ieee spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
08.08

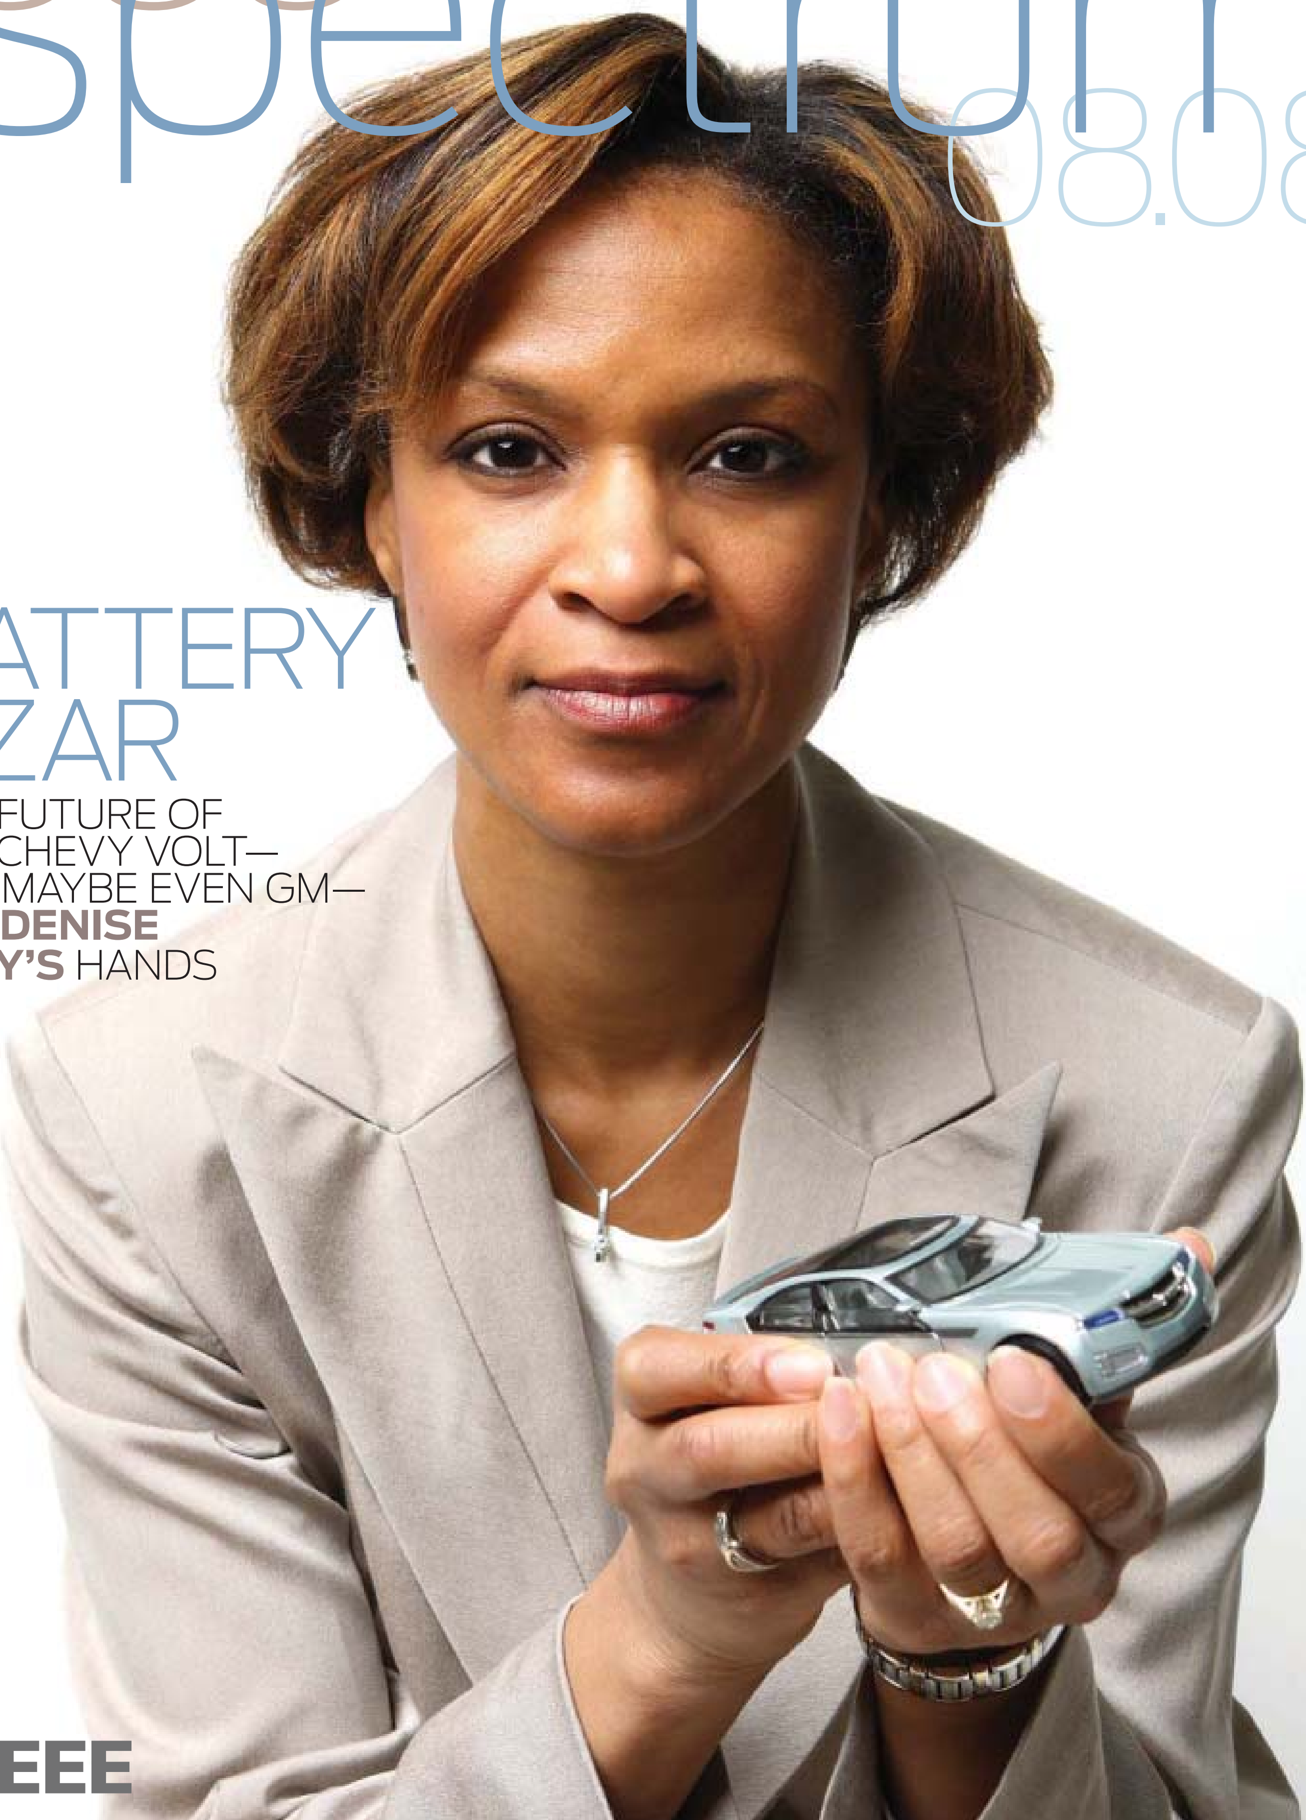

BATTERY CZAR

THE FUTURE OF THE CHEVY VOLT—AND MAYBE EVEN GM—IS IN **DENISE GRAY'S** HANDS

IEEE

volume 45 number 8 international

ieee spectrum 08.08

28

40

22

EARTH AND BEYOND: Denise Gray is helping reenergize GM; the Brazilian sci-fi game *Taikodom* runs on a mainframe; First Solar's photovoltaics are bigger and cheaper.

COVER: JOSHUA DALSIMER

THIS PAGE: CLOCKWISE FROM TOP LEFT: JOSHUA DALSIMER; HOPLON INFOTAINMENT; BLITZSTROM/BECK ENERGY

COVER STORY

UPDATE

OPINION

DEPARTMENTS

volume 45 number 8 international

ieee spectrum 08.08

LEFT: SHAUN CURRY/AFP/GETTY IMAGES; RIGHT: ROZ WOODWARD/GETTY IMAGES

WWW.SPECTRUM.IEEE.ORG

AVAILABLE 1 AUGUST ON SPECTRUM ONLINE

MOTORHEADS INVADE LONDON

When's the last time you bought a car based on how low its CO_2 was? At this year's London Motor Show, that statistic led the launches of several new, ultralow-consumption vehicles from major European carmakers. Practically every automaker touted its commitment to reducing vehicular carbon emissions, using approaches as varied as idle-stop equipment and tiny, high-powered diesels. Automotive editor John Voelcker takes us on a photographic tour of the good, the bad, and the breathtakingly bizarre at the UK's major new-car show.

ONLINE FEATURES:

Q&A WITH WILL WRIGHT: The legendary designer discusses his new game, *Spore*, with *IEEE Spectrum* contributing editor David Kushner.

THE ULTIMATE ENGINEERING LIFE-FORM: Enter *IEEE Spectrum*'s contest by using the Spore Creature Creator to design an organism with everything the ultimate engineer needs. Visit http://www.spectrum.ieee.org/creature for the official rules.

ALSO ONLINE:

- Webcasts
- Podcasts
- News
- Blogs
- Jobs
- Career Accelerator Forum
- IEEE Xplore® Digital Library
- Interviews
- Opinions
- More!

WWW.IEEE.ORG/THEINSTITUTE

AVAILABLE 8 AUGUST ON THE INSTITUTE ONLINE

IEEE-SUPPORTED GENETIC ACT IS NOW U.S. LAW

Read about a new law—supported by IEEE-USA—that prevents employers and health insurers from discriminating against people based on their genetic information.

TUTORIAL ON MENTORING PROGRAM

If you've flirted with the idea of participating in the IEEE Mentoring Connection program but weren't sure what it was all about, check out an online archive of a recent orientation session.

RUTGERS IS NEWEST EPP PARTICIPANT

Rutgers University, the newest IEEE Education Partner, offers members discounts on its advanced technology short courses in engineering, IT, and professional certificate programs.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2008 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS-IEEE Members: $21.40 included in dues. Libraries/institutions: $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.

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

back story

KATHARINE STEVENSON

Something New Under The Sun

WHAT FIRST caught Richard Stevenson's eye was the sheer scale of production. First Solar, a maker of solar cells based in Tempe, Ariz., had just announced it was "going to open four more factories in Malaysia—substantial numbers, really," he says.

That was a year ago. Stevenson wrote it up as a news item for *Compound Semiconductor*, the industry magazine he works for in Bristol, England. But when *IEEE Spectrum* invited him to write on the most compelling thing he could think of, this is the topic he chose: how did First Solar take cadmium telluride, which has been kicking around for decades, and fashion it into large solar panels more cost-effective than anything else on the market?

In trying to answer that question he came up against a still more puzzling one: why the company absolutely refused to talk. "I was really shocked that they were so incommunicative," he says.

The reticence only heightened his curiosity—and his determination to tell the story. So he studied a handful of published papers, interviewed experts outside the company, and pieced together how First Solar had cracked at least part of the manufacturing problem that blocks the industry's way to its ultimate goal—selling photovoltaic electricity at a rate comparable to that of fossil-fueled generation.

Stevenson's sleuthing called on both his journalistic and his technical expertise. He got a Ph.D. from the University of Cambridge by designing an optical scanning microscope and using it to study conductive polymers. The degree was in physics, he says, "but in the United States, this work would be done in a materials-science or an electrical engineering department." That argument—together with his love of vinyl records and contract bridge—makes him okay in our book.

Stevenson found his calling as a journalist after leaving Cambridge in 2001 and spending three years working as a process engineer for a supplier of semiconductor wafers. The company was buffeted by the dot-com bust, a lot of the fun had gone, and he decided he'd rather explain engineering than actually do it. In fact, he realized that that was what he'd always preferred. "I quite enjoyed writing my Ph.D. thesis, and most people around me hated it," he remembers.

He took a job at *Compound Semiconductor,* as features editor, and began writing and editing articles—including a few by his former colleagues. He hasn't looked back. ❑

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 Update page is in *IEEE Spectrum,* Vol. 45, no. 8 (INT), August 2008, p. 7, or in *IEEE Spectrum*, Vol. 45, no. 8 (NA), August 2008, p. 11.

contributors

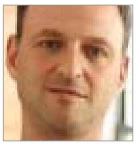

JOSH DALSIMER has taken on various assignments for *IEEE Spectrum*, including photographing an owl tracker and a bunch of little orange robots. For our cover story, "Battery Czar" [p. 28], he shot Denise Gray at GM's Vehicle Engineering Center, in Warren, Mich. Dalsimer says the battery-powered strobe packs he brought for the shoot looked like toys alongside the center's advanced technology: "We were definitely outdone in the battery department."

SUSAN KARLIN is a travel junkie who has visited every continent. Writing "Qatar University Opens EE Doors to Women" [p. 18] was especially satisfying to her because "women around the world don't get the same opportunity as women in Western countries," she says. "I've seen the frustration firsthand." A frequent contributor to *Spectrum,* Karlin covers science, technology, and entertainment.

JAMES TURNER, author of "Hacking the OLPC" [p. 20], says the One Laptop Per Child computer came through in a jam: he had forgotten to bring the charger for his regular laptop to a conference in Texas. "I ended up spending three days in Austin using just the OLPC, and it worked fine," he says. "I filed a story from that computer." Turner is a contributing editor for O'Reilly Media and a correspondent for the *Christian Science Monitor.*

JOHN VOELCKER wanted to profile Denise Gray from the first time he met her, but given Gray's intense work schedule, it took a year to arrange an extended interview. For "Battery Czar" [p. 28], they finally sat down together at a technical conference in San Diego. Right after that meeting, she climbed into a waiting sport utility and was driven to the airport, to fly back to Detroit on a GM corporate jet.

MICK WIGGINS, whose drawing appears in Technically Speaking [p. 16], has been an illustrator for more than 20 years. He starts with a pencil sketch, which he then scans into his computer; afterward, he adds color and refines the lines. His illustrations have appeared in such publications as the *Boston Globe, Smart Money,* and *Business Week,* as well as in theater posters and short stories.

LEI YANG and her colleagues explain how some clever software can replace costly hardware in "RAM for Free" [p. 34]. Yang is a research assistant at Northwestern University, where she studied with **ROBERT P. DICK**, an assistant professor of electrical engineering and computer science. For two summers, Yang interned at NEC Laboratories America, where she worked with **SRIMAT CHAKRADHAR**, head of the labs' systems architecture division, and with **HARIS LEKATSAS**, who is a consultant to designers of embedded systems.

CLOCKWISE FROM TOP LEFT: JOSHUA DALSIMER; RANDI SILBERMAN; MICK WIGGINS; JAMES TURNER; CARIN BAER spectral lines

Airlines: Got Fuel?

TOPPING OFF: An employee of Allied Aviation Fueling Co. of Houston gases up a Continental Boeing 767.
PHOTO: DAVID J. PHILLIP/AP PHOTO

High fuel prices are hammering the airline industry. I'm aware of that, *and* that all these annoying surcharges that have recently been announced—for luggage, for drinks—are attempts to somehow make up for the extra fuel costs.

What I didn't know was that airlines are also saving money on fuel by simply not filling up the tanks as high as they ought to. The U.S. Federal Aviation Administration requires that an aircraft carry enough fuel to reach its destination and the most distant alternate airport, plus an extra 45 minutes' worth. It's not to an airline's advantage to carry any more than the minimum requirement: more fuel means a heavier plane, and a heavier plane gets worse mileage.

Of course, the calculations are based on estimates—estimated passenger weights, estimated luggage, estimated speed. But after my last flying experience, in June, I'm thinking that the airlines are estimating a little low.

I flew from San Francisco to Newark, N.J., on a Continental Airlines Boeing 737. Great weather on both ends, only a little turbulence in between. A lovely flight, actually—the best I've had in a long time. The flight attendants were cheerful; passengers got two drinks and a bagel-and-egg sandwich without charge.

A little more than four hours into the flight, the pilot reported that we'd be landing about half an hour ahead of schedule, and the flight attendants began collecting trash in preparation for our approach. Perhaps 10 minutes later the pilot announced that we'd have to slow down a little to get in line for landing, but we'd still arrive well ahead of schedule. I was thrilled; this would be my first flight in at least a year that landed on time. Perhaps I could even call a friend for dinner.

And then, just a few minutes later, the pilot came on the public-address system again: "Uh, folks, we're going to make a quick stop for refueling." Huh? Passengers looked at each other in surprise. Flight attendants passed rapidly through the cabin checking seat backs and tray tables and strapped themselves in. Minutes later, we landed at Stewart Air National Guard Base, in Newburgh, N.Y., less than 100 kilometers from our destination. After a long ride on the ground past National Guard cargo planes, we parked and waited for the fuel trucks.

I was flabbergasted. I've logged a lot of airline miles over the years, but I've never been on a flight that ran out of gas. The pilot blamed the problem on air traffic delays, but from what he had said the delay seemed minimal—certainly not enough to eat up the 45-minute reserve the FAA requires. I wondered about the fact that we'd gotten so far ahead of schedule; likely we'd been flying a little faster than is optimal for fuel conservation. Or could the European tourists on board have gotten carried away by the cheap dollar and brought back a lot of extra luggage? Turns out, though, that the U.S. Department of Transportation recently singled out Continental for having the most of what it calls "minimum fuel declarations" on approach to Newark International Airport last year: 96, which is more than twice the amount it had the previous year. Declaring "minimum fuel" tells the air traffic controllers that if incoming planes need to be delayed they shouldn't delay this particular flight too much. The DOT also found that Continental has been pressuring pilots to cut back on the amount of fuel they carry, for ultimately the pilot makes that decision. Last October, the airline sent a memo to pilots pointing out that "adding fuel indiscriminately without critical thinking ultimately reduces profit sharing and possibly pension funding." Maybe that's why I ended up sitting on a deserted runway for an hour and a half?

Next time, Continental, forget the bagel. Just fill up the tank.

—Tekla S. Perry

A version of this column appeared in IEEE Spectrum*'s Tech Talk blog on 26 June.*

forum

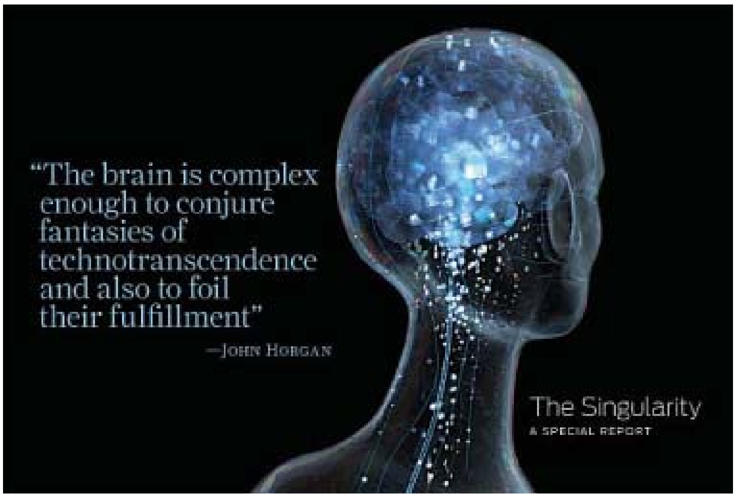

BRYAN CHRISTIE DESIGN

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.

SINGULARITY SKEPTICS

I HAVE READ "The Singularity: A Special Report" [June] and I am a skeptic. What Ray Kurzweil and others seem to neglect in their singularity theories is what I would call the damping term in a resonator: as soon as you have the slightest damping, your resonant quality factor goes from infinity to some smaller finite value. This damping effect is the same thing that keeps us from being massively productive: sickness, the need for sleep, hectic multitasking, faulty reasoning, and so on. In the same way, the humanoid machines at the center of these theories will spend at least 50 percent of their lives fighting computer viruses, broken wires, hackers, defects in semiconductors, and computer code "cancer," just as we do. The singularity visionaries also seem to ignore the human propensity for evil: create a new technology and someone will find a way to hack it and abuse it. Instead of the singularity being an infinitely better way of life, it may be only 2.718 times better and orders of magnitude more complex. Can mere humans possibly create a machine more rational and more reliable than themselves? Or are machines doomed to the same perplexities that we must endure?

Colin Joye
IEEE Graduate Student Member
Cambridge, Mass.

THE RAPID RISE in computer capability and the imminent singularity it will supposedly produce reminds me of a beer advertisement I saw some years ago. A new brewer had had phenomenal success. Its ad showed production figures with exponential growth. The ad concluded that within a few years Earth would be covered 3 feet deep with its beer. Unfortunately, that blessed event never took place. As most of us know, there are countervailing phenomena that keep the growth in check—no oversupply of beer and no oversupply of computer intelligence.

Maris Graube
IEEE Life Senior Member
Forest Grove, Ore.

THE SECRETS OF THEIR SUCCESS

I ENJOYED THE Christensen, King, Verlinden, and Yang article "The New Economics of Semiconductor Manufacturing" [May]. The 12 percent reduction in costs and 67 percent reduction in cycle time achieved in eight months by implementing the Toyota Production System are indeed impressive. But it is important to keep this in perspective: to achieve these results, it is critical to have an essentially fixed process (that is, a mature fab). On the other hand, a fab that constantly changes its process to keep up with Moore's Law has historically seen a 35 percent reduction in cost per transistor every year. Thus, the "new economics" promoted in the article are not very compelling compared to the "old economics" of Moore's Law—that is, until Moore's Law comes to an end.

Chris Mack
IEEE Senior Member
Austin, Texas

THE AUTHORS of this article deserve to be congratulated at translating the ideas of TPS into reality with such remarkable success. However, it is too early to be euphoric.

First, the TPS experiment at the unnamed facility has been running for about a year now. There are several other companies that have claimed to have cracked Toyota's secret and achieved similar successes, but only over a short term. Toyota, by contrast, has been doing it successfully for decades.

Second, there are reasons to fear that the experiment will fail in the long run. Humility, selflessness, and relentless self-discipline are required at all levels of an organization to find, accept, and correct mistakes—qualities scarce in Western societies spurred by short-term targets of promotions, raises, and stock options.

Vivek Sharma
Rapperswil, Switzerland

OF HITS AND HYPE

I WAS SKEPTICAL when I read Gordon Moore's claim that Google returns twice as many hits for "Moore's Law" as "Murphy's Law" ["Gordon Moore's Next Act," May], so I conducted the test myself. There were approximately 2.14 million hits for "Murphy's Law" versus 1.63 million for "Moore's Law." The extent of public awareness of Moore's Law is impressive, but let's not exaggerate it, please.

Blake Linton Wilfong
IEEE Member
Houston

more online at www.spectrum.ieee.org

update

A Computer for the Clouds

A proposed supercomputer would do just one job—model global climate—but consume far less electricity than a general-purpose machine

In May an IBM-built supercomputer called Roadrunner crunched through a quadrillion floating-point operations per second, officially becoming the first supercomputer to break the petaflop barrier. But unofficially, that barrier had fallen two years before, when MDGRAPE-3, a machine at Japan's Riken Institute, in Wako, powered up. Accepted benchmarking methods ruled out that performance because MDGRAPE-3 is a purpose-built computer, able to model molecular interactions and little else. Yet the machine cost Riken just one-tenth of Roadrunner's price—more than US $100 million—and consumes just one-tenth the power.

That power-saving potential is convincing many people who have belittled special-purpose machines to give them a second look. Electricity already accounts for more than half the lifetime cost of owning and operating a supercomputer—or any large server farm, for that matter—and power's share is expected to increase.

