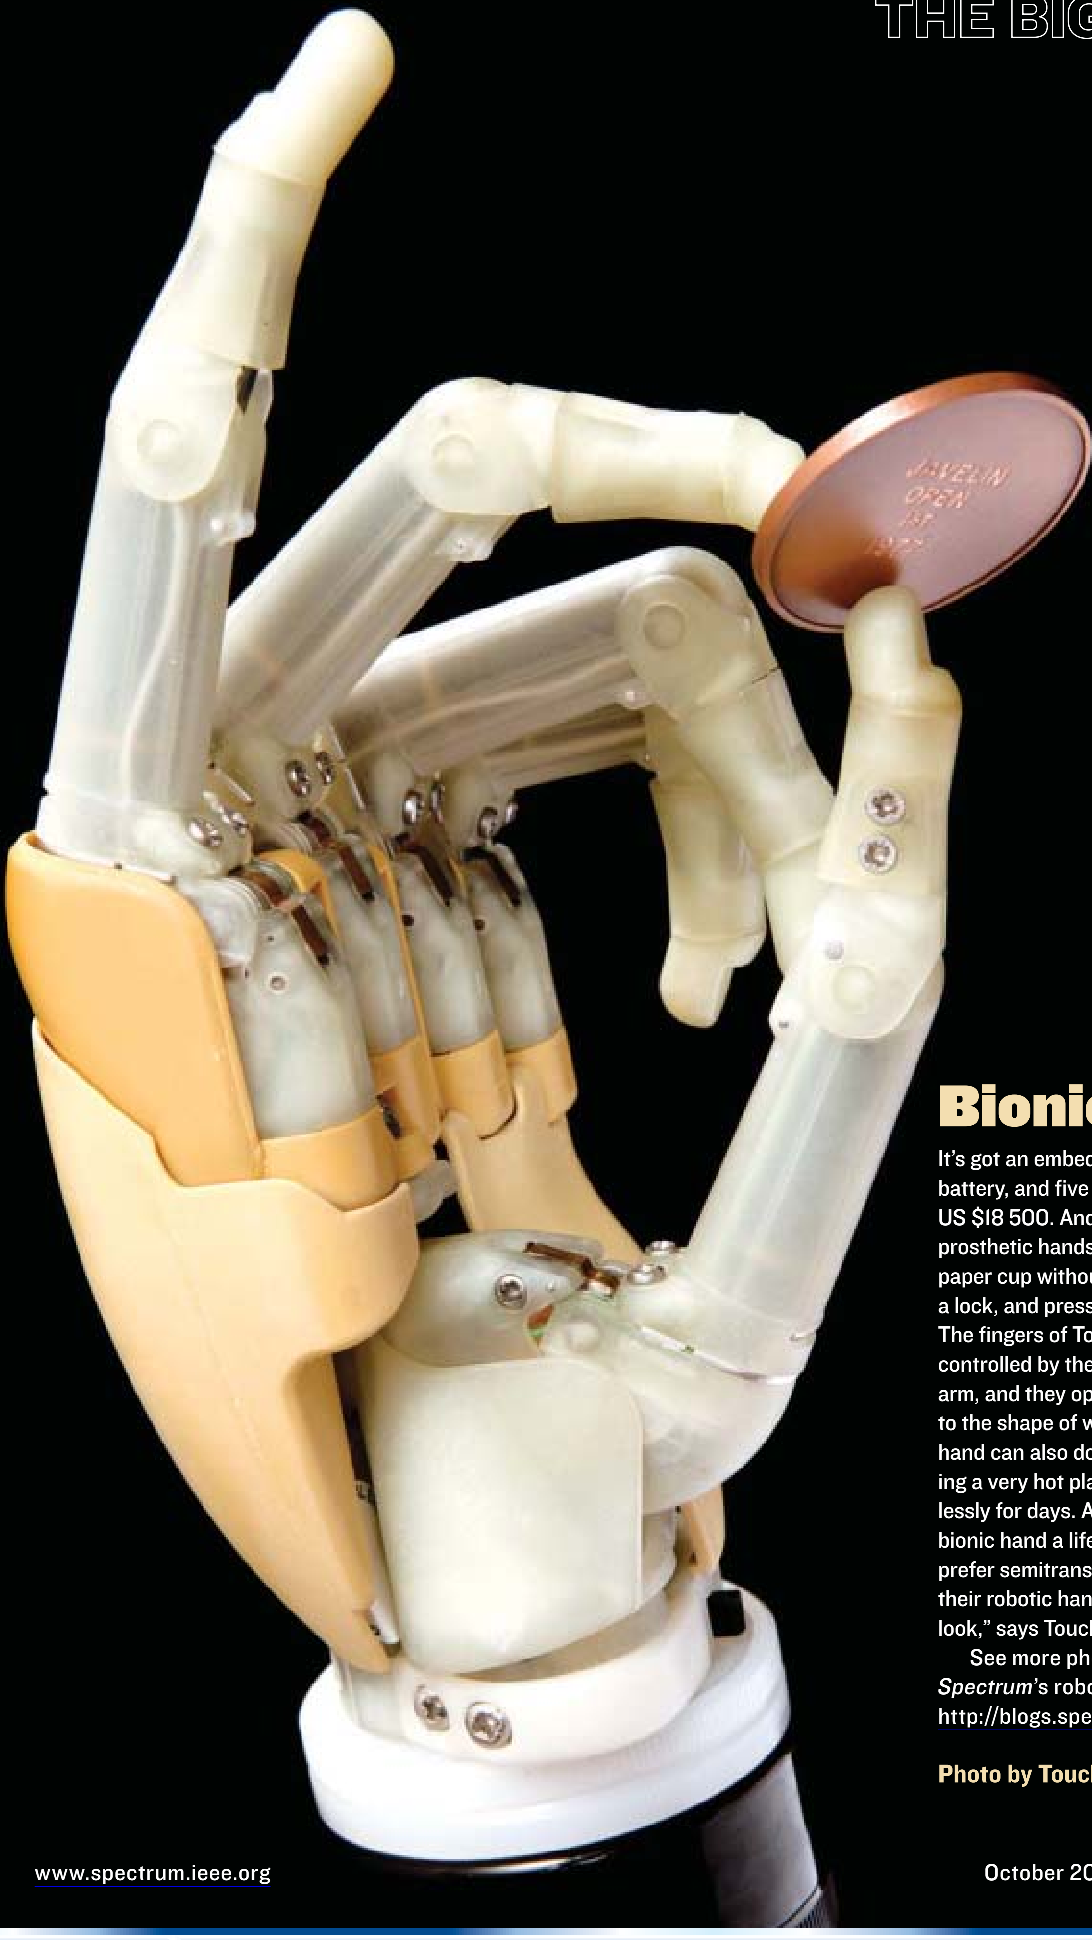

Bionic Hand

It's got an embedded computer, a rechargeable battery, and five small dc motors. It costs US $18 500. And it can do things most other prosthetic hands just can't, like grabbing a paper cup without crushing it, turning a key in a lock, and pressing buttons on a cellphone. The fingers of Touch Bionics' iLIMB Hand are controlled by the nerve impulses of the user's arm, and they operate independently, adapting to the shape of whatever they're grasping. The hand can also do superhuman tricks, like holding a very hot plate or gripping an object tirelessly for days. A skin-tone covering gives the bionic hand a lifelike look, but some customers prefer semitransparent models, to proudly flaunt their robotic hands. "They like the Terminator look," says Touch Bionics CEO Stuart Mead.

See more photos and videos at *IEEE Spectrum*'s robotics blog, Automaton, http://blogs.spectrum.ieee.org/automaton.

Photo by Touch Bionics

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

THE GREENING OF GO

Google's Robyn Beavers

ENERGY

OGLE

CORPORATE ROOFTOPS ARE THE LATEST FRONTIER IN SOLAR ENERGY GENERATION

By Sandra Upson

IT'S ANOTHER BRILLIANT DAY at the world headquarters of the hottest company on the planet. Some shirtless employees are playing a lunchtime game of volleyball while others stride across campus with laptops tucked under their arms. The place fairly crackles with energy, and in more ways than one.

Up here on a roof at Google's leafy and sprawling Mountain View, Calif., campus, with the shouts of the volleyball game just barely audible, sunlight glints off 9212 polysilicon solar panels stretching out toward the horizon. Amid the irregular jumble of angular roofs, a single south-facing wave stands out, a pitch and roll frozen in place against a backdrop of foothills.

Today, like most days, the panels will generate 9000 kilowatt-hours of electricity before the sun fades into a fat orange ball and disappears into the Pacific. All are connected to Mountain View's section of the electricity grid. The solar modules blanket virtually all the free roof space on the eight buildings at the center of the Googleplex. Even part of the parking lot is covered: two rows of carports, shaped like miniature gas stations, support yet more panels. When the last building is fully connected, by the end of this year, the panels will produce 1.6 megawatts of electricity. It'll be enough to satisfy 30 percent of the buildings' peak demand or power a thousand California homes.

Google's project is the largest corporate installation of solar panels in North America. It has grabbed headlines since Google announced it a year ago. That said, it isn't even in the worldwide top 10 of roof-mounted solar projects. A handful of factories in Germany and Japan take that honor, as well as a couple of roofs in Spain and the Netherlands. At the very least, the search giant's solar play adds one more country to the list of star performers in the world of commoditized sunshine. And it seems clear that Google's array won't be tops in North America for long.

After languishing through much of the 1990s, the market for photovoltaic installations in the United States and several other countries took off about five years ago, and it's now increasing by 40 percent annually in the United States alone. Spain's bullish market grew 100 percent in the past year. And percentages never tell the full story, as Noah Kaye, a spokesman for the Solar Energy Industries Association (SEIA), points out. "The German market was relatively flat in the past year, but Germany still installed more [photovoltaics] than the U.S. did," says Kaye, on behalf of the trade and lobbying group.

California has nonetheless become the second-fastest-growing solar market in the world, and that surge, especially in the United States, is being driven mainly by activity on corporate rooftops. Travis Bradford, president of the nonprofit Prometheus Institute for Sustainable Development, in Cambridge, Mass., calls corporate attention to solar power "an exploding interest." In 2006, the commercial sector accounted for 60 percent of newly installed capacity in the United States, up from 13.5 percent in 2001, according to data from the U.S. Department of Energy.

"We've stopped reporting the biggest systems," Bradford adds. "A new record is set every few months."

In March, Applied Materials of Santa Clara, Calif., announced a plan to install 1.9 MW of solar power on the rooftops of its Sunnyvale, Calif., complex. And it's not just high-tech titans retooling their roofs: Tesco, the British-based supermarket chain, says it intends to put up a 2-MW solar installation at an office complex in northern California. Wal-Mart, the world's largest retailer, intends to outshine all these companies with multipart plans to put more than 5.6 MW's worth of solar panels on the roofs of 22 stores in California and Hawaii. Two other discount-retailing giants, Target and Kohl's, have also begun

PHOTOGRAPHS: JOSON

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

Google Goes Green

Solar panels are just one of the many climate-change mitigation projects that Google is spearheading.

Plug-in hybrid cars. Google built carports in its parking lot that have solar panels on their roofs. Five outlets dangle down, all of which charge cars on a daily basis. The company is also accumulating a fleet of hybrids as part of RechargeIT, a Google program aimed at speeding the adoption of plug-in hybrid electric cars.

Climate-Savers Computing Initiative. Google partnered with Intel and 20 other companies to create a large industry coalition that will adopt strict energy-efficiency targets for IT equipment.

Upgrades to buildings. "We're changing lightbulbs, replacing air-conditioning equipment, and upgrading building systems to optimize when and how our electricity gets used," says Robyn Beavers, Google's corporate environmental programs manager.

transforming their roofs into tiny, independent utilities.

"It's not an illusion," says Craig Cornelius, program manager for the Department of Energy's solar division. "Corporate solar is really happening."

Amid the enthusiasm, it's important to keep this latest twist in the solar saga in perspective. Solar energy of all kinds fulfills less than 0.1 percent of electrical demand in the United States, and affordable, commercially available panels have hovered near 15 percent efficiency for years. Despite the recent burst of corporate enthusiasm, the prices of solar modules are expected to continue inching down at just 5 percent a year, and grid parity—the point at which solar panels can compete subsidy-free with utilities—isn't expected until 2015 at the earliest.

It's too soon to say whether these costly corporate installations will go down in history as the first of a limited series of impulsive, feel-good publicity moves by tech start-up billionaires, or as the beginning of a longer-term movement that will help sustain the market for solar photovoltaics during the next decade and enable solar to finally become cost-competitive. One thing is certain: the movement will flourish only to the extent it is nurtured by a complex patchwork of economic and bureaucratic conditions. Of the 9509 new grid-tied solar installations in the United States in 2006, which totaled 101 MW, 70 percent of them were in California. And that's not just because it's sunny. As it turns out, California subsidizes solar in a particularly generous way.

EVEN WITHOUT SUBSIDIES, solar panels may have found their logical home, at last, in the commercial world. The nice, flat roof design on most commercial buildings, unlike the pleasingly angular but less workable residential roof, is one obvious advantage. But some of the most compelling reasons are intangible.

"It's a key part of attracting and retaining employees," says Doreen Reid, a senior associate at The Climate Group, a London-based nonprofit that helps companies reduce their carbon emissions. "Students coming out of college are more conscious of a company's environmental image."

Robyn Beavers, Google's corporate environmental programs manager, confirms the transformative power of solar cells. "I've had so many people e-mail me and say, 'This is why I love working at Google' and, 'How can I install solar at home?'" says Beavers, who presided over the installation project. Google's solar enterprise is part of a larger mission to promote the growth of solar energy, she says. Google founders Sergey Brin and Larry Page have invested heavily in Nanosolar, a start-up that specializes in thin-film solar cells. (Both declined to be interviewed for this article.)

For its rooftops, Google chose Sharp modules capable of generating 208 watts each. The polycrystalline silicon cells average 12.8 percent module conversion efficiency. Because solar panels produce dc current, each system requires inverters to change the current into usable ac, and Google used a set of utility-grade inverters with an average of 96 percent conversion efficiency made by SatCon Technology Corp. of Boston. It partnered with EI Solutions, a solar project developer with headquarters in San Rafael, Calif., to do the electrical design work.

Google won't say how much the whole project costs, other than to indicate that it expects to recoup its investment in five to seven years. Nonetheless, experts estimate that a solar installation costs between US $3 and $5 per watt in California, and between $6 and $10 per watt in the rest of the United States after factoring in local and federal rebates for the cost of the system. (According to the Northern California Solar Energy Association, the average cost of installing large systems in the Bay Area in 2006 was $8.58 per watt before rebates, on par with national figures.) Using data from California's Solar Initiative program and based on a $2.80-per-watt incentive rate, Google likely retrieved about $4.5 million from California on a project that in total probably cost more than $13 million. Federal tax breaks through the Energy Policy Act of 2005 also help to burnish the appeal of what is still, for many, a prohibitively expensive system.

For other companies, an important piece has been added to the picture for solar. Where research and development have so far failed to slash the price of solar, clever financing schemes have filled the breach. Google's solar project, for all its trendy impact, was financed the old-fashioned way—with cash. But customers without billions of dollars in liquid reserves tend to shy away from such a move.

Rather than requiring that customers buy all the equipment for an installation, which can run into the millions of dollars, solar service providers are persuading customers to sign agreements that in effect turn those providers into miniature utilities.

The office-supply company Staples was among the first to pursue such a scheme—in 2004 with SunEdison, a prominent Beltsville, Md., solar electricity service provider, for a 280-kW installation on two of its California warehouses. The solar installation covered about 10 percent of the facilities' electric loads.

In this arrangement, SunEdison installs the solar modules on a customer's property and is responsible for maintaining them. But, crucially, it does not charge the customer for them. Instead, the customer signs a long-term agreement, usually lasting about 15 years, that locks the customer into buying back the electricity generated by those panels at a fixed rate. Typically, that rate is lower than retail utility prices. Prometheus's Bradford estimates that 40 percent of recent commercial installations have gone this route, and he says that it's likely to grow more popular as additional companies move into solar.

Through such long-term contracts, corporations are cushioning themselves from fluctuations in electricity prices, over both the long and short terms. In the short term, on-site generation lightens the customer's demand from the local utility precisely when the utility needs a break: during peak periods of demand, which are also, not coincidentally, the hours of the typical business day—yet another reason that corporate roofs make sense as hubs of solar activity. During these peaks, electricity prices can double, triple, or even quadruple. "Solar energy is generated when companies need it most, which is typically 10 a.m. to 5 p.m. on hot, sunny days," says Kaye, the SEIA spokesman.

Add to that executive-level nervousness about the prospect of restrictions on carbon emissions and the volatility of electricity prices, and some businesses are finding the case for solar compelling. "Companies can do all of the above," says Rick Whisman, the director of west region system sales at PowerLight, in Berkeley, Calif., a subsidiary of SunPower, one of the largest solar cell producers and installers in the United States. "They can make a wise decision on energy over the long term. They can reduce their footprint and also be prepared for regulations that may come into play in the future."

Whisman and others are quick to point out that solar makes sense only as a component of a larger plan. "What has made Google and Wal-Mart so noteworthy is the degree of thought that went into their planning," says Bradford. Indeed, pursuing electricity generation on-site is of limited value unless accompanied by a suite of energy-efficiency measures to reduce a company's overall demand [see "The Zero-Zero Hero," *IEEE Spectrum*, September]. Wal-Mart is the undisputed leader in driving the adoption of compact fluorescent lightbulbs, and the megaretailer hasn't stopped there. It is also modernizing its truck fleet to be more aerodynamic and fuel-efficient, and at a store in Texas, the company is testing sustainable design with an experiment that includes rooftop wind turbines and on-site recycling. Google, meanwhile, is pushing ahead with a broad package of ambitious environmental programs [see sidebar, "Google Goes Green"]. "This is just a first step for us," Beavers says.

WILL THE NEXT STEPS take Google beyond California? Beavers won't say. Google has 15 U.S. offices and several power-hungry data centers outside of California, as well as offices and facilities in 23 other countries. But few, if any, of those places offer the incentives California does.

Driven by those incentives, in the past five years Californians have edged above the threshold of 30-kW demonstration projects to larger systems, such as Google's, spawning a cottage industry of experienced local solar-installation service companies. The California Solar Initiative credits companies based on performance metrics that can amount to one-fourth the cost of the system, which—when combined with a federal tax credit on some solar equipment, and depending on the cost of the panels and the installation—can cover more than 50 percent of a system's total cost. In 2007, rebates in California evolved from being per-watt, based on system size, to a formula that takes into account details of the physical placement of the panels, so that systems that are expected to perform better will be reimbursed more generously. The new calculations factor in the panels' tilt and shading, as well as altitude and azimuth, which are the two coordinates commonly used to describe the sun's apparent position in the sky. The rebates are still part of a tiered system designed to reduce the incentives over time, and in 2008 energy-efficiency requirements will be tied to those rebate dollars. Only New Jersey, among the other 49 states, has shown a similar level of leadership in photovoltaics.

The question now is whether the movement can expand beyond a few isolated states and countries. Data from the Interstate Renewable Energy Council, a nonprofit that disseminates information on rules and incentives relating to renewables, suggest that Arizona, Colorado, Massachusetts, New York, and Texas are also promising markets. New Jersey, with its second-only-to-California inducements, has the second-highest installed solar capacity, with 18 MW in 2006. By comparison, Florida, the Sunshine State, installed a meager 170 kW of solar energy in 2006, the year that

UP ON THE ROOF: Solar panels blanket almost all of the available roof space at the heart of the Googleplex. New carports in the parking lot use solar energy to recharge employees' electric cars.

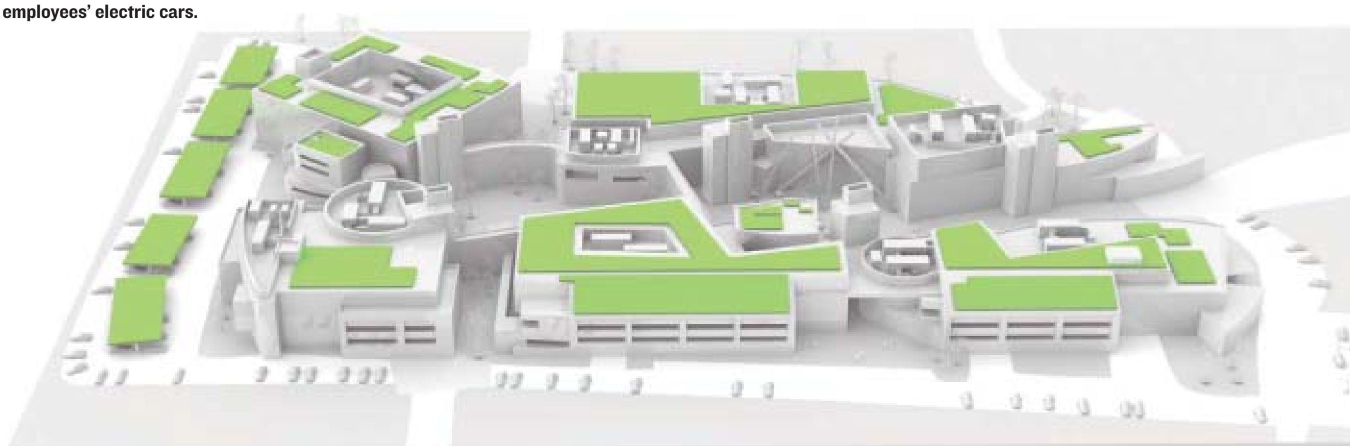

ILLUSTRATIONS: BRYAN CHRISTIE DESIGN

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

Photovoltaic Hot Spots

The countries where solar panels have fared well are not always in the sunniest places. In 2005, three countries accounted for 90 percent of the 3705 megawatts of installed photovoltaic capacity.

U.S. 479 MW (12.9% of world total) California contributed 30 percent of U.S. solar output.

GERMANY 1429 MW (38.6%) In 1999, a key program introduced large credits and, later, a feed-in tariff.

JAPAN 1422 MW (38.4%) In 1994, a subsidy for residential roofs covered 50 percent of total costs.

AUSTRALIA 60 MW (1.6%) The fourth largest producer, it's virtually tied with Spain (1.6%) and the Netherlands (1.4%).

SOURCES: IEA PHOTOVOLTAIC POWER SYSTEMS PROGRAMME, OKSOLAR.COM

its solar incentives program was launched. Its limited generation capacity speaks to the paltry nature of those incentives.

That disconnect between sunshine and solar output is even more pronounced outside the United States. The global leaders in solar energy, by virtually all metrics, are Germany and Japan. Both countries have sky-high electricity prices: on average 20 cents per kilowatt-hour in both Germany and Japan, double the average price of electricity in the United States, according to 2006 data from the International Energy Agency, in Paris. Starting in the mid-1990s, both governments began pouring money into renewable energy programs. As a result, today, in cloudy Germany, the renewable energy industry has become the country's second largest source of new jobs after the automotive sector. It employs some 200 000 people, according to Paul Runci, a senior scientist at Pacific Northwest National Laboratory, in Richland, Wash., who studies energy research and development trends.

But for corporations, the solar story depends on more factors than rates and rebates. In the United States, state-by-state rules on how to attach solar plants to the grid and how to compensate producers for electricity they export to the grid vary tremendously. The portion of the Google system that has the solar panels, for example, connects to a secondary grid, which does not accommodate excess power fed back into it, according to Johann Niehaus, the lead engineer on the project for EI Solutions. To account for that, the system was scaled both to fit the available roof space and to generate less than 50 percent of the buildings' minimum demand, so that the solar modules never come close to producing more electricity than the campus consumes. What they've installed approaches that 50 percent limit.

That is just one way that interconnection to the grid can be complicated. Other problems stem from what numerous experts have described as electric power companies' lack of familiarity with distributed generation. Some utilities have been reluctant to open up their grids to ever-larger quantities of electricity that they cannot manage. At times, a utility may declare that a generation system warrants an engineering study, which can cost up to $50 000, to analyze the impact of adding the system's electricity to the grid. That may not be prohibitive for Google, but for some prospective buyers it is. "They want to know, are you creating frequency disturbances? Voltage disturbances? How big are you in relation to the peak load on that circuit?" says Christopher Cook, a senior vice president for regulatory affairs and new markets at SunEdison. It's not unusual for a solar project to be killed because of the expense of commissioning an impact study. That's why, in one recent case described by Cook, a school in Virginia abandoned its plan to put solar panels on its roofs.

Some solar watchers have argued that those studies are sometimes unnecessary and redundant. The Interstate Renewable Energy Council and the Department of Energy, among others, are calling for guidelines to specify when they are needed.

"There are legitimate grid reliability concerns that can be addressed through technical means," Cornelius, of the Department of Energy, says. "And then there's the other reasons." Wal-Mart initially had hinted at a solar plan for its stores that would add up to 100 MW, but SEIA's Kaye suggested that the company had found it unfeasible in many states because of slow response rates from utilities. Such bureaucratic bunglings related to connecting to the grid led Wal-Mart to scale down the project to stores in just a few states, at least for now.

Cornelius points to Connecticut as an example of how things can go wrong for solar. Electricity prices there are among the highest in the country—15 cents per kilowatt-hour for the commercial sector in 2007, compared with a national average of 9 cents—and the state is full of congested distribution systems. Nevertheless, interconnection problems have not yet been formally ironed out. These have ranged from the painfully mundane, such as utilities not processing applications quickly, to unresolved concerns for the safety of a utility's distribution engineers. As a result, project developers are not prepared to invest in installations until they are confident the utility will agree to connect them hassle-free. "From the industry's perspective, we look at a state, and if it doesn't have interconnection rules, we say we can't do this project," Cook says. "It's not: let's go forward and see what happens. It's just simply: this state's not open for business."