"We think scientific computing is ripe for a change," says Michael Wehner, a climatologist at Lawrence Berkeley National Laboratory. "Instead of getting a big computer and saying, 'What can we do?' we want to do what particle physicists do and say, 'We want to do this kind of science—what kind of machine do we need to do it?' "

Wehner and two engineers, Lenny Oliker and John Shalf, also of Lawrence Berkeley, have proposed perhaps the most powerful special-purpose computer yet. It is intended to model changes in climatic patterns over periods as long as a century. Specifically, it should be able to remedy today's inability to model clouds well enough to tell whether their net effect is to warm the world or cool it. To solve the problem, climatologists figure they need to deal in chunks of the atmosphere mea-

COMPLEXITY: However pleasing to the eye, clouds are a computational headache because their numerous features can either accelerate or retard warming. These altocumulus clouds form in midatmosphere, at 2000 to 6000 meters.

PHOTO: DARRYL TORCKLER/GETTY IMAGES

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

update

suring 1 kilometer on a side—a job for an exaflop machine, one with 1000 times more computing power than even Roadrunner can provide.

Wehner, Oliker, and Shalf estimate that a general-purpose machine using today's technology would cost $1 billion to build and 200 megawatts to power—enough for a small city. By comparison, they estimate, a specialized machine would cost just $75 million and consume just 4 MW.

The researchers are now trying to validate their claims with a hardware mock-up, which they are building in collaboration with Tensilica, a custom-chip supplier in Santa Clara, Calif. The plan is to bench-test a single processor by November and a parallel array of processors by the middle of 2009. If the claims are vindicated, the researchers hope to get government funding for a full-size machine.

Critics of special-purpose machines say they've heard it all before. "The problem is that when we devise a new way to solve a problem, the machine designed for the old way will no longer be as good," says Jack Dongarra, a professor of electrical engineering and computer science at the University of Tennessee.

But according to Horst Simon, who heads the Lawrence Berkeley lab's research computing center, the proposed machine would not be so specialized that a new algorithm would render it instantly obsolete.

"We are building hardware that runs not just one algorithm but a large class of related algorithms," he says. "We are trying to eliminate unessential features of the architecture, much of it developed for desktop applications, and to optimize it for a class of applications that is scientifically focused."

Not that there wouldn't still be room for superspecialized machines. As *IEEE Spectrum* went to press, D.E. Shaw Research of New York City said that by the end of the year it will have a specialized machine, called Anton, that can simulate molecular interactions hundreds of times as fast as anything now available.

—Philip E. Ross

Efficiency of World's Top 10 Supercomputers:

Average power consumption
» 1.32 megawatts

Average power efficiency
» 248 floating-point operations per second per watt

Yearly electricity cost*
» US $1 029 124

*Assumes constant operation at $0.089 per kilowatt-hour.

Source: Consumption and efficiency from Top500.org.

A Better Camera Pill

A collaboration of engineers led by the Fraunhofer Institute for Biomedical Engineering, in St. Ingbert, Germany, has developed the first-ever control system for a camera pill. Natural motions propel today's camera pills through the esophagus too fast to take all the pictures that doctors would like to see. The Fraunhofer pill is steered using an external magnet held by a physician. "In the future, doctors will be able to stop the camera in the esophagus, move it up and down, and turn it," says Frank Volke, a researcher at Fraunhofer.

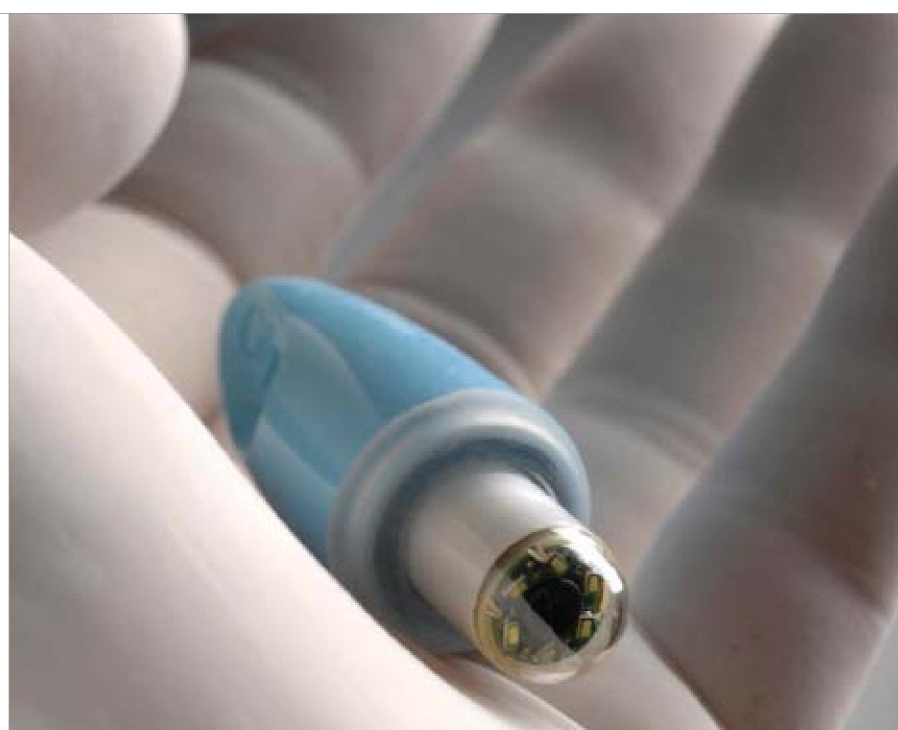

PHOTO: FRAUNHOFER INSTITUTE FOR BIOMEDICAL ENGINEERING

update

Wi-Fi Takes On Bluetooth

Ozmo hopes to replace Bluetooth with the Wi-Fi that computers already have

WIRING DIAGRAM
Neuroscientists in Switzerland and Massachusetts, using a variant of magnetic resonance imaging called diffusion spectrum imaging, have constructed the most comprehensive map ever made of the connections in the human brain. Unlike other techniques, diffusion imaging shows the orientation of the fibrous connections between neurons. The resulting map shows the location of millions of neural fibers in the outer layer of the brain.

IMAGE: INDIANA UNIVERSITY

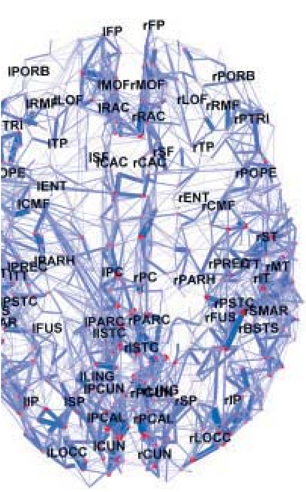

WHY DOES a laptop need both Wi-Fi and Bluetooth? A Bluetooth chip creates a personal area network, allowing peripherals like keyboards, mice, and headsets to communicate wirelessly with a laptop or mobile phone. But why can't that network be a subset of the Wi-Fi local area network? Why can't Wi-Fi manage all our wireless communications?

Two big problems have prevented convergence, but they've both been solved by Ozmo Devices, a Palo Alto, Calif., start-up—and at a data rate that's much faster than Bluetooth's.

ANNA DEMIAN

First, while a computer's Wi-Fi can communicate with other Wi-Fi devices, it cannot do so at the same time it's talking with a router. So for a Wi-Fi–enabled keyboard to interact with a laptop computer without taking it off-line, the keyboard would need to contact the router, which would then route the data back to the laptop. But this is so inefficient that users would notice delays in keystrokes showing up onscreen.

Ozmo has written a software patch that allows a laptop's Wi-Fi to communicate in tiny time slices, using some to talk with local peripheral devices and the rest to maintain its connection to the router. Company cofounder Roel Peeters says the strategy of letting the laptop talk directly to a peripheral keeps latency to about 10 milliseconds, the same as Bluetooth's.

Second, Wi-Fi is notorious for the way it sucks up power—typically cutting a laptop's battery time in half, making Wi-Fi unsuitable for a small-battery-powered headset or keyboard. Ozmo's chip set not only draws far less power than off-the-shelf chips but also outperforms Bluetooth, because it goes to sleep when not actively engaged in sending or receiving data. Based on in-house testing, Peeters estimates the power consumption to be one-third to one-fourth that of Bluetooth's. The tradeoff is the peripheral's data rate—9 megabits per second versus Wi-Fi's standard 54 Mb/s. Still, that's three times as fast as Bluetooth's top speed.

Intel had already been working to add personal area networking to Wi-Fi when Ozmo's chip set came along. Intel's ambitious research initiative, dubbed Cliffside, is expected to allow much higher data rates so that a computer could send a high-definition movie directly to a television screen, for example. Cliffside is far from commercial release, but in the meantime, Intel has adopted Ozmo's technology for the lower end of the data-rate continuum. New Centrino-based laptops will come with Ozmo's software patch.

Any reports of the impending death of Bluetooth may be exaggerated, however, or at least greatly premature. The first Wi-Fi peripherals with Ozmo's chip set won't ship until 2009. Meanwhile, Bluetooth is shipped in half the world's cellphones, one-third of its laptops, and, in 2007 alone, 7 million cars. The Bluetooth Special Interest Group, a trade association, claims the technology is embedded in more than a billion devices today.

Belkin International, in Compton, Calif., will be one of the first to come out with Wi-Fi–enabled keyboards, mice, and headsets. But the company expects to sell them only with new PCs, according to senior technologist Brian Van Harlingen. It can take years to displace an existing technology, no matter how technically superior the new one might be.

—STEVEN CHERRY

2010 The year worldwide investments for the production of photovoltaic cells will match those for semiconductor manufacturing, according to market research firm iSuppli, in El Segundo, Calif.

Earth-Size Radio Telescope Opens Its Eye

Seven telescopes act as one to produce finest radio images ever

ARECIBO, PUERTO RICO: The 1000-meter radio telescope dish at dawn.

PHOTO: ROGER RESSMEYER/CORBIS

THIS FALL, the world's largest telescope will begin its scientific mission. Made up of radio telescopes in Chile, Germany, Italy, the Netherlands, Puerto Rico, South Africa, and Sweden, the e-VLBI—for electronic very long baseline interferometer—creates in effect a telescope with a diameter of 11 000 kilometers; Earth's own diameter is about 12 750 km at the equator. Because a telescope's resolution is proportional to its size, the e-VLBI should see farther out in space and time and elucidate the finer structures of the most energetic phenomena in the universe, such as supernovas, pulsars, and black holes.

Although a smaller, Europe-wide e-VLBI has been in operation for more than a year, the full multicontinent version opened its eye only on 22 May 2008, when all seven sites were linked to a custom-built supercomputer, operated by the European VLBI Network (EVN), in a test observation.

VLBI increases the resolution of a pair of radio telescopes by using the time a particular radio wave arrives at each of them to estimate its frequency and pinpoint its origin. Although the technique has been in use since the mid-1980s, linking radio telescopes between other countries and continents in real time has not been possible until now.

Each of the radio telescopes used in the May test produced up to 1 gigabit per second of data. Until the EVN supercomputer was built, the only way to transport such large volumes of data across the globe for analysis was by recording them on magnetic media and then physically shipping the media to the Joint Institute for VLBI in Europe, located in Dwingeloo, Netherlands.

The new system transports data via a large number of network providers. Just getting the data to EVN "was an enormous technical hurdle to overcome," says Arpad Szomoru, head of technical operations and R&D at Dwingeloo. The supercomputer can handle flows of up to 100 terabytes per observation coming in from 16 radio telescopes.

Radio astronomers were initially skeptical that moving to real-time VLBI was worth the effort, says Huib van Langevelde, director at Dwingeloo. But the recent test showed that e-VLBI collapses processes that would have taken weeks without the supercomputer networking into a matter of hours.

For the new intercontinental e-VLBI system to work, all the telescopes must observe the same astronomical radio source at exactly the same wavelength. All the stations have atomic clocks, synchronized to within a few millionths of a second, so that a radio wave's arrival can be precisely marked at every telescope.

As Earth rotates, an observational target will drop beneath the horizon in relation to some telescopes and rise in relation to others. "The source is tracked by a changing set of telescopes," says Szomoru.

According to Szomoru, astronomers will get the most from the e-VLBI when they're tracking stellar explosions and other transient cosmic activity, including gamma-ray bursts and flaring microquasars. "It is now possible to observe a number of candidate cosmic sources and, if one of them is seen to come into an active state, do one or several follow-up observations, thus catching a flare in a microquasar at a very early moment, something which was not possible previously."

—BARRY E. DIGREGORIO

FACE OF EVIL

This sinister-looking gingerbread man is actually a cross section of some melted metal interconnects as seen with a focused ion beam. The focused ion beam is an instrument similar to the scanning electron microscope, except that it uses gallium ions to image a sample instead of electrons. The image came in sixth in the "Art of Failure Analysis" competition at IEEE's 15th International Symposium on the Physical and Failure Analysis of Integrated Circuits, in Singapore. See the top 10 images at http://spectrum.ieee.org/jul08/6433.

IMAGE: JACQUELINE KWA

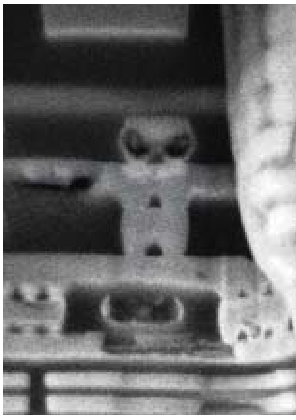

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

1/30TH The amount of energy required to boil water in a container lined with copper nanorods versus what's needed for an ordinary container, according to researchers at Rensselaer Polytechnic Institute, in Troy, N.Y.

update

A Connecticut Fuel Cell In South Korea's Grid

Breakthrough contracts revive hydrogen hopes in electric power

GOOD DIGESTION: A fuel cell plant powers a Kirin Brewery operation in Tokyo, which runs on its waste gases in a self-contained system. *PHOTO: FUELCELL ENERGY*

If you had asked a decade ago about fuel cells, you would have learned that two Connecticut companies dominated the U.S. market for small power plants, their only real application at that time. After a lot of ballyhoo in the meantime about putting fuel cells in cars and the "hydrogen economy," the same two Connecticut companies still reign in the fuel cell market today, selling essentially the same technology. Now one of them looks to be taking off, having received a big boost from some unlikely customers in Asia, beginning with a Japanese brewer. Fuel cells finally seem set to become a significant player in electric power.

FuelCell Energy (FCE), in Danbury, Conn., was founded in 1969 as Energy Research Corporation. It has focused single-mindedly on the molten carbonate cell, which operates at a relatively high temperature. Fuel cells cause hydrogen to react with oxygen, across an electrolyte and with the help of catalysts at the cathode and anode, to produce an electric current and water as a by-product. Generally, fuel cells run at temperatures that make it necessary for the catalytic elements in the electrodes to be made from expensive materials like platinum or palladium, and the catalysts tend to be sensitive to impurities like carbon monoxide. The Danbury cell, running at about 600 °C, consumes the carbon monoxide along with its hydrocarbon fuel. Its internal steam reforming system, which can take the hydrogen it needs from coal gas, natural gas, or waste processing and digester streams, distinguishes the FCE fuel cell from all the others.

FCE's big break came in 2003 with an order from the Japanese beer maker Kirin Brewery Co., in Tokyo, where a battery of cells would consume the brewery's digestive gases while producing electricity to run the plant. Further business came from a variety of customers in the United States and Asia, almost always organizations that wanted to brag about using an innovative green energy system.

Having gained some attention in Asia with the Kirin order, FCE took another big step up the path toward full commercial viability this year, when South Korea's leading independent power producer, POSCO Power Corp., in Seoul, placed an order for 25.6 megawatts of fuel cell power plants and entered into a 10-year manufacturing and distribution agreement with the company. That agreement affirmed POSCO's faith in the commercial promise of the technology. Its purchases doubled FCE's orders and unequivocally made FCE the world leader in stationary fuel cells. A second Connecticut company, United Technologies Corp.'s UTC Power, in South Windsor—maker of the fuel cells that powered the Apollo vehicles on their trips to the moon—is the next-most-important player.

Since obtaining the breakthrough Kirin order in 2003, FCE boasts that it has brought its system costs down by 70 percent. Tony Leo, the company's vice president of applications engineering, says it set up a team to drive down costs systematically, using "value engineering" techniques to assess every component of a plant. That led to a decision to separate the fuel cell stacks, the dc-ac converter, and the thermal system rather than put all three major components in one expensive enclosure.

FCE is aiming to get installation costs down to US $2 per watt by 2012, compared with $3.25 or more five years ago. Analysts at ThinkPanmure and Lazard Capital believe that FCE has reasonable prospects of achieving profitability within 18 to 24 months because of its successful cost cutting and the diversity of promising markets for its product.

—William Sweet

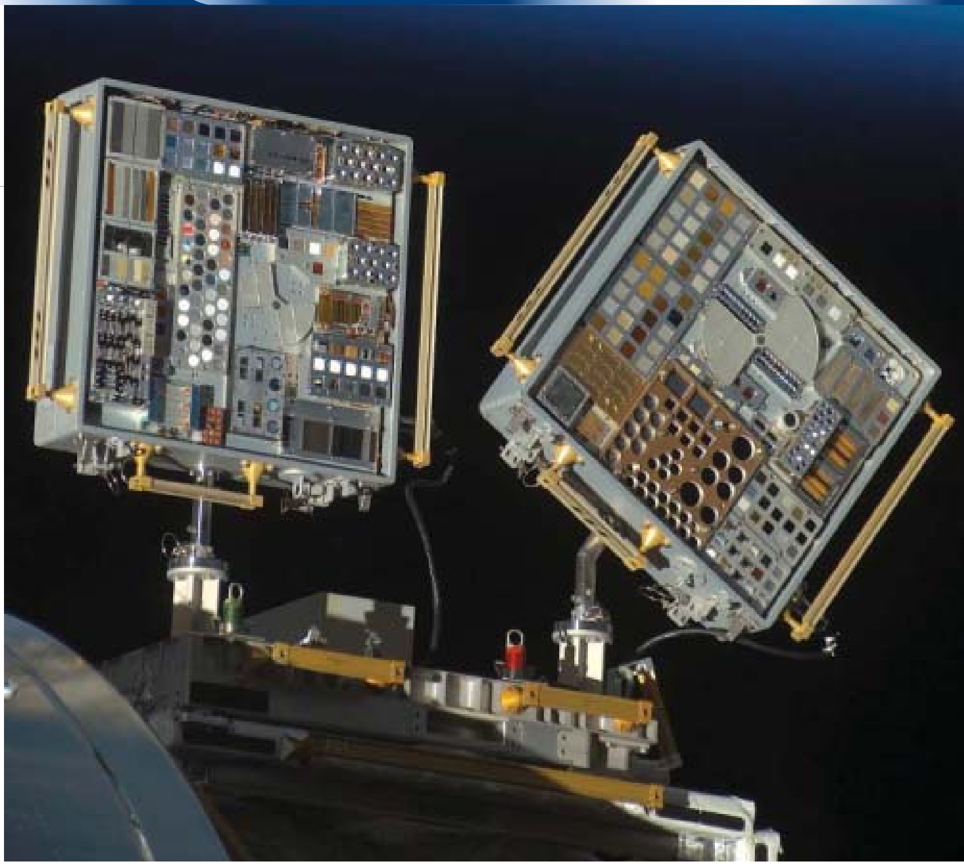

SPACE TEST: New radiation-resistant transistors, attached to the International Space Station for evaluation, orbit Earth.

PHOTO: NASA

Transistors in Space

NASA tests new radiation-proof chips

MISSION PLANNERS for future flights to the moon, Mars, and beyond rightly worry about how to keep astronauts safe from the hazards of high-energy protons, neutrons, and ions streaming through space. These particles can cause cancer and brain damage, but they can be just as damaging to a spacecraft's electronics.

NASA is now testing how well a new type of transistor, shown to be radiation-resistant on Earth, will hold up in space. The transistors went up on the space shuttle *Endeavor*, and astronauts placed them in a test setup on the outside of the International Space Station on 22 March. After a year, researchers will check how radiation affects the transistors' operating voltages and currents.

A typical transistor is made up of a gate separated from the channel between the source and the drain by an insulating dielectric, usually silicon dioxide. When the gate sees a voltage above a certain threshold, it creates a conductive path that lets charge flow between the source and the drain.

Radiation's effects on the silicon dioxide dielectric are what may cause a transistor to fail. Radiation ionizes the oxide, creating electron-hole pairs. The electrons flow out but the holes get stuck, building up a positive charge. This leads to two problems: it causes current to leak across the channel, draining batteries, and it decreases the transistor's threshold voltage, putting it out of sync with other transistors and disrupting the delicate timing of a circuit. "Every piece of circuitry is designed to operate in synchrony in a predictable manner," says Tobin Marks, materials science and engineering professor at Northwestern University, in Evanston, Ill., who led the work on the new transistor that NASA is testing. "You want [everything] to be rock stable."

Instead of silicon dioxide, the new transistors contain an organic dielectric, which assembles itself from three chemical solutions into a 15-nanometer-thick layer. The organic dielectric is more radiation-proof than silicon dioxide because it conducts holes well, says David Janes, an electrical and computer engineering professor at Purdue University, in West Lafayette, Ind. The cosmic radiation still generates electron-hole pairs, but the holes "bleed out" rather than building up, he says.

Marks and Janes have already shown that the transistors can withstand 2.85 kilograys (285 kilorads) of proton radiation. The space station test, scheduled to last 12 to 18 months, will show how the transistors hold up when exposed to ultraviolet radiation, as well as zero gravity and the corrosive effects of atomic oxygen. A one-way mission to Mars, including eight months of travel and three to six months on the planet, exposes electronics to about 10 kG of radiation—the limit of today's space electronics.

The custom processing required to make traditional radiation-hardened circuits ups their cost and keeps them several generations behind commercial computer chips in speed and power consumption. The electronics also require radiation shielding and backup circuits and devices, which add bulk. "We're always trying to save weight," says Geetha Dholakia, a research scientist at the NASA Ames Research Center, at Moffett Field in Mountain View, Calif. "If [the electronics are] inherently radiation-resistant, then one doesn't have to employ these other precautions." The new transistor, which contains a tiny zinc oxide nanowire as the semiconductor material, is lighter and could be cheaper to fabricate, claim Marks and Janes.

NASA might have other dielectric options for space electronics. To speed their transistors, Intel and other chip makers are turning to a new class of dielectrics, including oxides of hafnium and zirconium. But so far, there is little data about their radiation tolerance, Dholakia says.

—PRACHI PATEL-PREDD

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

the big picture

SPRAY OF LIGHT

The 90-meter-high piece of graffiti scrawled on the back of this office building in Rotterdam, Netherlands, was the product not of hundreds of aerosol cans but of Laser Tag, an open-source painting program from New York City–based Graffiti Research Lab. As an ordinary green laser pointer's beam was drawn across the building, a sensitive low-noise camera (developed for astronomers) captured its movements and relayed them to a laptop running the software. Then a movie projector "painted" the squiggles onto the multistory canvas in real time.

PHOTO: GRAFFITI RESEARCH LAB

 BY PAUL McFEDRIES

The Cloud Is The Computer

"You don't generate your own electricity. Why generate your own computing?"

—**Jeff Bezos**, CEO, Amazon

MICK WIGGINS

IN THE early 1990s, Sun Microsystems launched a new marketing campaign with a singularly perplexing slogan: "The network is the computer." What on earth did that mean? I knew about networks, of course, having worked in a Macintosh shop, but the wires between our computers were primarily for sending e-mail. My computer was my computer. After I became a freelance writer in 1991, my network was whatever floppy disk I had at hand (the infamous "sneakernet" solution). My computer was still my computer. Even after I cobbled together my first 10Base-T network, in 1993, the only consequence was that I used floppy disks less. My computer stayed resolutely in front of me, and I'm sure this was the case for most folks back then.

In the past few years we've seen Sun's slogan morph from perplexing to prophetic. As we do more and more online, we see that the network—that is, the Internet—is now an extension of our computers, to say the least. Particularly with wireless technologies, we see that a big chunk of our computing lives now sits out there in that haze of data and connections known as the **cloud**. In fact, we're on the verge of **cloud computing**, in which not just our data but even our software resides within the cloud, and we access everything not only through our PCs but also **cloud-friendly** devices, such as smart phones, PDAs, computing appliances, gaming consoles, even cars.

It's the dawn of the age of **pervasive computing**, when computation and information are ubiquitous and always available. It will be the age when the World Wide Web becomes the **World Wide Computer**, or simply (albeit ominously) the **megacomputer**.

The cloud is a powerful computer not only because so many of us now connect to it but also because many companies are spending billions to bulk it up. The basic unit of the new infrastructure is the **data center**, a typically massive building (think multiple football fields in size) housing enormous **clusters**, collections of computers (usually numbering in the thousands), networked together to power the online services we've come to rely on: Google's search engine, Amazon's e-commerce operations, YouTube videos, and so on. The stupendous size of these **server farms** is one reason even the most complex Google search yields results in a split second, and why even the longest YouTube video starts within seconds.

But the cloud doesn't exist just so you and I can get speedy search results and videos. It also offers companies more computer power. This began with Amazon's S3 (Simple Storage Service), which offered companies storage space on Amazon's thousands of servers for a fraction of the cost of purchasing physical hard drives. Now with EC2 (Elastic Compute Cloud), Amazon is letting companies use the spare cycles on Amazon's servers. (The virtual world of Second Life is generated by Amazon computers.)

This trend is called **hardware as a service** or simply **virtualization** because it offers companies the use of virtual computing equipment. A similar idea is **software as a service**, which gives companies access to virtual software packages such as payroll, accounting, and customer-relationship management. Both **HaaS** and **SaaS** let companies reduce their IT budgets by ordering only as much **computing on demand** as they need.

In his fascinating recent book *The Big Switch: Rewiring the World From Edison to Google*, which discusses the coming shift to cloud computing, Nicholas Carr points out that the cloud is rapidly becoming analogous to the electric grid: in the same way you can just plug a device into a wall socket and get electricity cheaply and instantly, you can now plug into the cloud and get data, storage space, and even processing power cheaply and instantly. This is **utility computing** powered by massive **utility data centers**—the cloud equivalents of power plants.

Those of us who've been around long enough to recall an analog-only world may resist this new age, but not so the **digital natives**, who are growing up in a world filled with computers, mobile phones, and other digital devices. **Digital nativism** is the belief that these **screenagers** have a major advantage over the rest of us, not only in using these newfangled gadgets but also in the increased cognitive skills that using those devices allegedly confers. Is that true? I don't know, but I bet the cloud does. ❑

"If you're a 20- to 30-year-old analog engineer, you're sitting pretty right now"
—JIGNASHA PATEL, Human Resources, Freescale Semiconductor

careers

U.S. Engineering Salaries Rise Sharply

EE starting salaries are up 13 percent; some other fields do even better

THE U.S. economy may be stuck in second gear, but for electrical engineers at all levels of experience, the job market just keeps getting better. Opportunities and salaries are increasing even more in Europe and Asia, where economies are healthy if not booming.

The field ranks third among bachelor's degrees and second among master's degrees on the list of majors that companies most want, says Andrea Koncz of the National Association for Colleges and Employers (NACE) in Bethlehem, Pa. Demand is particularly strong in defense, aerospace equipment, and medical and consumer electronics.

Even nontraditional employers are on the hunt. "We're losing people to places like Disneyland," says Boeing spokeswoman Cindy Wall.

The average offer this year was US $56 512, up 3.5 percent from last year. That's much higher than the $49 427 average for civil engineers, but modest compared with the chemical engineers' $63 749, buoyed by the soaring fortunes of the petroleum-products sector. "What's driving salaries in those industries can arguably be summed up in one word: energy," says Tim McCreight, marketing manager at the American Institute of Chemical Engineers.

EE salaries, on the other hand, are rising in part because the United States is not producing enough grads to replace retiring employees. Wall thinks that this might be because students have more career options now and might not consider math and engineering as appealing as they were back in the space age. "Look at the Sputnik era," she says. "There were so many engineers rolling into the workforce 40 years ago, and we don't have that now."

Engineers with analog, RF, and wireless skills, as well as those with postgraduate training, are becoming harder and harder to find. "If you're a 20- to 30-year-old analog engineer, you're sitting pretty right now," says Jignasha Patel, a human resources director at Freescale Semiconductor in Austin, Texas. "It's a buyer's market for you."

Telecom firms and chip makers are increasingly looking globally for new hires. Two-thirds of Freescale's employees live and work outside the United States. The company recruits heavily in China, India, and Romania, where it is also looking for analog talent. These hot markets are seeing high salary increases of 10 to 12 percent, Patel says.

U.S. aerospace and defense contractors can't cast their nets so broadly because EEs overseas can't easily obtain the necessary security clearances. "We think there might not be enough people to replace the baby boomers," Wall says.

Companies are also hiring more masters and Ph.D. grads to get the skills they need in analog and RF systems. At Texas Instruments, in Dallas, the percentage of new-grad hires with M.S. degrees or higher rose from 10 percent a decade ago to 54 percent last year and about 60 percent this year, says Diana Johnson, a campus recruiter for the company.

Recruiters are hiring more aggressively on campus, building stronger relationships with students and faculty. They are also starting earlier, often going out in the fall to look for the best candidates. Jeff Goodman of Raytheon says nowadays he will make firm job offers before Christmas. Companies are also hiring more interns; NACE says that almost 70 percent will get job offers, up from 57 percent in 2001.

More employers are offering signing bonuses, and the average bonus, now at $4450, is 25 percent higher than last year's. "Candidates at the top of their class are getting four or five offers, a blend of pretty aggressive base salary, sign-on bonus, and good relocation support," says TI's business recruiting manager, Scott Schmitt. "We're not hearing from candidates that they don't want to come to Texas Instruments due to salary as much as location or work lifestyle."

The very best new grads—those with internship experience, advanced coursework, and some extra effort in their classes—sit in the catbird seat these days. —PRACHI PATEL-PREDD

How do EE salaries fare against other fields over time? What are the best areas within EE? See The Data, last page of this issue.

GRADUATES: HANS NELEMAN/GETTY IMAGES; MONEY: DON FARRALL/GETTY IMAGES

Qatar University Opens EE Doors To Women

A small but oil-rich country needs all the electrical engineers it can produce

ENGINES OF CHANGE: Qatar University students conduct a lab experiment. Women entered the chemical engineering program in 2004; the EE program will see its first women students this fall. PHOTO: QATAR UNIVERSITY

NEXT MONTH, the electrical engineering program at Qatar University, in Doha, will open its doors to women.

Qatar is a peninsular country like Denmark but with one-quarter of the territory and a fifth of the population. Even so, it has a huge demand for engineers. Indeed, the entire Qatar economy is as hot as its Arabian sands, with a gross domestic product that's expected to grow 14 percent this year, to US $73 billion. To fuel yet more economic development and cultivate a culture of innovation, the government annually allocates an astonishing 2.8 percent of that GDP to technology R&D, ahead of the United States and just behind South Korea. It hopes as well to reduce dependence on foreign technology workers, and that's where Qatar University comes in.

"There is a great focus on education here that is unique in this part of the world," says Adnan Abu-Dayya, chairman of the university's electrical engineering department. "We want to make sure there are opportunities for local residents."

The plan calls for doubling the size of the EE class to allow for some 25 to 30 women, a 50:50 student ratio, which is higher than that of any Western country.

"There's a closer ratio of men to women engineers here than in the United States or Canada," says Abu-Dayya.

For employers, the influx of local female engineers to the workforce plugs a brain drain left by men who head to Europe and America for their education and never return. Abu-Dayya worked at AT&T Wireless in Seattle after earning a Ph.D. in electrical engineering from Queens University in Kingston, Ont., Canada.

"There's a cultural issue here. Men tend to go away to university, while the women stay home-based," says Ian Dowdeswell, incident manager of Q-CERT, the Qatar government's center for information security. "Engineers are in strong demand all over the world, especially in the Middle East," says John Challenger, CEO of the Chicago global outplacement firm Challenger, Gray & Christmas. "I see it as the world forcing the change, because the cultural norm is holding the country back."

But a big push for an EE major came from the women themselves. "Generally, women are serious, passionate, and dedicated to education," says Abu-Dayya. "They are hardworking, like to explore, and are as committed as the men, if not more."

The university's first engineering majors to open to women were industrial and computer engineering in 2001, followed by chemical engineering in 2004. Electrical engineering, critical to such fundamental elements of the infrastructure as power systems and telecommunications, comes next.

The program also helps place students in internships and postgraduate employment and engages in research collaborations with nearby branches of American universities. Eventually, the department plans to offer graduate engineering degrees.

"Right now, our economy is big in oil and gas. That may last 50 years, or it could last 300 years. Eventually, we have to depend on other sources of income," says Abu-Dayya. "This is part of a process to convert our society over time to a knowledge-based economy—to have human beings be our natural resource."

—SUSAN KARLIN

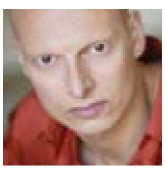

mini-profile
By Susan Karlin

JOSEPH GATT: MOTION-CAPTURE ACTOR
The demand for lifelike computer games has spawned a cottage industry of motion-capture actors. These anonymous performers wear dozens of reflective sensors so that special cameras can transmit their every movement to a computer, which assigns the actions to the CGI character. "Every neck crook and arm swing is picked up. It takes a while to get the hang of it," says Joseph Gatt, who modeled for the character of Kratos in Sony Computer Entertainment's award-winning *God of War II*. Rates run from US $900 to $1000 for a typical 10-hour day.

PHOTO: KEVIN MAJOR HOWARD

The Poor Man's Patent

Note to self: a note to yourself isn't worth the paper it's scrawled on

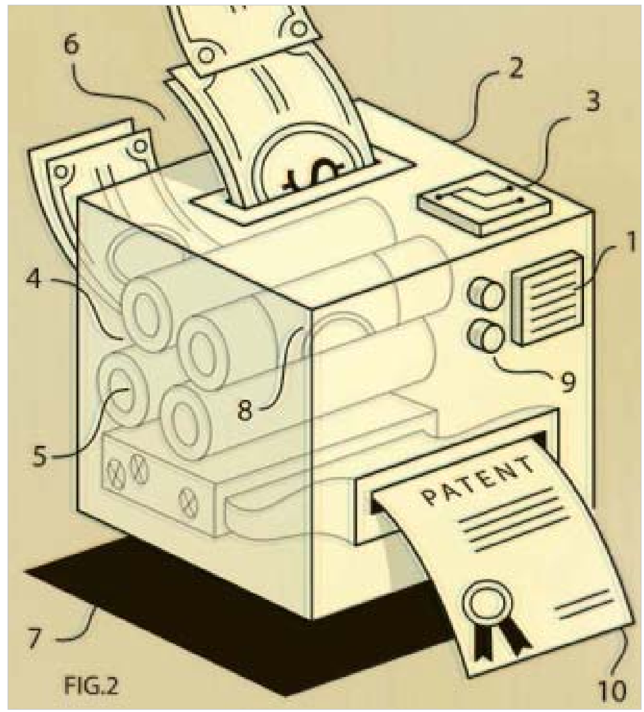

PATENT ATTORNEYS charge between US $7000 and $15 000 to prepare and file a patent application. If only there were a cheaper way, a kind of poor man's patent. But it just doesn't exist.

Some people think they can protect their invention by writing a patentlike description of it and mailing the document to themselves, but this is no substitute for patent pending. At best, the letter shows that you conceived an invention by a certain date, but you'll probably be able to prove that with engineering notebooks, e-mails, dated PowerPoint presentations, and the like. Moreover, evidence of an invention's conception date is useful only in a limited set of circumstances, most of which involve actually filing for a patent at some point in time. So save yourself the paper and the postage stamp.

There used to be a so-called disclosure document you could file with the U.S. Patent and Trademark Office for $10, but early last year the USPTO eliminated the option because it not only provided no real benefit to inventors but also misled them into thinking they had achieved patent-pending status.

Undoubtedly more useful is the provisional patent application. Once this application is filed, an invention can be disclosed or sold without fear of losing patent rights, so long as a full "utility" patent application is filed within a year of the provisional. The problem is that if you draft a provisional improperly, it can limit the scope of possible protection and might even be thrown out by a court, invalidating the utility patent you file later. Worthwhile provisionals may end up costing nearly as much as a real patent application, so they don't really count as a poor man's patent.

Unscrupulous operators regularly prey on naive inventors by touting inexpensive patent-pending status. The Federal Trade Commission and the Patent Office are onto such scams and have established programs to shut them down. In one notable case, a patent attorney working for an invention promotion firm saved money by filing only design patents for his clients. Design patents, which, generally speaking, protect the way an article looks, have limited usefulness for most inventions. The attorney was disbarred.

Nor can you save the cost of a patent by simply relying on copyright or trade secret protection. Copyright protection may prevent someone from photocopying your write-up of your invention, but by law it cannot stop anyone from implementing the ideas or functionality of your invention in a competing product. As for trade secret laws, they have a hard time protecting ideas or functionality, because they can't stop someone from reverse engineering your invention.

Suppose you invent a light saber toy (à la *Star Wars*) and document all its circuitry in schematics. A competitor buys your product, gains an understanding of its overall functionality, tears into it in order to reverse engineer the circuitry involved, and begins selling competing light sabers. There has been no copyright infringement, because your schematics were not copied, and there is no trade secret violation, because the minute the product was sold, its status as a secret ended. You say that you mailed the schematics to yourself? That does not give you grounds to sue anyone for anything. You have a design patent? That would be easy to get around. You have a provisional patent application? I hope it discloses all the relevant circuitry and is not just a bare-bones description of a mechanism producing a blade of contained energy. No version of a poor man's patent could possibly help you here.

Patents are expensive, no doubt about it, and the requirements are fairly strict. But as my grandmother used to say, you get what you pay for. —KIRK TESKA

HARRY CAMPBELL

CHILD'S PLAY: Journalist James Turner turned his One Laptop Per Child computer into a workhorse. PHOTO: STREETS OF BOSTON

Hacking the OLPC

With a little tinkering, the One Laptop Per Child XO can be a grown-up computer too

I GOT AN early look at the XO laptop, from the One Laptop Per Child Foundation, at the 2007 International Consumer Electronics Show. The friendly green-and-white computer, with a truly glare-free screen, appealed to me right away.

So last December I took advantage of the OLPC "Give One, Get One" program. For US $400 I scored a second-generation XO, a $200 tax deduction, the presumed gratitude of an anonymous child, and a year of T-Mobile Wi-Fi service. My hunch was that, with a little hacking, the XO could be as useful as an equally small Asus Eee PC, which also costs about $400 but comes with no charitable benefit.

I quickly found that for casual use, the XO's small size and built-in handle make it a great workplace or travel companion. It comes with Wi-Fi, and you can get a USB Ethernet interface if you prefer a wired connection. For all-day use, the keyboard is way too small for most adult fingers, and the touch pad is too twitchy, but a USB mouse and a folding USB keyboard solved those problems. If the built-in keyboard gets in the way, you can use the XO's "transformer hinge" to rotate the screen 180 degrees.

Getting the software side of things right took a little more effort. Sugar, the XO's preinstalled version of Linux, may be a uniquely great operating system for kids, but it really doesn't cut it if you're trying to do work-work. Luckily, plenty of XO owners before me seem to have felt the same way. I quickly found some great step-by-step instructions for installing alternative versions of Linux, including Ubuntu, my choice for this project.

The new operating system goes on your XO's secure digital card. To boot from an alternate device, you need to first access the OLPC's boot loader, and for that you need to get a developer key. Instructions are at http://wiki.laptop.org/go/Developer_Key.

With the key installed, you're ready to build a bootable SD-card image. Directions are at http://tinyurl.com/3xd8xb. I found that an 8-gigabyte SD card had more than enough room for applications and files, but physically placing the card in the XO wasn't easy. The SD reader is under the screen and awkward to get at.

For my operating system, I chose a distribution of Ubuntu known as Xubuntu. Given the XO's modest 1200-by-900-pixel screen resolution and 433-megahertz processor, I went with a lightweight window manager, Xfce.

With Xubuntu up and running, I used two common Linux tools for managing software, apt-get and the Synaptic Package Manager, to add new applications. One thing to note is that you'll get errors about not having an Xauthorization file unless you install Xauth using apt-get, then issue the command "xauth generate :0 . trusted." You'll also want to load up NetworkManager so that you can connect to secure networks and load CUPS (Common Unix Printing System) to print to network printers.

I also installed Gnome-RDP, a remote-desktop application. Remote Windows XP sessions were surprisingly usable, perfect for checking Outlook or taking a peek at other files on my desktop.

Within a week I had a pretty amazing little portable to carry around. Although the screen is smaller than a traditional laptop's, its resolution is higher. The XO is light and durable, has superb Wi-Fi reception, and is an eye-catcher wherever it goes. And nothing beats it as a portable Gmail reader at a conference.

Even with its limited memory and economy-class processor, the XO runs browsers and remote desktop sessions as well as regular business-class laptops do. I can't wait to try it out on a plane—no more laptop-in-the-belly syndrome, and the battery might even last the whole flight.

You can keep your Asus Eee PC. I love my Ubuntu-powered XO!

—James Turner

books

Patently Oblivious

IEEE Spectrum contributor Kirk Teska guides managers through the patent process

Even patent attorney and frequent *IEEE Spectrum* contributor Kirk Teska allows that patents can be a little dry. He recalls his first patent course in law school as "the most boring class I had."

But as I bounced through his latest book, *Patent Savvy for Managers*, I found it harder and harder to stay bored. I kept getting sucked in by quietly informative case studies, like the one about chewable dog toys that shows how vague language may contribute to a patent defense. While flipping through the pages I was distracted by a diagram of a Toro Corp. lawn aerator. Three pages later, I found I understood what happens when your patent claims are too broad or too narrow.

Throughout, *Patent Savvy*'s examples—like patentable inventions themselves—are useful, novel, and nonobvious. And entertaining. They include cubic zirconium (is it useful?), plastic leaf bags that look like jack-o'-lanterns when they're filled (isn't it obvious?), a collapsible escape ladder, laser eye surgery, and the aerogel used to keep crayons from melting.

There are detailed examinations of the famous Amazon one-click patent and of Gillette's patent for the three-blade razor, which may or may not have been infringed by rival Schick when it made a four-blade razor. (Oddly, Gillette claimed both that its patent was infringed and that the Schick razor was much better than its own.) There's also a chapter-long study of *Research in Motion v. NTP*, the case that decided the fate of the popular BlackBerry handheld—a US $612 million patent lawsuit with more ups and downs than Super Bowl XLII.

Useful sections show how to read a patent, file outside the United States, and document an invention while it's being invented; the difference between buying and licensing patents; and the hard economics that govern such decisions as whether to contest a patent rejection. There's even room for witty chapter titles, like "What to Do When Your Candy Bar Melts," and for a joke or two, like comedian Steven Wright's one-liner "I invented the cordless extension cord."

The key phrase in the book's title is "for managers": Teska faithfully focuses on the business side of patents. Most important, his book can help managers decide whether an innovation is worthy of a patent filing in the first place, by applying a hard-nosed cost-benefit analysis. If you manage engineers as they innovate or just want to manage your own innovations, *Patent Savvy for Managers* is a useful addition to your bookshelf portfolio.

—Steven Cherry

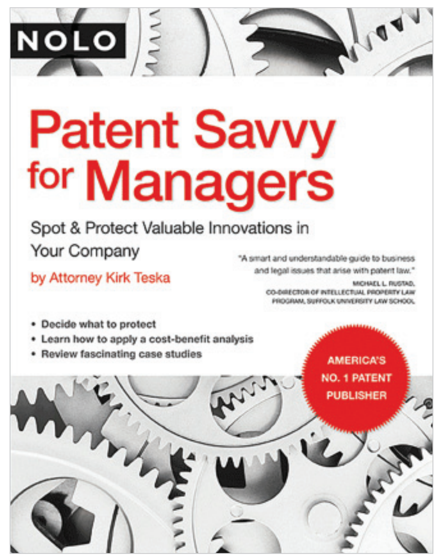

PATENT SAVVY FOR MANAGERS: SPOT & PROTECT VALUABLE INNOVATIONS IN YOUR COMPANY
By Kirk Teska; Nolo, 2007; 278 pp.; US $29.99 (paperback); ISBN: 978-1-4-133-0694-1