All these bureaucratic problems are surmountable, analysts say, largely by means of new industry standards. Approved in June 2003, IEEE Standard 1547 represents one step toward integrating the technical side of interconnection practices for distributed power sources. "We had something like 3000 utilities, each with their own interconnection requirements," says Richard DeBlasio, the technology manager for the National Renewable Energy Laboratory's Distributed Energy and Electric Reliability program, in Golden, Colo., who chaired the committee that drafted the standard. Since then, at least half of the 50 states have adopted the IEEE standard as well as another, UL 1741, as their minimum technical guidelines for equipment and safety requirements. If the political will is there, the environment should improve for photovoltaic installations. In fact, it's already happening: in June, Oregon's legislature approved a set of policies that can amount to a 50 percent tax credit for solar installations and manufacturing. And the historic materials crunch that has held the prices of panels aloft is likely to abate as new manufacturing capacity comes online starting in 2008.

Although California may be alone in the United States in swaddling itself in polysilicon panels, the leadership of the Googles and Wal-Marts of the world could cause corporate solar installations to pop up on rooftops in the rest of the country in almost no time at all. ■

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

THE HIGH-*k* SOLUTION

MICROPROCESSORS COMING OUT THIS FALL ARE THE RESULT OF THE FIRST BIG REDESIGN IN CMOS TRANSISTORS SINCE THE LATE 1960S

BY MARK T. BOHR, ROBERT S. CHAU, TAHIR GHANI & KAIZAD MISTRY

From left: Ghani, Mistry, Chau, and Bohr of Intel with a wafer of 45-nanometer microprocessors

As you read this, two of our most advanced fabs here at Intel are gearing up for the commercial production of the latest Core 2 microprocessors, code-named Penryn, due to start rolling off the lines before the year is up. The chips, based on our latest 45-nanometer CMOS process technology will have more transistors and run faster and cooler than microprocessors fabricated with the previous, 65-nm process generation. For compute-intensive music, video, and gaming applications, users will see a hefty performance increase over the best chips they are now using.

A welcome development but hardly big news, right? After all, the density of transistors on chips has been periodically doubling, as predicted by Moore's Law, for more than 40 years. The initial Penryn chips will be either dual-core processors with more than 400 million transistors or quad-core processors with more than 800 million transistors. You might think these chips don't represent anything other than yet another checkpoint in the inexorable march of Moore's Law.

But you'd be wrong. The chips would not have been possible without a major breakthrough in the way we construct a key component of the infinitesimal transistors on those chips, called the gate stack. The basic problem we had to overcome was that a few years ago we ran out of atoms. Literally.

To keep on the Moore's Law curve, we need to halve the size of our transistors every 24 months or so. The physics dictates that the smallest parts of those transistors have to be diminished by a factor of 0.7. But there's one critical part of the transistor that we found we couldn't shrink anymore. It's the thin layer of silicon dioxide (SiO_2) insulation that electrically isolates the transistor's gate from the channel through which current flows when the transistor is on. That insulating layer has been slimmed and shrunk with each new generation, about tenfold since the mid-1990s alone. Two generations before Penryn, that insulation had become a scant five atoms thick.

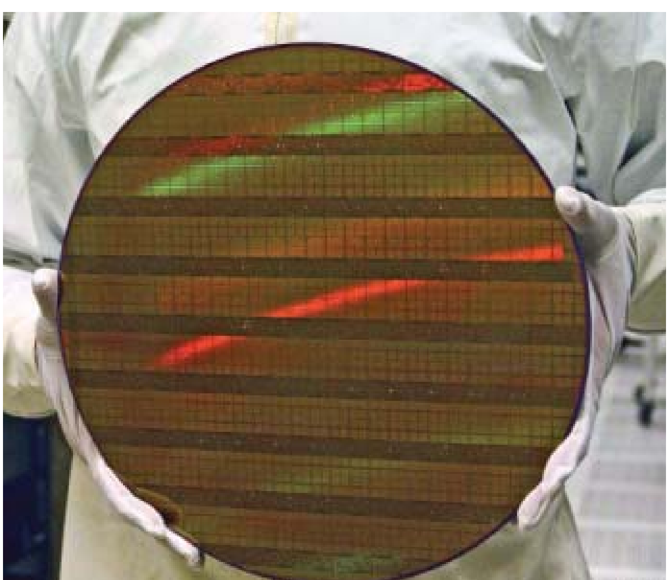

IN THE FAB: By the end of 2007, two fabs at Intel will be churning out the first commercial microprocessors made up of transistors fundamentally redesigned using new materials.

We couldn't shave off even one more tenth of a nanometer—a single silicon atom is 0.26 nm in diameter. More important, at a thickness of five atoms, the insulation was already a problem, wasting power by letting electrons rain through it. Without a significant innovation, the semiconductor industry was in danger of encountering the dreaded "showstopper," the long-awaited insurmountable problem that ends the Moore's Law era of periodic exponential performance gains in memories, microprocessors, and other chips—and the very good times that have gone with it.

The solution to this latest crisis involved thickening the insulator with more atoms, but of a different type, to give it better electrical properties. This new insulator works well enough to halt the power-sucking hail of electrons that's plagued advanced chips for the past four years. If Moore's Law crumbles in the foreseeable future, it won't be because of inadequate gate insulation. Intel cofounder Gordon Moore, of Moore's Law fame, called the alterations we made in introducing this latest generation of chips "the biggest change in transistor technology" since the late 1960s.

As difficult as finding the new insulator was, that was only half the battle. The point of the insulator is to separate the transistor's silicon gate from the rest of the device. The trouble is, a silicon gate didn't work with the new insulator material. The initial transistors made with them performed worse than older transistors. The answer was to add yet another new material to the mix, swapping the silicon gate for one made of metal.

It may not seem like such a big deal to change the materials used in a transistor, but it was. The industry went through a major upheaval several years ago when it switched from aluminum interconnects to copper ones and—at the same time—from SiO_2 cladding for those interconnects to chemically similar "low-*k*" dielectrics. And those changes had nothing to do with the transistor itself. A fundamental change to the composition of the transistor is pretty much unheard of. The combination of the gate and the insulator, the gate stack, hasn't changed significantly since Moore, Andrew S. Grove, and others described it in this magazine back in October 1969!

So when you boot up your next machine and you're surprised by how fast it rips through some video coding, remember: there's more new under its hood than in any computer you've ever owned.

The story of how we and our co-workers solved the gate-insulation problem may seem esoteric, and in a literal way it is. But it is also emblematic of how Moore's Law, the defining paradigm of the global semiconductor industry, is being sustained against often-daunting odds by the swift application of enormous intellectual and material resources to problems that, increasingly, are forcing engineers to struggle in realms until recently occupied only by physicists.

THE PROBLEM, ULTIMATELY, IS ONE OF POWER. At five atoms, that sliver of SiO_2 insulation was so thin that it had begun to lose its insulating properties. Starting with the generation of chips fabricated in 2001, electrons had begun to trickle through it. In the processors made just two years later, that trickle became some 100 times as intense.

All that current was a drain on power and a source of unwanted heat. Laptops were heating up too much and draining their batteries too quickly. Servers were driving up their owners' electric bills and taxing their air conditioners. Even before we

PREVIOUS PAGE: AARON HEWITT; LEFT: INTEL; ILLUSTRATIONS: BRYAN CHRISTIE DESIGN

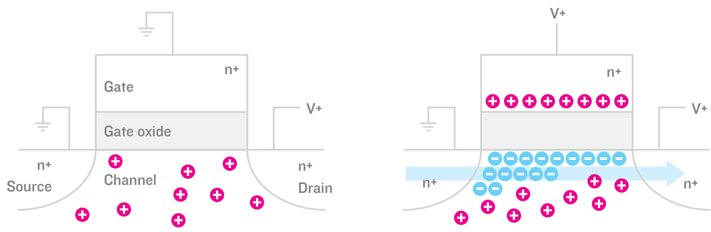

THE TRANSISTOR: **A positive voltage on the gate of an NMOS transistor drives positive charge in the channel away from the insulating gate oxide and attracts electrons, forming a path for electrons to flow.**

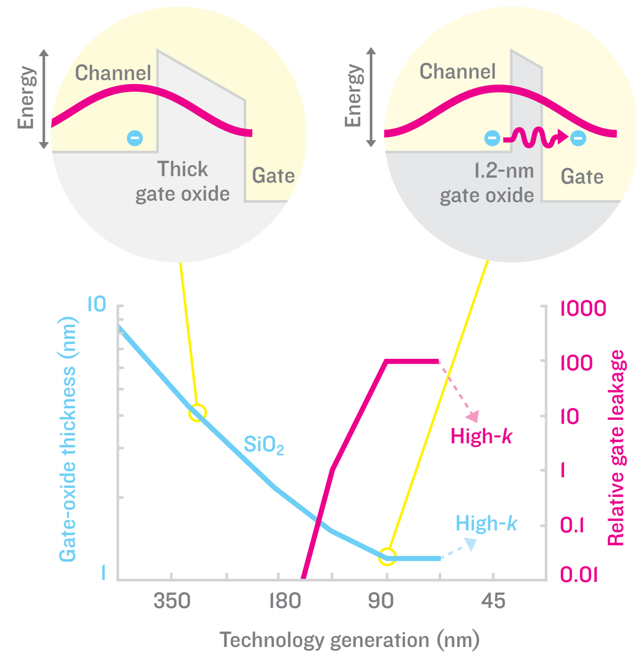

RUNNING OUT OF ATOMS: **The transistor's gate oxide thinned with each new technology generation until it reached just 5 atoms (1.2 nm) thick. At that scale, the wave describing the probable location of an electron [red curve, top] is broader than the gate oxide, and the electron can simply appear on the other side of the gate oxide, having tunneled through the insulation. This so-called gate leakage increased 100-fold in the last three generations of transistors. A switch to a new gate oxide, a high-*k* dielectric, was needed to plug the leak.**

ran out of atoms, designers had devised some tricks to throttle back on the power without losing speed. But without a way to stanch the unwanted flow of electrons through that sliver of insulation, the battle to make ever more powerful processors would soon be lost.

To understand why, you need a quick lesson (or refresher) in semiconductor basics. The type of transistor that is chained together by the hundreds of millions to make up today's microprocessors, memory, and other chips is called a metal-oxide-semiconductor field effect transistor, or MOSFET. Basically, it is a switch. A voltage on one terminal, known as the gate, turns on or off a flow of current between the two other terminals, the source and the drain [see illustration, "The Transistor"].

MOSFETs come in two varieties: N (for *n*-type) MOS and P (for *p*-type) MOS. The difference is in the chemical makeup of the source, drain, and gate. Integrated circuits contain both NMOS and PMOS transistors. The transistors are formed on single-crystal silicon wafers; the source and drain are built by doping the silicon with impurities such as arsenic, phosphorus, or boron. Doping with boron adds positive charge carriers, called holes, to the silicon crystal, making it *p*-type, while doping with arsenic or phosphorus adds electrons, making it *n*-type.

Taking an NMOS transistor as an example, the shallow source and drain regions are made of highly doped *n*-type silicon. Between them lies a lightly doped *p*-type region, called the transistor channel—where current flows. On top of the channel lies that thin layer of SiO_2 insulation, usually just called the gate oxide, which is the cause of the chip industry's most recent technological headaches.

Overlying that oxide layer is the gate electrode, which is made of partially ordered, or polycrystalline, silicon. In the case of an NMOS device it is also *n*-type. (The silicon gates replaced aluminum gates—the metal in "metal-oxide semiconductor"—in work described in the 1969 *IEEE Spectrum* article. But the "MOS" acronym has nevertheless lived on.)

The NMOS transistor works like this: a positive voltage on the gate sets up an electric field across the oxide layer. The electric field repels the holes and attracts electrons to form an electron-conducting channel between the source and the drain.

A PMOS transistor is just the complement of NMOS. The source and drain are *p*-type; the channel, *n*-type; and the gate, *p*-type. It works in the opposite manner as well: a positive voltage on the gate (as measured between the gate and source) cuts off the flow of current.

In logic devices, PMOS and NMOS transistors are arranged so that their actions complement each other, hence the term CMOS for complementary metal-oxide semiconductor. The arrangement of CMOS circuits is such that they are designed to draw power only when the transistors are switching on or off. That's the idea, anyway.

Although the basic features and materials of the MOS transistor have stayed pretty much the same since the late 1960s, the dimensions have scaled dramatically. The transistor's minimum layout dimensions were about 10 micrometers 40 years ago, and are less than 50 nm now, smaller by a factor of more than 200. Suppose a 1960s transistor was as big as a three-bedroom house and that it shrank by the same factor. You could hold the house in the palm of your hand today.

In the Penryn processors that we recently began fabricating, most of their transistors' features measure around 45 nm, though one is as small as 35 nm. It's the first commercial microprocessor to have features this small; all other top-of-the-line microprocessors in production as this article is being written have 65-nm features. In other words, Penryn is the first of the 45-nm generation of microprocessors. Many more will soon follow.

The thickness of the SiO_2 insulation on the transistor's gate has scaled from about 100 nm down to 1.2 nm on state-of-the-art microprocessors. The rate at which the thickness decreased was steady for years but started to slow at the 90-nm generation, which went into production in 2003. It was then that the oxide hit its five-atom limit. The insulator thickness shrank no further from the 90-nm to the 65-nm generation still common today.

The reason the gate oxide was shrunk no further is that it began to leak current [see illustration, "Running Out of Atoms"]. This leakage arises from quantum effects. At 1.2 nm, the quantum nature of particles starts to play a big role. We're used to thinking of electrons in terms of classical physics, and we like to imagine an electron as a ball and the insulation as a tall and narrow hill. The height of the hill represents how much energy you'd need to provide the electron to get it to the other side. Give it a sufficient push and—sure enough—you could get it over the hill, busting through the insulation in the process.

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

But when the hill (the oxide layer) is so narrow that you are counting individual atoms of thickness, the electron looks less like a ball and more like a wave. Specifically, it's a wave that defines the probability of finding the electron in a particular location. The trouble is that the wave is actually broader than the hill, extending all the way to the other side and beyond. That means there is a distinct probability that an electron that should be on the gate side of the oxide can simply appear on the channel side, having "tunneled" through the energy barrier posed by the insulation rather than going over it.

IN THE MID-1990S, WE AT INTEL and other major chip makers recognized that we were fast approaching the day when we would no longer be able to keep squeezing atoms out of the SiO_2 gate insulator. So we all launched research programs to come up with a better solution. The goal was to identify a gate dielectric material as a replacement for SiO_2 and also to demonstrate transistor prototypes that leaked less while at the same time driving plenty of current across the transistor channel. We needed a gate insulator that was thick enough to keep electrons from tunneling through it and yet permeable enough to let the gate's electric field into the channel so that it could turn on the transistor. In other words, the material had to be physically thick but electrically thin.

The technical term for such a material is a "high-*k*" dielectric; *k*, the dielectric constant, is a term that refers to a material's ability to concentrate an electric field. Having a higher dielectric constant means the insulator can provide increased capacitance between two conducting plates—storing more charge—for the same thickness of insulator. Or if you prefer, it can provide the same capacitance with a thicker insulator [see illustration, "The High-*k* Way"]. SiO_2 typically has a *k* of around 4, while air and a vacuum have values of about 1. The *k*-value is related to how much a material can be polarized. When placed in an electric field, the charges in a dielectric's atoms or molecules will reorient themselves in the direction of the field. These internal charges are more responsive in high-*k* dielectrics than in low-*k* ones.

Incidentally, back in 2000, leading semiconductor firms began to change the material used to insulate the metal wires that connect transistors to each other from SiO_2 to low-*k* dielectrics. In the case of interconnects, you do not want the electric field from one wire to be felt in other nearby wires, because it creates a capacitor between the wires and can interfere with or slow down the signals on them. A low-*k* dielectric prevents the problem.

We set about studying a veritable alphabet soup of high-*k* dielectric candidates, including aluminum oxide (Al_2O_3), titanium dioxide (TiO_2), tantalum pentoxide (Ta_2O_5), hafnium dioxide (HfO_2), hafnium silicate ($HfSiO_4$), zirconium oxide (ZrO_2), zirconium silicate ($ZrSiO_4$), and lanthanum oxide (La_2O_3). We were trying to identify such things as the material's dielectric constant, how electrically stable it was, and its compatibility with silicon. For quick turnaround, we experimented with simple capacitor structures, building a sandwich consisting of titanium nitride electrodes, the high-*k* dielectric, and a silicon gate electrode. We then charged them up and discharged them again and again, watching to see how much the relationship between capacitance and voltage changed with each cycle.

But for the first two years, all the dielectrics we tried worked poorly. We found that charges got trapped at the interface between the gate electrode and the dielectric. This accumulated charge within the capacitor altered the voltage level needed to store the same amount of energy in the capacitor from one charge-discharge cycle to the next. You want a transistor to operate exactly the same way every time it switches, but these gate-stack structures behaved differently each time they were charged up. The results were very discouraging, but eventually our team got an important break.

It turned out that the problem lay in how we constructed the test capacitor. To make the dielectric layer, we were using one of two different semiconductor-manufacturing techniques: reactive sputtering and metal organic chemical vapor deposition. Unfortunately, both processes produce surfaces that, though remarkably smooth by most standards, were nevertheless uneven enough to leave some gaps and pockets in which charges could get stuck.

We needed something even smoother—as smooth as a single layer of atoms, actually. So we turned to a technology called atomic layer deposition, so new that its debut in CMOS chip production comes only this year with our new high-*k* chips. Atomic layer deposition lets you build up a material one layer of atoms at a time. In this process, you introduce a gas that reacts with the surface of the silicon wafer, leaving the whole substrate coated in a single layer of atoms. Then, because there is no more surface to react with, the deposition stops. The gas is evacuated from the chamber and replaced with a second gas, one that chemically reacts with the layer of atoms just deposited. It too lays down one layer of atoms and then stops. You can repeat the process as many times as you want, to produce layered materials whose total thickness is controllable down to the width of a single atom.

Deposited in this manner, both the hafnium- and zirconium-based high-*k* dielectrics we studied showed much more stable

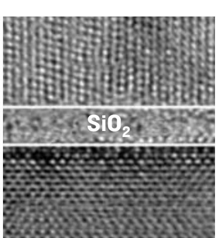

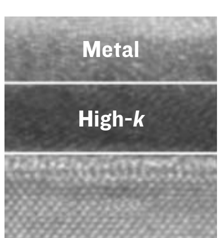

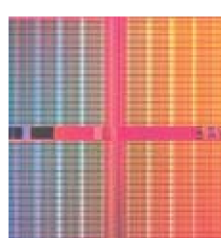

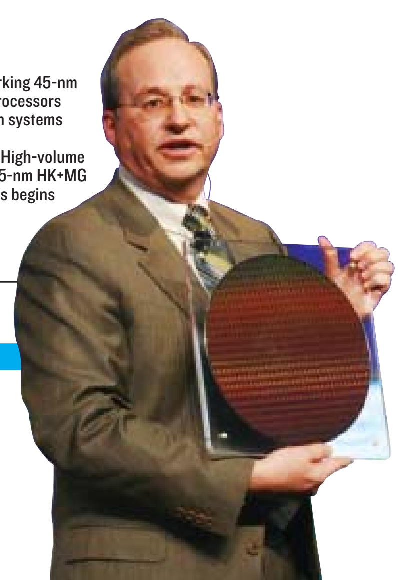

Jan. 2007 Working 45-nm HK+MG microprocessors demonstrated in systems

2nd half 2007 High-volume production of 45-nm HK+MG microprocessors begins

1990 | 2003 | 2004 | 2005 | 2006 | 2007

Late 1990s SiO_2's limits realized. Search begins for a high-*k* dielectric to replace it

Nov. 2003 Working high-*k* plus metal gate (HK+MG) research transistors

Jan. 2006 Fully functional 1-billion-transistor 45-nm HK+MG memory chips

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

electrical characteristics in comparison with the ones formed by sputtering or chemical vapors. The trapped-charge problem seemed to have been smoothed out.

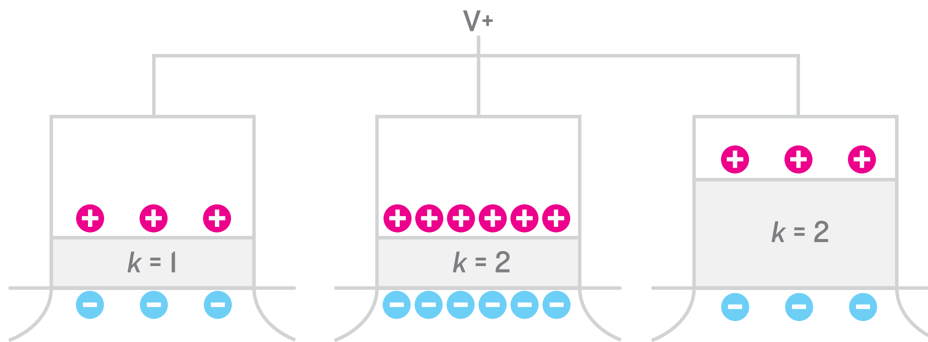

THE HIGH-*K* WAY: The dielectric constant, *k*, is a measure of an insulator's ability to concentrate an electric field. If one gate oxide has twice the dielectric constant of another, a given voltage will draw twice as much charge into the transistor channel. Or, the same amount of charge will accumulate, if the higher-*k* dielectric is made twice as thick.

WITH TWO CANDIDATE MATERIALS IDENTIFIED, we started making NMOS and PMOS transistors out of them. Then came the next snag. These transistors, pretty much identical to our existing transistors except for the different dielectric, had a few problems. For one thing, it took more voltage to turn them on than it should have—what's called Fermi-level pinning. For another, once the transistors were on, the charges moved sluggishly through them—slowing the device's switching speed. This problem is known as low charge-carrier mobility.

We weren't the only ones encountering these problems; just about everybody else was struggling with them, too. With the countdown in progress for the next generation predicted by Moore's Law, understanding why the high-*k* dielectric transistors performed so poorly and finding a solution became an urgent task. Using a combination of experimental work and physics-based models, we began to figure out what had gone wrong. The source of the trouble, ultimately, came down to the interaction between the polysilicon gate electrode and the new high-*k* dielectrics.

Why this is so has a complicated explanation. The dielectric layer is made up of dipoles—objects with a positive pole and a negative one. This is the very aspect that gives the high-*k* dielectric such a high dielectric constant. These dipoles vibrate like a taut rubber band and lead to strong vibrations in a semiconductor's crystal lattice, called phonons [see illustration, "Bumpy Ride"]. These phonons knock around passing electrons, slowing them down and reducing the speed at which the transistor can switch. But theoretical studies and computer simulations performed by us and others showed a way out. The simulations indicated that the influence of dipole vibrations on the channel electrons can be screened out by significantly increasing the density of electrons in the gate electrode. One way to do that would be to switch from a polysilicon gate to a metal one. As a conductor, metal can pack in hundreds of times more electrons than silicon. Experiments and further computer simulations confirmed that metal gates would do the trick, screening out the phonons and letting current flow smoothly through the transistor channel.

What's more, the bond between the high-*k* dielectric and the metal gate would be so much better than that between the dielectric and the silicon gate that our other problem, Fermi-level pinning, would be solved by a metal gate as well.

NOW OUR ENGINEERS HAD A NEW MAJOR CHORE: find a metal they could use for the gate electrode that would combine well with the new high-*k* dielectric. Because the electrical characteristics of the gates of NMOS and PMOS transistors are different, they actually needed not one metal but two—one for NMOS and one for PMOS.

Just as standard MOS transistors use *n*-type and *p*-type polysilicon gates for NMOS and PMOS transistors, high-*k* transistors would need metal gate electrode materials with a key property similar to polysilicon's. This key property is known as the work function. In this context, work function refers to the energy of an electron in the gate electrode relative to that of an electron in the lightly doped silicon channel. The energy difference sets up an electric field that can modulate to the amount of voltage needed to begin to turn the transistor on, the threshold voltage. Unless the gate's work function is chosen well, the threshold voltage will be too high, and the transistor will not turn on easily enough.

We analyzed, modeled, and experimented with many types of metals, some with work functions that more closely matched highly doped silicon than others. But by themselves, none had exactly the work function of the doped silicon, so we had to learn to change the work function of metals to suit our needs. Eventually, the research group identified NMOS and PMOS metals by first building capacitors out of them and then transistors. We cannot disclose the exact makeup of our metal layers, because after all, the IC industry is very competitive!

We built our first NMOS and PMOS high-*k* and metal gate transistors in mid-2003 in Intel's Hillsboro, Ore., development fab. We started out using Intel's 130-nm technology, which was about three years old at the time and was used in high-volume production. The transistors, with a hafnium-based oxide and metal gate electrodes, had everything we needed: they turned on at the right voltage, leaked little current through the gate oxide, and passed a large amount of current through the channel for a given voltage. And that current moved quickly. In fact, for a given off-state current, these first transistors drove more current than any transistor reported at the time.

OF COURSE, WE WEREN'T ALONE. And there were still plenty of unknowns. By 2003, researchers in university labs and other semiconductor firms around the world had zeroed in on hafnium-based materials as the gate dielectric. A variety of them were under earnest study: hafnium oxides, hafnium silicates, and hafnium oxides containing nitrogen. The method of forming the high-*k* film, too, was unsettled, with different groups trying sputtering, chemical vapor deposition, and atomic layer deposition, which we eventually settled on. But the biggest unknowns at the time were what metal gate materials to use and how to fit them into the transistor-manufacturing process.

The normal fabrication method is known as "gate first." As the name implies, the gate dielectric and gate electrodes are constructed first. Then the dopants for the source and drain are implanted into the silicon on either side of the gate. Finally, the silicon is annealed to repair the damage from the implantation process. That procedure requires that the gate electrode material

INTEL (4)

be able to withstand the high temperatures used in the annealing step—not a problem for polycrystalline silicon but potentially a big one for some metals.