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

FIRST SOLAR: QUEST FOR THE $1 WATT

WITHIN FIVE YEARS, THIS COMPANY'S THIN-FILM SOLAR CELLS COULD COMPETE WITH COAL

By Richard Stevenson

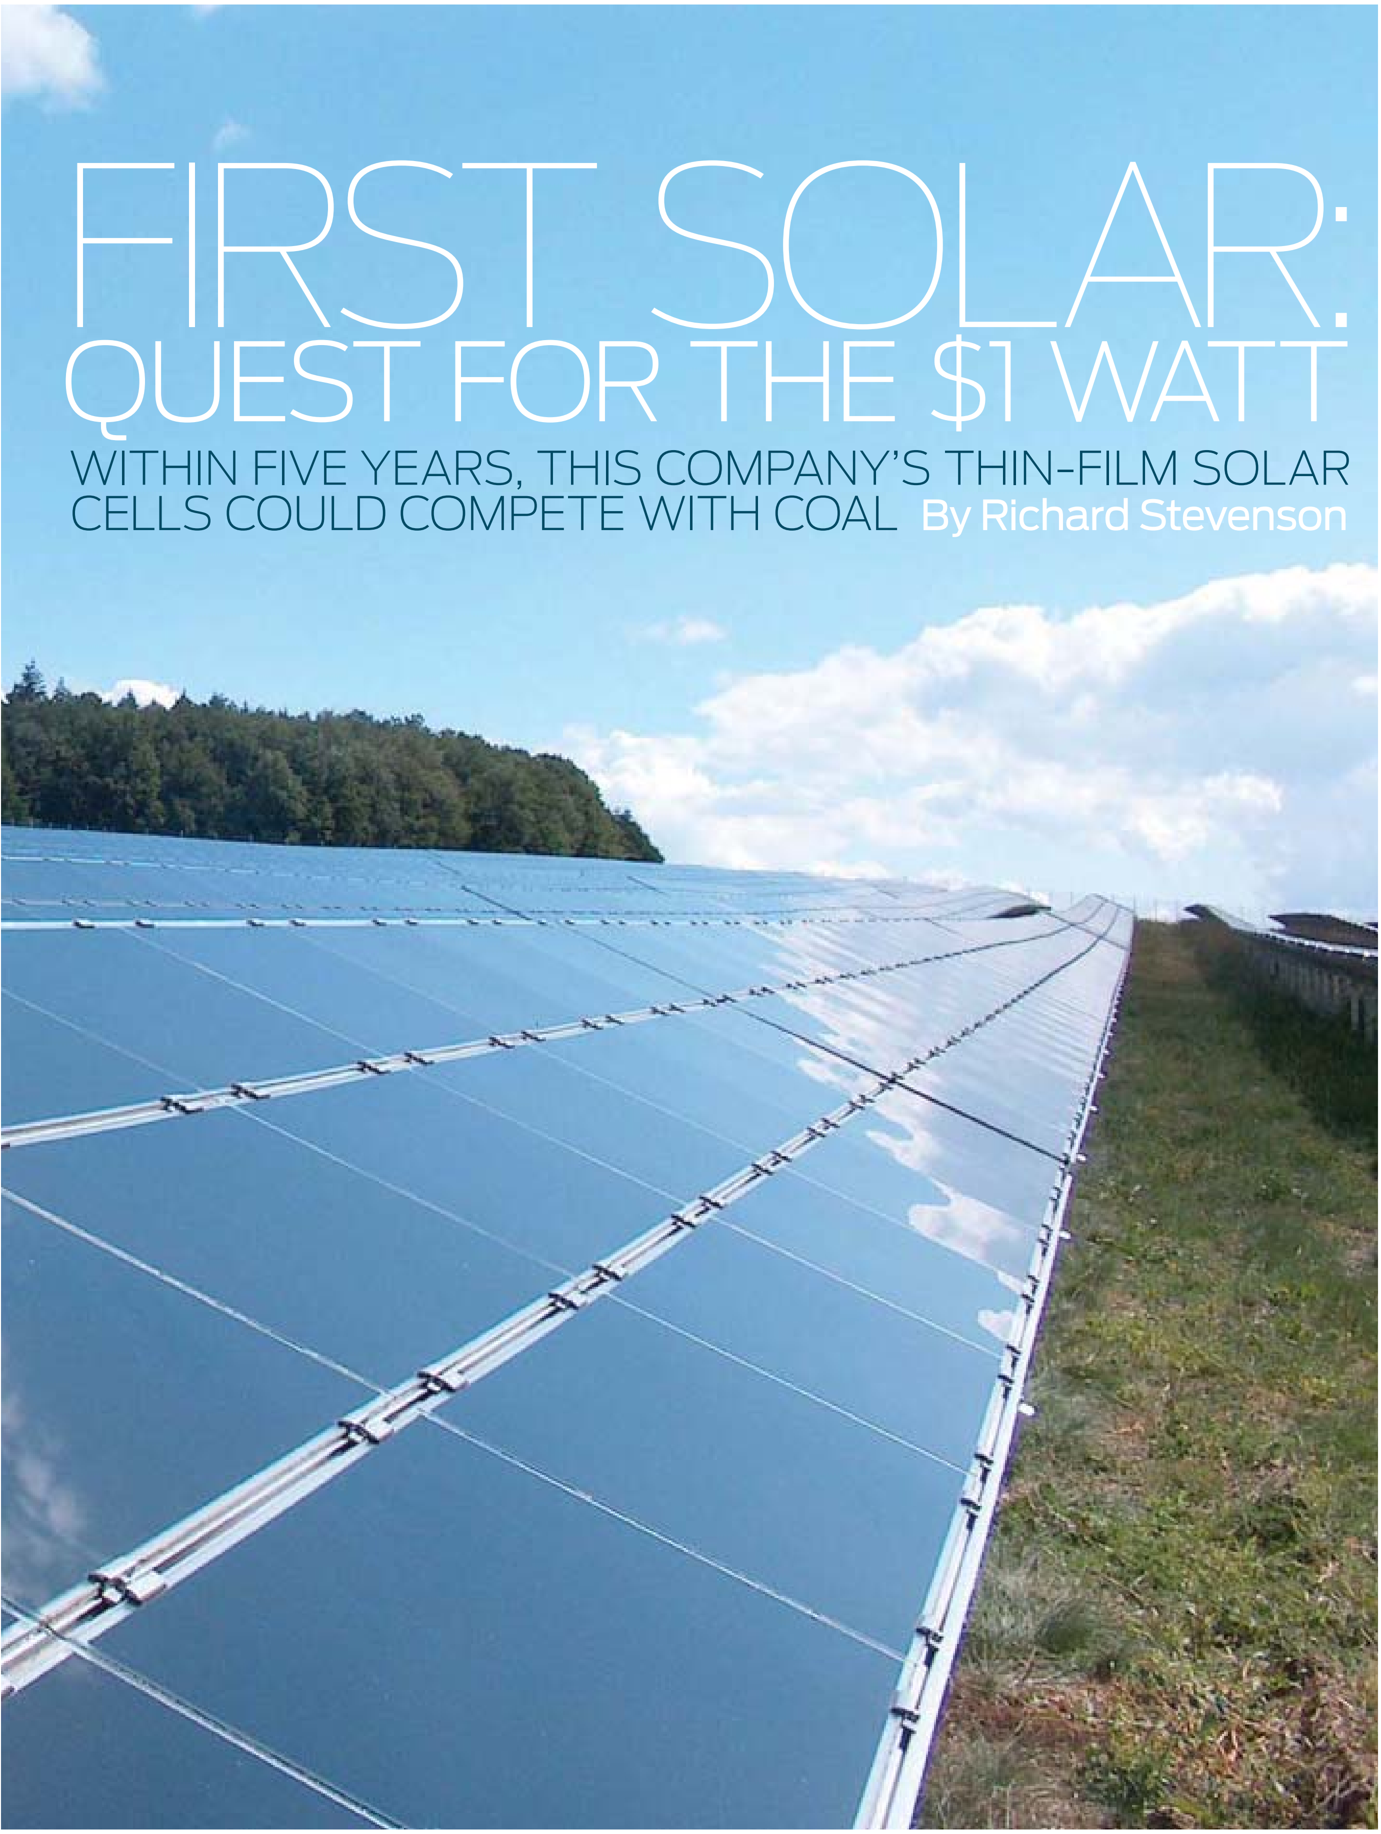

IT'S EASY TO MAKE A SMALL PILE OF MONEY OFF photovoltaic cells but very hard to make a big one. The reason is one of the most fundamental in free-market economics: the larger the market you aim for, the more competitors you'll have to face.

THE HUGE SCALE of First Solar's manufacturing prowess is evident in single panels [previous spread] and in arrays like this 40-megawatt solar farm in Brandis, Germany.

IF YOU JUST WANT TO POWER a billion-dollar space probe, almost any price per watt is acceptable. If you are selling to lonely farmhouses, you just have to charge less than the cost of running a power line to the boondocks. In some parts of the world, competing with grid electricity itself may be an easy game during peak consumption hours. But if you want the off-peak market, you'll have to price your cells at about US $1 per watt. That price is called grid parity, and it's the holy grail of the photovoltaic industry. At least 80 firms around the world, from Austin to Osaka, are in the chase.

Surprisingly, at the moment no company is closer to that grail than a little start-up called First Solar, which until very recently had been known only to specialists. It's located in Tempe, Ariz., and analysts agree that it will very likely meet typical grid-parity prices in developed countries in just two to four years. It's got a multibillion-dollar order book, it's selling all the cells it can make, it's adding production capacity as fast as it can, and its stock price has rocketed from $25 to more than $250 in just 18 months.

The most tantalizing fact about First Solar? The company will not talk to reporters. At all.

The company's coyness seems to be related to the nature of its industrial secrets. These have less to do with First Solar's device—a decades-old design based on a thin film of cadmium telluride—than with the way the company manufactures it. Somehow, First Solar has scaled up the light-catching area from postage-stamp to traffic-sign dimensions. What the company does reveal is that its product has three massive cost benefits. Its active element is just a hundredth the thickness of the old standby, silicon; it is built on a glass substrate, which enables the production of large panels; and manufacturing takes just two and a half hours—about a tenth the time it takes for silicon equivalents.

Of course, it's not enough that First Solar match the costs of fossil-fuel generation on the grid; it must also maintain its economic edge over other photovoltaics. There are additional nascent technologies, including cells based on copper indium gallium diselenide (CIGS), silicon on glass, and the combination of germanium, gallium arsenide, and gallium indium phosphide. Even conventional silicon technology, which has dominated the market since its commercial launch in the 1950s, seems to have a lot of kick left in it. Currently, though, it's suffering from its own success, as an insatiable demand for silicon cells has led to a scarcity of raw material. However, if the silicon shortage disappears by the end of the decade, as expected, the sale price should drop substantially from recent levels, which have fluctuated between $3 and $4 per watt.

PREVIOUS SPREAD: GEHRLICHER SOLAR; THIS PAGE: JUWI

Right now, First Solar depends mainly on a government-subsidized program in Germany, where it has contracts worth more than $6 billion through 2012. Other markets with the same type of subsidies (known as feed-in tariffs, which spread the cost of alternative energy among all customers) include France, Italy, Spain, South Korea, and Ontario, Canada. To fill these orders, the company is undergoing a massive expansion of its manufacturing facilities that should boost annual production capacity to just over 1 gigawatt by 2009. This capacity could supply one-sixth of that year's estimated global solar-cell business, which is currently growing at 50 percent per year.

This rapid ramp-up is impressive for a company founded only in 1999, after it acquired its cadmium telluride (CdTe) technology from the purchase of Solar Cells Inc. (SCI). Cash for the launch came from the equity firm JWMA, whose president, Michael Ahearn, became First Solar's CEO and is still running the company.

First Solar began by developing its manufacturing technology at its Perrysburg, Ohio, facility. Commercial operations started in January 2002 with a 25-megawatt base plant, which began high-volume production a couple of years later. Since then the company has replicated its manufacturing line at the Ohio site, built four more lines in Germany, and begun constructing a fourth plant in Malaysia, which will bring the total number of production lines in that country to 16. Ahearn recently told investors that the first Malaysian plant has just started to produce cells and that it should be operating at full capacity by the end of next year. Line capacity has risen also, to 45 MW.

THE CELLS ARE MANUFACTURED on 0.6-by-1.2-meter sheets of glass, which are cleaned and cut on an angle to produce the strong, defect-free edges required for processing. The glass has already been coated with a transparent tin oxide that provides electrical contact to the device. This starting platform is radically different from that for silicon cells, which are made from far smaller monocrystalline and polycrystalline wafers.

Next, the device layers are deposited onto the sheets. This is the stage at which First Solar's secret surely applies, says John Hardy, an analyst at American Technology Research, in Greenwich, Conn. In his view, keeping this secret is one of the main reasons that First Solar refuses to talk to the media.

Nevertheless, it is still possible to uncover some of the details of First Solar's growth process. Dieter Bonnet, a coinventor of the CdTe cell and the chairman of Solarpact, a research consortium in Germany, says that First Solar's process is just a refined version of that used by its predecessor, SCI, which released a report about its manufacturing technology in 1993. Interestingly, this document was coauthored by James Nolan, a current director of First Solar, the person responsible for designing and building prototype equipment for the pilot manufacturing line.

The report from SCI describes an elemental vapor deposition process that takes place in four chambers. Glass is placed on rollers and fed into the first chamber, where it is heated to 600 °C. Then it is transferred into the second chamber, which is full of cadmium sulfide vapor, formed by heating solid CdS to 700 °C. The vapor forms a submicrometer deposit on the glass as it moves through this cloud, after which a similar process in a third chamber adds a layer of micrometers-thick CdTe in about 40 seconds. Then a gust of nitrogen gas rapidly cools the panels to 300 °C in a fourth chamber, strengthening the material so that it can withstand hail and high winds.

The two layers—CdS and CdTe—are critical because they constitute the electronic junction that converts light into electricity. Most of the sunlight entering the glass passes through the thin CdS layer before being absorbed by the much thicker CdTe film. Here the light transfers energy to electrons in CdTe, freeing them from their normal bound state so that they can move through the material. To get them moving, however, you need an internal electric field.

In silicon cells, that field is created internally by constructing two adjacent layers with different electronic properties. One layer consists of silicon doped with small amounts of phosphorus, which has one more electron in its outer orbital than silicon does. When a phosphorus atom is inserted in place of a silicon atom, that extra electron is transferred to the crystal lattice. Because these electrons move about freely and carry a negative charge, this material is known as *n*-type silicon. *P*-type silicon, on the other hand, gets its corresponding positively charged particles from tiny amounts of boron, an element that has one less electron than silicon in its outer shell. In this case there are not enough electrons to form all the covalent bonds required, so the electrons move around to try to fill this deficiency, which is called a hole. Holes act like free, positively charged particles.

When *p*-type and *n*-type materials are placed together, they form a *p*-*n* junction. The electrons and holes attract one another, congregate by the interface, and leave the *p*-type and *n*-type regions with negative and positive charges, respectively, creating the required electric field.

CdTe beats silicon in this respect because it can form a *p*-*n* junction more simply. CdTe and CdS are made by bonding two elements with different numbers of electrons in their outer shells, and any deviation from an exact 50:50 balance between these elements produces doped material. In fact, a slight imbalance naturally occurs in both materials, and that makes it very easy to make *p*-type CdTe and *n*-type CdS. In silicon, the two halves of the junction require more steps to manufacture.

A further advantage results from the position of CdTe's absorption edge. Calculations have shown that the ideal solar cell would start absorbing sunlight at a wavelength of 910 nanometers. CdTe is close to this sweet spot, with absorption kicking in at 850 nm, while silicon starts to absorb only at 1.1 micrometers.

After forming the junction between CdTe and CdS in the four-chamber tool, First Solar heats the panels to improve the efficiency with which it converts light into electricity. Bonnet says that this process takes place in the presence of some form of chloride,

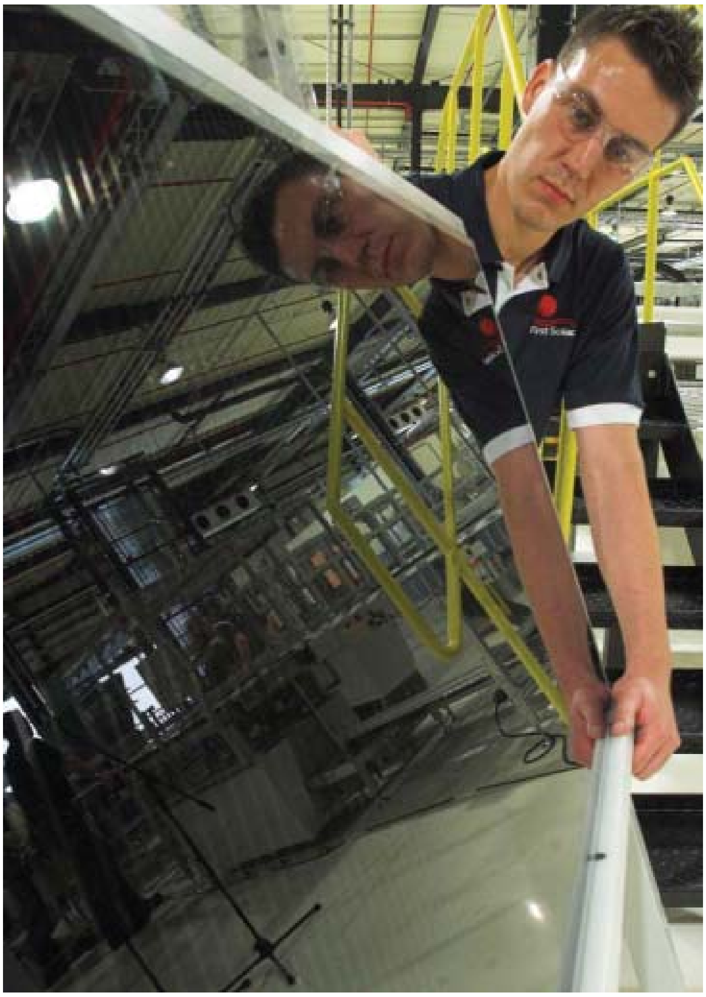

A RARE LOOK inside a First Solar plant—here, in Frankfurt/Oder, eastern Germany—gives away none of the company's secrets, which involve manufacturing very large panels. The cell structure itself is known [see diagram below]; as in all thin-film designs, the active layers are built on an inert glass wafer.
PHOTO :SVEN KAESTNER/AP PHOTO; DIAGRAM: EMILY COOPER

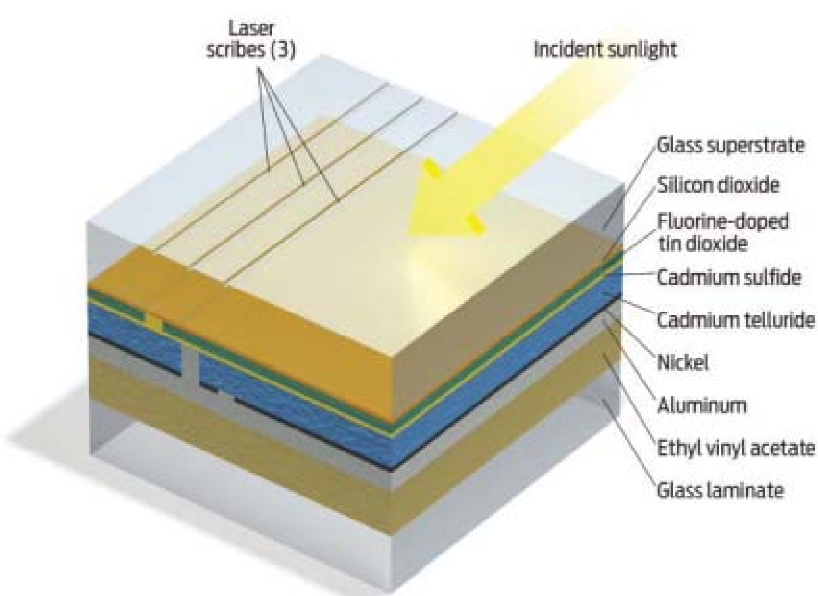

because it boosts efficiency, although the mechanism is not yet understood. The benefits are significant, and lab results have shown that cell efficiencies more than double as a result, to better than 10 percent.

Once the heating is over, a laser patterns the CdTe sheet into an array of smaller, rectangular solar cells connected in series. By stringing together the right number of cells, this process tailors the panel's output to produce 70-volt modules, delivering a current that ranges from 0.97 to 1.08 amperes. Finally, First Solar deposits a metal contact onto the CdTe and adds a laminate, a rear glass cover sheet, and termination wires.

Today's modules deliver up to 75 W at a conversion efficiency of 10.6 percent and have a manufacturing cost of $1.14/W. This is way below the selling price of $2.45/W, so the company enjoys a healthy profit margin. However, to compete against fossil-fuel sources on the free market and pick up a tidy profit, the company will have to get manufacturing costs down to between $0.65/W and $0.70/W. To do so, it has told investors that it needs to reduce manufacturing costs and increase conversion efficiency to 12 percent. Getting there is entirely feasible, as CdTe cells have a theoretical maximum of well over 20 percent; the National Renewable Energy Laboratory, in Golden, Colo., has already produced cells with 16.5 percent efficiency.

At first glance you might think that conversion efficiency shouldn't matter, so long as the price per watt is low enough. That argument applies only at the level of an individual module, however, not an entire installation. Because First Solar's panels are less efficient than silicon designs, they need more space to soak up enough sunlight, and that raises the cost both for real estate and installation. The company is aiming to reduce the installation costs through its recent cash purchase of the U.S. firm Turner Renewable Energy, which designs and deploys commercial solar projects.

Total cost must also reflect the expected life of the modules. First Solar says its product will last 25 years, after which the materials will be returned to the company for recycling. Besides helping the environment, recycling would provide First Solar with material, albeit after a long wait.

Cadmium is plentiful as a by-product of mining, but some critics doubt the long-term global availability of tellurium. Company president Bruce Sohn dismisses this notion. In a conference call in May he said, "We are not seeing any supply issue for tellurium. We have a couple of sources, and we have locked down our long-term contracts for raw materials. That has helped us maintain the supply as well as the price."

Although the modules would be great at providing solar electricity to homes, for the time being First Solar isn't selling to the public. Ahearn has told investors that the company is having a hard enough time supplying the demand for solar farms—some 55 percent of its business—and commercial rooftop installations. Panels for solar farms can be mounted on low-cost trellis systems. Most rooftop installations are for business premises. Such deployments could be seen as hazardous, because in a fire the panels could give off potentially fatal cadmium fumes. To fight these fears, First Solar cites an experiment done at Brookhaven National Laboratory, in Upton, N.Y., in which cells heated to 1100 °C lost just 0.04 percent of their cadmium, an insignificant amount.

More serious is the threat posed by rival thin-film technologies, which together have been receiving very high levels of investment since the silicon shortage began.

Of the alternative technologies, CIGS has been grabbing most of the headlines, thanks to its claims for maximum efficiencies of up to 20 percent. Another advantage is the ease with which CIGS can be deposited as a thin film. This technology has yet to live up to its billing, however.

"It's an awful lot of hype as opposed to a lot of reality," says Robert Castellano, president of the Information Network, based in New Tripoli, Penn. "No one has come up with a full-blown production setup, and that has soured all the venture capital and private equity companies."

He says that investors were lured by promises of simple, quick production processes for panels having an efficiency of 12 percent. In fact, though, efficiencies have been lower, and manufacturing has been delayed. Heliovolt Corp., in Austin, Texas, and Nanosolar, in San Jose, Calif., for example, have each raised over $100 million of investment but are only on the fringes of manufacture after more than five years of development.

Another thin-film technology, called amorphous silicon on glass, is already making an impact on the solar market. It has efficiencies of around 7 percent, and because it uses only tiny quantities of silicon, it has been largely unscathed by the silicon shortage. Also, because manufacturing equipment is more readily available than for CdTe technology, there is a lower barrier to entry for would-be manufacturers. "The customer can get everything from us," says Juerg Steinmann, head of marketing communications at Oerlikon Solar, in Truebbach, Switzerland. Its services are proving to be popular, and customers are currently inquiring about production systems for the manufacture of 40 MW or 60 MW per year. "Within a relatively short time we will have gigawatt factories," says Steinmann.

Oerlikon's tools produce 85-W solar panels covering 1.4 square meters, using a 0.3-μm layer of amorphous silicon that strongly absorbs visible light. However, output can be increased by nearly half with the addition of a 1.5-μm microcrystalline layer of silicon that absorbs infrared radiation too. Late last year Oerlikon introduced modified process equipment for microcrystalline growth. Even with this additional growth step, manufacturing throughput is fast, and the Swiss company contends that a manufacturer wielding its tools could make a solar module about as quickly as First Solar can. Nevertheless, the all-important cost-per-watt ratio is slightly inferior. "Today we're looking at $1.50 per watt, and by 2010 our goal is going to be $0.70 per watt," says Steinmann.

The idea of using more than one material to capture a higher proportion of the sun's radiation has also been pursued by Emcore Corp. This Albuquerque-based company uses layers of germanium, gallium arsenide, and gallium indium phosphide to manufacture cells that are roughly three times as efficient as First Solar's. But the production costs are astronomical because the technique requires relatively slow growth rates and because deposition occurs on small germanium substrates. Even so, these drawbacks did not prevent Emcore from enjoying success in its initial target market, aerospace applications, in which high efficiency and reliability are paramount.

The high cell costs can be offset in terrestrial solar-cell systems that use large lenses or mirrors to focus sunlight by a factor of several hundred, boosting conversion efficiency to almost 40 percent. But this strategy makes sense in only 10 to 20 percent of the world market. "Where we play is in the very sunny, high-solar-resource areas, where it's also very warm," says David Danzilio, the company's vice president in charge of photovoltaics. For that reason, sales of Emcore's product are unlikely to take much of a bite out of sales of First Solar's systems, which can be used in all climates.

All this means is that in the short term, First Solar's main competition will continue to come from conventional silicon cells. This mature technology is unlikely to deliver any major hike in efficiency from today's figure of around 16 percent, but if the silicon shortage disappears in a year or two, lower material costs will propel a major reduction in the cost-per-watt figure. "However, even if polysilicon came down in price significantly, First Solar could cut their prices and still see the same margins that the traditional module makers do," says Hardy of American Technology Research. "It would definitely hurt them on pricing, but they would still be extremely profitable at lower prices."

With no strong challenger in sight, First Solar is well placed to continue its quest for grid parity. Getting there would substantially reduce greenhouse-gas emissions. That achievement would be a great legacy and would make a really great story, but will First Solar be willing to tell it? ❑

TO PROBE FURTHER Solar Cells Inc. describes the details of its elemental vapor process in its 1993 report "Fabrication of Stable Large-Area Thin-Film CdTe Photovoltaic Modules," which is available at http://www.osti.gov/bridge/product.biblio.jsp?osti_id=10181903.