To make a long story short, the search for gate electrode materials with both the right work function and tolerance to high-temperature processing was very difficult and full of dead ends. Especially for the PMOS transistor.

Another transistor process sequence, dubbed "gate last," circumvents the thermal annealing requirement by depositing the gate electrode materials after the source and drain are formed. However, many of our peers saw the gate-last process, which we ultimately adopted, as too much of a departure and too challenging.

Meanwhile, a third approach remarkable in its simplicity emerged. Called fully silicided gates, it lets you follow the normal gate-first process but then lets you turn the polysilicon gate into a metal-silicide gate, essentially replacing every other silicon atom with metal (usually nickel). Then, by doping the nickel silicide, you can alter its work function for use in either an NMOS device or a PMOS one. By late 2006, though, nearly everyone, including us, had given up on the fully silicided gates approach. No one could move the silicide's work function quite close enough to where it needed to be.

Nevertheless, the search goes on at other major chip makers to find the materials with the right work function that could survive high temperatures and enable the industry standard gate-first process flow.

HAVING BUILT WELL-FUNCTIONING TRANSISTORS using old technology, in the second half of 2003 it was time to move from research to development of high-*k* dielectric plus metal gate transistors, as we called them. Engineers began working to determine whether these early transistors could be scaled to the upcoming 45-nm dimensions and still meet the rigorous performance, reliability, and manufacturability requirements of an advanced microprocessor technology.

It was no cakewalk. The research group engineers had provided a critical lead in identifying promising high-*k* and metal gate materials, but the NMOS and PMOS transistors had not yet been combined on one wafer as they would be in a microprocessor, using a manufacturing process that could make both. What's more, there were hard questions still to be answered about how many good chips we could expect for every bad one (yield) and how reliable those chips would be.

During the months that followed, the team cracked one problem after another—making changes to materials, chemical recipes, and manufacturing processes. It wasn't until late 2004 that the team felt it had enough convincing data that the new transistors could be made to work on our 45-nm technology. At that point, there was no turning back. Intel was now committed to making a high-*k* dielectric plus metal gate transistor structure using the gate-last process flow. It was a gutsy call. Our team knew it was committing all of Intel's next generation of microprocessors to the biggest change in transistor technology in 40 years.

The next key milestone was to demonstrate working test chips using the final scaled dimensions combined with the new transistor features. The traditional chip to test a new technology on is static random access memory, or SRAM, which is the type of memory collocated on the same chip with the microprocessor. Typically, microprocessor makers have designs for SRAM that are a year or more ahead of their processor designs. SRAM is a very regular array of memory cells, each of which consists of six densely packed and interconnected transistors. Because of their density and regularity, SRAM chips provide good data on how many defects a manufacturing process produces.

Our first fully functional test SRAM chips with the new transistors came off the line in January 2006. They were of a 153-megabit design consisting of more than 1 billion transistors. Each six-transistor memory cell in the chip occupied little more than one-third of a square micrometer. This test chip had all the features needed to build a 45-nm microprocessor, including the high-*k* plus metal gate transistors and nine layers of copper interconnects. Considering how new and radically different the transistor and manufacturing process were, it was a surprise even to some of the engineers in the development group that it all worked together so well. Even so, the development team still had a lot ahead of it to bring the performance, reliability, and yield of the process up to the level needed for manufacturing microprocessors.

The new gate stack worked wonders in battling leakage through the gate, reducing it by more than a factor of 10. But the gate oxide is not the only source of transistor leakage chip makers have to worry about. The other significant leak is called source-to-drain or subthreshold leakage. It's a trickle of current seen even when the transistor is intended to be in the "off" state. Making transistors smaller has also meant steadily lowering the amount of voltage needed to turn them on, the threshold voltage. Unfortunately, steadily lowering the threshold voltage lets more current slip through. For many years, each new generation of transistor would increase drive current (and improve performance) by about 30 percent but would pay the price of about a threefold increase in subthreshold leakage. Leakage currents have reached levels high enough to be a noticeable portion of total microprocessor power consumption.

The industry is now in an era where power efficiency and low leakage are more important than raw speed increases. But a transistor can be designed to operate to favor either priority by adjusting the channel length or adjusting the threshold voltage. A shorter channel leaks more but allows for a higher drive current. A higher threshold voltage pinches off the leak but also throttles the drive current. Adjusting the threshold

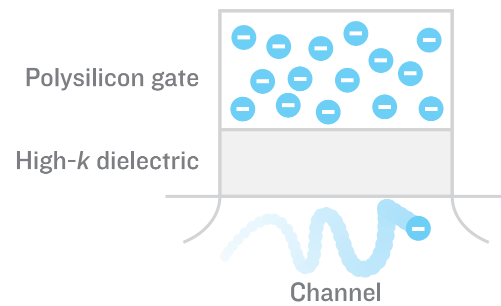

BUMPY RIDE: The particular density of electrons in a traditional polysilicon gate allowed inherent vibrations in the high-*k* dielectric to move into the transistor channel and disrupt the flow of current.

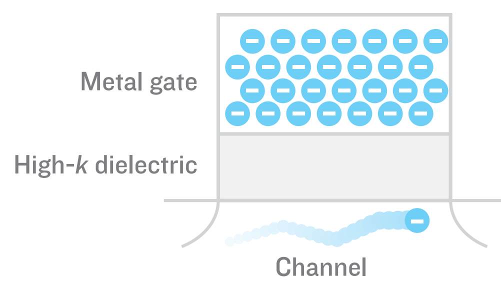

The higher density of electrons in a metal screened out the vibrations, allowing current to flow more smoothly.

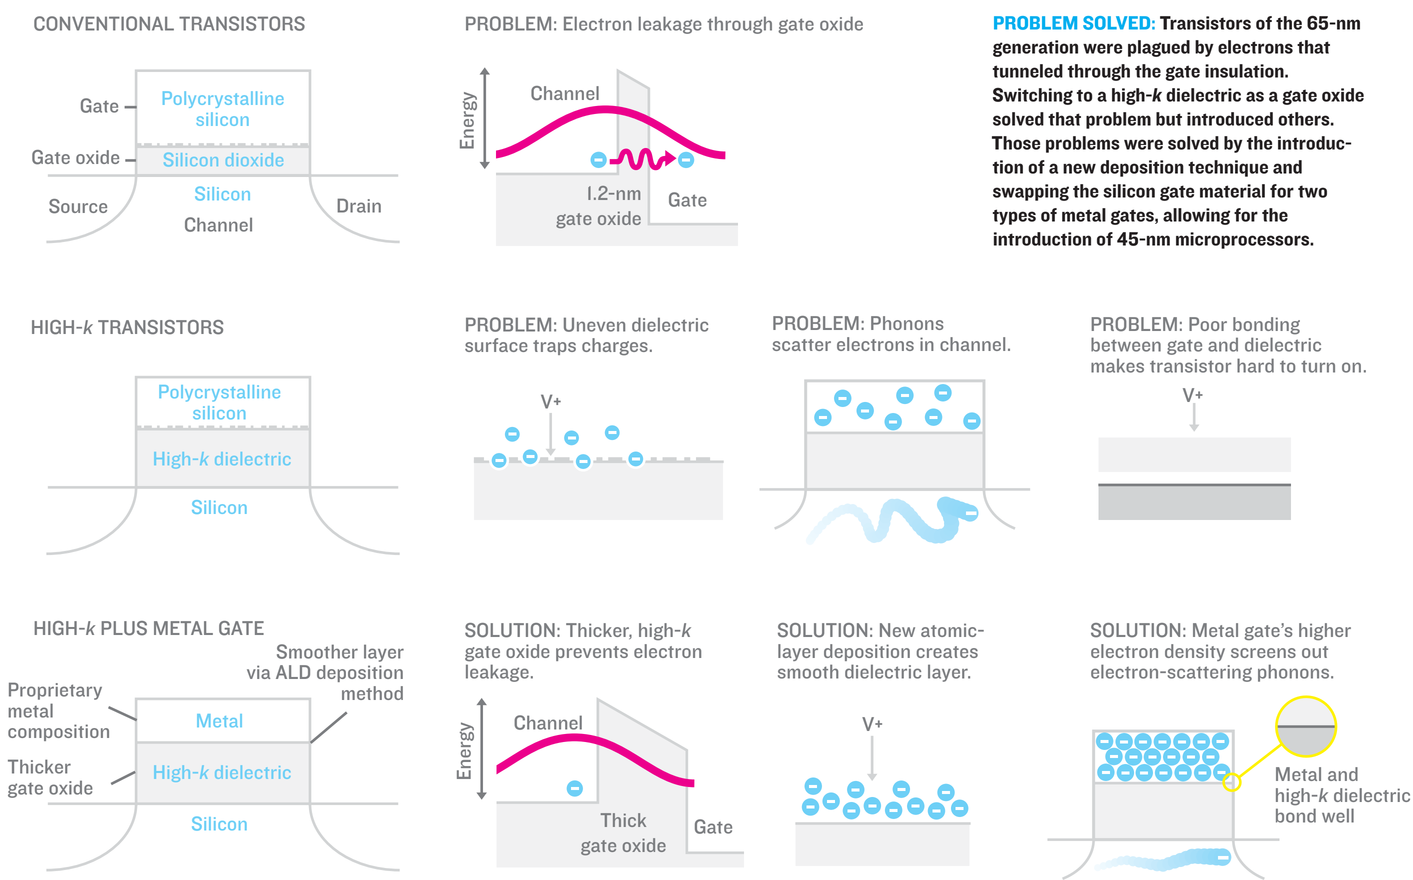

PROBLEM SOLVED: Transistors of the 65-nm generation were plagued by electrons that tunneled through the gate insulation. Switching to a high-*k* dielectric as a gate oxide solved that problem but introduced others. Those problems were solved by the introduction of a new deposition technique and swapping the silicon gate material for two types of metal gates, allowing for the introduction of 45-nm microprocessors.

voltage is where the high-*k* dielectric comes into play. A thicker dielectric reduces the gate's ability to open a conductive channel between the source and the drain, increasing the threshold voltage. A thinner dielectric layer has the opposite effect. Compared with the previous 65-nm transistors, 45-nm high-*k* plus metal gate transistors provide either a 25 percent increase in drive current at the same subthreshold leakage or more than a fivefold reduction in leakage at the same drive current, or anywhere between those values. We can make the choice on a product-by-product basis, or different circuits on the same microprocessor chip can use different transistors to optimize for performance or power.

IN JANUARY 2007, Intel made the first working 45-nm microprocessors using these revolutionary high-*k* plus metal gate transistors. One was the Penryn dual-core microprocessor, which has 410 million transistors. Different versions of Penryn will be optimized for mobile, desktop, workstation, and server applications. The quad-core version of this product will have 820 million transistors. Penryn was followed a few months later by Silverthorne, a single-core microprocessor with 47 million transistors that is designed for low-power applications, including mobile Internet devices and ultramobile PCs. There are more than 15 new chips under development at Intel using our new technology. Production of Penryn and Silverthorne will start later this year at Intel plants in Oregon and Arizona. Next year, we'll start up the process at two other high-volume manufacturing fabs, in New Mexico and Israel.

The invention of high-*k* plus metal gate transistors was an important breakthrough. Although we could have continued to shrink transistors to fit the dimensions needed for the 45-nm generation without this breakthrough, those transistors would not have worked much better than their predecessors, and they certainly would have expended more watts. We're confident this new transistor can be scaled further, and development is already well under way on our next-generation 32-nm transistors using an improved version of high-*k* plus metal gate technology. Whether this type of transistor structure will continue to scale to the next two generations—22 nm and 16 nm—is a question for the future. Will we need new materials and new structures again?

Nobody knows for sure. But that is what makes integrated circuit research and development so exciting. ■

ABOUT THE AUTHORS

MARK T. BOHR, an IEEE Fellow, is the director of process architecture and integration at Intel. **ROBERT S. CHAU**, an IEEE Fellow, is the director of transistor research and nanotechnology. **TAHIR GHANI**, an IEEE member, is the director of transistor technology and integration. **KAIZAD MISTRY**, an IEEE senior member, manages the development of Intel's 45-nanometer CMOS technology in the logic and technology development group.

TO PROBE FURTHER

Robert S. Chau and colleagues explained the problem that led to the use of a metal gate in detail in "High-*k*/Metal-Gate Stack and Its MOSFET Characteristics," *IEEE Electron Device Letters*, June 2004.

Intel and others will be presenting the latest high-*k* dielectric and metal gate transistor research at IEEE's 2007 International Electron Devices Meeting, in Washington, D.C., from 10 to 12 December.

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

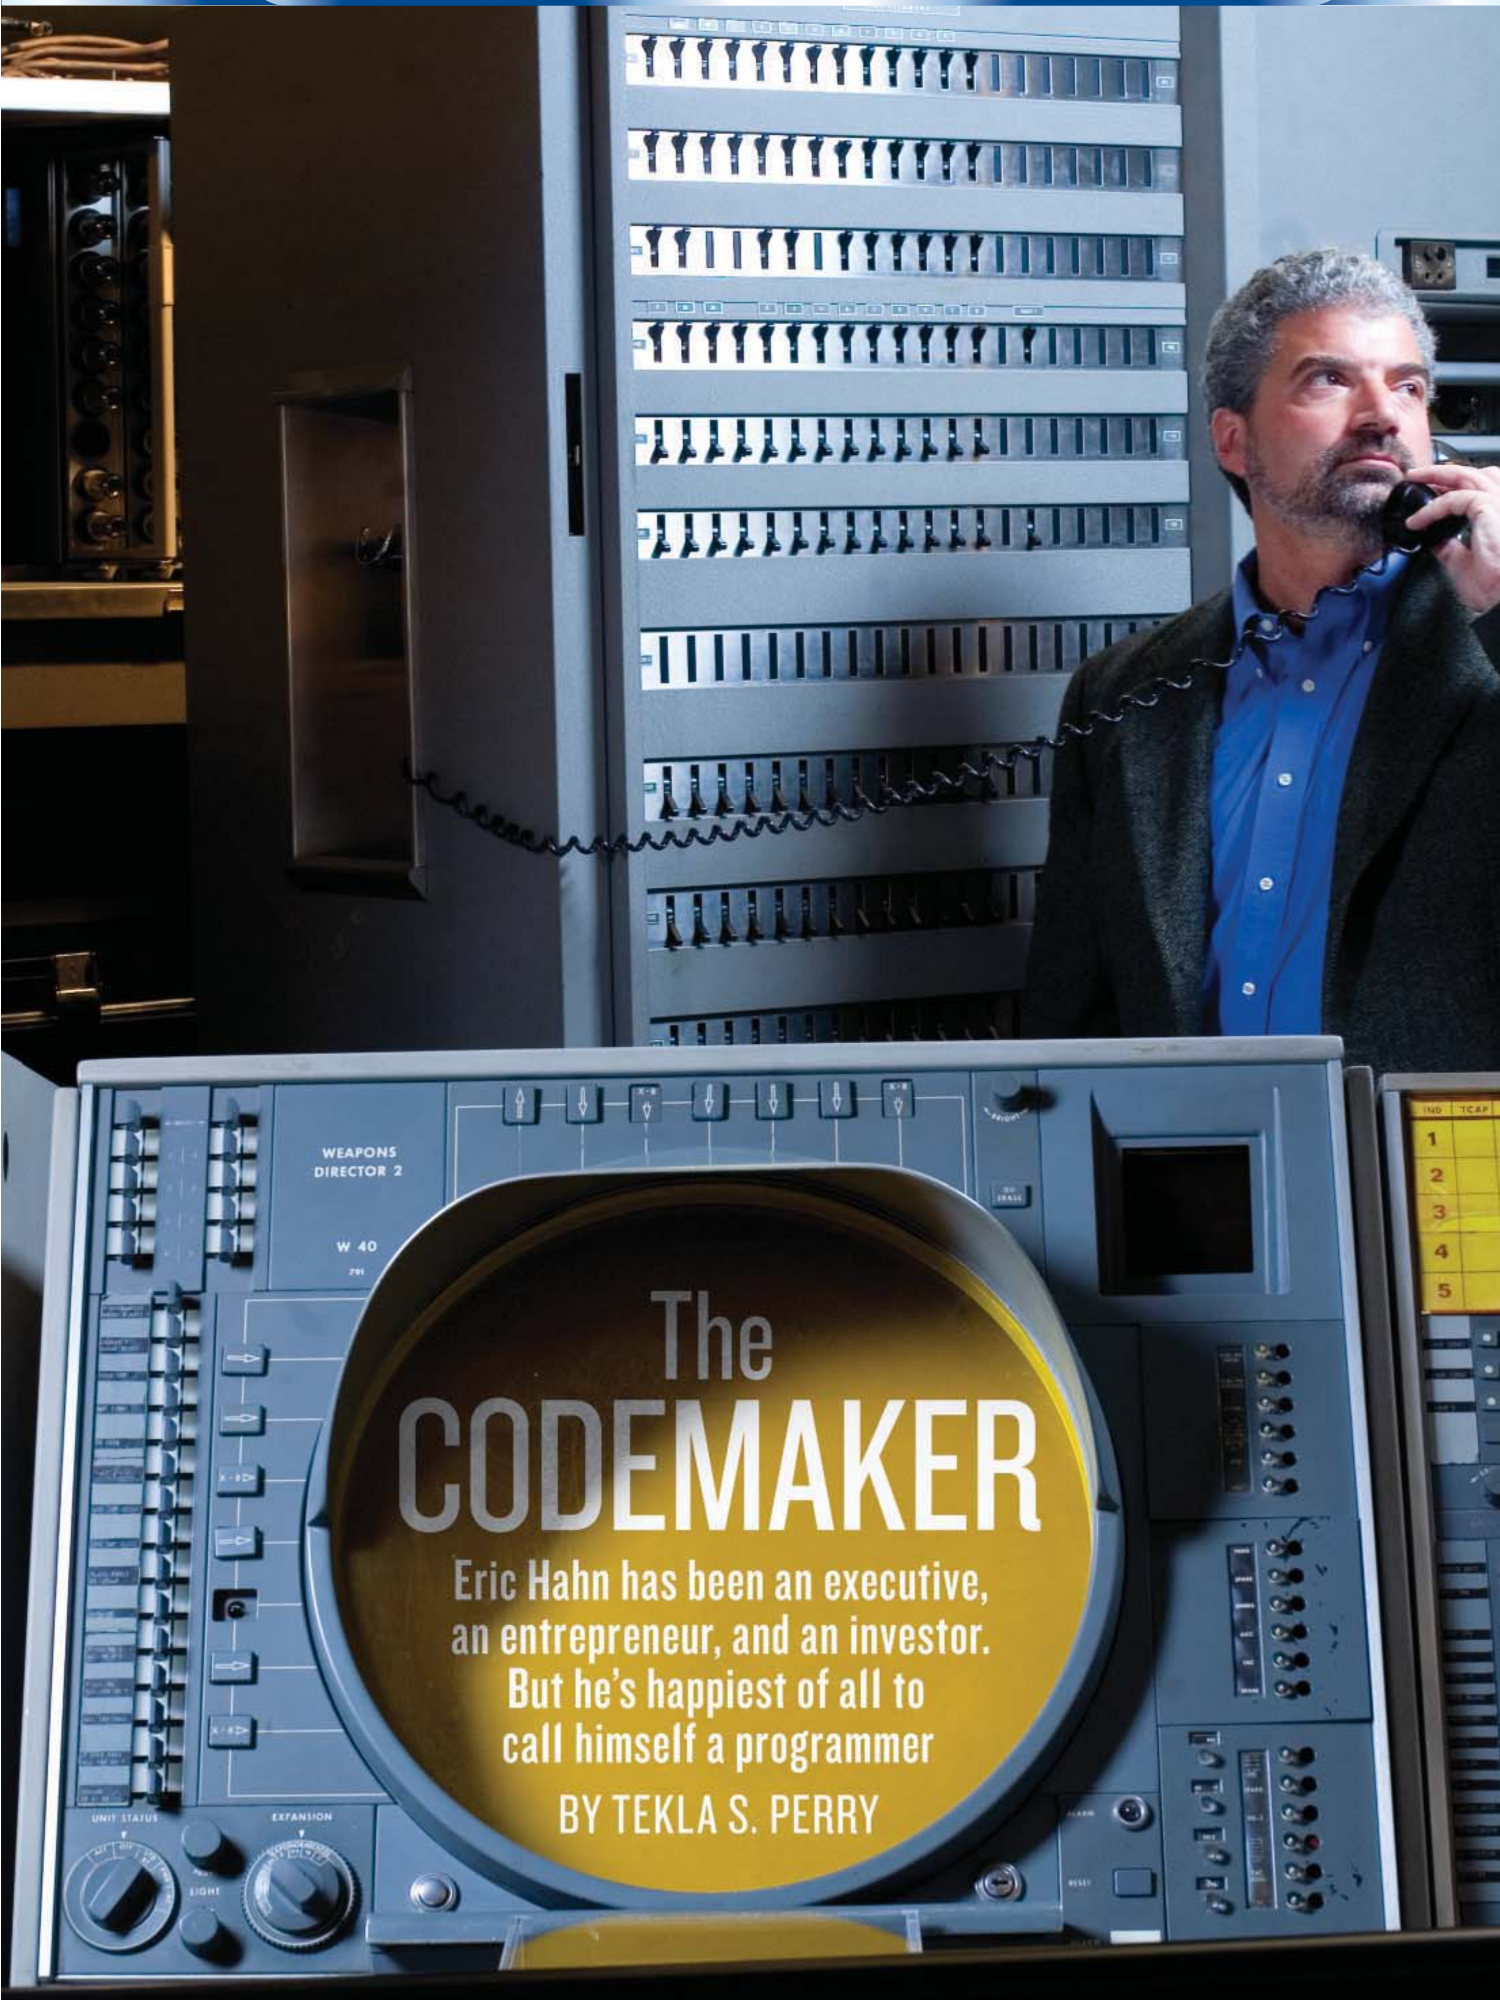

The CODEMAKER

Eric Hahn has been an executive, an entrepreneur, and an investor. But he's happiest of all to call himself a programmer

BY TEKLA S. PERRY

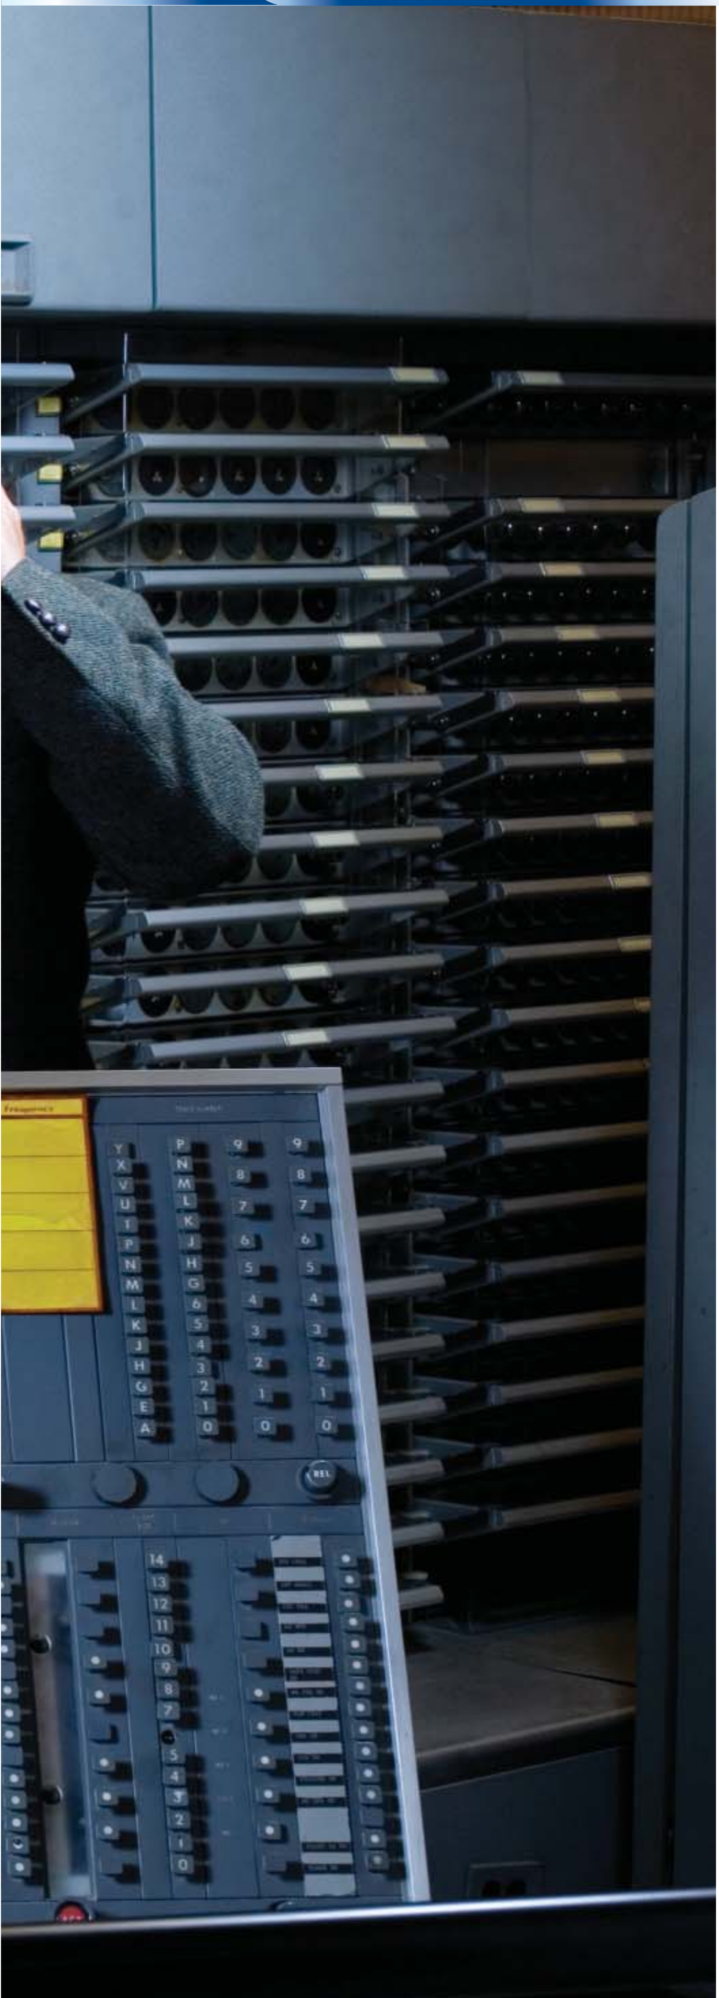

Eric Hahn, a geeky 12-year-old in middle school in East Hampton, N.Y., wanted a computer in the worst way. It was the early 1970s, and computers were owned by corporations and schools, not by kids, but Hahn had to try to get one. He wrote a letter to C. Gordon Bell, then the brash vice president of research and development for the mighty Digital Equipment Corp., at the time the world's largest maker of minicomputers.

The object of Hahn's desire was a Digital PDP-8/a minicomputer. It may be hard to remember what it was like to get excited about a computer the size of a microwave oven with 4 kilobytes of main memory and a 12-bit word length. But this was at a time when men's sideburns were big, women's shoes were high, and Donny Osmond and the Carpenters ruled the airwaves.

Hahn didn't want charity—just a price break. A PDP-8/a, then two years past its introduction, could be had for around US $1000—in quantities of 100. All Hahn wanted from Bell was the 100-quantity price. He'd already saved up close to $1000 by soldering circuit boards for his father, who had a small electronics company that did one-off projects.

Bell, already a minor legend for having led the design of the time-shared PDP-6, knew a publicity opportunity when he saw it. So a few months later Hahn and Bell met in an office in DEC's Maynard, Mass., headquarters, Bell in a dress shirt, Hahn in a sweater, and posed over the computer gear [see photo, "At the Keyboard"]. Hahn gave his $1000 to Bell, and Bell handed over a PDP-8/m, a much faster and more expensive machine than Hahn had sought. "It was a no-brainer to get him a computer at a price he could afford," Bell says, "and it turned out to be one of the better investments Digital ever made!"

And then, Hahn recalls, came "one of the high points of my young life. I spent two or three hours debating with Gordon Bell, who had personally designed the instruction set used on the PDP-8, about the foibles of programming the machine."

Hahn carried his new computer home and began writing code for it. Within four years—at age 16—his work on the PDP-8/m became the basis of his first successful software company, Amide Software, which sold an emulation program that enabled Intel 8080–based personal computers to run PDP-8 software.

Touch a child's life, they say, and you never know what other lives might be touched in turn. But in Hahn's case, you can make a darn good guess. He has helped start about a dozen companies, including the e-mail and collaborative software company Collabra, which was acquired by Netscape in 1995. He ran technical initiatives at Bolt, Beranek & Newman (now BBN Technologies), Convergent Technologies, Lotus (now IBM Lotus Software), and Netscape. Along the way he became a millionaire.

Nine years ago he started the Inventures Group, a tiny, early-stage investment business in Palo Alto, Calif., that is similar in some respects to a venture capital firm but invests mainly Hahn's money. It has stakes in some highly touted start-ups, including Linux pioneer Red Hat and Opsware (formerly Loudcloud), purchased by Hewlett-Packard in July for $1.6 billion.

But what has Hahn really excited these days is Zimbra, a

San Mateo, Calif., company that is replacing traditional low-function e-mail software with much more versatile software that uses the power of the browser. Hahn calls Zimbra his quintessential project. He invested in it in 2004 and, shortly thereafter, joined its board of directors. "I eat, drink, and sleep Zimbra these days," he says. Zimbra's key technical insight was that the browser itself could be used to deliver fully functional e-mail to users without installing any software. People knew about Web e-mail, but they had never seen a system like this. Zimbra was one of the first browser-based e-mail systems to have more functionality and a better user experience than traditional desktop packages.

After Hahn took Zimbra under his wing, his first step was a product architecture review. During weeks of intense meetings, Hahn grilled the key developers. "Why is this in a single database table? What happens if there is a corruption?" Hundreds of such questions were hashed out in the next few weeks. "I've made many, many product mistakes over the years," Hahn says. "I should at least help make sure we make new mistakes this time around!"

Hahn got hooked on software early. Hobbyist computers running on the Intel 8080 microprocessor came out in 1975. A friend, Howard Cannon, got an Imsai. The boys, in high school, quickly discovered that while lots of free software was available from user groups for Hahn's PDP-8, little software existed for the Imsai. So the two wrote a program that let the 8080 emulate a PDP-8, thereby opening up the PDP-8's vast software library for the 8080. They distributed the program on paper tape, selling it by mail for $35 a copy.

"We sold hundreds of copies," Hahn recalls. "It was a significant amount of money." The boys wanted to spend the cash on computer gear; their parents insisted they save it for college.

AT THE KEYBOARD: Eric Hahn [right] was only a middle schooler when he met C. Gordon Bell in the 1970s.

Cannon studied artificial intelligence at MIT and went on to become a key contributor at newly founded Symbolics, in Cambridge, Mass. Hahn, rejected by MIT, went to Worcester Polytechnic Institute. WPI had close ties with DEC, so given Hahn's passion for PDPs and his connection with Gordon Bell, it was an easy choice. He blasted through in less than three years and graduated at age 19.

Computer classes at WPI were easy for Hahn. It was a new field, and the knowledge bank wasn't yet immense. In many cases he and his fellow students knew nearly as much as their professors.

At the time, many users accessed the WPI computers by dialing through ordinary phone lines. The campus phone system wasn't particularly reliable, and users would regularly lose connections, at which point the computer would cancel their work in progress. So Hahn wrote software that would take everything the computer was doing when a call got disconnected and would save it to a file. He called the program "Freeze and Thaw" because users who were disconnected could come back, "thaw" their work, and start from where they had stopped.

His popularity soared. "Because this was an engineering school, and everybody used the computer, you affected everyone. What you did was probably right up there with changing the menu in the cafeteria," he says.

Upon graduation, Hahn had three job offers—one from DEC, another from DEC's up-and-coming competitor Prime Computer, in Natick, Mass., and a third at a much lower salary from Bolt, Beranek & Newman, in Cambridge. He took the job at BBN, and he swears it is the smartest choice he has ever made.

Hahn liked BBN because it was the prime contractor behind the ARPANET, a high-speed data network that connected the scattered laboratories and contractors of the U.S. Defense Department's Advanced Research Projects Agency (DARPA). One of the Internet's key forerunners, the ARPANET then had no more than a few hundred connected hosts. But it was booming, adding as many as two hosts each month. (Today's Internet grows at a rate of a few million sites per month.)

"At DEC, I would have been working on the PDP-10 operating system," Hahn says. "It was perfectly wonderful stuff, but there was nothing mystical about it.

"People working on the ARPANET, to me, were light years more evolved in thinking about computers and networking than anyone in the traditional minicomputer world."

But he wasn't quite through with WPI. For his undergraduate thesis, he had created a paper design for a campus network that used 8-bit Zilog Z-80 microprocessors as switches, connecting users to the school's PDP-10 computer. It would let more users connect to the computer than could do so using direct phone connections. After he left, WPI decided to build the network, and Hahn offered to help. He'd drive his little green Dodge Dart the 60 kilometers to WPI and spend most of the weekend in a windowless office in the computer center, writing Z-80 assembly-language code. On Saturday night he'd crash at the house of Allan Johannesen, who worked with him on the project, or with another friend.

"He didn't get a penny for it; he didn't get aggrandizement," says Johannesen, who ran the computer center at the time and is now VP of technology infrastructure for WPI. "He just did it because it was something he knew how to do."

Says Debby Meredith, an executive consultant who has worked with Hahn at several companies, "If you give him free time, he'll do what makes him happy, and what makes him happy is programming."

At BBN, Hahn was the most junior programmer in a group that worked on the interface message processor (IMP), a specialized computer that handled the comings and goings of packets of information in and out of the network. IMPs gathered packets of information coming in, performed error-checking routines, and then forwarded them on to their destinations. Today we call this a router, and companies like Cisco Systems and Juniper Networks churn them out by the millions. But in 1980 it was a work in progress. A decade earlier, BBN had built the original IMPs out of Honeywell 316 computers. Then BBN made another version of them using its own hardware, the BBN-C/30, but that system simply emulated the Honeywell 316.

Hahn was horrified. For someone who looks for beauty in code, what he found was anything but. And some things were downright ridiculous.

The Honeywell 316 was a computer without a stack—a data structure that lets program functions be queued and lets multiple functions use the same subroutines. "It was probably the last one ever built that way," Hahn notes. "So you couldn't write a subroutine and test it and know that it would function reliably, because its behavior would change, depending on what instruction invoked it." The BBN-C/30 copied this frustrating feature of the Honeywell exactly.

A PLACE IN HISTORY: Computers are Eric Hahn's vocation and avocation. Hahn is a major backer of the Computer History Museum in Mountain View, Calif., where this Digital PDP-1 from the 1960s is exhibited.

With the energy and optimism of youth, Hahn, then 20, asked his boss, Jim Herman, if he could create a new instruction set, reprogram the C/30 microcode to implement it, and then rewrite the IMP program to take advantage of the new instructions. Herman told him to go ahead. Says Hahn, "I'm guessing that he must have thought that I was nuts, but they were only paying $20 000 a year for me, so it wouldn't cost them much. And maybe something would come from it, even if they didn't actually ever put the software I wrote on the ARPANET."

Herman recalls that he did agree with Hahn that there was a lot of stuff in the IMP software that was old and had been fussed with for way too long. And he was willing to let Hahn tackle the redesign because Hahn was a wunderkind, immensely productive, and would work it out much faster than anyone else possibly could.

Hahn worked feverishly for six months. He built a test network in the laboratory and spent virtually every waking hour there coding like crazy, then testing the program, writing down failure points, and figuring out how to fix them.

Herman remembers walking in on Hahn one day during this period and seeing him sitting with his keyboard on his lap. "He looked like he was playing the piano, like a musician, he was so fluent and so fast," Herman recalls.

At the end of the six months, in 1980, BBN began rolling out a version of the IMP that contained Hahn's new microcode and new program. Besides being easier to work with than the old software, the new system ran about 10 times as fast as the old one. It could accept double the number of connections, so it could process much more message traffic, invaluable during this time of ARPANET growth. The BBN IMPs continued to route ARPANET traffic through most of the 1990s.

Then California beckoned. In 1982 Hahn flew to the Bay Area for a few days of meetings with officials at the University of California at Berkeley, in connection with a BBN project to build a network for the university's library. He called an old friend and scheduled what he thought was to be a casual lunch; it turned out to be a job interview for Convergent Technologies, in Santa Clara, a company building workstations then sold by Sperry, Burroughs, and others.

"The weather was beautiful, there were girls everywhere, and California was just kind of mythical," Hahn recalls. And he was getting frustrated with the research orientation of BBN, which specialized in building limited numbers of custom products and shied away from standard, off-the-shelf microprocessors and languages. Hahn wanted to build things for the masses.

Within months he was in California, first working on communications processors for Convergent computers, which were then migrating into the company's operating system group. Four years later, in 1986, he was 26 and managing the server products division, with 120 people under him. His group released a powerful server based on the Intel 386 chip that acquired some serious geek chic because Intel used it for benchmarking.

"I was general manager of a successful division," Hahn says, "and I hated it. I guess it was good for me, but my heart wasn't in it. Managing a lot of people, particularly a lot of hardware people when I'm a software guy, is not what I wake up every morning wanting to do."

In 1988 Unisys Corp., in Blue Bell, Pa., acquired Convergent. Hahn hadn't been entirely thrilled with the Convergent culture, which he says emphasized deadlines above all else and rewarded employees for how late their cars were spotted in the parking lot. But Unisys was worse, with its time-clock, defense-contractor

PREVIOUS PAGES AND RIGHT: TIMOTHY ARCHIBALD. LEFT: ERIC HAHN/DIGITAL EQUIPMENT CORP.

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

mentality. So he left to join another start-up—cc:Mail, one of the first developers of e-mail applications for computers.

Hahn became vice president of engineering for cc:Mail in 1988. He stayed until shortly after Lotus acquired the company in 1991; he made enough money on his stock options to buy the Palo Alto property on which he later built a five-bedroom brown-shingled house. He and most of the cc:Mail management left Lotus soon after, in disappointment after discovering it really had no interest in the cc:Mail product. Lotus had simply acquired the company for its 21 million users, whom it intended to convert into Lotus Notes customers.

Then came Collabra. And he never could have seen it coming.

After leaving Lotus, Hahn moped around, with no ideas about what to do next. But after a few weeks, partners from Merrill, Pickard, Anderson, & Eyre, a venture capital firm, offered him a spot as entrepreneur in residence. It was pretty much the standard entrepreneur-in-residence deal offered by California VC firms: a salary, a title, desk space, business cards, and a year or so to come up with an idea for a company. The firm might or might not decide to invest in the idea.

For Hahn, the title was the most valuable part of the deal; it enabled him to call executives in Silicon Valley and get answers to questions. He came up with an idea for a company within three months, though he didn't incorporate it until early 1993.

That company, Collabra Software, built a software package that worked as an add-on to existing e-mail programs, enabling users to share files easily and communicate with each other via discussion boards. It basically made a practical system out of some fairly abstract and complicated concepts floating around until then under the general rubric of groupware. Collabra's product, called Collabra Share, shipped in 1994; it was a finalist for *PC Magazine*'s Technical Excellence Award for Systems Software that year, losing the crown to Microsoft Windows NT Workstation 3.5. (As Hahn likes to say, "That's a quality problem to have.")

Netscape acquired Collabra for $100 million in Netscape stock in late 1995, shortly after Netscape's mammoth and legendary initial public offering. By the time the lockup period ended six months later and the Collabra founders could sell their Netscape stock, its value had jumped to more than $250 million. "Every one of the engineers who had joined us at the beginning made over $1 million," Hahn recalls.

The sale to Netscape was a relief to Hahn, and not just for financial reasons, says Meredith, Collabra's vice president of engineering. "As CEO," she says, "he had the weight of the world on his shoulders. He's a responsible guy. He wants to always do the right thing, and he wanted a great outcome for everyone who worked for him. But he's got a little bit of the Andy Grove paranoia in him. He would imagine horrible things that would happen. The acquisition was a relief to him."

Hahn stayed at Netscape three years, running the server products division and then becoming chief technology officer. He says his biggest contribution at Netscape was in the decision to make the browser an open-source system—a highly controversial move at the time. When Microsoft famously cut off Netscape's air supply by releasing its free Internet Explorer in 1995, Netscape found itself in crisis. Hahn's "heretical" view—notably shared by Netscape developer Jamie Zawinski and sales engineer Frank Hecker—was that the only way out of this box was to turn the "lemon into lemonade." They would rally the goodwill of the Internet and its army of developers around Netscape by positioning it as an underdog. Meeting discreetly at CEO Jim Barksdale's house in Palo Alto, Hahn persuaded fellow executives Marc Andreessen, Mike Homer, and Peter Currie to pursue the plan.

During the browser wars, Netscape was really fighting two battles: the Microsoft battle was public and oft-reported, but less known were the company's internal struggles. Netscape had experienced years of hypergrowth. It had created an environment where speed was valued—users expected new products, and investors expected new revenues and profits. But, Hahn recalls, that also meant that developers rarely had time to do things the ideal way. The developers couldn't reliably add new features without destabilizing the existing, fragile code. Whenever he could, Hahn would ask each of a product's engineers to "sign the box" before a product shipped—in essence to publicly declare personal pride in the work. "If we couldn't do that, it probably was too soon to ship the code," he says. And in fact they weren't able to do it often. "We were in a fight for our lives," he adds, "and speed usually won out over other concerns. But I'm not complaining: war is hell."

Not every feature was unstable. In fact, one of Hahn's major triumphs was created in the heat of battle: he turned the Lightweight Directory Access Protocol (LDAP) into the Internet's de facto directory standard, allowing users to search white pages and directories for the first time. "Looking up e-mail and Web addresses was very easy to do after LDAP," Hahn says, "and very hard to do before."

Before 1996, the wilds of the Internet provided no standards for directories. Commercial users found that especially troublesome. At a company with 50 000 employees, there was no good way for one employee to look up another employee's e-mail address. Although commercial vendors existed—Novell Netware, for example, had a proprietary directory program—it was by no means Internet-wide, and in fact usually not even companywide. "One group would have Novell, another would have Microsoft, another had Unix, and it was a mess," Hahn recalls. "People forget, but in the early '90s every company had a different proprietary e-mail system—you couldn't even reliably send e-mail outside your company." That problem was fixed early with SMTP and POP, but the directory problem lingered for another decade.

Then Hahn heard about Tim Howes, a graduate student at the University of Michigan. Howes was solving the problem as part of his dissertation, but he was limited by what he could do as a student. Hahn plucked him out of Michigan and dropped him into Netscape, which put its full force behind the protocol. In 1997, LDAP, Version 3, won *PC Magazine*'s Technical Excellence Award for Networking.

In the end, Netscape was no match for Microsoft's deep pockets and ultimate weapon—the ability to bundle its browser with operating system software. This practice was later slammed by the U.S. Department of Justice as anticompetitive, but the resolution of the antitrust action against Microsoft came too late to help Netscape. AOL acquired Netscape in 1998 mostly, analysts said at the time, for the Netscape brand, Web site, and other nonbrowser software.

Hahn left Netscape just before AOL came in. He had promised to stay two years after the Collabra acquisition, and he fulfilled that promise; he had never planned to stay longer. He wanted to spend more time with his young sons—which had been impossible with the rigorous travel schedule Netscape demanded. He was hoping to slow down a bit after cc:Mail and Collabra. At the time, most friends and colleagues said he should join a venture capital firm.

Hahn turned down offers from a number of VCs. "I was worried about all those breakfasts and luncheons and that I'd have to wear a suit and that I wouldn't get to be hands-on with the

TURNABOUT: Eric Hahn helped C. Gordon Bell realize his dream of building the Computer History Museum in Silicon Valley: he and his wife, Elaine, donated more than US $10 million to the effort. The museum's collection includes this IBM 1620 from the 1960s.

TIMOTHY ARCHIBALD

technology," he says. He also decided against doing a start-up as a CEO, realizing that would be all-consuming. He looked into angel investing but decided that was too hands-off a relationship.

So he did something pretty interesting: he began "mentor investing." Hahn finds people who have an idea for a company and have started to pitch it to venture capitalists but are still trying to fill in the nitty-gritty. He writes a check, typically between $100 000 and $1 million, which buys him a piece of the eventual company and the right to be involved with developing the idea, hiring people, figuring out the marketplace—and even writing code. He often gets an office at the company and spends time there on a regular schedule. He holds the company retreats and holiday parties at his house.

Bill Lee, who founded one of Hahn's investments—RemarQ, a company that developed and operated Web-based discussion boards for companies including Amazon, eBay, and Novell—recalls that Hahn committed to spending a certain amount of time each month at the company but always exceeded that. "He got down to the level of critiquing the code," Lee says.

So far, more than 10 of his dozen-plus investments (a few are not public knowledge) are successful or on track. "I think of myself as the blue-haired lady in Atlantic City sitting at the slot machine with a box of quarters," Hahn says. "Success for me is simply not running out of quarters so I can keep playing."

Besides Red Hat and the recently sold Opsware, his hits include Good, which builds technology to keep travelers continuously connected wirelessly to corporate computers; Proofpoint, which makes an e-mail security product; and RemarQ. Four of his investments have already been sold; three others have gone public. And he's had two blatant failures: Disappearing Inc., an e-mail add-on, and Bing, a peer-to-peer project. The Bing experience is still fresh; he started the company in early 2006 to help people back up their computers, using free disk space on friends' computers. Hahn, three engineers, and a marketing executive worked feverishly for nine months to develop the product, but by the fall of 2006, Hahn recalls, "it became painfully clear that our thinking was flawed." People just wouldn't trust their friends' PCs with their data. In early 2007 he admitted defeat and returned the investors' capital, consoling himself that "with start-ups, if you don't have a few failures you're not taking enough risk."

Because in his current role he can do just about whatever he wants, Hahn's coding again. "I see the same aesthetic in code as I do in music or art or poetry," he says. "For me, coding has always been a creative and aesthetic joy, not just a technical joy. That creative outlet has always been the draw." Lookout, one of the companies that was sold, started as a personal project in 2002, when Hahn was looking for an excuse to program. He decided he could really use a search engine for his e-mail, and so he wrote one in C# for Microsoft's popular Outlook e-mail package. Today desktop search engines built into operating systems search e-mail as well as standard documents, but in 2002 that wasn't the case.

From the moment Lookout sprang to life, Hahn says, he "could no longer imagine not having it." He posted it on the Web for his friends to download and use. The software spread rapidly, with thousands and then tens of thousands and soon hundreds of thousands of users. And they all were turning to Hahn for support.

"I'd get bug reports from Norway saying, 'Hey, this weird thing happens when I run it in the Norwegian character set,'" Hahn recalls. "And I'm thinking, Norwegian character set? Huh?"

And then Hahn would stay up late trying to fix the bugs or adding features users requested. "When there's end-user passion, it's intoxicating," he says. "Franz from somewhere would send me a message, and I would find I really wanted to make him happy."

Hahn ran into Lookout users everywhere. On a family trip to Panama, visiting a coffee plantation that got power only a few hours a day, he spotted Lookout software running on the plantation's single PC. Within a year, though, it got to be too much. "I was hemorrhaging under the love," Hahn says. He hired Mike Belshe to work on the project full time; other friends helped when they could. He had no business plan. There was no income—he paid Belshe and the others out of his own pocket.

But for Hahn it was still emotionally rewarding. And more important, it made him a programmer again. "It was all about the code," he says. "I wrote 60 000 lines of code, and it was a product, and people liked it."

In 2004 Microsoft offered to buy Lookout, making, says Hahn, a generous offer for "what was basically two guys and a dog." Hahn and Belshe, who was by then a partner, sold it.

But once coding was back in his life, he wasn't going to give it up again. These days, though busy with Zimbra, he's still trying to solve the computer backup problem that he tackled with Bing.

"I wonder," he says, "how many programmers are trapped in the bodies of Silicon Valley executives. We tend to leave programming jobs because they just don't pay enough to support kids and mortgages here in Silicon Valley. But increasingly, when people have some material independence, they revert.

"It's a lot better than buying a football team," he says. ■

TO PROBE FURTHER

To find out more about Eric Hahn's mentor investing efforts or to submit a business plan, see http://www.ingroup.com.