Get information about inverted triple-junction technology in "High-Efficiency GaInP/GaAs/InGaAs Triple-Junction Solar Cells Grown Inverted With a Metamorphic Bottom Junction," Applied Physics Letters, *91 023502, 2007.*

For comparisons between CdTe and CIGS manufacturing, see Michael Powalla and Dieter Bonnet's paper "Thin-Film Solar Cells Based on the Polycrystalline Compound Semiconductors CIS and CdTe" published in Advances in OptoElectronics, *2007, available free of charge at http://www.hindawi.com/GetArticle.aspx?doi=10.1155/2007/97545.*

DENISE GRAY HAD TO TAKE A CITY BUS TO HER FIRST JOB AT GM, BUT NOW SHE HOLDS THE COMPANY'S GREEN-CAR FUTURE IN HER HANDS

By John Voelcker

Photography by Joshua Dalsimer

BATTER

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

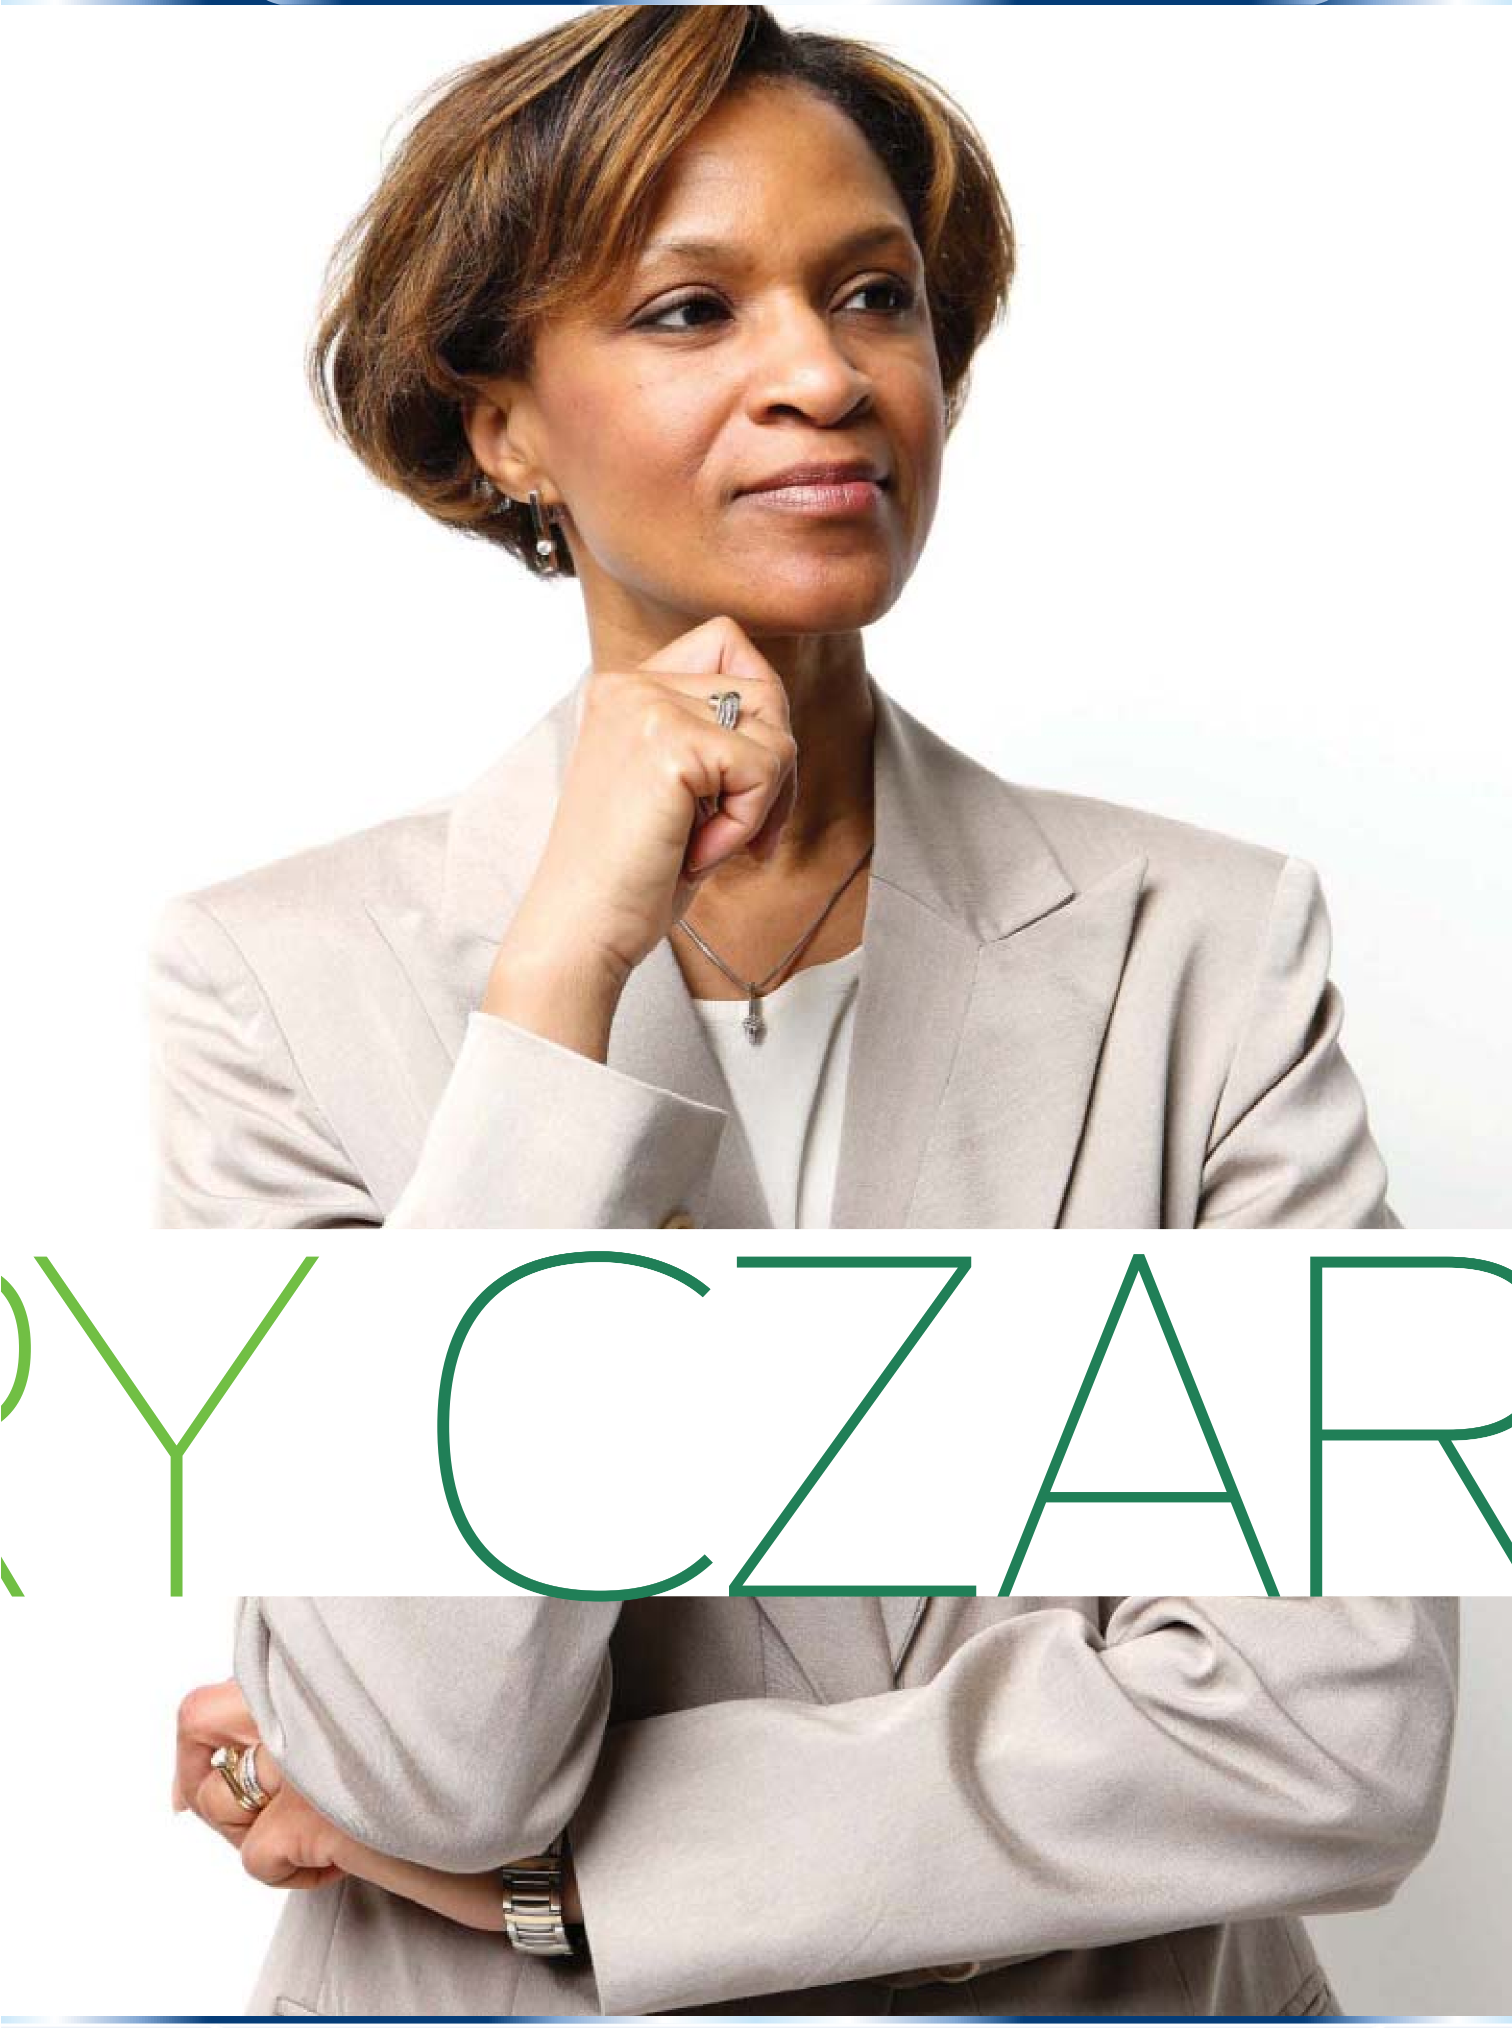
RY CZAR

HOT SPOT: Denise Gray checks her ever-present BlackBerry while leaning against a row of thermal chambers used for testing batteries in different climate conditions.

CARMAKERS really, *really* don't like to take chances. Not since the earliest days of the industry have they tried to develop a new body and chassis *and* a new energy storage and power delivery system, all at the same time and all for the same car. Yet that's exactly what General Motors is doing with its upcoming Chevrolet Volt extended-range electric car. And the stakes couldn't be higher.

Caught off guard a decade ago as rival Toyota launched and then refined its hybrid-electric-drive cars, GM has struggled to catch up. Within the last year, Toyota has equaled GM's global production—*and* announced the manufacture of its millionth Prius hybrid. But with the Volt and various other hybrid vehicles of its own, GM has mapped out a bold strategy that will pit it squarely against Toyota. The arena? Technology.

The executive at the center of this battle is an African-American engineer in her 40s—one of very few black women in the auto industry's upper ranks—who didn't even own a car when she took her first job at GM 28 years ago. Denise Gray, unofficially GM's "battery czar," is the company's director of energy-storage devices. Her job is nothing less than overseeing GM's efforts to develop a new generation of batteries that will give it an edge in electric vehicles.

At the top of her to-do list is testing and approving the battery pack for the much-touted Volt, which GM is working feverishly to release in November 2010. The clock began ticking when the first battery pack was delivered last year on 31 October and won't stop until the first Volt rolls into a dealer's showroom. Until then, Denise Gray will be the executive to watch in the U.S. automotive industry.

IT'S HARD TO OVERSTATE how much GM has riding on the Volt. Decades of downsizing and dwindling North American sales have the company locked in a neck-and-neck race with surging Toyota for the title of the world's largest car company, which GM held for 70 years. With the radical Volt, GM hopes to leapfrog its rival's decade of experience in hybrid electric vehicles.

Technically, the two firms' approaches to electric drives could scarcely be more different. Toyota has built more than 1 million "power-split" hybrids, which use a battery with a storage capacity of 1 to 2 kilowatt-hours to assist the gasoline engine and store energy regenerated while braking. These cars travel only 1 or 2 kilometers on pure electric power.

By contrast, the Volt is an "extended-range electric vehicle," which will take advantage of the new large-format lithium-ion batteries just now entering the market. Its 16-kWh battery pack will give the car a pure-electric range of up to 65 km, or 40 miles, with a small gasoline engine providing another 480 km, or about 300 miles, on a single tank of fuel. That gas engine, however, won't drive the wheels directly; it will power only a generator that recharges the batteries, which drive the electric motor that spins the wheels. The Volt will also have a plug that will let its batteries be recharged from any outlet; that's why this kind of vehicle is also called a plug-in hybrid.

With sales of big, gas-thirsty vehicles in free fall in the United States, GM has called the Volt its most important new car program. And success or failure will hinge on those lithium-ion battery packs; if they don't prove robust enough to last 10 years and 240 000 km (about 150 000 miles), the car will flop.

All Gray has to do is make sure GM picks the right lithium-ion system and then ensure that those battery packs are exhaustively tested. She'll have to navigate government safety-certification requirements in multiple countries and have the packs manufactured in large quantities and to very high quality standards. She's got two years to get it done.

Car talk surrounded Gray during her childhood in Detroit in the 1960s and 1970s. "My entire family worked in the industry," she says. "My aunt assembled

V6 engines at GM's Livonia engine plant; my other aunt and my uncles assembled axles at GM Gear & Axle. Working for auto companies was just a part of our life." Her mother worked at GM too, making stabilizer bars and other parts in a General Motors forge plant—some of the hottest and most grueling work in the industry. Still, even though cars were in her blood, Gray didn't own one when she first started working at GM in 1980; she rode the city bus to her position in Warren, Mich., at the GM Technical Center.

Today, when time is tight, she flies on a GM corporate jet. And back on earth, she gets a brand-new company car every few months. These days, she's driving a sporty Cadillac SRX crossover, after recently trading in a low-slung, V8-powered Corvette. Her smile flickers when she says, quietly but firmly, "I like to merge with *pride*."

Her mother, Vernice, had moved to Detroit from Arkansas in the 1950s, seeking a job in the booming hub of the U.S. auto industry. Denise's father left when she was five; after that Vernice worked full-time to support her children.

Gray attributes much of her own success to Vernice. "My mom is the epitome of determination to succeed," she says. "She probably didn't have a lot of good cards dealt to her, but she's made the best of it. She's an extremely hard worker and was extremely focused on taking care of her family."

The family lived in an old frame house in Detroit, and Gray attended that city's public schools. In middle school, a mathematics teacher impressed by her math and science abilities suggested she apply to Lewis Cass Technical High School, then Detroit's only magnet school. The young Denise hadn't given it much thought, but she followed the teacher's advice and easily passed the entrance exam for admission.

At Cass, she was one of only two girls who focused on electronics. "I began quickly recognizing there's a gender issue here," she says. "It's not something that girls usually go into. But it really did keep my interest."

This being Detroit 30 years ago, autos were a core part of the curriculum. Gray reminisces fondly about learning to weld. "I remember being in the garage, learning how to use a torch," she says. "It was neat." She compiled a 3.86 grade point average and caught the eye of GM coordinators looking for bright students with good grades and self-starter attitudes to work half-day jobs.

"You worked with the mechanics, so you wore a uniform," she recalls. "When I would catch the bus home, I'd have on my uniform. And the kids knew I had gone to Cass, and they were all wondering, What happened to this smart girl? 'Cause now she's wearing a gas station attendant's uniform!"

In 1981, after finishing high school, she went on to attend the General Motors Institute, now Kettering University, in Flint. For five years she alternated 12 weeks of studies with 12 weeks in a GM job, rotating among the divisions. She spent time in the machine shop, the foundry, the prototyping studio, and other operations—nine different assignments altogether.

IN 1986, HER LAST YEAR at General Motors Institute, Gray met Michael Steel, an African-American design engineer a few years older than her, who worked in GM's electric group. Steel would become her friend, boss, and mentor. "There weren't a lot of African-American engineers then," he recalls, "let alone female engineers, so I introduced myself." He ended up advising her on her fifth-year thesis, on instrument clusters.

When she graduated, Steel hired her for the midsize-car division as a lead engineer. She would follow in his wake for more than a decade. For the first six years she designed electrical components for cars' instrument panels, including such innovations as radio controls integrated into the hub of the steering wheel. She designed interior and exterior lighting, air bags, all kinds of audio equipment, and interfaces between the electrical systems of the vehicle and its engine and transmission.

Meanwhile, in 1989, Steel left the instruments group to manage the Corvette team's electronics development lab. Gray followed him to the new team, as she did again in 1992, when Steel moved to GM's power-train group to help develop a new generation of the company's iconic "small block" V8 engine.

Those were the days when software-controlled electronics was supplanting older ignition systems. Ignition systems had traditionally been mechanical: a rotor inside the distributor closed a succession of circuits that caused the engine's coil to deliver enough power to each spark plug to make it produce a spark, igniting the gasoline-air mixture already delivered into the cylinder. But with increasingly stringent emissions limits in the 1980s, only electronics could manage engines precisely enough to control the combustion adequately. The combination of controllers, sensors, and software for those systems had to be tested in a new way.

Steel put Gray in charge of developing and implementing the tests for what became the 1997 Corvette. Her team had to define and test various operating conditions the engine might face. For instance, to produce the right torque, the software algorithm that controlled the firing of the spark plugs had to "know" the position of every cylinder to fire at the optimal time. The same engine was used in three different classes of vehicles: not only the Corvette two-seat sports car but also the four-seat Camaro coupe and a range of trucks.

Once every type of use had been written down, Gray's team developed tests that would expose the system to the extremes of each, then calculated what a satisfactory performance should be for each test. Once every test was run, the results had to be verified (were the right tests used? were the results correct? did the whole test cycle follow our procedures?) and the testing team had to officially sign off—all before the software could be released to the larger engine team.

"SHE WAS NOT SHY about taking verification reports to the platform manager," Steel says, grinning. Translated from the corporate jargon, that means that despite constant pressure to stay on schedule, Gray simply refused to release any software that hadn't met the list of requirements put together by the vehicle teams. "She would offer [the manager] a conscious decision if there was a 'Part Not Validated' report, saying, 'I'm not backing down, I'm going to write the report' " to indicate that the system had failed its tests, says Steel. Given the challenges of integrating new and often costly electronics into the complex mechanical systems of a modern engine, it was the right decision—but hardly a popular one at the time.

Slowly but surely, she built a record of technical and managerial success. And in 1995, when Steel became software manager for all the engine groups, Gray was able to take the procedures she'd developed for the small-block team and extend them throughout the power-train organization. It was her first director-level job. In GM-speak, that means she was fully responsible for all the work of both engineers and their managers toward delivering a product—in this case, engine-control software.

That engineers must rigorously validate software is hardly a revolutionary thought nowadays. Ten years ago, though, a lot of people in auto manufacturing yawned—or sneered—at the kinds of formal procedures that standardized software testing and validation require. In the car business, electronics engineers were traditionally the stepchildren of the power-train group, compared with the mechanical engineers who designed the moving parts of the engine. The EEs' world of computer screens was foreign to the engine guys. After all, you could tell if an engine was running properly by listening to it, right? Reading through reams of printouts to figure out if something worked correctly wasn't a popular activity.

Nonetheless, Gray persevered in her quest to integrate modern software practices into the macho world of engine development. Over and over, says Steel, she took the time to educate other managers about what her team was doing and why it benefited them. But, he says, "she was clear that it was a discussion of *how* they would buy in on the method—not on the fact that she was doing this" in the first place.

Gray's reasoning and persuasiveness did the trick. "She has the tenacity of a Chihuahua," Steel says, grinning. "She would have been a pit bull, but she's so small..."

That inner tenacity has helped Gray overcome any opposition sparked by her sex and race. She says, diplomatically, that she tends not to notice such attitudes. "You probably have to hit me over the head with a two-by-four to tell me you don't respect my presence there because of my being a female" or being black, she says.

Steel is less circumspect. "As a female African-American engineer and manager in a historically white male company," he says, bias "was there, for sure. But she put her nose to the grindstone—and she gained the respect of a *lot* of people." And it worked.

In 1996, employed full-time and caring for two children—she was married in 1987—she decided she needed a master's degree to get to the next level in her career, even though she intended to stay on a technical track. She took four years to earn her master's in engineering science remotely from Rensselaer Polytechnic Institute, in Troy, N.Y., and ended up being inspired to switch from a technical track to a managerial one. She calls the coursework "pivotal" and says, "It truly gave me insight, and I think confidence, that I could really be in—that I *needed* to be in—management."

In 2000, GM promoted Gray to her first executive-level job, director of software engineering, soon followed by director of controller integration and applications. These steps on the corporate ladder took her higher than Steel, her mentor. She was responsible for all the controller systems that the group engineered for any vehicle, both the hardware and the software. These included systems for fuel delivery, ignition timing, valve opening, overall engine control, transmission shifting, and a host of others.

In 2003, Steel became manager of software at the electronics integration and software group, reporting to Gray. "She was nervous about it," Steel recalls. "When she got the promotion," he says, "I told her, 'It's been a great ride. Now it's time to unhook your car from my train.' "

In 2004, Gray moved to a position overseeing all software development for the transmission group, with the same mandate: to weave systematic and consistent software development and testing procedures through the organization. Her engineering teams designed, tested, and calibrated all the software for five of GM's range of four-, five-, and six-speed automatic transmissions.

In 2006, she began to think about her next assignment. She knew she wanted to stay in a position where she was responsible for "getting product out the door," but she wasn't sure where her next focus should be. And then the pieces came together. She heard that the company had created a new director position to oversee advanced battery work. It was perfect.

"I was an electrical engineer who loved chemistry in college," she explains. "I had the vehicle systems background, power-train background; I'd worked on leading-edge software algorithms, meeting fuel economy and emissions requirements, all of that.

"I wanted it, hands down," she says. And on 12 October 2006, she got it.

IT'S GOOD TO BE BATTERY CZAR. The job is highly visible—as part of the core Volt team, Gray meets often with GM CEO Rick Wagoner, product guru Bob Lutz, and other top GM executives. And there's no doubt that the job is vital to the company's future.

That kind of responsibility can feel overwhelming at times, she says, especially considering the task she's been given: to pack 10 years' worth of user testing into a couple of years.

"I recognize to the *n*th degree the criticality of making this happen," Gray says, as only an engineer could. "I understand what it means to the company and, quite frankly, to the industry, to our country. I understand its significance from an energy-security perspective." And she calls the company's support for the project nothing less than "awesome": "They've said, 'Whatever you need, Denise.' And it's been provided."

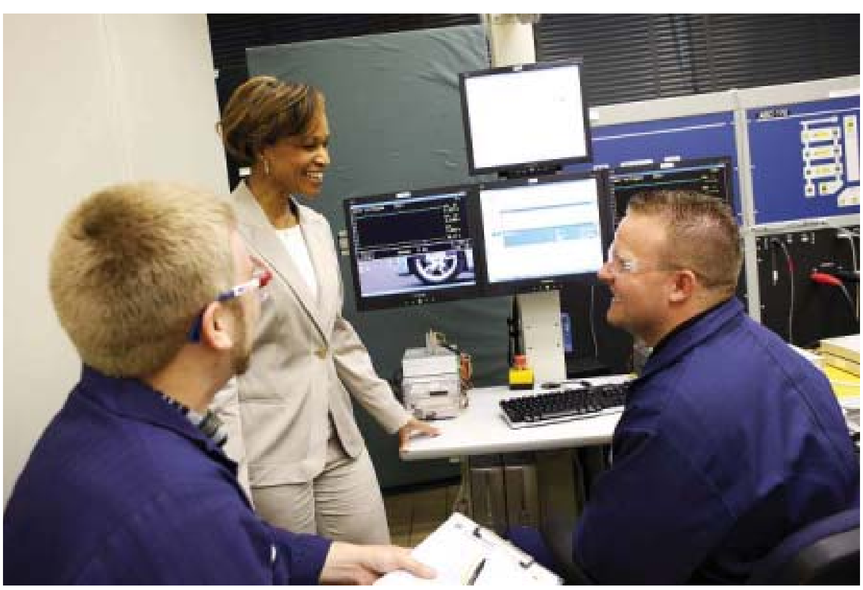

PROGRESS CHECK: Denise Gray meets with battery engineering specialist Alan Martin [left] and test engineer Scott Herz at GM's Vehicle Engineering Center in Warren, Mich.