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

SAFETY

Anatomy of a Crash-Test Dummy

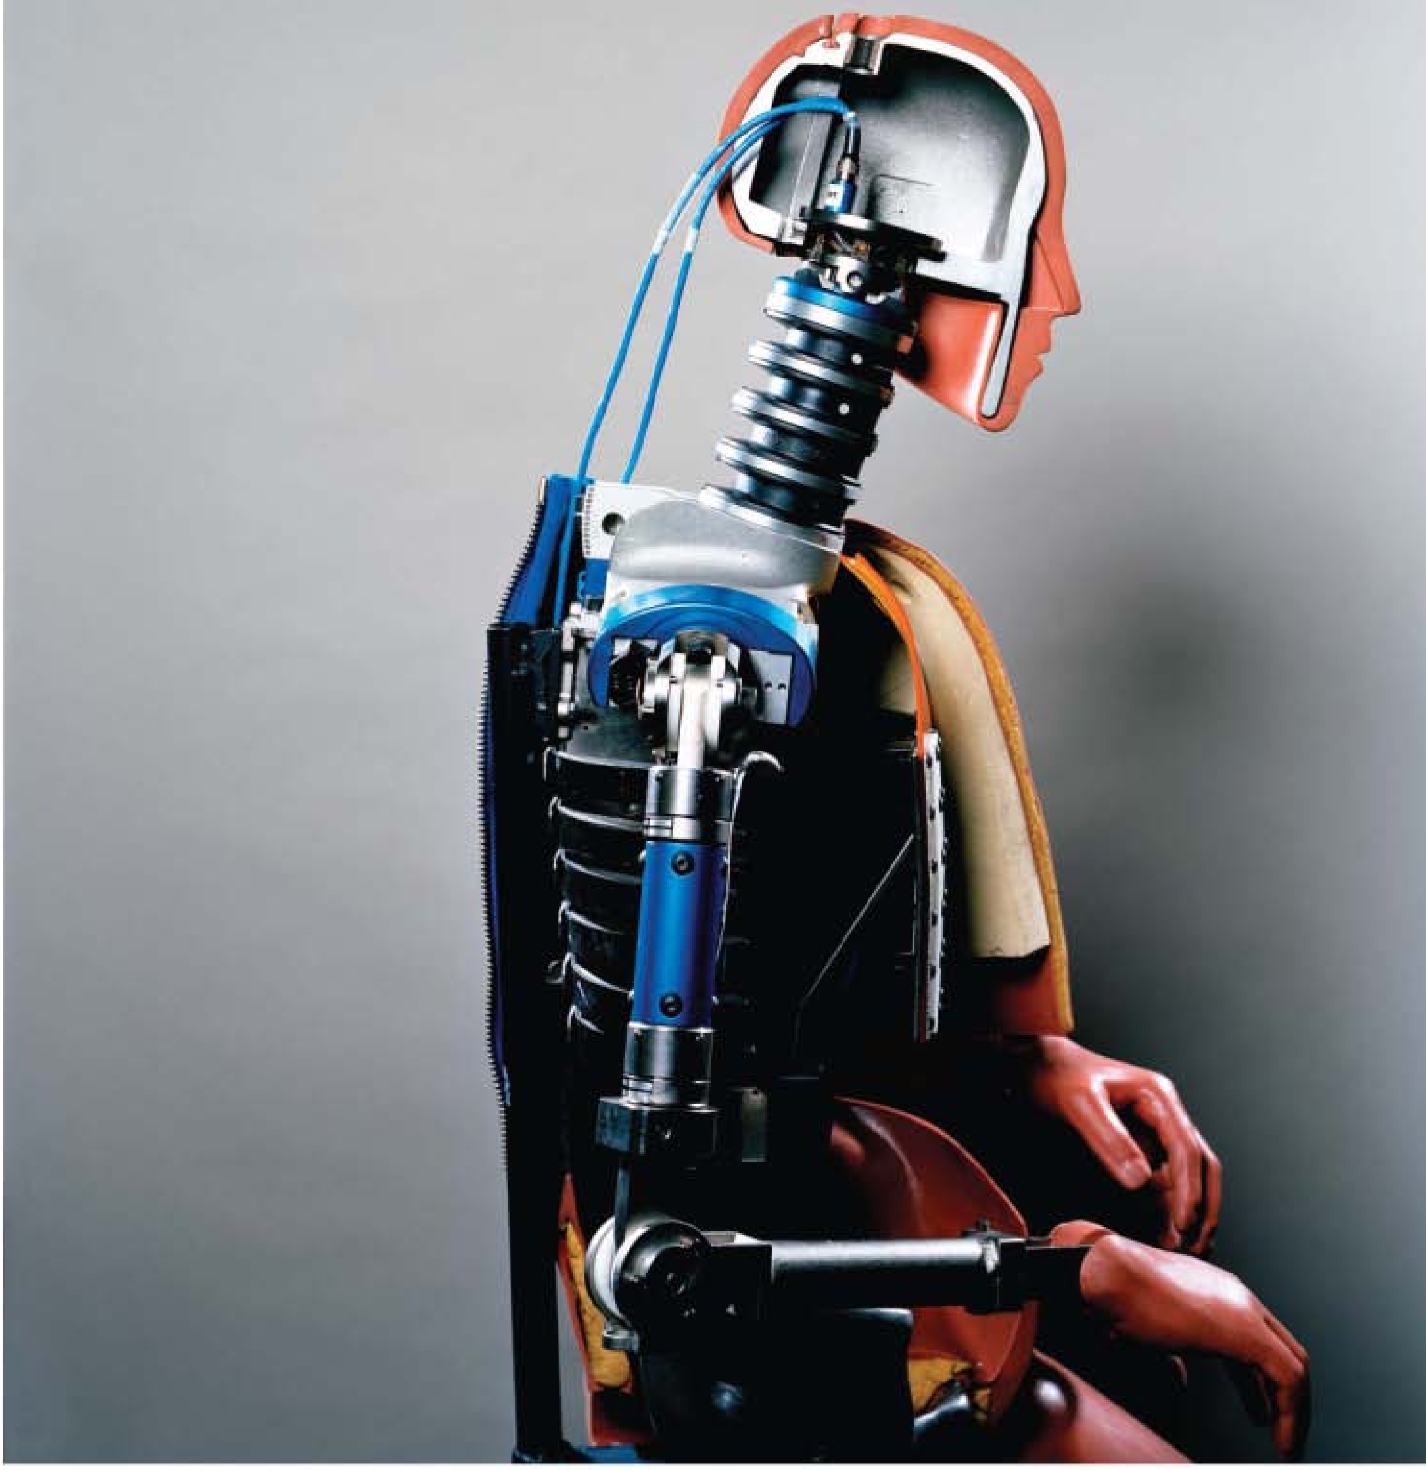

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

BY ERICO GUIZZO • PHOTOS BY CHRIS MUELLER

WE'VE BEEN LEARNING A LOT FROM DUMMIES

HALF A DUMMY: The vinyl flesh was removed from half of this dummy to expose its skeleton and sensors. Its maker, Denton ATD, normally assigns only serial numbers to its dummies, but this special unit, used in trade shows all over the world, was nicknamed Fred.

His bare buttocks rest on the cold steel shelf; the smooth, hairless skin has a ghastly pinkish-orange hue. His toeless feet lie nearby, alongside his head, rib cage, arms, hands, and legs. Could this be the grisly scene of some ritualistic slaying?

Not quite. A white-coated technician enters the room and transports the body parts to a wooden workbench. He takes an Allen wrench and screws the feet to the legs, the hands to the arms, and then the limbs and head to the torso. When he's finished, another Hybrid III midsize adult male anthropomorphic test device has begun to come to life. Or at least what passes for life for a crash-test dummy.

This Hybrid III is the handiwork of Denton ATD, a 170-employee company with facilities in Michigan and Ohio that manufactures some of today's most advanced crash-test dummies. These human surrogates simulate how a real person's body would respond in a car crash and help ensure that a new car's seat belts, air bags, head and armrests, structural frame, interior padding, and other elements provide good protection.

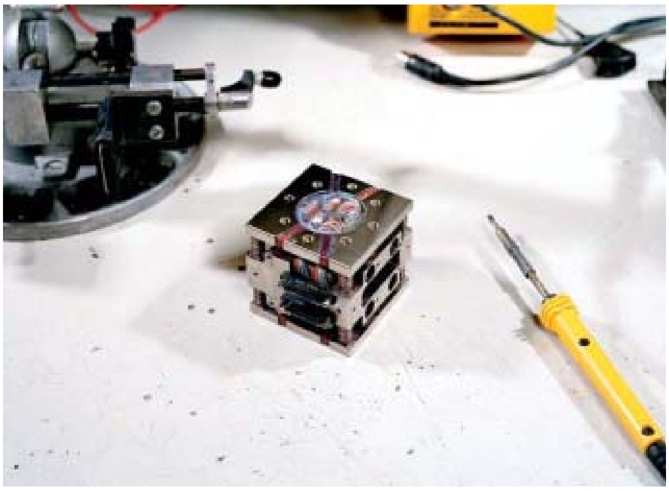

In a few days this new Hybrid III unit will be instrumented with force, torque, and acceleration sensors and then shipped to an undisclosed automaker in the Detroit area. There he'll be placed into brand-new cars and endure a torturous range of injury and insult: head-on collisions, rollovers, rear and side crashes—all to certify that the carmaker's vehicles can protect their human occupants in the event of an accident.

The dummy's ordeal, in other words, could someday save your life. But you'll probably never get to meet this electromechanical marvel. He doesn't even have a name. In the records that register the dummy's parts, his crash-test history and, ultimately, his retirement date, he'll simply be known as No. 0200-137.

THIS PAST SUMMER, I visited Denton to see how the company makes its extraordinary dummies. Denton's assembly plant sits amid cornfields just outside the picturesque town of Milan, Ohio (population 1445), birthplace of Thomas Edison.

When I step inside the company's unassuming building, the first things I see are body parts—everywhere. "Here's a thorax," says Mike Beebe, a senior vice president at Denton and one of the world's leading experts on the art and science of making dummies. "There's a spine box, with all the different pieces. That's an abdomen. Those are arms. Legs. Heads." I try to mentally arrange a full body out of the disordered parts, but what springs to mind is something alarmingly Picasso-esque.

Beebe points to a photo showing a group of dummies. "Family portrait," he quips. The family includes the most widely used dummy, the Hybrid III 50th-percentile male, meant to represent the average North American man. He weighs 78 kilograms and is 1.75 meters tall—or would reach that height if he could stand, which he can't, because he's in permanent sitting mode. Hybrid III has a petite wife (Hybrid III 5th-percentile female), three kids (Hybrid III 10-year-old, 6-year-old, and 3-year-old), and an oversized cousin (Hybrid III 95th-percentile male), who tips the scales at 100 kg—the "big guy," as Beebe puts it.

This family of dummies is designed for use in crash tests simulating frontal impacts: cars running squarely into other cars, trees, walls—that kind of thing. Also in Denton's catalog of 40 dummies are models for testing side impacts, rear impacts, accidents involving pedestrians, and air-bag blows on small children. Denton's customers include Chrysler, Ford, Honda, Hyundai, Nissan, Porsche, Volkswagen, India's Tata Motors, and China's FAW. The company ships 20 to 25 dummies a month.

Beebe explains that in a crash test, a dummy's sensors register a range of parameters: the force of a blow to the thigh, the torque on the neck during sudden deceleration, the compression of the chest against a seat belt. These measurements are then converted into injury criteria, which reveal the harm—anything from minor concussion to death—that would have been done to the vehicle's occupants had they been human. Such injury knowledge comes mostly from researchers studying how Newtonian mechanics applies to the human body, usually by performing impact and deceleration tests on cadavers, pig carcasses, or eager graduate students.

SMART PARTS: Denton's most complex force sensor [this page], the pelvis load cell of a WorldSID dummy, can record 12 data variables. FOCUS, an enhanced dummy face [opposite, top], has multiple load cells behind its eyes and facial bones. Waiting to receive new parts, a Hybrid III family sits in the Dummy Emergency Room [bottom, left] at Denton's assembly plant in Ohio. A dummy's vinyl and metal parts converge in the assembly area [bottom, right].

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

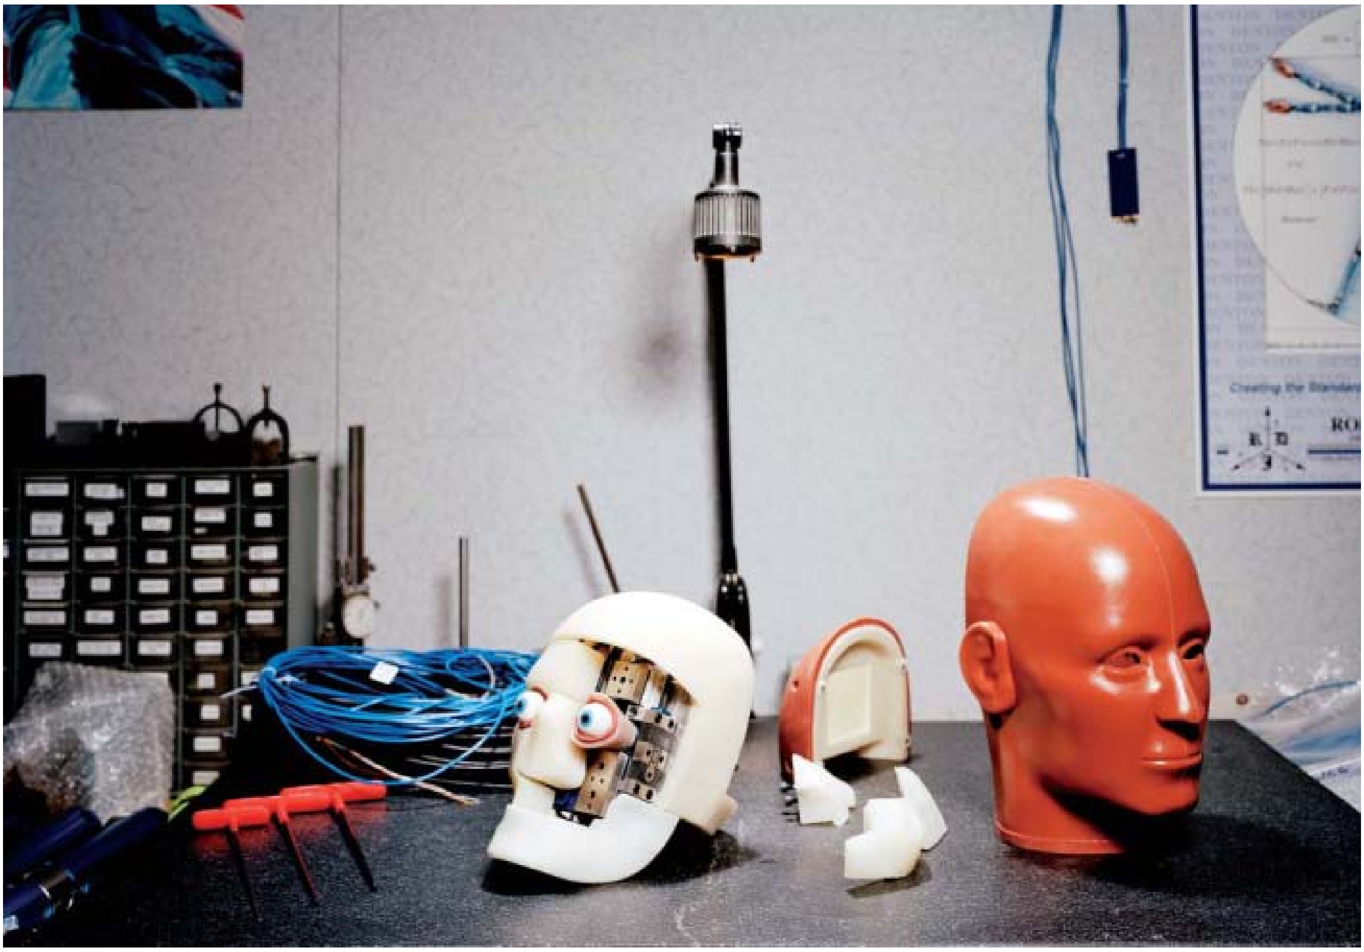

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

But Denton's dummies do more than car crashes. They've been used in roller coaster tests in Iowa and in simulated train wrecks in India. They've been dropped out of airplanes, strapped into crashing helicopters, and shot out of cannons. They've checked out school-bus seats, motorcycle air bags, and ski-slope protection nets. An Australian clothing company ordered a "perfect size 10" dummy to try on its new styles. And in a TV show, a Denton Hybrid III was punched in the face by a professional boxer and held in a neck lock by a Brazilian jujitsu fighter.

"There were some applications where we had no clue what they were doing," Beebe says. "It was proprietary or government related. The dummies left brand-new. They came back in parts."

A DUMMY LIKE NO. 0200-137 consists of 350 metal and plastic parts. Denton fabricates most of them itself, and just about everything is done by hand.

First comes the skin, the salmon-colored flexible plastic that covers a dummy's body. To make the feet, for example, a worker pours a milkshake-like substance—liquid vinyl—into an aluminum mold the size of a brick. The mold's interior is shaped like a foot (that is, a dummy's toeless foot), and the vinyl will solidify, or cure, when it goes into an oven. The skin for the dummy's head, upper arms, lower arms, hands, thighs, and shins is made the same way.

In another part of the factory, a group of workers fabricates steel and aluminum parts for the dummy's skeleton. One technician loads some specs into a computer numerical control, or CNC, machine, which automatically cuts, drills, and mills a steel part—in this case, an intricate disk for the dummy's shoulder. Over in another corner, a worker bends long strips of steel that will form No. 0200-137's ribs. Co-workers call this guy "Rib Man."

As the workers weld the smaller parts to the larger ones, pieces of the skeleton begin to take shape: skull, spine, hips, ankles, knees, elbows. The dummy's neck is more intricate. A large fraction of car accidents, especially rear collisions, result in severe neck injuries. To create a structure that can mimic the movement of a human neck, a worker mixes together natural rubber, polyacrylates, nitriles, neoprene, and butyl to obtain precise damping characteristics. He injects the mixture into a press that molds a handful of disk-shaped pieces, which will be alternated with metal rings to form the Hybrid III's characteristic segmented neck.

From different sectors of the plant, No. 0200-137's vinyl and metal parts converge in the assembly area, where they wait for the white-coated technician. The technician starts by measuring and weighing the head, limbs, and torso, and with a special scale he determines each section's center of gravity, which has to match that of a real person. To assemble the dummy, all it takes is a bunch of hex screws and a wrench. But No. 0200-137 is not ready yet. He needs some sensors.

ON MY SECOND DAY at Denton, I head out to its headquarters in Rochester Hills, Mich. From Beebe's remarks, I already have an inkling that the company appreciates the humor in the otherwise serious work it does. My suspicions are soon confirmed: dummy bobbleheads greet visitors at the reception desk; a poster of a dummy posing as Rodin's *The Thinker* hangs in a corridor.

When I walk into the corner office of Denton's president and CEO, David Stein, I get still more. "I spend my day with dummies!" is one of his favorite tension breakers. Stein, an electrical engineer turned dummy-industry executive, shows me his collection of crash-test miniature toys and dummy dolls bought on eBay.

The goal of my visit is to find out how the company is pushing the envelope of dummy design, and I had anticipated that Stein wouldn't be forthcoming with details. A dummy's specs, I reasoned, are probably like the formula for Coca-Cola or the blueprints for the Boeing 787—secreted away in a locked vault under heavy guard.

Not so. Dummy specifications in the United States are public, Stein tells me. You can walk into the offices of the National Highway Traffic Safety Administration in Washington, D.C., and request docket 74-14; in it you'll find schematic drawings, assembly descriptions, performance requirements, and a 16-page parts list for the Hybrid III 50th percentile.

The rationale behind keeping dummy specs open is so that automakers, safety equipment suppliers, dummy makers, and the NHTSA, which crash-tests most new models of vehicles before they can go on the U.S. market, are all on the same page. (The European Union and other countries have similar regulations.)

So in principle, anyone can get the specs and build a Hybrid III dummy. The challenge, Stein says, is consistency—meeting the requirements while making your next dummy indistinguishable from the previous one. Fashioning dummies that are like clones takes a lot of expertise. In fact, only two companies in the world have the know-how to build the most advanced dummies: Denton and First Technology Safety Systems, in Plymouth, Mich.

Making the dummy specs public and official also has a down-

DUMMY MAKING: [From left] Fabricating a dummy's parts is mostly done by hand: workers pour liquid vinyl into a hot metal mold to make a dummy's chest skin. Long strips of steel become dummy ribs in the hands of "Rib Man." At the Dummy Body Shop, excess vinyl is trimmed and parts are polished. A technician glues tiny electrical gauges to make a load cell.

side: you can't improve a given dummy model once the government freezes its design. The Federal Motor Vehicle Safety Standard No. 208, Occupant Crash Protection, which dictates how dummies should be built, among other things, was promulgated in the United States back in 1972 and hasn't changed substantially since then. So dummy makers have to stick with some 30-year-old designs and materials.

In fact, certain fabrication techniques date back to just after World War II, when the U.S. military developed the first modern dummies to test ejection seats in airplanes. Colonel John Stapp, a U.S. Air Force medical officer, pioneered the field of biomechanics with research that involved subjecting volunteers, including himself, to death-defying deceleration runs on sleds. He later realized it was more productive to develop and use crash-test dummies. An annual meeting named after him, the Stapp Car Crash Conference, is still the key gathering of car-crash testers, biomechanics researchers, and other safety industry experts.

In the late 1960s and early 1970s, with traffic accidents killing more than 50 000 people on U.S. roads each year, demand for safer vehicles grew stronger. In 1971, General Motors, which had been developing some dummy prototypes, decided to combine elements from two competing designs, one from Alderson Research Laboratories and the other from Sierra Engineering. GM named the resulting hybrid dummy, aptly, Hybrid I.

All the while, different makers created a myriad of other dummies, some now long retired, others still in active duty. The aerospace variety includes Model T Parachute Dummy, Torso, and Dynamic Dan. The medical profession has Rando for Radiotherapy, Dexter Dental Dummy, and the Phantom family—Cardiac Chest Phantom, Nuclear Phantom, and Organ Scanning Phantom.

The automotive sector raised its own dummy families: the now-retired Sierra Family (Sierra Sam, Sierra Stan, Sierra Susie, Sierra Saul, little Sierra Sammy, and Sierra Toddler) and the Hybrid clan (Hybrid I, II, and III—this last variety created by GM for the NHTSA, which made the design official in the late 1970s). More recent are the side-impact units: SID, EuroSID, BioSID, and WorldSID. These dummies, like the Hybrids, have specifications set by government agencies, and Denton and other makers follow such designs to manufacture and sell them to customers.

Today, WorldSID is by far the most advanced dummy. It was the first to be designed by a worldwide consortium of industry, government, and academic experts, with the goal of harmonizing test protocols and safety standards, which can vary widely from country to country. Loaded with sensors, it can record 258 different measurements in a single crash test. One unit can run close to US $350 000—more than twice the price of a Hybrid III.

DENTON DIDN'T START OUT as a dummy maker. The company originally focused on making force and torque sensors, called load cells, which are used in crash-test dummies but also in many other pieces of equipment, such as power tools and digital scales. Robert A. Denton founded the company in 1969. A quiet and creative engineer, he produced his first load cells at his home near Troy, Mich., using his kitchen oven to cure the components. Denton and his engineers went on to design many of today's most widely used automotive load cells—including the ones used in the Hybrid III.

After its mechanical parts and skin have been assembled, No. 0200-137 is shipped from Denton's Ohio plant to the company's load cell unit, adjacent to its Michigan headquarters.

In a vast, high-ceilinged hall, half a dozen milling machines hum away. Occasionally, a machine stops and a worker retrieves the newly milled part. These parts, made of aluminum and steel, come in all shapes and sizes and will form the structural elements of the load cells.

No. 0200-137's femur load cell, for example, is encased in a cylinder. Inside, a small metal beam traverses the cylinder's length. If you compress the cylinder, the beam will deform. To measure the extent of deformation, the load cell uses a thin zigzagging wire of brittle metal such as a titanium alloy. The electrical resistance of this wire, or gauge, changes when you compress or stretch it. When the gauge is glued to the beam, it converts the beam's deformation into a variable voltage. Add more beams and gauges, and your load cell can measure force in additional directions, and because you know the dimensions of the beams, you can also measure torque.

The metal parts are carted to Denton's electronics lab, where they will be outfitted with gauges. These need to be precisely glued to the center of the beam or the measurements will be distorted. Whereas the making of the dummy so far has required lots of heavy lifting, hammering, and milling, the alignment and gluing tasks require great finesse and hand-eye coordination.

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

Fourteen women are on the job. Sitting in cubicles decorated with flowers and children's drawings, they peer intently into microscopes and nimbly maneuver tiny tweezers and wire cutters.

Like other Hybrid III dummies, No. 0200-137 will receive two other types of sensors: accelerometers and potentiometers. To measure accelerations experienced by the dummy's head, for example, three uniaxial accelerometers are installed inside the skull at the center of gravity. The potentiometers measure deflection by translating their movement into voltage; the dummy receives one of these units behind the sternum to measure forces exerted by the seat belt or other object against the chest.

Now fully instrumented, No. 0200-137 has one last stop before it can leave the factory. It must be thoroughly checked out in the certification lab. Considering what goes on here, you could just as easily call this place the dummy torture chamber.

First comes the Head Drop Test—and the name pretty much says it all. A technician detaches No. 0200-137's head from the neck and hangs it from shafts at a precise height of 37.6 cm above a heavy block of steel. A magnetic release mechanism drops the head, which hits the block with a thud. This test ensures that the head has the right weight and damping properties.

Next, the technician reattaches the head to the body and places the dummy on a platform, positioning him so that his chest is in line with a 23-kg steel cylinder—the thorax impactor probe—hanging from above. *Three, two, one!* The probe swings down and connects with the dummy's sternum, sending No. 0200-137 flying backward into a net. The impact deflects the potentiometer, and the technicians check that the data fall within a certain range.

Minutes later, the dummy is getting his neck bent in an evil-looking contraption, and his knees whacked by another swinging probe. When it's all done, No. 0200-137 receives a certification stamp on his records and is ready to meet his future owner. Which means he gets packed into a cardboard box and trucked away.

THE DUMMIES OF THE FUTURE—what will they look like? Answering that question would require visiting the many organizations that in one way or another are advancing dummy technology—automakers such as GM and Ford, the NHTSA and its counterparts in other countries, research groups at places such as Wayne State University, the University of Virginia, and the University of Michigan, and the dummy manufacturers, of course.

But since I'm here at Denton, I pose the question to Randy Kelly, the company's sales vice president and a dummy spokesman of sorts. Kelly has been on television a bunch of times talking about dummies. The dummy of the future, he declares, is already here—in Rochester Hills. He's got two examples to show me.

One of the long-standing goals of biomechanics has been to find out just what goes on inside the rib cage during an accident, Kelly says. That's especially important in a side-impact collision, where the armrest, a door, or an SUV fender can hit passengers on the side, snapping their ribs. In current side-impact dummies, potentiometers attached to each rib register the rib's movement. But the device is tracking deflections only in one direction, so it's somewhat crude. Denton engineers came up with a better way to capture all that rib action.

"This is RibEye," Kelly says, pointing to a dummy's torso. Nothing appears unusual, but he explains that instead of potentiometers, each of the dummy's 12 ribs is equipped with an LED, and two light-angle sensors are mounted on the spine. "The sensors track the position of each LED or of a point on the rib," Kelly says. "It's just like celestial navigation that the sailors did back then."

The advantage of RibEye over existing methods is that it measures movement in all three dimensions, with an accuracy of 1 millimeter. Denton, which partnered with Boxboro Systems of Boxborough, Mass., to develop the system, has installed it in several dummies. Customers are now testing RibEye in R&D programs.

The second project Kelly mentions is FOCUS (facial and ocular countermeasure for safety headform), developed by Denton with the U.S. Army Aeromedical Research Laboratory and the Center for Injury Biomechanics, run jointly by Virginia Tech College of Engineering, in Blacksburg, and Wake Forest University School of Medicine, in Winston-Salem, N.C. FOCUS consists of an enhanced dummy face, with synthetic eyeballs

FACE IT: Denton and its partners designed the FOCUS system [opposite page] to study eye and face injuries and evaluate protective gear. In Denton's certification lab [this page], a dummy must endure a series of gruesome tests, including the Head Drop Test, required to ensure the head has the right damping properties.

made of a silicone-like material to register penetrating injuries and load cells at the back of each eye socket to measure nonpenetrating impacts. The face also has custom-made multiaxis load cells behind the frontal bone above the eyes, the zygomatic bones on each side of the eyes, the nasal bone, and the upper and lower jawbones.

What's FOCUS good for? Eye injuries to soldiers have increased dramatically since World War II, Kelly explains, so the Army plans to use FOCUS to evaluate helmets, goggles, and the protective features of its vehicles. The sensor-packed face could also be used to study air-bag face impacts, motorcycle-related injuries, and sports injuries. Oh, and popping corks. Turns out they account for about 10 percent of eye-related hospital admissions in Europe.

I ENCOUNTER NO. 0200-137 on a muggy afternoon in June at Autoliv, in Auburn Hills, Mich. Autoliv, one of the world's largest suppliers of air bags, seat belts, and other car safety systems, also performs crash tests for customers that don't have their own crash-test facilities. One such customer—No. 0200-137's owner—is conducting a frontal crash test today. The customer allows me to observe the test as long as I omit certain details that might reveal the automaker's identity and "stay away from the vehicle." In exchange, I get to watch a brand-new $40 000 car get totaled. Sounds like a fair trade to me.

The test takes place in a temperature- and humidity-controlled hangar. Eight other men and I sit inside a viewing room above the test track, kind of like a VIP box at a Formula One race. Except this race lasts about 10 seconds, and the track is just 200 meters long.

The car, a greenish-gray four-door sedan, sits at the head of the track. To the bottom of the vehicle, technicians attach a cable that will tow it down the center of the course. At the opposite end is a 45-metric-ton barrier made of reinforced concrete. It's 1.8 meters high, 1.8 meters thick, and 3.6 meters wide.

No. 0200-137 sits in the driver's seat, his expressionless face registering none of the last-minute preparations going on around him. The technicians carefully adjust the angle of his head, the space between his torso and the steering wheel, and the inclination of his thighs. The dummy wears a form-fitting cotton short-sleeve shirt, above-the-knee shorts, and a pair of $125 size 11 black oxfords. It's all to reproduce real driving conditions (people usually drive with their clothes on, after all).

Suddenly a horn goes off. An orange light flashes. Huge light panels flood the test bay. Fifteen high-speed video cameras begin rolling. We hear the noise of the tow cable dragging down the track. In just 3 seconds, the sedan accelerates from 0 to 48 km/h, and for another 7 seconds it maintains that exact speed. An instant before impact, the tow cable is released, and in the next instant, the car crashes into the barrier. *Boom!* Headlight fragments fly off in all directions. The back wheels almost jump from the ground. Then—just silence.

Measurements from the dummy will be transferred to computers, processed, and translated into an injury criteria report, which will tell the carmaker how real passengers would have fared in such a crash. I'm not allowed to see No. 0200-137 afterward. But odds are he's still in good shape. He may need a replacement part or two or a realignment or maybe some new clothes—but little more than that. For the next decade or so, this will be his routine. Until the day a new dummy takes his place. ■

TO PROBE FURTHER

To see more photos of dummies along with a video of a rib-impact test, visit http://www.spectrum.ieee.org/oct07/dummies.

For technical details about Denton's dummies, go to http://www.dentonatd.com.

Accidental Injury: Biomechanics and Prevention, edited by Alan M. Nahum and John W. Melvin (Springer-Verlag, 2001), has several chapters on how dummies are used in crash tests.

For an entertaining view of the field of biomechanics, see *Stiff: The Curious Lives of Human Cadavers*, by Mary Roach (W.W. Norton and Co., 2003).

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

PHOTOGRAPHS: MARK LEONG

ARTIFICIAL INTELLIGENCE

Cracking GO

Brute-force computation has eclipsed humans in chess, and it could soon do the same in this ancient Asian game

BY FENG-HSIUNG HSU

IN 1957, HERBERT A. SIMON, a pioneer in artificial intelligence and later a Nobel Laureate in economics, predicted that in 10 years a computer would surpass humans in what was then regarded as the premier battleground of wits: the game of chess. Though the project took four times as long as he expected, in 1997 my colleagues and I at IBM fielded a computer called Deep Blue that defeated Garry Kasparov, the highest-rated chess player ever.

You might have thought that we had finally put the question to rest—but no. Many people argued that we had tailored our methods to solve just this one, narrowly defined problem, and that it could never handle the manifold tasks that serve as better touchstones for human intelligence. These critics pointed to *weiqi*, an ancient Chinese board game, better known in the West by the Japanese name of Go, whose combinatorial complexity was many orders of magnitude greater than that of chess. Noting that the best Go programs could not even handle the typical novice, they predicted that none would ever trouble the very best players.

Ten years later, the best Go programs still can't beat good human players. Nevertheless, I believe that a world-champion-level Go machine can be built within 10 years, based on the same method of intensive analysis—brute force, basically—that Deep Blue employed for chess. I've got more than a small personal stake in this quest. At my lab at Microsoft Research Asia, in Beijing, I am organizing a graduate student project to design the hardware and software elements that will test the ideas outlined here. If they prove out, then the way will be clear for a full-scale project to dethrone the best human players.

Such a result would further vindicate brute force as a general approach to computing problems, if further vindication were needed. Even now, the method is being applied to such forbidding challenges as protein folding, scheduling, and the many-body problem.

MANY OF THE EARLY computer-chess researchers hailed from the fields of psychology or artificial intelligence and believed that chess programs should mimic human thinking. Specifically, they wanted computers to examine only playing sequences that were meaningful according to some human reasoning process. In computer chess this policy, known as selective search, never really

CHESS VS. GO

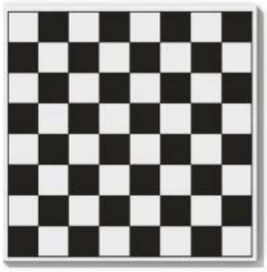

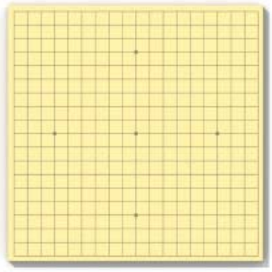

	Chess	Go
GRID SIZE	8 x 8	19 x 19
AVERAGE NUMBER OF MOVE CHOICES PER TURN	35	200–300
LENGTH OF TYPICAL GAME	60 moves	200 moves
NUMBER OF POSSIBLE GAME POSITIONS	10^{120}	10^{170}

EXPLOSION OF CHOICES
(starting from average game position)

Chess		Go
35	Move 1	200
1225	Move 2	40 000
42 875	Move 3	8 000 000
1 500 625	Move 4	1 600 000 000

THE GOAL OF GO

The object of Go is to enlarge your territory at your opponent's expense. One way is by surrounding your opponent's stones by putting your own stones on the adjacent points, which are known as "liberties." Once surrounded, stones are removed from the board and become your prisoners, each worth a point.

Another way to claim territory is by surrounding empty space—that is, unoccupied intersections, each of which is also worth a point. Here, for instance, the two groups in the corner each enclose nine spaces, worth as many points. Obviously, it takes fewer stones to enclose territory at the corners than in the middle of the board.

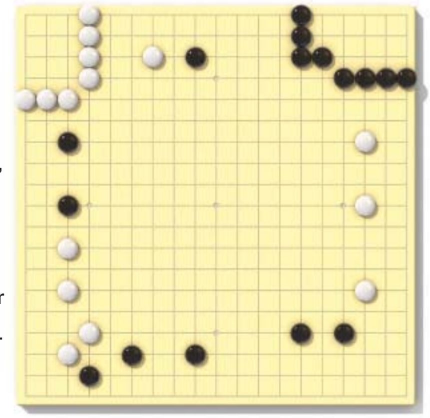

CAPTURING STONES

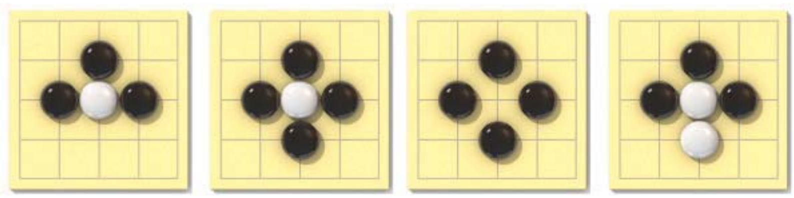

Here, with one more move the white stone will be surrounded...

...leaving it "dead"...

...so that it may be taken off the board.

If it were white's turn to move, however, the player could block the above maneuver, thus.

made progress. The reason is that humans are extremely good at recognizing patterns; it is one of the things that we do best.

It was only in the late 1970s, with the success of Northwestern University's Chess 4.*x* program, written by David Slate and Larry Atkins, that the engineering school of thought became dominant. The idea was to let computers do what *they* do best, namely, calculate. A simple legal-move generator finds all the permissible moves in a position, considers all the possible responses, and then repeats the cycle. Each cycle is called a ply, each generation of new possibilities is called a node—that is, a branching point in a rapidly widening tree of analysis. The branches terminate in "leaf," or end positions.

Carried to its logical extreme, the tree would grow until it exhausted every legal continuation, leaving the program nothing to do but examine the end positions to see which of them were wins—that is, checkmates—and which were draws, then work backward along the branching structure to choose the line that led to the best outcome, assuming that both sides play perfectly. Such exhaustive analysis is impractical, though, because it would produce a tree containing about 10^{60} positions. That's about a thousand times the number of hydrogen atoms in the sun.

There is, however, a course midway between selectivity and exhaustiveness. Instead of analyzing to the end, the program can merely look a few moves further ahead than a human could manage. Deep Blue typically looked 12 plies ahead in all variations (and 40 or more plies in selective lines), generating around 170 million leaf nodes per second. Next, the program would evaluate each of these positions by counting "material," that is, the standard values of the chess pieces. For example, a pawn is worth one point, a knight or bishop three, and so on. Then it added points for a range of positional factors, chosen with the help of human grandmasters.

The resulting evaluation function probably was no better than a middling amateur's ability to grade a single position. But by grading 200 million of them, it was able to do very well indeed. Just ask Kasparov.

This substitution of *search* for *judgment* is the essence of the brute-force method, and it turned out to have two critical advantages over selective search. To begin with, the program became easier to write, had far fewer bugs, and did not have so many blind spots. And crucially, the program played significantly and measurably better as the processing power increased, once the switch to brute force had been made.

Slate and Atkins believed their program was playing at only Class C level—that is, about the level of the typical avid tournament player, who is rated between 1400 and 1600 on the U.S. Chess Federation's rating scale. However, when they moved their program to a supercomputer, it shocked everyone by winning a tournament among Class A players, with ratings between 1800 and 2000. A Class A player is good enough to beat a Class C player 9 times out of 10, on average.

Moving to a supercomputer made this enormous difference because it allowed the program to look just a little further ahead. Detailed measurements later showed that when a brute-force program searched just one ply deeper, its strength improved by between 200 and 300 rating points. When two players are separated by that big a gap, the higher-rated player will win, on average, 4 out of 5 games.

It was this almost linear relationship between search depth and playing strength that first made me believe chess could be solved. I wondered whether the relationship would continue

DIAGRAMS: BRYAN CHRISTIE DESIGN

all the way up to the World Champion level—about 2900 on the Chess Federation's scale. In the end, this conjecture proved to be partly true. That is, the program did continue to play better as search depth increased, but additional gains in rating could also be achieved by improving the evaluation function and the selectivity of its search.

DEAD OR ALIVE?

To play any board game well you have to assess the situation on the board astutely, over and over again. For Go, doing this involves determining whether a group of connected, like-colored stones (yours or your opponent's) is "alive" or "dead." Stones that cannot be captured are alive; spaces that are surrounded by living groups of the first side that cannot sustain living groups of the second side belong to the first side. The game ends when both sides agree on the final disposition of territory.

The challenge of Go comes from the fact that analyzing whether a group of like-colored stones is likely to live or die can be a hugely tricky affair. In the figure below, the black stones are

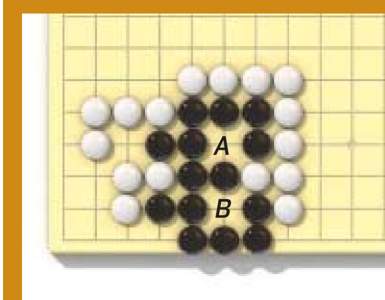

unconditionally alive as they have two "eyes," indicated by *A* and *B*. The liberties *A* and *B* cannot be occupied by white (for a white stone placed in either spot would itself be dead), and therefore black stones cannot be captured no matter what.

However, the situation is usually more complicated, and to judge whether a group is alive or dead the program must search many moves ahead. In the figure below, if black places a stone on *C*, then

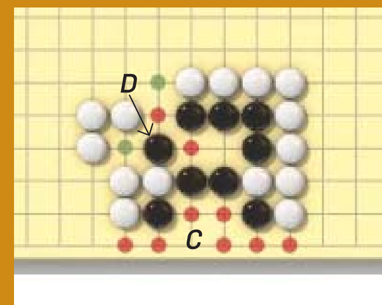

the black stones will live, but not unconditionally. That is, black may have to make additional moves to keep the stones alive if white plays nearby. For instance, if white plays on any of the points marked in red, then black will have to respond appropriately—and immediately—to keep the black group alive.

Consider the black stone labeled *D*. If the nearby green spaces are occupied by white and black doesn't react right away, then white can kill black by playing on the space to the right of stone *D*. If black captures the new stone, then white plays on the space above stone *D* and destroys the eye. If black connects by playing on the space above stone *D*, then white captures the four newly connected stones, killing the entire group.

A program would have to follow these branching possibilities to see whether the black stones are alive or dead, a far more arduous job than simply counting up men and positional features, as in chess. *—F.H.*

GO IS PLAYED on a board crisscrossed by 19 vertical and 19 horizontal lines whose 361 points of intersection constitute the playing field. The object is to conquer those intersection points.

A player makes a move by placing a lozenge-shaped "stone" on an intersection, then the other player counters, and the two alternate moves. Players capture enemy stones by surrounding them, that is, by removing their "liberties," which consist of either the vacant points adjacent to a stone itself or to friendly stones to which it is itself connected (see illustration, "The Goal of Go"). When no more moves are possible, the players count up the intersection points they control, and the player with the most points wins.

All the leading Go programmers today belittle brute force. In this they resemble the computer chess experts of 40 years ago. Selective search dominated thinking on computer chess from the late 1940s to the late 1970s, and that mind-set prevented any program from advancing beyond the level of a Class C player.

Go does, however, present two real problems, both having to do with the amount of searching the program must perform.

The first problem is the tree of analysis. Because Go offers more possibilities at every turn, the tree is far bigger for Go than for chess. At the start of the game, the first player can place a stone on any one of 361 positions, the second player has 360 choices, and so on. A typical game lasts about 200 moves, so it averages at least 200 move choices per turn—nearly 10 times as many as in the average chess position.

The second problem is the evaluation of the end positions. In Go you can't just count up stones, because you have to know which stones are worth counting. Conquered territory is defined as board space occupied or surrounded by "living" stones—stones the opponent cannot capture by removing their liberties. Before you can count a stone as live, you have to calculate several moves ahead just to satisfy yourself that it is really there in the first place.

Put these two problems together and you get a computational problem that at first glance seems intractable. But there are ways to engineer around it.

LET'S START with the problem of the exploding tree of analysis. If we assume that the program must consider every possible continuation that could arise 12 plies into the future, as Deep Blue did in chess, you might expect to have to search a million times as fast. But we don't really need to pay that high a price, because there are ways to prune the tree of analysis.

One old standby, implemented in all chess programs, is called alpha-beta pruning, and it works by curtailing the examination of a move the moment it becomes clear that another move must be better. Let's say the program is comparing move A with move B, and it already knows that A leads to an advantage. If it finds, early on, that move B allows the other side to obtain a draw at the very least, then the program can cut off its analysis, saving a lot of time.

Alpha-beta pruning reduces the effective branching factor to about the square root of the number of move choices. For example, to look ahead 12 plies in pure brute-force mode, you would need to search only about 4 billion positions, or 4×10^9, instead of 38^{12}—or 10^{19}—positions.

A newer way to cut back the overgrowth—null-move pruning—was not implemented in Deep Blue, even though one of its inventors, Murray Campbell, was a key member of the Deep Blue team. The algorithm performs a kind of thought experiment, asking how the position would look if you were to give up the right to move for one turn, thus allowing your opponent to make two moves in a row. If after that enormous sacrifice you still have a good position after a relatively shallow search, then the algorithm can stop

CAN MONTE CARLO WORK ON GO?

Some of the best Go programs today employ Monte Carlo methods, which play out move possibilities internally, in random games, then select the move with the best win/loss index. It can be considered a brute-force technique.

Monte Carlo has long been known to work reasonably for games of "imperfect information" such as backgammon, in which the rolling of dice introduces an element of chance. The method is also a good option for games of perfect information that are too complex to crack by more straightforward means. Had it been applied to computer chess back in the 1950s, before today's search algorithms were perfected, it might well have raised the standard of play.

Monte Carlo techniques have recently had success in Go played on a restricted 9-by-9 board. My hunch, however, is that they won't play a significant role in creating a machine that can top the best human players in the 19-by-19 game. Even so, Monte Carlo is worth keeping in mind for games and gamelike computing challenges of truly daunting complexity. *—F.H.*

its analysis right there. It has identified a cutoff point—a point at which the branch can be pruned, thus saving the labor of going through all the other possible responses.

Imagine that the program examines a line in which it has just won the opponent's queen—giving it an enormous material advantage—and the opponent has responded. Now the program asks: If I do nothing, can my opponent hurt me after, say, $n-2$ plies—where n is the number of plies I would have searched after a legal instead of a null move? If the answer is no, the program concludes that it has indeed won an entire queen for nothing, that its position is likely won, and that no further analysis is necessary. This dividend is well worth the shallow "$n-2$ ply" search the computer has invested.

In computer chess, the main risk in null-move pruning comes from the null move (or pass) itself, which is illegal in chess. Because it is illegal, certain positions that could be defended by a pass must lose; the null-move trick can cause a program to ignore this condition. In Go it doesn't matter, though, because players are allowed to make passes.

Null-move pruning was first proposed as a fairly conservative technique, curtailing the depth of search only by $n-1$ plies, but experimenters soon found that $n-2$ or even $n-3$ reductions sometimes gave good results. Even better performance comes from applying null-move pruning *inside* the reduced-depth search itself. Such "recursive null-move pruning," when coupled with standard alpha-beta pruning, appears to reduce the branching factor to about the square root of the square root of the number of move choices. This means that recursive null-move pruning can keep the analysis tree from growing any faster in a Go program than it would in a chess program that did not use null-move pruning.

The upshot is that a machine searching no faster than Deep Blue did 10 years ago could go 12 brute-force plies deep in Go (with additional selective search extensions). It does so, however, without making a full and proper evaluation of the resulting positions, as it could do for chess.

YET ANOTHER TIME-SAVING TECHNIQUE emulates human thought (for a change). When human players search through the Go game tree, they generally check the live-or-dead status of each stone only once, then in effect *cache* the result in their memories. In

other words, they don't check again unless they have good reasons to do so. The point of caching is to fetch often-used information rather than recalculate again and again.

The idea has been tried in computer chess, in the "method of analogy" algorithm, which reuses conclusions reached in one branch of analysis when similar positions arise in other branches. It reduces the search tree by a factor of three or four, but unfortunately the operations needed to cache, retrieve, and apply the conclusions slows the program by the same proportion. To wring a net gain out of the method, therefore, the slowdown must be contained, for instance, by using special-purpose hardware to do the computation or by finding new ways to chop the search tree even further.

Think of the tree again. What the method of analogy basically does is to take an entire branch from one part of the tree and graft it to another. Suppose that early on the program discovers a sequence in which white can win in just one ply, by capturing black's queen with a bishop. The program will then cache that sequence and apply it to latter parts of the search tree, provided that nothing major has happened in the meantime (like losing a piece) and that the bishop can still capture the queen.

In chess, this method of analogy works only for very short branches or for branches that contain mostly "forced" moves, that is, checks, check evasions, and captures. However, if the branches contain more than a ply or two of nonforcing moves (which present far more possibilities for calculation), then the program's accounting system breaks down.

The reason has to do with the nature of the accounting system, which consists of a map of on/off bits that tracks the "to" and "from" squares of each chess piece. The program uses this bitmap to decide whether anything has happened to invalidate the graft—for instance, by making the winning move in a grafted branch illegal or providing the losing side with a way out of a sequence of forced moves. It turns out in chess that if grafted branches contain more than one ply of nonforcing moves, the bitmaps will quickly cover most of the board, the accounting system will become unmanageable, and the grafting operation will fail.

In Go, however, the method of analogy should be much more useful. Because the board is so large (19 by 19 versus 8 by 8 in chess), a battle typically occurs in a relatively small part of it, so the bitmaps will mostly have "off" bits, making it more likely for them to be useful. Also, the program can generally reuse the maps many more times than in chess, because each of the many local battles tends to be unaffected by battles elsewhere. Therefore, the program should be able to graft deep branches—the kind needed to decide life-and-death questions—from one part of the game tree to another.

This ability to answer life-and-death questions cheaply is vital if the brute-force approach is to work. To determine whether a group of pieces will live or die the program may have to search from 1000 to 1 000 000 positions. That wouldn't be so bad, really, if it were the extent of the problem. It isn't. In a typical game, we may easily have more than 10 such problems on the board at the same time, and the status of one group can affect that of its neighbors—like a cowboy who points a revolver at another cowboy only to find himself covered by a rifleman on a roof. Such interactions can complicate the problem by something on the order of taking 1 million to the 10th power—enough to stretch a calculation lasting a microsecond into one vastly dwarfing the age of the universe.

This is where the bitmaps we mentioned earlier come to the rescue. They make it easy to tell when maps do and do not intersect and also allow caching to work, thereby drastically reducing the cost of dynamic search required for proper evaluation of positions.

It is conceivable that with caching techniques, including but not limited to the method of analogy, it may take no more than 1000 to 1 000 000 nodes (or one individual life-and-death decision tree) of dynamic search to properly evaluate an end position. Although that's more expensive than in the case of chess, it's manageable.

WHAT, THEN, CAN WE EXPECT FROM THE HARDWARE? Deep Blue used 0.6-micrometer CMOS technology, kind of creaky even in 1997. Each of its 480 custom-designed processors searched up to 2.5 million positions per second. The theoretical peak speed was more than 1 billion positions per second, but the sustained speed was only 200 million positions per second because of communication overhead, load-balancing issues, and implementation inefficiency.

Today 45-nanometer process technology is just getting into production. With it, a machine searching as fast as Deep Blue could easily fit on a single chip. In fact, with gains expected from technology and from optimization of chip architecture, a single-chip machine could actually be more than 100 times as fast as Deep Blue. If we then made 480 copies of that monster chip and integrated them all in a parallel architecture, we could get at least another 100-fold increase in computational power. On top of that, in 10 years Moore's law is expected to present us with still another 100-fold speedup.

Put it all together and you should be able to build a machine that searches more than 100 trillion positions per second—easily a *million times* as fast as Deep Blue.

That would be enough to build a tree of analysis for Go as big as Deep Blue's was for chess and to evaluate all its end positions properly. If we assume the top Go players calculate about as deeply as the top chess players do, the result should be a machine that plays Go as well as Deep Blue played chess.

Well enough, that is, to beat any human player.

MY GUT FEELING is that with some optimization a machine that can search a trillion positions per second would be enough to play Go at the very highest level. It would then be cheaper to build the machine out of FPGAs (field-programmable gate arrays) instead of the much more expensive and highly unwieldy full-custom chips. That way, university students could easily take on the challenge.

At Microsoft Research Asia we are seeding university efforts in China with the goal of solving some of the basic problems. Whether these efforts lead to a world-champion Go machine in the next decade remains to be seen. I certainly wouldn't bet against it. ■

ABOUT THE AUTHOR

FENG-HSIUNG HSU earned a Ph.D. in computer science at Carnegie Mellon University, Pittsburgh, where he and fellow students designed the first grandmaster-level chess machine. Then he moved to IBM to develop its successor, Deep Blue, which beat World Champion Garry Kasparov in 1997. Hsu now manages the platforms and devices center of Microsoft Research Asia, in Beijing.

TO PROBE FURTHER

For a full account of the IBM project to build a chess machine, see *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion*, by Feng-hsiung Hsu, Princeton University Press, 2004.

To experiment with a Go program, readers can download GNU Go at http://www.gnu.org/software/gnugo. Offered by the Free Software Foundation, in Boston, this free program has performed well in recent computer Go events.

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

50th ANNUAL IEEE GLOBECOM

2nd ANNUAL COMMUNICATIONS INDUSTRY FORUM & EXPO

26 – 30 November 2007 • Hilton Washington Hotel • Washington, DC

PLAN TO ATTEND THE PREMIER TELECOMMUNICATIONS EVENT

Where industry meets to network, exchange ideas, and learn the latest developments

50th ANNIVERSARY CELEBRATION

Dr. Irwin Jacobs
Co-Founder & Chairman
QUALCOMM, Inc.