She is always aware of the pressure on this highest of high-profile projects. But her preferred approach remains collaborative and calm. "I can be in command and be the quietest person in the room," she says. Her style is to support people working together to solve problems on their own. "We'll be better off if everyone learns how to get from point A to point B, as opposed to being *told* how to get from point A to point B."

But that's not her only mode. She recalled that recently a fellow executive criticized one of her team members for doing something he wasn't allowed to. "I closed my book," she says, "and said, 'I gave that person authority to do it. You weren't in the meeting; you have no authority to speak on this particular subject,' and I grabbed my coat with lots of emphasis and stormed out of the room. I usually don't do that, but the person had done that kind of thing a couple of times, so I had to come at it from a different angle to get the point across.

"The point," she says—quietly—"was well-taken."

GRAY REMAINS as diplomatic as she is driven. She assiduously shares information among vendors wherever possible and works hard to make sure that GM's internal procedures don't gum up the works, says David Vieau, CEO of A123Systems, one of two teams vying to provide advanced lithium-ion packs for the Volt battery pack. "She's down to earth," he says, "and direct—but reasonably so."

Prabhakar Patil, CEO of Compact Power—part of the competing team—says Gray has insisted that any unexpected hitches or shortfalls be discussed openly, without finger-pointing or corporate self-protection. He recalls the challenges that the company faced in designing the internal cooling for the Volt pack. Other vendors might have minimized the problems to avoid scaring their client. But, emboldened by Gray's edict, Compact Power, a U.S. subsidiary of South Korea's LG Chem, was quick to detail its difficulties. Gray responded immediately by making several of GM's experienced cooling-system engineers available to Compact Power. With their help, the company solved the problems and kept on schedule.

On a typical day, Gray wakes before dawn, dispatches a night's worth of e-mail on her BlackBerry, and drives her sons Taylor, 17, and Nathan, 13, to school. She gets to the office at 7:30 a.m. and starts making phone calls to Europe. She tries to be back at home by 7:30 p.m. for family dinner. After that, more e-mail.

It's a grueling schedule, but not so grueling as it was 18 months ago, when Gray had to spend at least one week out of four visiting lithium-ion cell makers in South Korea and Japan, touring their plants, explaining GM's needs and plans, and educating herself. Now the competition to build the Volt battery pack is down to just two teams—Germany's Continental, using cells from A123Systems, versus Compact Power, using cells from its parent company, LG Chem. So Gray spends more time these days with GM teams throughout Michigan or in Mainz-Kastel, Germany.

NEXT GENERATION: The battery pack that will be used for the Chevrolet Volt [right] poses next to its older sibling, the original EV1 battery pack. PHOTO: JOHN F. MARTIN/GENERAL MOTORS

FAMILY FIRST, church second, work third." Not every executive would relegate a top job at one of the world's biggest companies to third place. But Gray is comfortable with her choices and works hard to balance the three pillars of her life. Rob Peterson, head of communications for the Volt and E-Flex programs, recalls reminding her of a presentation on the first media day of the Detroit auto show. "I'll be there," she said, "but it'll have to wait until after church."

Similarly, she works hard to attend all of her son Taylor's high-school football games. She may be using her BlackBerry in the stands, but she's there—and she's probably shouting louder than anybody else.

The Gray family lives in a suburb half an hour west of Warren, Mich., where the GM Technical Center is located. Husband Kevin, whom she met while she was a General Motors Institute student in the early 1980s, works for GM as well, as a director overseeing cooling-system development. Dinner with the family is a must. Her sons do their homework at the kitchen table—and with Mom and Dad plugging away, sometimes there are four laptops going at the same time.

WHAT'S NEXT? "I told my boss, hopefully there are no plans to move me, because we're just getting started," Gray says. "We've got to keep thinking about what's next, and then what's next"—for future batteries, for her company's role in the industry, and for increased research into advanced energy-storage systems.

GM's surprising openness about the progress of the Volt's development may be a good indicator of her team's success thus far. Global automakers tend to guard technical details closely. The Volt team, by contrast, invited 80 reporters to its labs this spring and detailed the test algorithms that let them simulate 10 years of usage in only two calendar years.

And while Gray finds herself increasingly in the spotlight, it doesn't interest her much. "In my mind, it doesn't matter," she says. "I haven't earned it. I'm flattered, but it doesn't really mean a lot until after I've delivered"—when that first Volt rolls off the line, probably amid more hoopla that any new GM car in 20 or 30 years.

Occasionally she'll let herself marvel at her life and her job, with its improbable arcing ascent: "When what you're doing is that important to the leaders of your company, it makes you feel awfully good as an engineer."

And it is important; in fact, it's vital. "The advanced battery team is the group that the entire Volt project hinges upon," says Aaron Bragman, of industry analyst Global Insight, in Troy, Mich. If GM is to survive in a market that's clamoring for high-tech mileage champs, Bragman says, "then beating Toyota to market with a plug-in hybrid is absolutely critical."

No one is more aware of that than Denise Gray. ❑

RAM FOR FREE

NOVEL DATA COMPRESSION DOUBLES THE MEMORY IN EMBEDDED SYSTEMS WHILE HARDLY SLOWING THEM DOWN

BY LEI YANG, ROBERT P. DICK, HARIS LEKATSAS & SRIMAT CHAKRADHAR

"GIVE US A READING on the 1202 program alarm," radioed Neil Armstrong to Mission Control in July 1969, seemingly about to lose his famously cool demeanor. He was busy trying to steer his spacecraft to the first-ever manned landing on the moon and was worried that this error message from his guidance computer meant serious trouble.

Fortunately, a young computer engineer at Mission Control had the insight to realize that this error was not as ominous as it seemed, and on his signal the Apollo 11 landing went forward. Within hours of the astronauts' safe touchdown, it became clear what had happened: the lunar module's rendezvous radar had remained switched on during descent, when only the landing radar was needed, and the craft's navigation computer had become overtaxed trying to process radar data from the two sources. The system's programmers had, however, built in a fail-safe mechanism that would shed the less critical tasks so that the computer could do its main job.

The Apollo Guidance Computer was arguably the first example of an embedded system, one that incorporates a special-purpose computer dedicated to a single function. Embedded systems have become the predominant form of computing, exemplified by the microcomputer-operated brains inside your microwave oven, your MP3 player, your cellphone, even your refrigerator, to name just a few. Although many such devices now have 10 000 times as much random-access memory (RAM) as the mere 4 kilobytes of their Apollo-era great-granddaddy, memory constraints continue to dog their designers.

Cost is one reason: although the prices of RAM have plummeted fast, the need for memory has expanded faster still. Another concern is all the energy that RAM requires. Manufacturers of mass-market products, especially portable devices like music players and mobile phones, must therefore take care to add no more memory than the software needs to operate. Hewing to that line is no easy trick. To speed a new product

EMILY COOPER

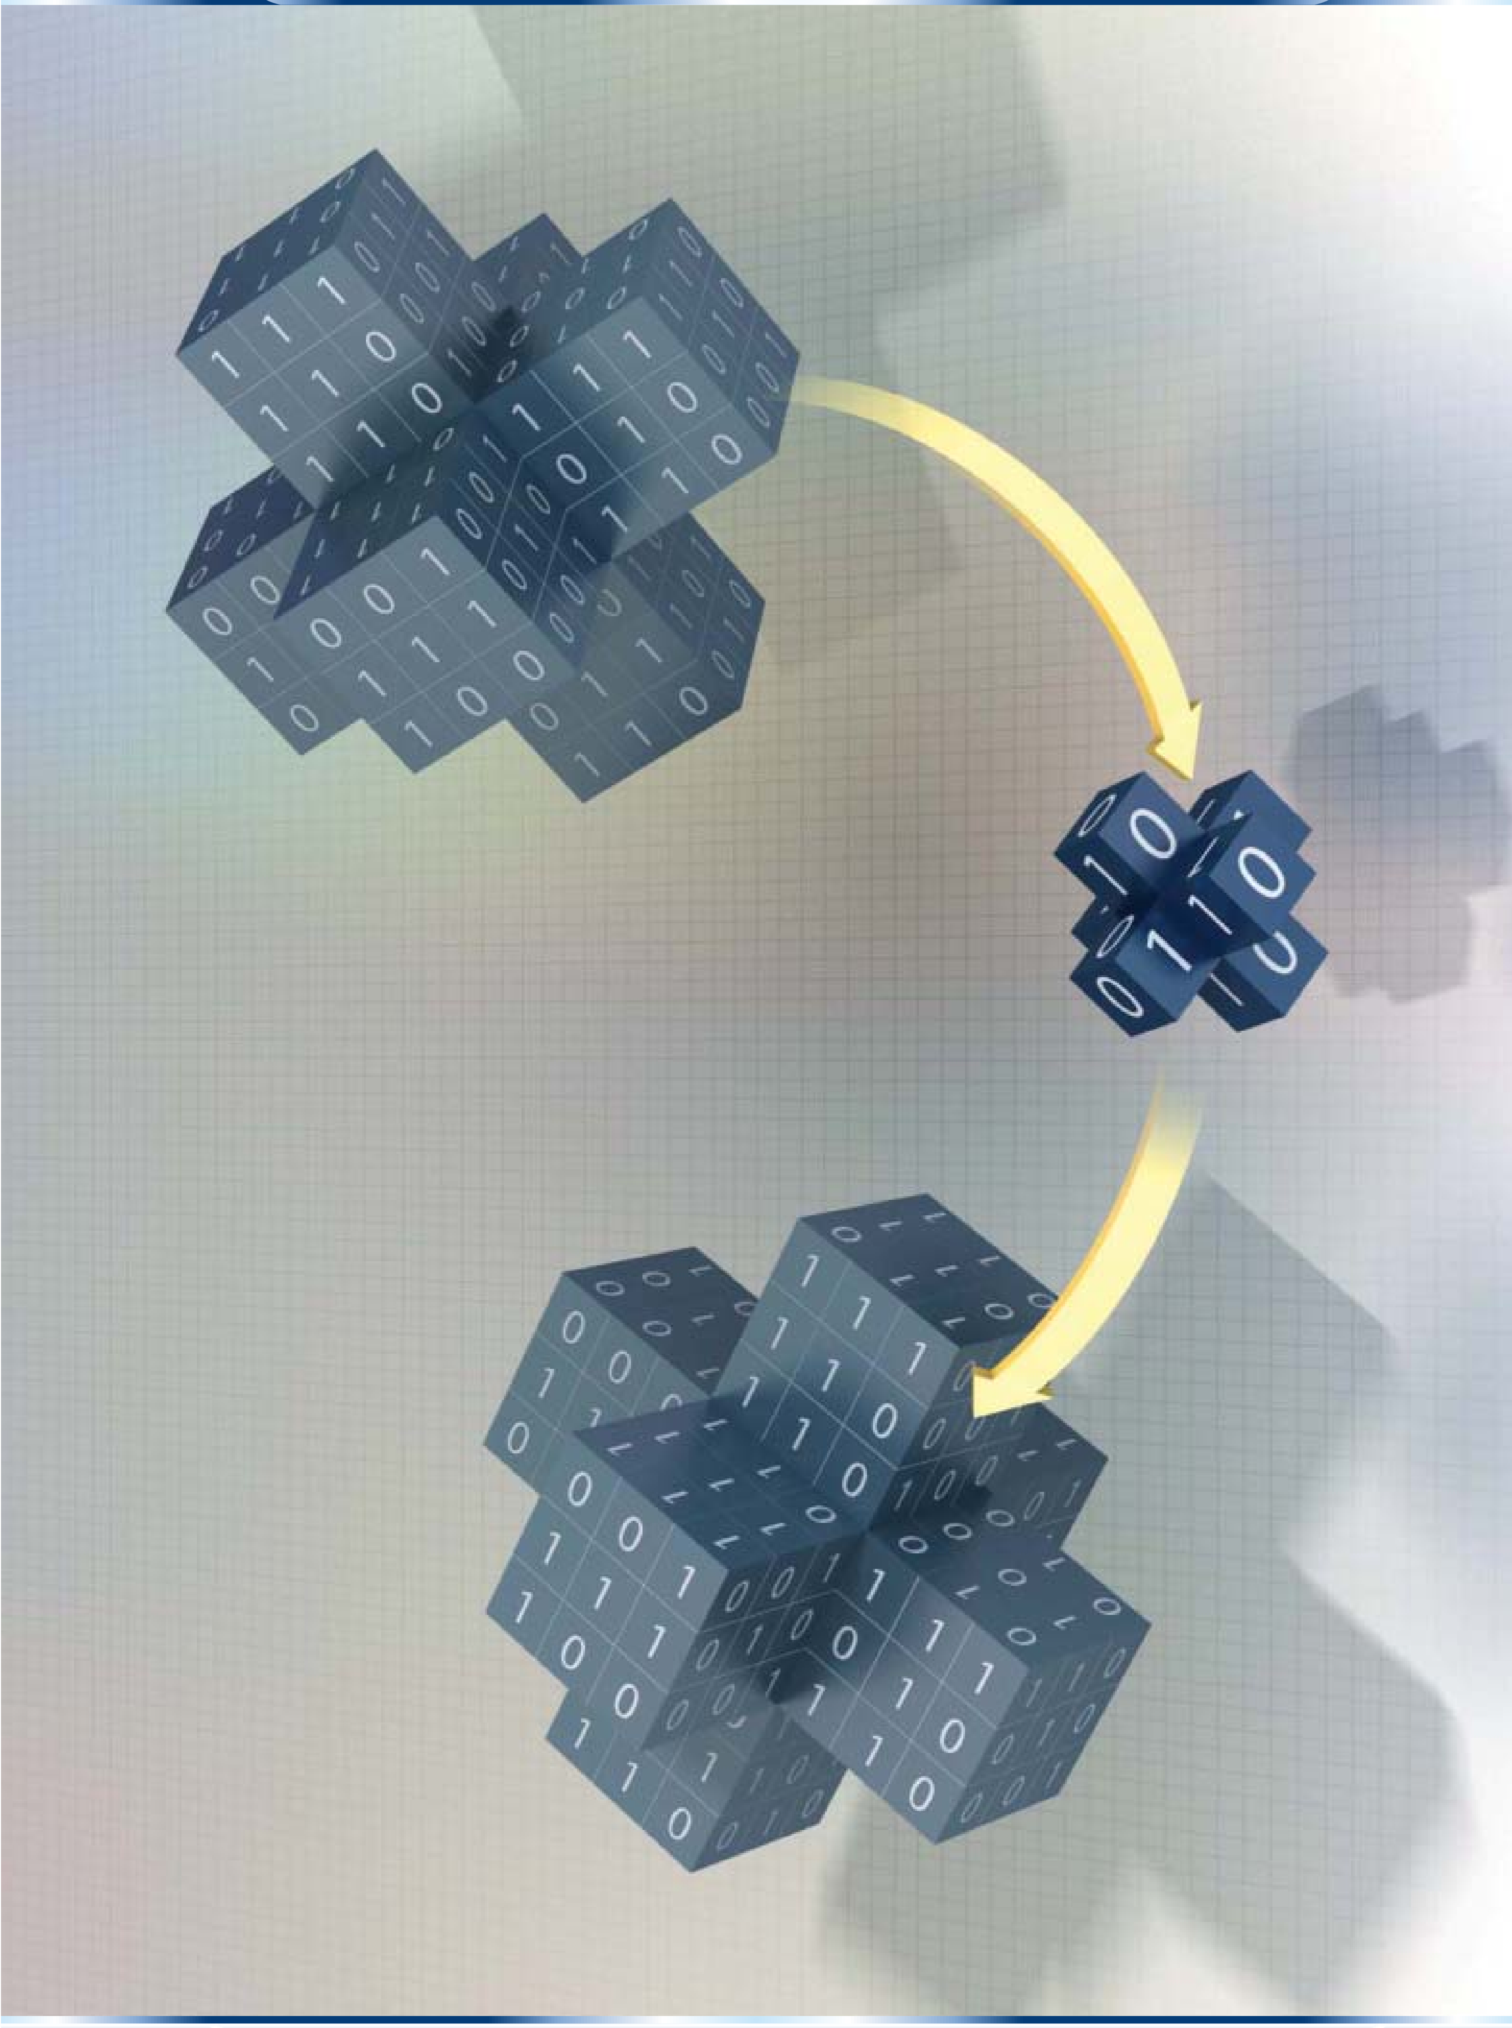

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

to market, hardware and software engineers have to work along parallel tracks, which means that neither side can know quite what the other has in store.

If software designers yield to temptation and ask for more memory than they could possibly need, they risk wasting a lot of money—even pennies matter when you're producing millions of units. Or the product could end up being too power hungry. Yet if they skimp on RAM, they may prevent the unit from running some new killer app that would allow the gadget to beat the competition. Such mistakes sometimes force companies to redesign their hardware, a process that is enormously costly and time-consuming.

We have spent the better part of three years trying to give designers of embedded systems a third option: to increase effective memory by compressing the data stored in RAM using just software.

DATA COMPRESSION is standard fare in other parts of the computer business, of course. DiskDoubler was a top-selling software utility during the early days of the Apple Macintosh, for example. By encoding regular patterns in the data in more compact form, it gave users the feeling they'd doubled the space on their disks—something well appreciated at a time when hard drives held only a few tens of megabytes.

Unlike the "lossy" compression employed, say, for encoding JPEG images, the compression used to store programs and associated data has to be "lossless": it can't drop a single bit. Conventional wisdom held that additional special-purpose chips were required for this (or the functions of those chips had to be designed into the processor, at great expense). Otherwise, the thinking went, operations would be too slow for memory compression, and in portable systems the energy usage would become prohibitive.

Speed and power are indeed often critical for embedded computers, as is consistency of response time. So it's no wonder that their designers have shied away from data compression and just added RAM as needed. But it occurred to us that if you could use data-compression software to control the way embedded systems store information in RAM, and do it in a way that didn't sap performance appreciably, the payoff would be enormous.

We began our investigations by taking a long, hard look at existing compression techniques. The most promising appeared to be something called Lempel-Ziv-Oberhumer, or LZO, one of the family of widely used Lempel-Ziv algorithms invented in the late 1970s by the Israeli computer scientists Abraham Lempel and Jacob Ziv. In the mid-1990s, an Austrian programmer named Markus Oberhumer wrote the LZO variant in ANSI C, designing it expressly for speed.

Once we had selected the compression algorithm, we needed to devise ways to determine which of the data in memory should be compressed and when. We also had to come up with the means to expand the compressed data when they were needed. While this might seem nightmarishly difficult, in fact, getting LZO compression to work reliably in our test-bed system—a PDA—proved straightforward. We were able to make use of the virtual memory feature of the PDA's Linux operating system, which swaps infrequently used data between RAM and disk storage. We modified this mechanism to compress such data when memory requirements exceed physical RAM capacity. But instead of writing the compressed data to disk, the system stores it on a virtual device—a portion of the available RAM.

When an application requires the previously compressed information, our software locates and decompresses it, then moves it back to the uncompressed part of the PDA's memory so that the application can continue to work normally. What's more, our scheme allots space for the compressed data as needed, so that applications that work fine without extra RAM aren't slowed down at all.

Programmed in this way, our PDA could use far less memory to run various applications—games, office tools, even media apps, whose compressed sound and image data are normally expanded when they are in RAM. Attempts to operate these same applications using reduced physical RAM but without data compression generally crashed the system—or made it grind along at an intolerably slow pace. To ensure that the compar-

CONVENTIONAL WISDOM HELD THAT SUCH DATA COMPRESSION REQUIRED SPECIALIZED HARDWARE

COMPRESSED-RAM STATISTICS

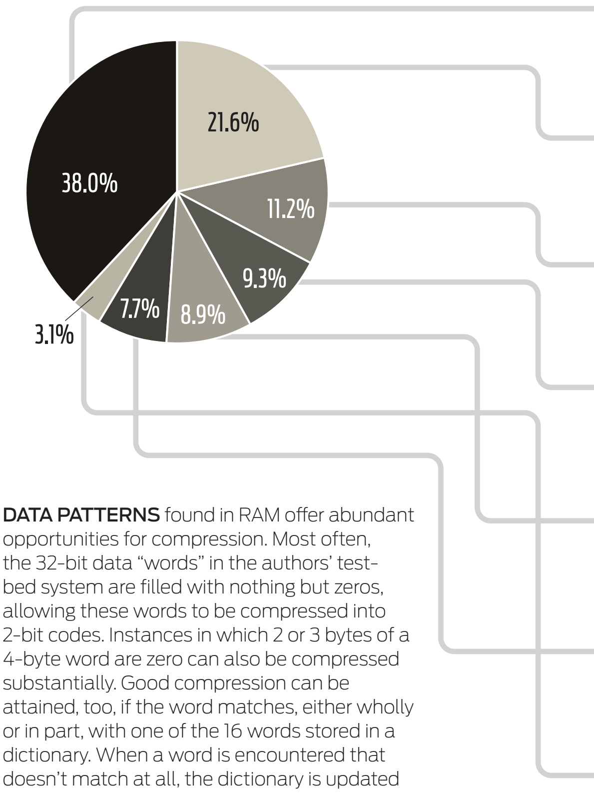

DATA PATTERNS found in RAM offer abundant opportunities for compression. Most often, the 32-bit data "words" in the authors' test-bed system are filled with nothing but zeros, allowing these words to be compressed into 2-bit codes. Instances in which 2 or 3 bytes of a 4-byte word are zero can also be compressed substantially. Good compression can be attained, too, if the word matches, either wholly or in part, with one of the 16 words stored in a dictionary. When a word is encountered that doesn't match at all, the dictionary is updated to include it.

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

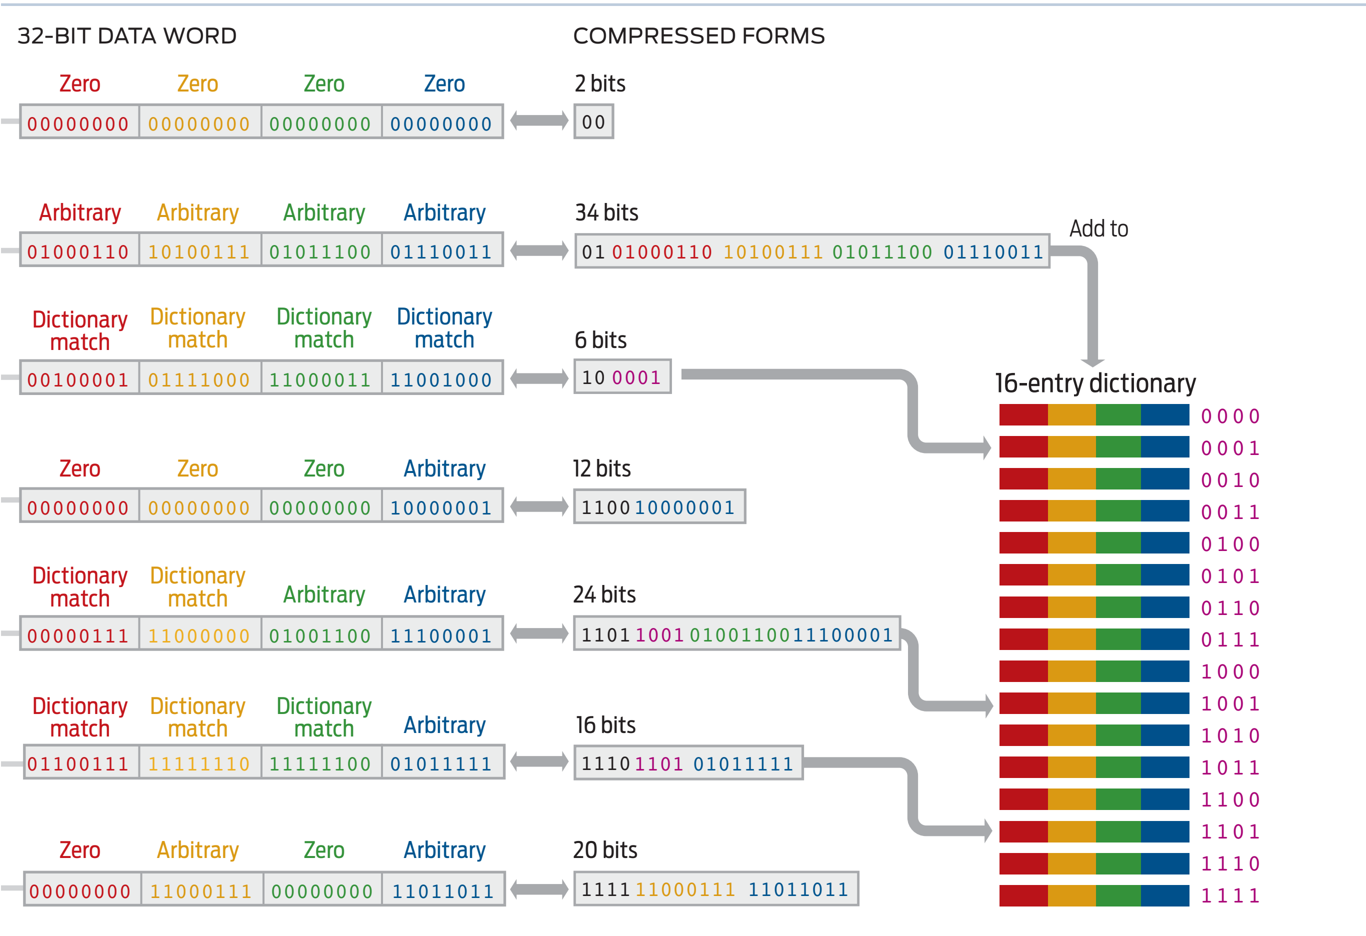

ison was valid, we wrote special software to monitor and then replay user input with identical timing properties, with and without data compression. And to fully quantify how much the compression degraded performance, we ran some benchmark tests that didn't require any user interaction at all. Execution time went up by almost 10 percent on average and by nearly 30 percent in the worst case. Recognizing that some situations might not tolerate such lethargy, we sought to do better.

THE CHALLENGE, then, was to come up with an algorithm that would rival LZO's degree of compression but be much quicker. We managed to satisfy both requirements by exploiting regularities in the kinds of data that are typically stored in RAM, a scheme we dubbed pattern-based partial match. It resembles many other compression techniques in that it replaces frequently used patterns with short codes and rarely used patterns with longer codes. The basic strategy is not unlike what Samuel Morse adopted when he translated the alphabet into a series of dots and dashes—which explains why the Morse code for the commonly used letter *E* is just a single dot, while the rarer *J* requires a dot and three dashes.

Pattern-based partial match takes advantage of the fact that much of the RAM in embedded systems is wasted. For example, in a system with 32-bit (4-byte) data "words," numerical values often demand just 4, 8, or perhaps 16 of these bits. The rest of the bits are zeros. An integer variable, for example, is normally just 16 bits wide. And when a particular integer is a small number—say, the hour or minute value you set on your cellphone alarm—most of those 16 bits, too, will be zeros. And even when all 32 bits are really required (to index an arbitrary spot in memory, for instance), nearby words often hold similar values, because they point to adjacent memory locations. So such a word can be encoded compactly merely by keeping track of how much it differs from a neighbor.

The extent to which you can compress data depends on the patterns in the input. For example, a 32-bit word whose 3 most significant bytes are zeros is packed into just 12 bits in our scheme. What's more, our system maintains a dictionary of frequently used data words. If a word of the input exactly matches something currently stored in the dictionary, those 32 bits get squeezed down into just 6 bits.

The dictionary is constructed on the fly, so it can adjust if the statistical properties of the data change—say, when a pattern that was rare becomes common. You might think this would make decompression impossible, but in fact, this neat trick really works. The key is that the dictionary is not saved for later lookups. Rather, when it comes time to decompress things, the compressed data themselves are used to re-create the dictionary—and to update it in a manner that keeps it matched with what was used to compress the data in the first place.

DICTIONARY ENTRIES change on the fly. When presented with a data pattern (shown here as a letter, which may represent 32 or even 64 bits, for example), the compressor checks whether the pattern can be found in the dictionary, which in this illustration has just two entries. If not already present, the pattern is put in. If it is there already, the system outputs just the 1-bit dictionary index, which provides substantial compression.

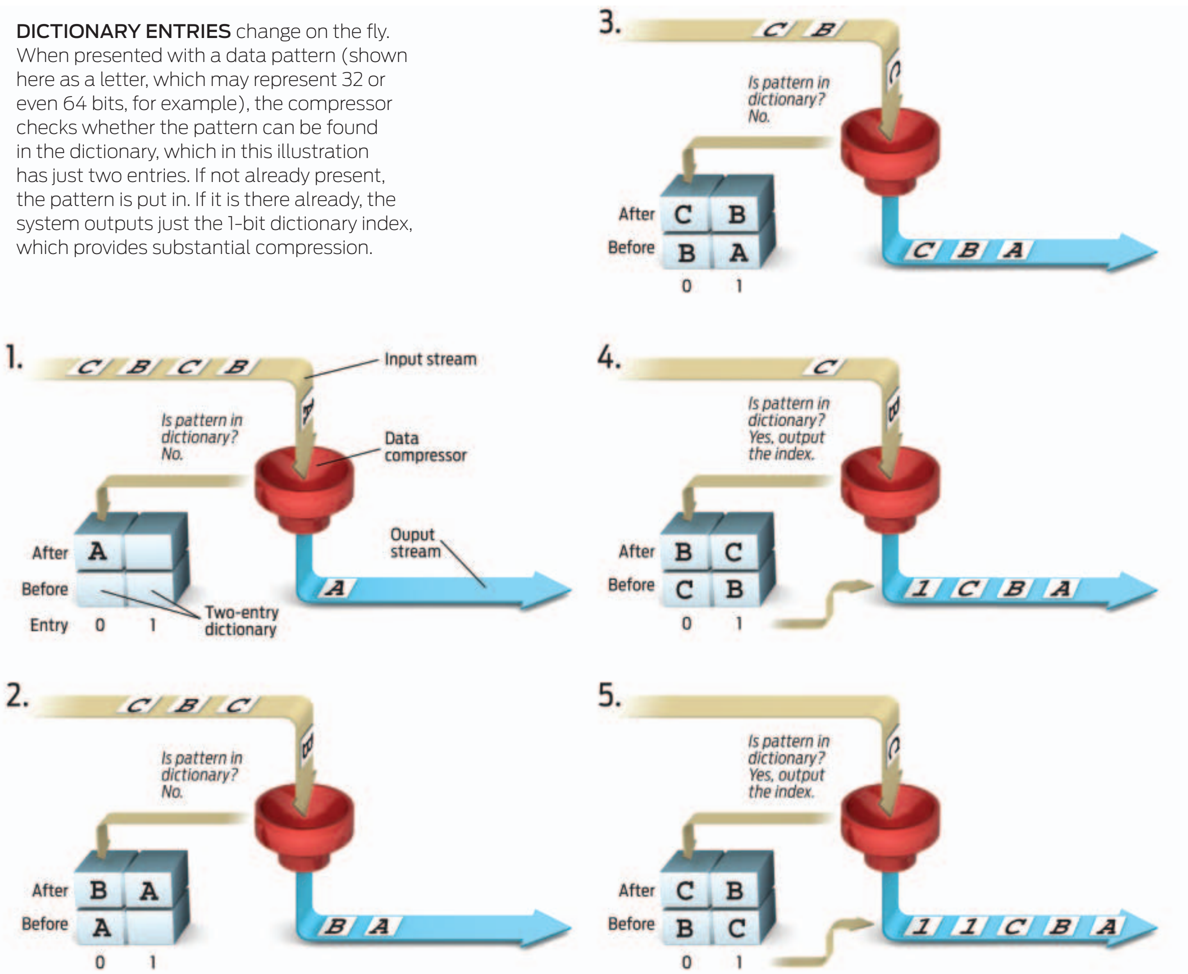

Jon Louis Bentley of Bell Labs, along with three colleagues, invented this technique for data compression in the mid-1980s, an approach that has since become known as move-to-front coding, because newly encountered patterns are placed at the front of the dictionary. Our version allows for 16 possible matches (or partial matches) to the contents of eight separate two-entry dictionaries, a somewhat odd arrangement that allows the lookups to be made at blazing speed.

Our investigations of the data held in RAM showed plenty of opportunities to squeeze things down. The algorithm we devised for this can't match LZO or some other well-known techniques for general-purpose data compression—say, for compacting the contents of an image file—because most files don't have the same tendencies as the data typically found in an embedded system's RAM. A 32-bit word used to encode the color of a particular pixel, for example, is not likely to be chock-full of zeros. But for compressing RAM, our system excels. It reduces the space needed by about 60 percent, and it's startlingly fast. Indeed, our testing revealed that ripping out this much memory results in a performance penalty of just 0.2 percent on average and 9.2 percent in the worst case. That is, this software effectively gives an embedded system more than twice the memory it had originally—essentially for free.

Translating these gains from the lab bench to the marketplace has not been a trivial undertaking, however. In January 2007 we filed for patents on the process, which we dubbed CRAMES, for Compressed RAM for Embedded Systems. We were keen to use it to address a real-world problem: industry's seemingly continual need to redesign embedded systems like mobile phones so that they can run ever more complex applications. (Who would have thought that high-school students would be using their phones to study for the SAT?)

NEC, which together with the U.S. National Science Foundation sponsored our project, wanted to reuse existing hardware for its next-generation cellphone applications, and our compression scheme allowed it to do just that: the company's Foma N904i smart phone, released in September 2007, uses CRAMES. But plenty of practical hurdles sprang up along the way.

For example, in real life an application sometimes terminates when its data are still in the compressed part of memory. Because we initially designed our software not to know the "owner" of any particular piece of compressed data, it was not possible to find out which parts of memory belonged to defunct software processes. So over time, useless data would pile up.

We licked this problem by keeping track of the compressed data and the applications that created them so that the system could free up orphaned chunks of memory. This change required slight modifications to the operating system, however. The production ver-

EMILY COOPER

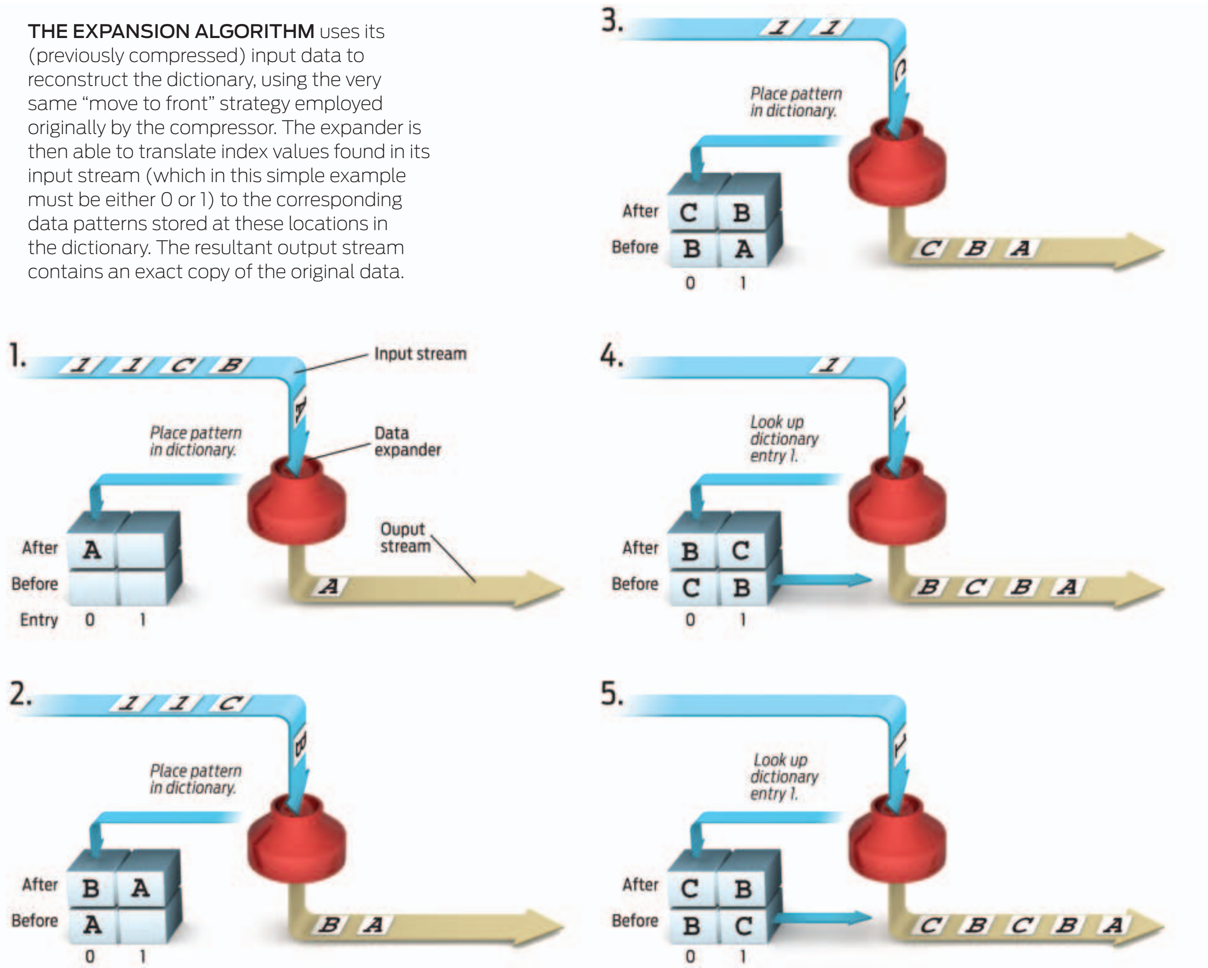

THE EXPANSION ALGORITHM uses its (previously compressed) input data to reconstruct the dictionary, using the very same "move to front" strategy employed originally by the compressor. The expander is then able to translate index values found in its input stream (which in this simple example must be either 0 or 1) to the corresponding data patterns stored at these locations in the dictionary. The resultant output stream contains an exact copy of the original data.

sion of CRAMES remains highly modular, but it is not quite as elegant as the stand-alone operating-system add-on we first envisioned it to be.

Initially, we tested CRAMES on a cellphone prototyping board that ran applications that didn't require access to the airwaves. We measured performance with and without CRAMES switched on, while either cutting back on the amount of memory available or starting more applications than is normally possible.

Later, NEC engineers in Japan gave CRAMES the acid test by running it on hardware connected to a telecom network. Because there was very little time before the first product was to ship, they couldn't test our system of pattern-based partial match completely. So they chose to install a version of CRAMES that uses older, slower, but more thoroughly proven Lempel-Ziv compression code.

WE WERE THINKING only of embedded systems when we first engineered CRAMES. But since then we've also been working on some exciting possibilities involving general-purpose computers (which might allow us to compress our acronym for this data compression, fittingly enough, to CRAM).

Consider the recent trend of putting more than one processor core on a chip. Doing so boosts computing power, sure, but it sacrifices the space on the chip that's available for cache memory, which can hamstring an application whenever it requires frequent access to main memory. We're now designing pattern-matching hardware to compress the data in cache memory to speed up applications.

Another exciting prospect on the horizon involves computers at the other end of the spectrum: tiny ones found in the least-expensive, lowest-cost systems, like those in kitchen appliances. Some of these systems have only 4 KB of RAM—matching what Armstrong and Aldrin had at their disposal in 1969. We've developed software that manages the memory of such tiny computer systems and compresses data when space gets tight.

Although we don't envision our memory-compression technique being used for anything as dramatic as a lunar landing, it's ready for use in equipment that's every bit as critical to safety: sensor networks that monitor the structural integrity of buildings or bridges. These applications will require a great deal of testing, of course, because the software must be able to handle unexpected situations without failing catastrophically. People accept that their desktop computers will crash every once in a while, but such embedded systems must be robust enough not to cause real crashes—as the designers of the very first one, the Apollo Guidance Computer, knew well. ❑

TO PROBE FURTHER For a listing of the authors' related articles, see http://robertdick.org/tools.html.

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

FOUR BRAZILIAN GEEKS DREAM OF CREATING THE NEXT BIG HIT IN MASSIVE ONLINE GAMES—WITH CREATIVITY, A BIT OF LUCK, AND AN IBM MAINFRAME

TEXT & PHOTOS BY ERICO GUIZZO

THE GAME-FRAME GUILD

"THIS ARROWY THING IS A BARRACUDA SPACESHIP, and these little spheres are asteroids," says Tarqüínio "Tarq" Teles, pointing to a jumble of wire-frame outlines floating on a computer screen. "This is how our server sees the world. It keeps track of planets, ships, stations, everything. Shut it down and you shut down the universe."

This is the virtual universe of *Taikodom*, a science-fiction online game that Hoplon Infotainment, a small company in Brazil, plans to launch next month. Teles, Hoplon's hyperenergetic 37-year-old CEO, says that *Taikodom* will allow tens of thousands of people to play together in a sprawling virtual galaxy. The game is Brazil's first incursion into the rapidly expanding market of massively multiplayer online games, or MMOGs. It's a genre made famous by powerful franchises such as Sony Online Entertainment's *EverQuest* and Blizzard Entertainment's *World of Warcraft*. But Teles and his team are not intimidated.

"Our secret weapon is fun," he says, "and we Brazilians know a thing or two about fun, right?"

A visit to Hoplon's office helps confirm Teles's claim. The start-up is based in a lush island city named Florianópolis that's just off Brazil's southern coast and only an hour's flight from São Paulo, the country's business center. With pristine beaches and a bustling town, Floripa—as the locals call it—is known for its laid-back surfing scene and vibrant nightlife. If you're lucky, you may run into supermodel Gisele Bündchen in one of the trendy local clubs.

On this scorching November afternoon, sipping a *caipirinha* on the popular Brava Beach is an enticing idea. But over at Hoplon, nestled atop a wooded hill on the island's north bay, the window shades are down, and the only tan the T-shirt-clad staffers are getting is from their computer screens. They hunker down in cubicles cluttered with *Star Wars* bric-a-brac and Frank Herbert novels, churning out Java code and

BRAZIL'S MASSIVE ATTACK: Tarqüínio "Tarq" Teles [top right], CEO of Hoplon Infotainment, a Brazilian start-up, is leading a team of 72 artists and programmers in developing *Taikodom*, a massively multiplayer online game, or MMOG, that will allow thousands of people to pilot spaceships and engage in large-scale battles in a sprawling virtual universe.

ALL GAME IMAGES: HOPLON INFOTAINMENT

three-dimensional models in hopes of meeting their tight launch schedule.

Just four years ago, when Teles and three friends—Tiago Luz Pinto, Carlos Eduardo Knippschild, and Cristóvão "Cristo" Buzzarello—started developing *Taikodom*, they worked out of Teles's one-bedroom apartment. Now, after recruiting 72 employees and raising over US $10 million in venture capital funding, they're hoping to capture a fraction of the booming MMOG market, which reached $3 billion last year and should double by 2012, according to game industry research firm DFC Intelligence, in San Diego.

But if there's one factor that really sets Hoplon apart from other independent game developers, it's the unusual game server the company adopted. Clusters of PCs are the workhorses behind most MMOGs. Hoplon, however, plans to run *Taikodom* on an IBM System z machine. For the uninitiated, that's a mainframe. The Brazilians have partnered with Big Blue to develop a server optimized for large-scale multiplayer games by fitting a System z with Cell processors, the nine-core parallel-processing chips that power Sony's PlayStation 3. IBM calls this bit-crunching beast the gameframe.

The two companies claim that their machine can do a better job than conventional servers at handling the two most demanding computing tasks in an MMOG: transactions and simulations. In *Taikodom*, the first job involves things like keeping track of each user's spaceships, weapons, and virtual money. The second involves things like calculating the trajectory of objects and checking for collisions. Today's typical game clusters handle both types of tasks, and they can host a few thousand users at a time. Add more and the servers slow down, and users experience lags in the game.

Games with lots of players like *World of Warcraft* have gotten around this problem by splitting the work among multiple clusters, creating duplicate worlds that don't communicate. That means one subscriber can't rendezvous with another unless they arrange to play on the same group of servers. IBM and Hoplon think they have a better idea: the gameframe's hybrid hardware can divide the workload more effectively and hold all the users in a single universe.

"For Hoplon, it's like the perfect machine," says Joe Clabby of IT consultancy Clabby Analytics, in Yarmouth, Maine. "Mainframes are the best transaction engines, and when combined with the Cell to handle math-intensive tasks, the result is probably the sleekest, densest, most scalable architecture they could ever get."

Hoplon and IBM are also creating a middleware layer called bitverse, a software platform that makes the underlying hardware transparent for programmers. For IBM, the gameframe and bitverse are products it can offer to customers interested not only in MMOGs but also in scientific simulations, virtual worlds, and a future 3-D Internet. It's a market that has caught the attention of other big companies like Microsoft, Sony, and Sun Microsystems, as well as start-ups like BigWorld and Forterra.

For the four Brazilian gamers, the partnership with IBM is a dream deal: they get unusual access to high-end technology, and they save money on equipment because they rent all their mainframe resources from IBM and can request upgrades as their subscriber base grows. Their ultimate ambition, of course, is to see *Taikodom*, which is currently undergoing a public beta test, competing with the big titles—a goal they know will require much more than heavy-duty hardware.

IT'S FRIDAY, 1:30 P.M., and a dozen Hoplon staffers cram inside a tiny conference room for the daily 15-minute status meeting. There are no chairs or even a table. Everyone has to stand. A board filled with yellow Post-its hangs on one of the walls.

"The Post-its are our task-management system," says Knippschild, a collected, detail-oriented 29-year-old in shorts and sneakers. "It's based on a method called Scrum, used for agile software development." The term *scrum* comes from rugby, he explains, and refers to a play in which several team members try as a group to gain control of the ball. The idea is to break from project management models that function as relay races, in which members work separately and in sequence.

Knippschild is in charge of keeping *Taikodom*'s development on track. Call him the Scrum Master. In the meeting, representatives from Hoplon's art, software, and other departments report the tasks they have finished and those they plan to do next. They control the flow of assignments—fixing a bug with the game's weapon interface, creating a new space station—using the colored notes, so that everyone can keep track of things at a glance.

After the meeting, Knippschild shows me how some of the groups work. In the game-design department, where a team of nine conceives the mechanics and visual elements of *Taikodom*, two long-haired guys dressed all in black are locked in a heated argument that has something to do with a "portable cyclotron." At the other end of the office, in a room decorated with drawings of insectlike spaceships and bright movie posters, art director Samara Sena is busy generating, as she puts it, "textures and explosions."

But what catches my attention is a quieter bunch nearby, their headphones on, lines of code flashing on their screens. The software team, Hoplon's largest, is led by Pinto, the company's chief technology officer. The bespectacled 31-year-old wrote his first Basic program at age 9, discovered Smalltalk as a teenager, and became an engineering student before dropping out to start Hoplon.