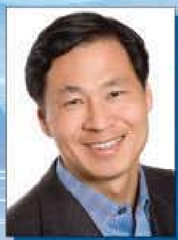

Dr. Jeong Kim
President
Bell Labs at Alcatel-Lucent

Prof. Leonard Kleinrock
UCLA

PLENARY SPEAKERS AS OF SEPTEMBER 2007

Matt Bross
Chief Technology Officer
BT Group

Scott McGregor
CEO
Broadcom

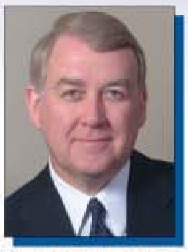

Mark A. Wegleitner
Senior Vice President, Technology
& Chief Technology Officer
Verizon Communications

Prof. Wu Hequan
Vice President
Chinese Academy of Engineering

COMMUNICATIONS INDUSTRY FORUM & EXPO

> ACCESS '07 EXECUTIVE BUSINESS FORUM
12 Sessions

> DESIGN & DEVELOPERS FORUM
9 Sessions

> TUTORIALS & WORKSHOPS
20 Tutorials & 9 Workshops

EXPO EXHIBITORS AS OF SEPTEMBER 2007

- Auerbach Publications - Taylor & Francis Group, LLC
- Cambridge University Press
- CANTO
- Elsevier, Inc.
- Institute for Information Industry, Taiwan
- The Institute of Engineering and Technology
- NIKSUN
- Now Publishers, Inc.
- OPNET Technologies
- Springer
- Telcordia
- University of Maryland - Engineering
- Wiley Blackwell

IEEE GLOBECOM

> GENERAL SYMPOSIUM
17 Sessions
100+ Technical Presentations

> 9 TECHNICAL SYMPOSIA
164 Sessions
900+ Technical Presentations

> 50th ANNIVERSARY CELEBRATION
IEEE GLOBECOM Panel
Trivia Contest with prizes

> EntNet @ IEEE GLOBECOM
Keynote address
8 Panel Sessions & 2 Tutorials

CONFERENCE PATRONS • AT&T • Broadcom • Institute for Information Industry, Taiwan • Telcordia

IEEE GLOBECOM 2007
GLOBAL COMMUNICATIONS CONFERENCE
EXHIBITION & INDUSTRY FORUM

REGISTER by 31 OCTOBER and SAVE!
www.ieee-globecom.org/2007

IEEE
IEEE COMMUNICATIONS SOCIETY

RESOURCES

From Nerd to Wonk

Tired of designing devices that policy-makers will misuse? Go back to school and train to become a policy maven yourself BY PRACHI PATEL-PREDD

OUR MAN IN D.C.: Paul Parfomak parlayed his engineering smarts into a policy job in Washington.

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

RESOURCES

THE DOCTOR IS IN: A Ph.D. in engineering and public policy puts Parfomak "in the sweet spot."

Lined up end to end, the pipelines that carry oil and gas around the United States could reach the moon and back.

CAREERS

Should the United States need to import liquefied natural gas from Canada or distribute ethanol coast-to-coast, the country will need still more pipelines, all of them vulnerable to mishap or malice. It's all grist for Paul Parfomak's mill.

Parfomak analyzes pipeline security for the Congressional Research Service, a nonpartisan agency sometimes referred to as "the brains of Congress." He estimates the risks of accidents and terrorist attacks and draws up mitigating policies for members of Congress to consider. The job requires an understanding of a wider range of issues than most professionals normally cover.

"You simply would have a much more difficult time understanding the policy dimensions if you didn't understand the engineering," Parfomak says. However, you also need to understand the effects of the price of oil and gas, as well as the manifold ways in which the governmental and legislative processes work.

Fortunately for Parfomak, his Ph.D. in engineering and public policy from Carnegie Mellon University, in Pittsburgh, has trained him to think not just as an engineer but also as a risk analyst, economist, and social scientist. His degree, he says, puts him "in the sweet spot."

ENGINEERING AND POLICY programs originated in the United States in the 1970s at Carnegie Mellon, MIT, and Stanford. Others were later established at Washington University and the University of Maryland. In Europe, the Netherlands leads, with technology and policy programs at Delft University of Technology and Eindhoven University of Technology, and emerging programs at Utrecht University and a few other schools. Other programs are found at the Instituto Superior Técnico, in Lisbon, and at the University of Cambridge, in England.

Graduates work mainly in local and national government agencies, such as the U.S. Environmental Protection Agency, the Department of Energy, and NASA, and in international organizations such as the World Bank and the United Nations Development Program. Some also work for corporate giants such as AT&T and Lockheed Martin.

All the programs train engineers to solve problems that lie at the junction of technology, society, and politics—say, where best to string new power lines, or how to estimate the environmental impact of public transportation, or how to evaluate the safety of cellphones on airplanes. IEEE Fellow Patrick O'Shea, chairman of the electrical and computer engineering department at the University of Maryland, in College Park, says that the standard engineering curriculum teaches students to ask whether something is possible, and if so, whether it is practical. That may be fine for computers and iPods, he says, but it won't work for energy and transportation, where "things may be feasible and even practical, but there is tremendous resistance to them for various reasons that have nothing to do with technology."

Dava Newman, head of MIT's program, says that the idea is to coax engineers out of their discipline's black-and-white realm to engage in the gray areas of policymaking, where debate is a way of life. "We firmly believe we're trying to train engineers with a difference, so that they might go out and lead."

In addition to technical courses, the curriculum covers the methodologies of policy, economics, decision analysis, risk analysis and assessment, and management. Students learn to take big, messy, unstructured problems and then identify the most important pieces and ask the right questions.

To round out the qualifying exam for a doctorate at Carnegie Mellon, several faculty members spend the better part of a month creating a fictional scenario that the students must work out in five days. This year, students had to decide what a real estate company in Florida's Miami-Dade County should do to prepare for a 2-meter rise in sea level over the coming century. "They had a bunch of options. Do we abandon the city, do we turn it into the

THIS PAGE AND PREVIOUS: CHAD DOWLING

Venice of the West, or do we try to dike the whole area?" says IEEE Fellow Granger Morgan, the head of the program.

The correct answer? There isn't one, but "there are typically lots of wrong answers," Morgan says. The students needed to recognize, for instance, that just one dike would be a disaster should it fail, and that the city would need multiple dikes, which would carry a hefty cost. Also, moving out might be safe in the long run, but it might be a public-relations disaster.

These programs are popular choices for engineers with a bit of work experience. "On average, [students] entering the program are about 25, with a few years of experience," says MIT's Newman. According to Morgan, having such people makes for a richer school environment for the entire class.

The interdisciplinary nature of technology and policy attracts a lot of women. At Carnegie Mellon, the engineering and public-policy program has a higher proportion of women than any other engineering program on campus does.

The popularity of engineering and policy programs should increase as society faces ever more challenges falling at the intersection of technology and policy. The environment, energy, and transportation have been central issues at least since the 1973 oil crisis. Now climate change and environmental issues are more in the forefront in the United States than they have ever been before, Newman says. According to the University of Maryland's O'Shea, "Energy is the most important technopolitical problem we face." And then there are newer challenges in national security, information security, and biotechnology. "The great thing," Newman says, "is that we're in vogue now."

UNDERGRADUATE ENGINEERS who decide they want to go into the field can try to land a policy-oriented internship. Matthew Ezovski, an IEEE student member at Rensselaer Polytechnic Institute, in Troy, N.Y., got such a gig through the Washington Internships for Students of Engineering, a program focusing on the greater Washington, D.C., area.

"It's important that those who influence policy understand the scientific implications of everything they do, whether it's telecom regulation or funding basic science research," Ezovski says. The interns don't work for particular agencies but pursue their projects through the applicable offices, agencies, and leaders. Ezovski's research on the security and effectiveness of biometric passports led him to work with the State Department and the House Judiciary Committee.

Very few undergraduate programs let students combine engineering and policy courses. But Susan Bailey, vice president of AT&T's global network operations planning, did just that, double-majoring in electrical engineering and public policy at Carnegie Mellon. "I wasn't satisfied to do problems and labs and come up with the answers to the technical question without asking, 'Why do I care about how this works?' " she says.

Although her job is mostly technical, Bailey says she does better because of her policy training, which gave her a way to understand how the world works, with models of economics, behavior, and decision making. "If you don't have exposure to the models, you are more limited by the number of tools in your tool kit that you can bring to bear," she says.

Of course, there are other options for engineers who want to view the world through a wider-angle lens than the one they got in college. There's always business or law school. But if you want to shape technology's impact on people or, in O'Shea's words, "want to play a more direct role in doing good for society," then policy just might be right for you. ■

ABOUT THE AUTHOR

PRACHI PATEL-PREDD, a regular contributor to *IEEE Spectrum*, is a freelance writer who covers technology, energy, and the environment.

HOW HIGH IS THE PATENT BAR *NOW*?

Has the Supreme Court's ruling eviscerated the patents for hosts of products—even Apple's iPod?

BY KIRK TESKA

INVENTION

This spring, the U.S. Supreme Court made it harder to patent things by raising a standard known as "obviousness." A lot of people had figured that this bar had fallen so low you could practically step over it. Now some people are talking as if there's no longer any point in getting a patent—but that's going too far. Still, some existing patents probably will fall by the wayside.

I procure patents much more than I litigate them, but within only a month of the decision, even I found myself in court, using the Supreme Court case as ammunition against a patent. I am not alone: within days of the ruling, Vonage, based in Holmdel, N.J., asked to retry a US $58 million lawsuit it had lost in March for infringing on New York City–based Verizon's patents. The real question for engineers, though, seems to be this: Just how high is the obviousness bar now?

To answer that question, you have to take a good, hard look at the case that induced the court to redefine the word "obvious."

Teleflex of Limerick, Pa., had sued KSR International of Ridgetown, Ont., over the adjustable gas pedals that KSR was supplying to General Motors. The court invalidated Teleflex's patent because the claimed innovation consisted of a combination of known components—an adjustable gas pedal equipped with an electronic sensor that controlled the throttle—that all functioned just as they were designed to function. That's "ordinary innovation," the court decided on 30 April 2007, the kind anyone schooled in the art could have foreseen without the slightest effort. It is therefore too obvious to be patentable.

This ruling contradicted that of the U. S. Court of Appeals for the Federal Circuit, the usual last chance in patent matters, given that the Supreme Court hears only one or two patent cases a year. The lower court had reasoned that because no one had written out or taught the idea of combining an adjustable gas pedal with a sensor, the combination constituted real innovation. The Supreme Court rejected this legal test, arguing that although the presence of such a "teaching" would indeed make the combination obvious, the lack of it does not necessarily make a combination unobvious.

Immediately following the Supreme Court's KSR ruling, the Federal Circuit then added another category of obvious, and thus not patentable, ideas: electronic versions of previously known mechanical devices.

RESOURCES

Leapfrog, a retailer of children's learning toys in Emeryville, Calif., held a patent for an electronic interactive learning device that sounds out the letters in a word as a child selects them. Leapfrog sued Fisher-Price of East Aurora, N.Y., alleging that the latter's similar PowerTouch product infringed on the Leapfrog patent. The Federal Circuit disagreed, arguing that Leapfrog's patent was obvious, based on a 1973 patent describing an electromechanical device with a phonograph record actuated by puzzle pieces. Pressing on a puzzle piece imprinted with a letter caused the record to play the sound of that letter in a word. "Applying modern electronics to older mechanical devices has been commonplace in recent years," the court held.

Does that mean the first MP3 player is obvious in light of a 1950s-era Wurlitzer jukebox? Not so fast, says the Supreme Court. Recognizing that "inventions in most, if not all, instances rely on building blocks long since uncovered" and that "discoveries almost of necessity will be combinations of what, in some sense, is already known," the Supreme Court said engineering would count as unobvious if it went against the conventional thinking in the field. Also unobvious is any combination of old elements that work together in an unexpected manner.

THE iPOD'S KEY PARTS were known when Apple patented them as a combo. Will the patent stand up?

ROBYN TWOMEY

Prove that a product is commercially successful, that it addresses a long-felt need, that it has been copied by or is licensed to others, and you'll have an easier time convincing the patent office that you deserve a patent or, if you already have one, that it's valid. The more unexpected the success of an innovation, the stronger will be the claim of being unobvious.

Most patents will survive the court's latest ruling, because it leaves several other doctrines regarding unobviousness intact. One such doctrine deals with a solution that appears obvious only after the problem it addresses has been identified. In one old case involving a mixing vial with two compartments separated by a seal, everyone believed that leakage from one compartment into the other occurred around the seal. One inventor discovered that leakage occurred through the seal and won a patent, even though his solution to this problem—using a different material for the seal—was fairly obvious.

Another way to prove unobviousness is by showing that something that seems obvious after a problem has been simplified may well not have been so beforehand. This is best explained through the example of an engineer who found a way to keep all parts of a mold at a set temperature by using a single temperature sensor to control several valves that allowed a cooling or heating medium to enter—a job that prior machines had done by assigning each valve a sensor of its own. The courts ruled that the new simplified system was unobvious.

Still another doctrine holds that something pieced together from prior art is not obvious if the construction of the pieces had to be changed to make them fit. In one old case, a patent for a blood filter put the inlet and outlet on the top, whereas an earlier patent for a fuel filter had them on the bottom. But the court decided that this apparent resemblance didn't stop the blood-filter patent from being unobvious. Reason: if you turned the fuel filter upside down so that it resembled the blood filter, aspects of its internal construction would have kept it from working. That was difference enough.

That's not to minimize the importance of the court's latest ruling. Let's take a look again at the case of an MP3 player. Imagine a company came along and tried to patent a media player with three key components—a housing that encloses various electrical components that perform computations, a touch pad based on polar coordinates "including angular input areas for providing input from a swirling finger motion," and "a button disposed at a central portion of the touch pad" that provides another kind of input.

Of course, we're talking about Apple's iPod.

Separately, all three of these components were known when Apple filed its patent application in 2001, but Apple convinced the patent office that the three components had never before been combined in one product. But don't each of the components function as expected? And aren't polar coordinate touch pads designed to process swirling motion input? And aren't buttons to be pressed so as to provide an input? Would the notion of a central button surrounded by a polar coordinate touch pad have surprised any consumer electronics designer? I doubt it.

What happens if the Apple patent is re-examined under the court's new standard? Time will tell. At the end of the day, what is obvious is still highly subjective. Following the KSR ruling, for better or worse, obviousness is also a legal question decided by a judge rather than a jury. And that has some patent owners truly worried. ■

ABOUT THE AUTHOR

KIRK TESKA is an adjunct law professor at Suffolk University Law School, in Boston, and the managing partner of Iandiorio & Teska, an intellectual-property law firm in Waltham, Mass. His book, *Patent Savvy*, has just been published by Nolo Press of Berkeley, Calif.

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

TOOLS & TOYS

BLAST FROM THE PAST

Hewlett-Packard is offering a calculator that looks like the first one it ever sold. Big deal BY KENNETH R. FOSTER

In 1971, Bill Hewlett, cofounder of Hewlett-Packard, took a good long look at the HP 9100A, a 40-pound electronic calculator that his company had introduced just three years before. Then he asked his engineers a question: Why can't we make it fit in my shirt pocket?

The marketing people said that it wouldn't sell, because slide rules—which could also calculate logarithms and other math functions—were much cheaper than this gadget would ever be. Hewlett ignored them, and a year later the HP 35 appeared at an initial cost of US $395, or nearly $2000 in today's dollars. Engineers and scientists lined up for it, and some 100 000 units were sold in its first year, making it one of the company's most successful products ever. The slide rule soon became landfill.

As a new Ph.D. graduate in physics, I lined up, too, but for a cheaper, four-function calculator from Texas Instruments that appeared soon afterward. I treasured my calculator deeply and briefly: someone stole it from my office a few weeks after it arrived.

The HP 35, named for its 35 keys, was discontinued after three years, replaced by more advanced models. Now, 35 years after its introduction, it's back in a commemorative edition called the HP 35s. The color and case design are reminiscent of the original, although this one has eight more keys and is a bit thinner and lighter.

The HP 35s is well made, the two-line LCD is clear, with adjustable contrast, and the keys have an inviting feel. It comes with a well-produced 359-page user's guide, supplemented by extensive online training modules.

Although the company's marketing pitch takes a loving backward glance, it also portrays the retro model as the "ultimate scientific calculator," crammed with far more power and complexity than the original. Its more than 100 math and programming functions include numerical integration, two-variable statistics, and regression. There are 800 storage locations, 42 built-in physical constants, unit conversions, and lots of other goodies. The user can choose between Reverse Polish Notation (RPN), an efficient way to manage calculations that HP has long favored, or algebraic entry logic, used by nearly all the other calculators.

In a clever twist, a user can enter an equation and solve it numerically for any of its variables without rearranging the equation. However, the calculator lacks some features often found in high-end calculators, such as graphics and symbolic math capabilities, removable storage, and the ability to interface with other computers.

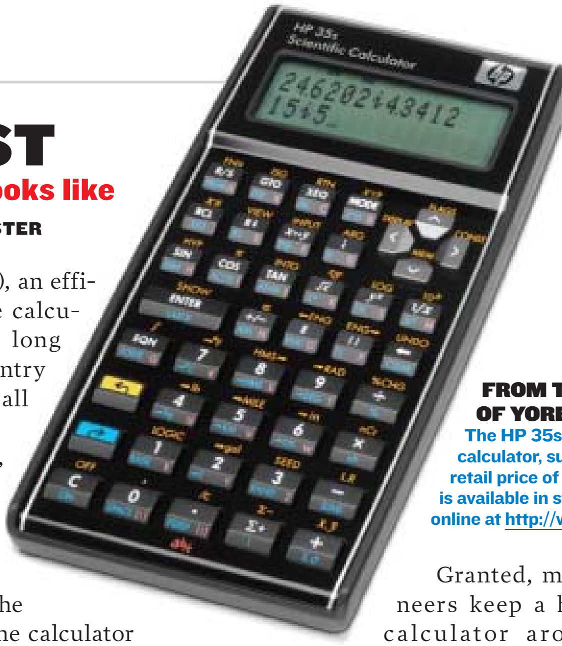

FROM TIMES OF YORE
The HP 35s scientific calculator, suggested retail price of US $60, is available in stores and online at http://www.hp.com.

HEWLETT-PACKARD

BUT WHO NEEDS all these features? Like other high-end scientific calculators, the HP 35s suffers from a bad case of feature creep. That comes from the irresistible urge of designers to stuff in ever more features as a product evolves, pushed by the need to stay ahead of the competition or to entice users to "upgrade" to new versions. This proliferation of features can introduce unexpected failure modes. In at least one HP "all in one" printer/scanner/fax machine, for example, the user cannot scan a page if the printer cartridge has run dry.

Feature creep can also lead to overly complex but unimaginative products—and the HP 35s is a case in point. How many users would need to integrate an equation on a handheld calculator? Or know whether the true result of a calculation is slightly above or slightly below the value indicated on the display when the calculator is set to display fractions instead of decimal numbers? Or need to choose between RPN and algebraic-entry systems, with their very different approaches to doing calculations? Perhaps very advanced users might benefit from these functions, but I suspect that they would have abandoned the calculator for a computer long before they reached that point.

Even my old Pickett slide rule suffered from feature creep. It had 34 scales, only a few of which I ever bothered to learn to use.

Granted, many engineers keep a handheld calculator around for occasional use, and others rely on software versions. I myself have a few scientific calculators loaded onto my PalmPilot. But would any engineer go out and buy an HP 35s for the sake of its large grab bag of features? I doubt it. Maybe that's why HP is also plucking the heartstrings of nostalgia.

Nowadays, the real market for handheld scientific calculators is in education, from middle school through college. The gadgets are commodity items sold in college bookstores, discount department stores, and office-supply centers. My own university bookstore has a rack filled with them. In its present display the HP 35s would occupy the high end of the Hewlett-Packard section, at a price of $59.99. Still, that is only about half as much as the top-of-the-line models from Texas Instruments, which are even more bloated with features. However, many middle school and high school teachers require students to buy graphics calculators, which would rule out the HP 35s.

Nevertheless, the HP 35s is a highly competent product with capabilities that would have been unimaginable 35 years ago, and the price is right. I can't wait to see what Microsoft will do for the 35th anniversary of Microsoft Windows. ■

ABOUT THE AUTHOR

KENNETH R. FOSTER, an IEEE Fellow, is a professor in the department of bioengineering at the University of Pennsylvania, in Philadelphia. His e-mail address is kfoster@seas.upenn.edu.

TO PROBE FURTHER

See the virtual HP calculator museum at http://www.hpmuseum.org.

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

IEEE INFORMATION

EDITORIAL CORRESPONDENCE
IEEE Spectrum, 3 Park Avenue, 17th Floor, New York, NY 10016 Attn: Editorial Dept. Responsibility for the substance of articles rests upon the authors, not the IEEE or its members. Articles published do not represent official positions of the IEEE. Letters to the editor may be excerpted for publication.

Institute of Microelectronics

A*STAR

R&D at IME covers the semiconductor technology chain, viz integrated circuit design, wafer fabrication process technology, packaging and assembly, and reliability testing and analysis.

These R&D are classified under three laboratories, namely:

- Integrated Circuits and Systems
- Microsystems, Modules and Components
- Semiconductor Process Technologies

IME has also established three application-driven programmes to focus its R&D efforts in new technologies and open up new opportunities for the semiconductor industry:

- Bioelectronics Programme
- Nano Electronics and Photonics Programme
- Micro-Electro-Mechanical-Systems (MEMS) Programme

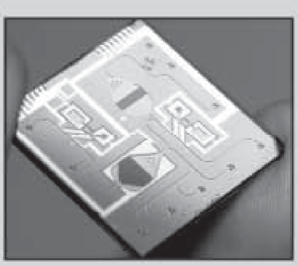

We would like to invite applications for:

Programme Director, MEMS

With the increased demand of MEMS technology for new and innovative applications, IME targets to accelerate this program further and add value to the Semiconductor Industry. IME houses an advanced 200 mm Si CMOS fabrication facility with MEMS specific process technology and advanced packaging capabilities, and analog and mixed IC design. Combining on-chip electronics and advanced MEMS technology, several projects are currently being carried out internally and in collaboration with the industry and academia on topics such as optical, inertial and RF MEMS for applications towards biometrics, medical technology, consumer electronics, display and communication. The R&D programs are driven substantially by commercial applications as the end goal. As Programme Director, you are expected to set strategic technology directions and develop technology roadmaps together with industry partners, build up a strong patent and publication portfolio and establish high visibility and reputation for the organization. You are also expected to work closely with industry in developing commercial applications of IME's research outcomes.

Programme Director, Bioelectronics

Microelectronics provides key and striking opportunities for advancing bio-medical research and creating new markets. IME is in a position to leverage on its expertise across the entire value chain of microelectronics and also on the collaborations of A*STAR's biomedical research institutes to make intellectual and economic impact in the bioelectronics area. The Bioelectronics Programme aims to explore the use of Si technology to build nanoscale sensors, devices and Lab On a Chip platforms for biomedical applications. The programme is of multi-disciplinary nature involving aspects of Si technology, bio-chemistry, bio-microsystem packaging and CMOS circuit design for bio-sensor signal processing and transmission. Currently, the program includes projects targeting "sample to answer microsystem" for clinical and molecular diganostics, devices for drug discovery applications and water quality / environmental monitoring. The successful candidate should direct the program towards transforming lab scale demonstrations to industrial applications.

Requirements

We seek candidates with vision, drive, commitment and strong leadership qualities. You should posses a relevant PhD with an outstanding record of publications in top-tier international journals and conferences. As these positions interact extensively with industry, the scientific community, lay audiences and staff at all levels, and oversee the research and human capital development of the Programmes, you should have excellent management and communication skills. Industry experience is an advantage.

For more information on the Institute and other job openings, please visit our website at www.ime.a-star.edu.sg

Our compensation packages are globally competitive with benefits such as comprehensive medical & life insurances, relocation assistance and vacation leave. Please email or fax us your detailed resume, including post applied, publications and patents list, expected and current salaries, and contact information.

Institute of Microelectronics
HR & Admin Department
11 Science Park Road,
Singapore Science Park II,
Singapore 117685

Email: personnel@ime.a-star.edu.sg
Fax : (65) 67731912

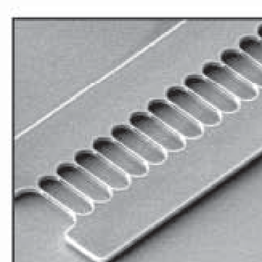

Universität Karlsruhe (TH)

Forschungsuniversität · gegründet 1825

KIT Karlsruhe Institute of Technology

The **Computer Science Department** at the Universität Karlsruhe (TH) invites applications for the position of

Full Professor in "Design and Computing in the Nano Era"

The professorship **(New Field Group)** which is to be filled as soon as possible is expected to cover research and teaching in the area of design methods and future computing architectures based on circuit technologies when migrating from micro to nano technology.

Candidates should have an international reputation and excellent research record as well as experience in leading research groups plus excellent teaching credentials.

Research expertise is expected in several of the following fields:

- Novel computing architectures and communication architectures when deploying nanotechnology based components (e.g. fault tolerant, adaptive, reconfigurable, parallel etc.)
- Design and synthesis methods for nano-technology based components and system architectures
- Analysis, classification characterization modeling and simulation of nano-technology based components

Engagement in teaching of new and existing courses at the CS department is expected. Excellent on-site research capabilities exist beyond the CS department in form of a desired collaboration with the **Center for Functional Nanostructures (CFN).**

Candidates should have a PhD degree, a habilitation or equivalent scientific credentials.

The Universität Karlsruhe (TH) is strongly committed to diversity. It therefore especially welcomes applications from women, handicapped etc. which will be preferred in case of otherwise equivalent qualifications.

In case of a first-time employment as professor, a candidate is subject to evaluation after an appropriate time period. Exceptions to this rule are possible.

Applications including up to 5 relevant publications are to be sent **before Oct 31st 2007** to the **Universität Karlsruhe (TH), Dekan der Fakultät für Informatik (dekanat@ira.uka.de), Kaiserstr. 12, 76131 Karlsruhe, Germany.**

University of Waterloo: The Department of Electrical and Computer Engineering invites applications for faculty positions in most areas of computer engineering, software engineering, and nanotechnology engineering, and in VLSI/circuits, information security, photonics, MEMS, control/mechatronics, signal/image processing, and quantum computing. Please visit https://eceadmin.uwaterloo.ca/DACA for more information and to apply online.

The Department of Electrical and Computer Engineering, in the faculty of Engineering and Architecture (FEA), at the American University of Beirut (AUB) in Beirut, Lebanon: invites outstanding candidates to apply for positions at the rank of Assistant Professor, starting September, 2008. Applicants should possess a Ph.D. degree in Electrical Engineering, Computer Engineering, Computer Science, or a closely related field. Applicants should have experience in performing research and teaching, in one of the following areas; computer architecture and parallel processing systems, applied electromagnetics and radio-frequency circuits and systems, software engineering, industrial and power electronics. Applicants for visiting positions will be considered. Applications must include a letter of interest, research, teaching, and service statements, complete curriculum vitae, and the names, e-mails, and addresses of at least three professional references. The complete application should be received by December 14, 2007, and must be addressed to the Dean of Faculty for Engineering and Architecture, American University of Beirut, P. O. Box: 11-0236, Riad El Solh, Beirut 1107 2020, Lebanon. An electronic copy must also be sent by email to fea@aub.edu.lb. The American University of Beirut is an Affirmative Action, Equal Opportunity Employer.

NANYANG TECHNOLOGICAL UNIVERSITY

Faculty Positions in Computer Engineering

The School of Computer Engineering (SCE) at Nanyang Technological University, Singapore offers undergraduate programmes leading to Bachelor of Engineering (Honours) in Computer Engineering or Computer Science, as well as full-time and part-time graduate studies leading to Master of Science and PhD. SCE's strengths lie in constantly maintaining industrial relevance in its programmes, as well as pioneering innovative cutting-edge research. Further information about the School can be obtained at **http://www.ntu.edu.sg/sce**

Applications are invited for appointment as **Associate Professor** or **Assistant Professor**. High-calibre applicants who possess a PhD and with a proven track record in research and teaching at university level are invited to apply for suitable appointments in the **Embedded Systems, Computer Networking and Ubiquitous Computing** areas.

Job responsibilities, application instructions and more details of available faculty openings can be viewed at **http://www.ntu.edu.sg/sce/job-acad.asp**

Submission of application and informal enquiries can be made to **VD-SCE-ACAD@ntu.edu.sg**

KAIST Korea Advanced Institute of Science and Technology

Faculty Positions in the fields of Entertainment Engineering at KAIST

KAIST has established a major research institute for interdisciplinary research in the field of entertainment engineering. The newly established Institute aspires to positions of international leadership at the intersection of Computer Science, Electrical Engineering, Mechanical Engineering, and Culture Technology. Successful candidates should have experience in some of the following topics: moving imagery technology (movie, animation, and broadcasting), interactive technology (computer games and ubiquitous life), and experiential technology (entertainment robots, exhibition, digital performance, media arts, digital sound and music). The successful candidates will be appointed as a faculty member at one of the departments at KAIST (primary), and at the same time at KAIST Institute for Entertainment Engineering (secondary).

Benefits

- Off-campus university housing
- Seed money (up to US$300,000)
- Immediate access to the graduate students (Both Masters and Ph.D. levels)
- Teaching load is limited to one course per semester

Professor Kwangyun Wohn,
Director, KAIST Institute for Entertainment Engineering
Korea Advanced Institute of Science and Technology (KAIST)
Telephone: +82-42-869-2901
E-mail: wohn@kaist.ac.kr
Website: http://www.kaist.edu/

Technische Universität Berlin

The Faculty of Elecrical Engeniering and Computer Science invites applications of outstanding scientists for two tenured faculty positions (Full professorships W 3):

Full professorship (W 3)

in "Microsensors & Actuators"

Reference number: IV-682

Topic: Research on sensors and actuators for automotive, medicine and automation electronics, with a focus on techniques like modelling, simulation, self test and recalibration as well as operation in harsh environment. Educational tasks are measurement of non-electrical values, sensors and actuators, technologies of microsystems and participation in the existing bachelor, master and diploma curriculums.

Full professorship (W 3)

in "Integrated Systems (Nano Architectures)"

Reference number: IV-681

Topic: Research on integrated digital systems with focus on wired and wireless communication technology based on Nano-CMOS technologies. The candidate shold have profound knowledge in the area of design, circuitry and architectures of integrated circuits.

The holder of the positions are expected to conduct research and teaching in the above fields, both in the Bachelor program and in the Master program of the Faculty. The successful candidate is expected to teach courses in German after two years.

Requirements: Compliance with Berlin appointment laws in § 100 BerlHG (further information can be provided on request), exceptional scientific and didactical qualification; perennial career in industry or research institutes with successful experience of fundraising.

To ensure equal opportunity for men and women, applications from female candidates with the advertised qualifications are explicitly solicited. Severely disabled applicants with equivalent qualifications will be given preferential treatment.

Applications should be sent within 4 weeks after publication to: **Dekan der Fakultaet IV, Sekr. FR 5-1, Franklinstrasse 28/29, D 10587 Berlin, Germany**

SHANGHAI JIAO TONG UNIVERSITY
School of Microelectronics

Applications are invited for full-time faculty positions at the ranks of Assistant, Associate or full Professor. The candidates should possess a Ph.D. degree in microelectronics, EE, or CE related areas and be able to demonstrate strong potential and abilities in both teaching and research. Desirable areas of research include but are not limited to IC design, semiconductor device and fabrication, electronic packaging/assembly and testing, microprocessor architecture, and embedded system. The successful candidates are expected to teach both undergraduate and graduate courses, supervise graduate students, and maintain and develop leading research programs. Detailed information can be found at http://ic.sjtu.edu.cn/recruit/chmyc.asp.

Applications should be forwarded to
Mr. Larry Luan,
School of Microelectronics,
Shanghai Jiao Tong University,
7F Haoran Hitech Building,
1954 Huashan Road,
Shanghai 200030, China
or by email to some-hr@ic.sjtu.edu.cn

Inquiries are welcome by telephone to 86-21-62839407.

INSTITUTION: Masdar Institute of Science and Technology, Abu Dhabi, UAE
POSITION: Professors, Associate Professors, Assistant Professors

THE MASDAR INSTITUTE OF SCIENCE AND TECHNOLOGY (MASDAR INSTITUTE), with the assistance and advice of the Massachusetts Institute of Technology (MIT), is being founded as a new and independent non-profit, tax-exempt research and educational institution (initially at the graduate level) dedicated to premier engineering research and education. The goal of the Institute is to develop, over a period of years, indigenous R&D capacity in Abu Dhabi, addressing issues of importance to the region in critical areas such as: renewable energy, sustainability, environment, water resources, and systems engineering and management, as well as to provide qualified men and women in the region with the opportunity to obtain graduate degrees in these technical fields.

JOB RESPONSIBILITIES: Teach graduate courses, supervise master and doctoral students, develop a research program and seek external funding for such research, and participate in the Institute's service and outreach activities.

QUALIFICATIONS: The Masdar Institute is seeking applicants for Professor, Associate Professor, and Assistant Professor in the fields of Water Resources and the Environment, Engineering Systems and Management, and Information Technology. Applicants should have a strong record of published research, experience in supervising graduate students, and relevant teaching experience. The applicant must be fluent in English. An earned doctorate in the relevant field is required. Relevant non-academic work experience would be an advantage.

POSITIONS IN WATER RESOURCES AND THE ENVIRONMENT: Faculty will provide research leadership and educational activities to engineering and science students concerned about the future of water supply and use, as well as an understanding of environmental systems. Specialists from Civil, Environmental, Mechanical and Chemical Engineering are encouraged to apply, as well as faculty from other disciplines such as Chemistry, Materials Science, Biological Engineering, Electrical Engineering, Biology, and Nanotechnology, whose advanced scientific work may have implications for water and the environment, even if they have not directly worked in the water and environmental area. Candidates specializing in the area of desalination and advanced water reuse in such areas as polymer membranes, ceramic membranes, biological filtration and new oxidation techniques will be considered, as well as experts in environmental systems methods such as life cycle analysis and in environmental sciences focused on fate and transport of pollutants in the environment and the mitigation and remediation of the impacts of new water and environment technologies.

POSITIONS IN ENGINEERING SYSTEMS AND MANAGEMENT: The Engineering Systems and Management program will provide intellectual research leadership and educational activities to students interested in applying a systems approach to the engineering and management of renewable and sustainable energy technologies. The Institute is looking for candidates in one of the following areas: Operations Research (Stochastic Modeling, Optimization, Decision Science, Simulation), Operations Management (Manufacturing, Technology Innovation) and Industrial Economics (Industrial organization, Technology policy, Economic Development).

POSITIONS IN INFORMATION TECHNOLOGY: Faculty candidates must be capable of teaching and conducting research in such IT areas as: Software Engineering; Information Processing for Engineering Systems; Database, Internet, and Systems Integration Technologies; Data Mining; Artificial Intelligence and Semantic Web; Communications and Connectivity among Information Systems; Information Management; Intelligent Systems; and Pervasive Computing. In addition, faculty are expected to effectively collaborate with other faculty areas such as Water Resource and the Environment, Materials Science and Engineering, Mechanical Engineering and Engineering Systems and Management.

APPLICATION SUBMITTAL INFORMATION: The Massachusetts Institute of Technology is assisting the Masdar Institute in the search. Initial screening of applications will begin immediately. Application deadline is November 30, 2007. Application materials should include your name, address, telephone numbers, curriculum vitae and what specific position you are applying for, your current position, a description of how your experience matches the position requirements, and e-mail contact information for three references.

Materials should be submitted electronically as a MS Word attachment to:
Dr. Russel Jones, President
Masdar Institute of Science and Technology, Abu Dhabi, United Arab Emirates
Co-Chair, Search Committee for Masdar Institute of Science and Technology
e-mail: rjones@masdar.ae

Copy to:
Dr. Fred Moavenzadeh, Co-Chair
Search Committee for Masdar Institute of Science and Technology
c/o Technology and Development Program
Massachusetts Institute of Technology
e-mail: tdpmail@mit.edu

TECHNICALLY SPEAKING

By Paul McFedries

All A-Twitter

The blog search engine Technorati now tracks nearly 100 million blogs, with nearly 200 000 new ones added every day. (As the Technorati tagline has it: some of them *have* to be good.) Folks from all walks of life use blogs to opine to the world, and it's a rare day when some celebrity doesn't start a **celeblog**. But the dirty secret is the massive number of abandoned blogs. The market research firm Gartner, in Stamford, Conn., recently put the number at more than 200 million. Clearly most people who start a blog soon give it up for dead.

Why do so many blogs go belly-up? Probably because blogging is *hard*. Unless you love to write, churning out even remotely interesting mini-essays every day of the week is a tough slog, particularly when you're not even sure anyone outside of your immediate family is actually reading your musings. Perhaps this explains the recent explosion of interest in **microblogging**, posting short thoughts and ideas to a personal blog, particularly by using instant messaging software or a mobile phone. Jaiku (http://jaiku.com) lets you create a **miniblog** to which you post short messages—called **jaikus**—either via its site or by texting the messages through your mobile phone. Fotolog.com enables members to exchange short messages about posted photos.

But the major buzz in microblogging centers around Twitter (http://twitter.com) a site that combines **social networking** and microblogging. It periodically asks members a simple question: "What are you doing?" Members respond, or **twitter**—via e-mail, instant messaging, short message service (SMS), third-party programs, or the Twitter site itself—with text-based posts no more than 140 characters long. (When Twitter recently won an award, its representative's acceptance speech was apropos: "We'd like to thank you in 140 characters or less. And we just did!") This is why services such as Twitter and the similar Dodgeball.com are known as **notification tools** or **quick-ping media**.

What happens to all those posts? They get displayed on the user's Twitter home page, of course—that's the microblogging part—but they also get sent out to the user's circle of **Twitter friends**. These friends can receive the updates via e-mail, IM, SMS, or an RSS feed. Because of this, Twitter-like sites are also known as **constant-contact media**.

The goal of all this **twittering** seems to be to enhance one's cyberspace **presence**, an elusive concept that seems to refer to being "out there" (wherever "there" is) as much as possible. Peel back the layers of a typical Twitter user and you'll probably find that he or she also maintains a regular blog, a Facebook or MySpace account, a Second Life avatar, and so on. The dream is to achieve a sort of **virtual omnipresence**. Such people are said to be **ultraconnected**, although sometimes it's possible to be *too* connected. For example, if someone twitters that a particularly interesting event is occurring at a nearby location, the site can become overwhelmed by the unruly **Twittermob** that materializes.

Not that the entire world is in love with Twitter. Most people just don't see the point, and others dismiss it as a massive **time-suck**. (Almost everyone who gets into Twitter calls it "addictive," which may explain why there are so many **Twitterholics** out there.) For some people, however, Twitter bemusement has turned into outright **Twitter hate**. The biggest complaint is the unremitting triviality of most people's updates, particularly **dinner Twittering**, posting updates about what you are making or eating for dinner (or lunch or breakfast). One wag described Twitter as "the 'Seinfeld' of the Internet—a Web site about nothing." Other users complain about **Twitter storms** or **Twitterrhea**, update deluges consisting of dozens of messages per day from people who can't seem to stop themselves from posting.

Others deride Twitter and its ilk as **hipster narcissism**, a charge they say is confirmed by Twitter users' insistence on creating their own Twitter-based lingo. For example, a Twitter update isn't a post, it's a **tweet** (and posting is called **tweeting**); people who sign up with Twitter aren't called members or users but **twitterers** or, inevitably, the **twitterati**; the nonfriends who read a person's tweets are called **followers** (some of these folks probably deserve a less euphemistic name: *stalkers*); and the subset of cyberspace that consists of Twitter and its tweets, twitterers, and followers is called the **twitosphere**. Hmm. They might want to rethink that last one. ■

PAUL MCFEDRIES is a technical and language writer with more than 40 books to his credit. He also runs Word Spy, a Web site and mailing list that tracks new words and phrases (http://www.wordspy.com).

GREG MABLY

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

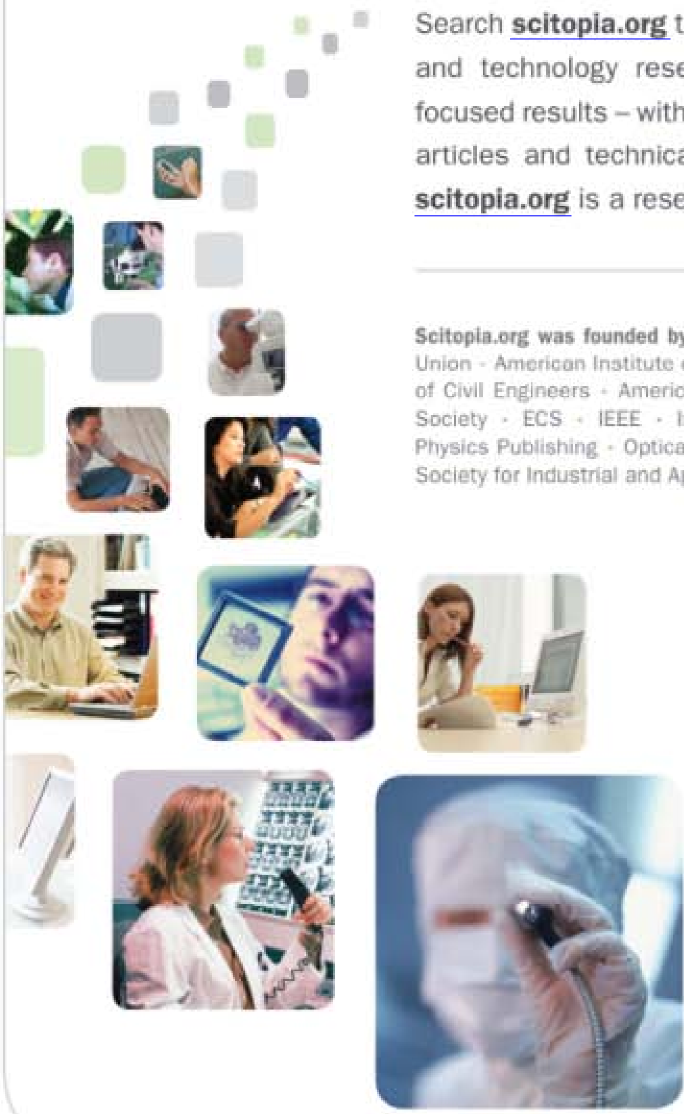

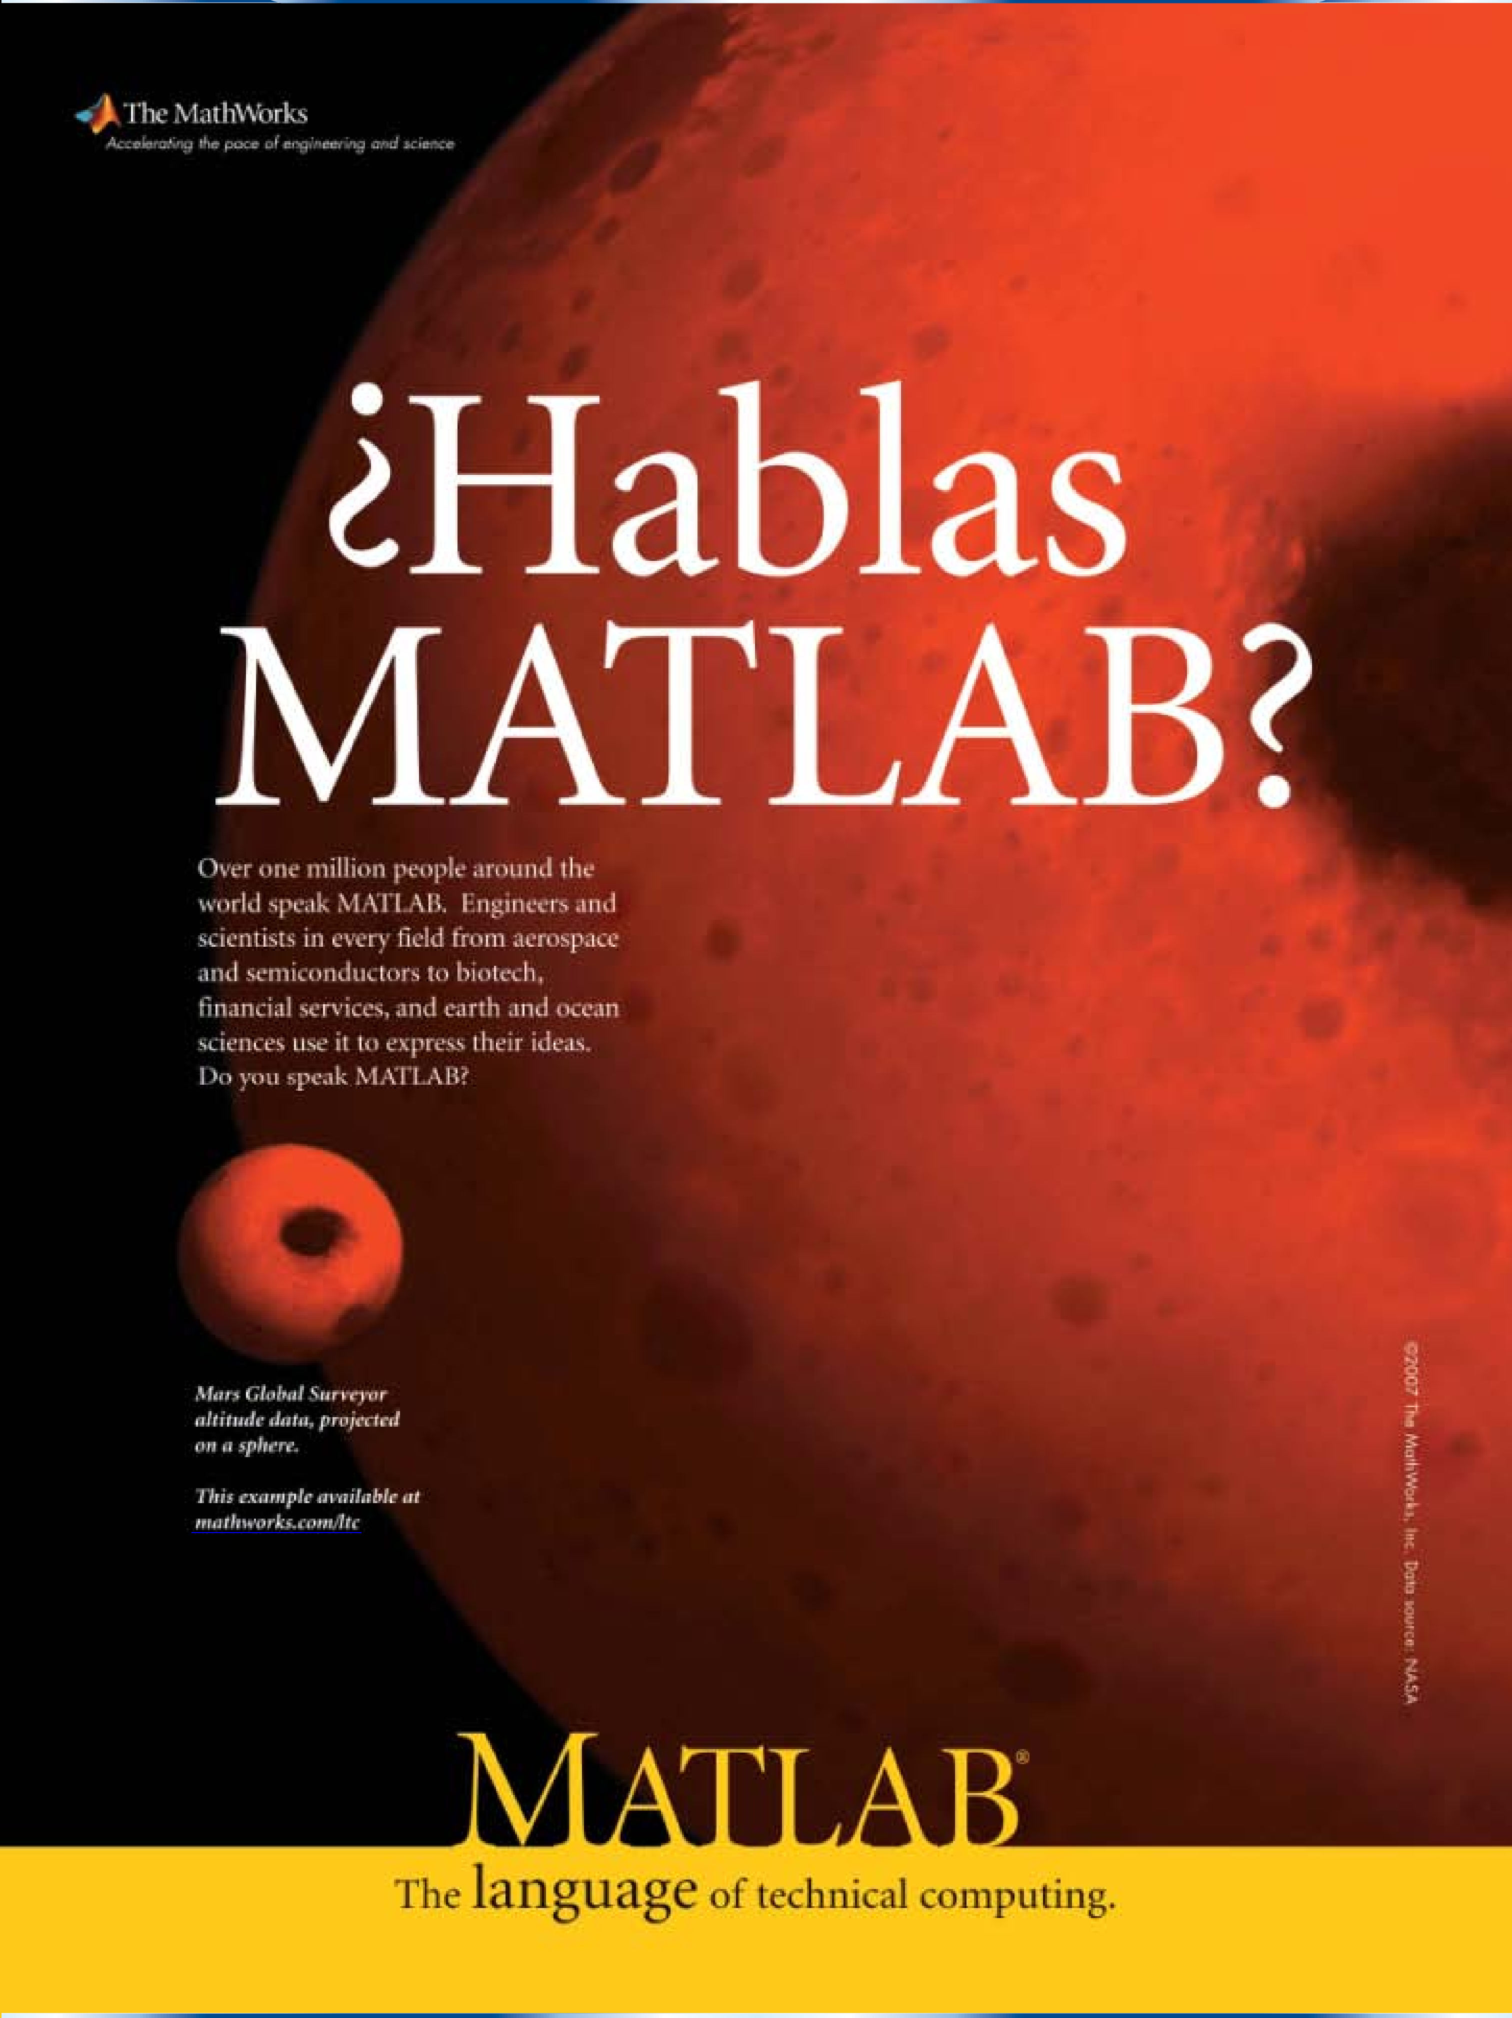
The MathWorks
Accelerating the pace of engineering and science
¿Hablas MATLAB?
Over one million people around the world speak MATLAB. Engineers and scientists in every field from aerospace and semiconductors to biotech, financial services, and earth and ocean sciences use it to express their ideas. Do you speak MATLAB?
Mars Global Surveyor altitude data, projected on a sphere.
This example available at mathworks.com/ltc
©2007 The MathWorks, Inc. Data source: NASA
MATLAB®
The language of technical computing.