Pinto says he can't show me the gameframe; it resides at an IBM data center in Hortolândia, outside São Paulo. But he offers to show me Hoplon's server farm, where his team tests the software before uploading it to the remote site. The servers are in a wooden closet with sliding doors, packed with PCs of different sizes and colors. "We call it our server barn," he says.

Like so many other business deals, the partnership with IBM began at a bar. Back in 2003, Teles and Buzzarello attended an IBM briefing on grid computing in São Paulo. Later that day, waiting at the airport for a flight back to Florianópolis, they met the IBM engineer, Marlon Machado, who had just given the briefing. With time to kill, the three sat at a bar, ordered beer, and began to sketch out an MMOG system architecture on a napkin.

They kept in touch, and in 2005, when they met again, Teles told Machado that they were having problems with their databases: whenever many users tried to log on at the same time, all attempting to retrieve masses of data with which to generate their view of the universe, the databases would crash.

"That's when the idea of using a mainframe—designed for concurrent transactions—came about," Machado recalls. He told Teles he would see what he could do. Two days later he called

COSMOS CREATORS: Carlos Eduardo Knippschild [top right] keeps *Taikodom*'s progress on track using a software-development method called Scrum, with Post-it notes on a board controlling the flow of tasks; Tiago Luz Pinto [left] leads a team creating a software infrastructure for the game server that Hoplon Infotainment and IBM call the gameframe, a hybrid of a mainframe and Cell processor–based servers.

SPACE FACES: Hoplon Infotainment's game designers and artists created an elaborate universe for *Taikodom*, with an intricate story line and a host of characters with whom players can interact. The company says that players' actions will have an impact on how the story unfolds and that in a future release users will be able to create their own avatars and wander inside space stations.

with good news. "Tarq, I got a server for you," he said. "It's a z9."

With an architecture descended from the famed IBM System/390 of the 1980s, the z9, a refrigerator-size black monolith, is one of the world's most powerful mainframes. It can have up to 512 gigabytes of memory and 54 central processors—some general purpose, others dedicated to cryptography or Java applications.

But the greatest thing about this high-powered, supersecure machine, Machado says, is that you can create hundreds of logical partitions, each with its own operating system and a share of the processors, memory, and storage. These partitions work as virtual servers, exchanging data directly through the mainframe's memory at much faster rates than if they were connected through conventional networking interfaces.

"The mainframe saves you from having to have massive farms of *x*86 servers, consuming lots of power and generating lots of heat," he says. "It's a cluster in a box."

PACING AROUND THE OFFICE, Teles looks tired after returning from a weeklong game conference. But when I ask about *Taikodom*'s fictional universe, he becomes animated. He says the game is different from other space-themed MMOGs, including *EVE Online*, which hit the market in 2003 and now has 230 000 subscribers. Not only will players be able to form alliances and participate in large-scale battles, they'll also be using social networking tools to create new kinds of interactions. One user may coach other players to help them obtain fighter-pilot certificates; another may use an in-game camera to broadcast battles.

"You can be a pilot, patroller, trader, miner, merchant. Or a spy, pirate, mercenary—maybe even a journalist," Teles tells me.

The story, he continues, takes place in the 23rd century. Humanity has been living in space ever since a mysterious energy shield engulfed Earth. Nations don't exist, and societies center around megacorporations. Some seek to build large industrial empires. Others dedicate themselves to hedonistic pursuits, like making beer. Human communities have a genetic pool that artificial intelligences use to generate new breeds. These AIs raise the kids as well. "Sex still exists, but no one wants to have babies," Teles says. "The only ones who want to are the Atavic people, who believe in the Galactic Spirit. They are a crazy sect!"

The idea for *Taikodom* began to take shape when Teles met Pinto and Buzzarello at the Federal University of Santa Catarina, in Florianópolis. They enjoyed playing multiuser dungeon games, or MUDs, the early computer role-playing games based on text, running on the university's mainframe. In the late 1990s, when they saw *EverQuest*, with its 3-D graphics and thousands of users around the world, they thought the market for such games was going to explode and decided they wanted to be part of it.

Their idea was to create a game without using the concept of levels and skills typical of existing MMOGs. "That's when you choose a character—a dwarf, wizard, warrior—with certain skills, and you need to kill a monster to increase your level so you can kill a bigger monster and so on," Teles says. "It's called the treadmill model." They wanted to give users more freedom and emphasize social interactions over mechanical, repetitive game play.

In early 2000, the trio put together a business plan and went after investors. But they came knocking just as the dot-com bubble was bursting, and all the doors closed in their faces. To bring in some cash, they took small jobs writing software for the Web and for scientific simulations. Just when the group was starting to run out of money, Knippschild, then finishing his engineering degree, joined as a Java programmer and mentioned that his uncle, Erich Muschellack, might be able to help.

"The greatest thing I saw was the opportunity of developing not only a game but an industry of games in Brazil," say Muschellack, the founder of venture capital firm Idee Tecnologia, in São Paulo. "Their plan is very ambitious. But that's good: they know they have to work 30 hours a day."

In January 2004, after eight months of negotiating, Hoplon and Idee closed a first round of funding. The friends became partners—and their game, business.

When Teles told his staff that *Taikodom* was going to run on a mainframe, they were worried they'd have to learn PL/I, BAL, JCL, or some other esoteric mainframe programming language. But Teles told them that the mainframe ran not only specialized software but also Linux and Java. For the programmers, that was like handing them the keys to a Ferrari.

And not just any Ferrari. The z9 easily handled the transactional workload of the game but didn't do so well with the floating-point math operations. So IBM decided to bring in its calculation turbocharger, the Cell processor.

"The technology vision here is to create a hybrid that has the best of both worlds," says Roland Seiffert, a senior engineer at IBM's Böblingen lab, in Germany, which designed the System z architecture and ported Linux to run on the mainframes. "It puts computations where they belong."

Seiffert says that today's multicore *x*86 processors perform parallel calculations by chopping a stream of instructions into threads. But the Cell beats the *x*86 because in addition to multiple threads, it has parallelism *inside* the threads. For example, to determine if a missile has hit a spaceship, you must compute the coordinates of the missile over time and check to see whether they coincide with those of the spaceship. You take a matrix containing the coordinates of the missile and use kinematics equations to calculate the future coordinates. Whereas an *x*86's threads calculate the new coordinates sequentially, the Cell's take sets of coordinates and perform mathematical operations on them all at once.

IBM used a Gigabit Ethernet interface to connect the z9 mainframe to the Cell chips, which are housed in blade servers, thin computers that sit on a rack like books on a shelf. Hoplon now has 10 virtual servers running SUSE Linux on the z9. In one test, the system was able to handle 700 spaceships in one section of the universe. That's still very far from the company's goal of supporting some 40 000 users at a time. But the Brazilians say the solution—more mainframe resources and faster networking equipment—is on its way.

IBM is replacing the z9 with the recently unveiled z10, a system the company says is 50 percent faster than its predecessor and provides the computing power of a cluster of 1500 *x*86 servers. And Hoplon, which had been sharing a mainframe with other IBM customers, will now have the new machine entirely to itself. The two companies are also relocating the z10-based gameframe to another data center with network switches and backbone connections better tuned to keep latency in the tens of milliseconds, as online games demand.

Still, if IBM wants to make the case that the gameframe is the ideal platform for MMOGs and virtual worlds, it has yet to demonstrate how it compares with other approaches. One alternative, also developed by IBM, is the system running *EVE Online*—a cluster of IBM blade servers with AMD processors and arrays of solid-state disks—which the game's developer, CCP Games, says can host 41 000 simultaneous users in a single universe.

Another approach that challenges the gameframe concept head-on is Sun's Project Darkstar, a Java-based open-source game server platform. Chris Melissinos, chief gaming officer at Sun, says that IBM is "throwing hardware at a software problem." The gameframe, he says, forces developers to conform their games to its hybrid hardware. By contrast, Sun is building a game-agnostic platform that handles all computing tasks using conventional servers. "To the developer, it looks like he's talking to one CPU on the back end," says Melissinos. "In reality, it could be thousands of CPUs."

Hoplon's answer is that its bitverse middleware is a key piece of the gameframe and works much like Darkstar, except that it has more powerful hardware at its disposal. "The Cell can give MMOGs more realism, more advanced physics and AI, which has long been a challenge," Machado of IBM says.

IT'S FRIDAY NIGHT in Florianópolis. I'm still peppering Teles and Pinto with questions about *Taikodom*, so we head out to Por do Sol Açoriano, one of the best seafood joints in town, where Hoplon's founders have delivered many pitches to investors, partners, and of course, journalists. "The only reason Hoplon has survived to this day is because of the quality of life here," Pinto says, gesturing from the spacious veranda to the sea.

A key factor that will distinguish *Taikodom* from other MMOGs, Teles says, is its elaborate story. Although the four founders had their ideas about creating fictional universes, they decided to get professional help, by hiring Gerson Lodi-Ribeiro, a famous Brazilian sci-fi author. The idea is that in *Taikodom*, players' actions will have an impact on the story. "A single player may affect the universe's fate," Teles says.

Within the Brazilian game community, opinions are divided. Some say that *Taikodom* will show the world what Brazil can do with enough resources. But others have accused Hoplon of hyping its game, some even calling it vaporware. Flávia Gasi, a game journalist, says the problem is that Hoplon has announced the game's launch on several occasions, only to pull back each time. "It seems that whenever the game is ready, the market has already advanced, so they start all over again," she says.

Outside Brazil, *Taikodom* is still barely known. A small number of English-speaking players are testing the beta version. Some of them have banded together within the game to fight against Brazilian alliances. This could be just friendly rivalry, or it may be a sign of a bigger problem: a massive Brazilian presence could pose an entry barrier for non-Portuguese speakers.

I ask Teles why it's taking so long to release the game. He sighs. "Just because we have an investor and a partner like IBM, people think we live in paradise," he says. "They have no idea what we're going through. It's been work, work, work."

Challenges, he says, have varied from expected hurdles—difficulty in finding good game programmers—to surprising snags, like the flame war that erupted in *Taikodom*'s discussion forum after a player reported seeing a gigantic bee flying in outer space (it was a software glitch), or the time a water pipe burst in the office, soaking computers and destroying drawings.

The next month will be crucial for Hoplon, because this may be its only chance to get it right. The company plans to publicize *Taikodom* in Brazil and eventually amass thousands of subscribers. Then it will seek partners in North America and Europe for a late 2009 release.

"As they say, if you knew how hard it was going to be, you wouldn't have decided to do it in the first place," Teles says. "Then again, you wouldn't have had a shot at making something successful." ❏

TO PROBE FURTHER To see videos and screen shots of Taikodom, *go to http://www.spectrum.ieee.org/aug08/taikodom.*

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

Sharp minds wanted for cutting-edge research.

Institute for Telecommunications Research

Under the leadership of newly appointed Director, Professor Alex Grant, the Institute for Telecommunications Research is broadening its research horizons and entering a new and exciting period of growth.

We are seeking to strengthen our existing research program and to expand into new areas of wireless communications networks. Of particular interest are fourth generation networks, mobile broadband, high-speed indoor wireless, infrastructure-less networks, cognitive radio and cooperative communications.

The Institute is already Australia's largest University based telecommunications technology research organisation. Founded in 1994, the Institute is internationally recognised for excellence in the areas of information theory, error control coding, mobile broadband, cellular, wireless and satellite communications. The Institute is home to a broad range of research activities, from basic theory and commercially funded applied projects, right through to the development of high-quality software and hardware products for delivery to industry.

Current Research Career Opportunities

Research Professor: Wireless Communications Networks

(2 positions available)

This is a unique and exciting opportunity for two Research Professors to build significant research programs in complementary areas of wireless communications, adding to the existing vibrant research environment.

ITR Director's Fellowship

This is a competitive fellowship supporting early career researchers in areas of strategic importance to the Institute. The inaugural fellowship is offered in the area of Information Theory. The successful candidate will pursue high-quality, fundamental research in the mathematical foundations of information transmission, processing and utilisation.

For more information on these positions visit: www.unisa.edu.au/hrm/careers

Future Research Career Opportunities

In addition to these specific opportunities, we are seeking expressions of interest from other highly qualified candidates.

Mid-Career and Senior Researchers
Outstanding mid-career and senior researchers are invited to submit an expression of interest in a wide range of wireless communication fields to discuss opportunities within the Institute for Telecommunications Research.

Post Doctoral Fellowships
Contact us about opportunities to further your career and research interests.

PhD Scholarship
Scholarships up to $30,000 per year are available.

Visiting Scholars
We welcome national and international scholars of high standing to spend time at the Institute.

For further detail on any of these opportunities, please contact: Professor Alex Grant on: +61 8 8302 5219 or email: alex.grant@unisa.edu.au

0704HR_IEEE CRICOS PROVIDER NO 00121B

www.unisa.edu.au/hrm/careers

Gold Award Winner – 2005 ACCI/BCA National Work and Family Awards

EOWA EMPLOYER OF CHOICE FOR WOMEN

The University is an Equal Opportunity Employer • Employer of Choice for Women • People of Aboriginal Australian and Torres Strait Islander cultural and linguistic background are encouraged to apply

Experience. The Difference.

University of South Australia

ieee spectrum classifieds

KIT Karlsruhe Institute of Technology

KIT – The cooperation of Forschungszentrum Karlsruhe GmbH and Universität Karlsruhe (TH)

Universität Karlsruhe (TH) Research University · founded 1825

Forschungszentrum Karlsruhe in der Helmholtz-Gemeinschaft

Presently, Forschungszentrum Karlsruhe (FZK) and Universität Karlsruhe (TH) are merging their activities in the Karlsruhe Institute of Technology, KIT. Within the framework of KIT, applications are requested for the position of

Scientific Director

of the Institute for Data Processing and Electronics
of the Forschungszentrum Karlsruhe in association with a

Professorship (W3)

at the Universität Karlsruhe (TH)

Applicants should have an excellent scientific qualification and international reputation in at least one of the following areas:

- Real-time data processing
- Microprocessor systems and analogue electronics
- Embedded systems
- Software engineering for embedded systems

We expect from the candidate new impulses for the development of data recording systems for applications with high data rates and high time resolution. These systems shall be applied in the areas of astro particle physics and atmosphere and climate.

The personality to be appointed is expected to be highly capable of managing interdisciplinary projects and heading a large institute with diverse research tasks and is also expected to represent the above scope of topics also by teaching at one of the faculties of physics, electrical engineering and information technology or computer sciences. Candidates should have a university lecturing qualification or equivalent scientific degree.

Applications of women are strongly encouraged, as we wish to increase the proportion of females at the management level. Handicapped persons with equal qualification will be preferred.

Applications, including a CV, a list of publications, documents about previous research and teaching work as well as offprints of the five most important publications shall, up to **September 30, 2008** be addressed to:

Forschungszentrum Karlsruhe GmbH , Dr. Peter Fritz, Director
P. O. Box 3640, 76021 Karlsruhe, Germany

In addition, we would appreciate the submission of your application documents electronically.
E-mail: peter.fritz@vorstand.fzk.de

The University of Michigan and Shanghai Jiao Tong University Joint Institute (UM-SJTU JI): invites applications for tenure-track positions at assistant, associate or full professor levels in all emerging engineering fields. Successful candidates will be expected to establish research programs in cutting-edge areas, to mentor Ph.D. students and to teach courses at both the graduate and undergraduate levels in either mechanical engineering, or electrical and computer engineering. Salary will be highly competitive and commensurate with qualifications and experience. For more information, please visit (http://www.umji.sjtu.edu.cn/job/faculty_hiring_ad.htm).

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Professorship in High Power Electronics and the Technology of Electrical Power Systems

The Department of Information Technology and Electrical Engineering (www.ee.ethz.ch/en/home.html) at ETH Zurich invites applications for a full professor position in High Power Electronics and the Technology of Electrical Power Systems. The responsibilities of this position include research and teaching in the area of power electronics converter systems and technologies with the long-term goal of a fundamental transformation of the existing infrastructure for transmission and distribution of electrical energy into a highly stable, highly efficient and extremely secure system, integrating a large share of renewable energy sources and storage capabilities. The focus will be on the theoretical and experimental analysis of high power converter systems and on the physically well-founded, multi-domain modelling and simulation of the circuits, thermal and electromagnetic properties at different levels of abstraction. Further research topics include the technological basics of new concepts for the transmission of very large quantities of electrical power as well as the simultaneous distribution of different energy carriers, i.e. multi-carrier systems.

Candidates must hold a PhD degree and should have established an internationally recognized research record in the area. Furthermore, they should present evidence of their leadership in a research team and of their ability to successfully colla-borate with other faculties as well as with the industry. He or she will be expected to teach undergraduate level courses (German or English) and graduate level courses (English).

Applications should include a curriculum vitae, a list of publications, a list of research activities, a research statement, and the names of at least three referees and should be submitted to the **President of ETH Zurich, Prof. Dr. Ralph Eichler, Raemistrasse 101, 8092 Zurich, Switzerland, no later than October 15, 2008.** With a view toward increasing the number of female professors, ETH Zurich specifically encourages qualified female candidates to apply.

L The northernmost University of Technology in Scandinavia
World-class research and education

GRADUATE SCHOOL OF SPACE TECHNOLOGY SEEKS FOR

Four Postdoctoral Research Fellows

More information: www.ltu.se/eng

Luleå University of Technology has an annual turnover of more than EUR 130 million. Today the University has 1,600 employees and 12,000 students. Research is carried out in close cooperation with companies such as Shell, Ericsson, Ford/Volvo, LKAB, Airbus, Volvo Aero and IBM and with the University's international partners such as Monash University and Stanford University. This research has a turnover of EUR 66 million.

LULEÅ TEKNISKA UNIVERSITET

Sign me up!

What's New @ IEEE

10 FREE NEWSLETTERS TO CHOOSE FROM

SUBSCRIBE TODAY:
www.ieee.org/whats-new

IEEE

World Class University Project

The Korean Ministry of Education, Science and Technology (MEST) is inviting excellent scholars and researchers from around the world to establish new academic projects and conduct joint research at Korean universities in the field of emerging technologies

WHAT IS THE WORLD CLASS UNIVERSITY (WCU) PROJECT? The WCU project is a higher education subsidy program of the Korean government, which invites international scholars who possess advanced research capacities to collaborate with Korean faculty members and establish new academic programs in key growth-generating fields.

WHICH FIELDS DOES THE WCU PROJECT SUPPORT? Focus is placed on supporting new growth-generating technologies that will spearhead national development. The ministry will give priority to inter-disciplinary studies that consolidate the fields of basic sciences and humanities & social sciences which will contribute to national, social and academic development.

HOW DOES THE WCU PROJECT WORK?
TYPE 1 ***Establishing new academic departments or specialized majors***
Under this type, high-quality foreign scholars are employed at Korean universities as full-time faculty members, on a contract of three years minimum, to establish new academic departments or specialized majors at the universities. Undergraduate degree programs should be established and opened by the spring semester of 2010 at the latest, and graduate programs should be in operation by the fall semester of 2009 at the latest.

TYPE 2 ***Recruiting foreign scholars to existing academic projects***
Under this type, foreign scholars are employed as full-time faculty members at existing departments of Korean universities to conduct joint research with Korean academics. Recruited foreign scholars are expected to be capable of developing new growth-generating technologies and also creating interdisciplinary studies.

TYPE 3 ***Inviting distinguished world-class scholars***
The third type invites distinguished scholars (including pioneering high-tech engineers) as part-time faculty members to conduct academic or research activities in a Korean university for a period of at least two months per academic year.

WHO IS ELIGIBLE FOR WCU PARTICIPATION (FOREIGN SCHOLARS)? All faculty members/researchers employed at a university, research institute or enterprise outside Korea are eligible to apply, including scholars of foreign nationality, ethnic Koreans who hold foreign nationality or citizenship, and scholars of Korean nationality.

WHERE CAN APPLICANTS FIND FURTHER DETAILS? Interested scholars are invited to view details and post questions about the WCU project on KOSEF's website (***http://www.kosef.re.kr/english_new***)

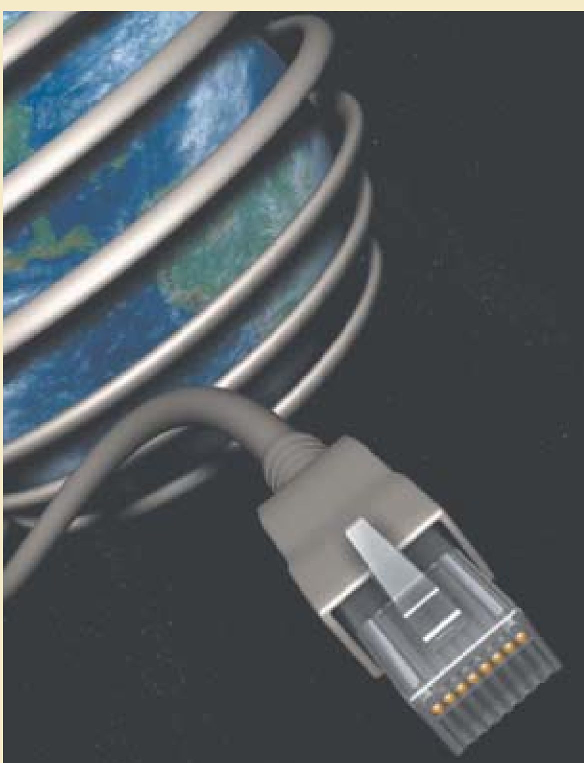

IEEE

Stay Connected

Spend less time worrying and more time communicating with an IEEE Email Alias

An @ieee.org email address shows you're part of a community of innovators. The IEEE Email Alias gives you an @ieee.org address while forwarding mail to your real account — and scanning for viruses at the same time.

It's free with your IEEE membership. Sign up today at www.ieee.org/alias

the data

BY PRACHI PATEL-PREDD

Engineers Are Doing Well by Doing Good

SALARY OFFERS for new U.S. electrical engineering grads over the past five years may not have increased as much as they have for those majoring in other areas of engineering, but demand is still strong. That should come as no surprise, given that technologists are behind the scenes in solar energy and search engines, cellphones and fuel cells, DNA sequencing and Hollywood blockbusters.

This rise in starting salaries would be even higher were companies not able to get young talent from such places as India, China, and Romania. The U.S. Bureau of Labor Statistics projects that over the next decade, EE employment will grow much more slowly than other engineering areas, because of the job outflux to other countries.

The consumer electronics industry, telecommunications, and aerospace and defense are sustaining the U.S. EE job market. And companies are willing to pay larger premiums for higher talent and education levels. ❏

U.S. Engineering Salaries

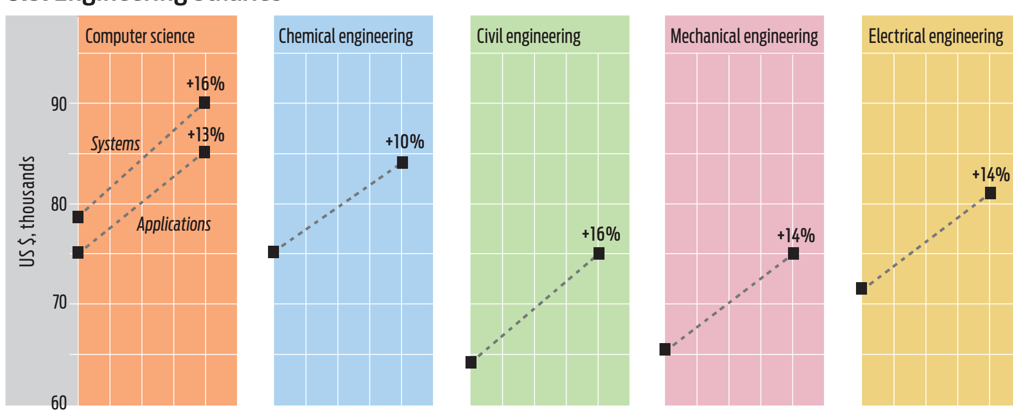

U.S. Starting Salaries for Bachelor's Degree Graduates

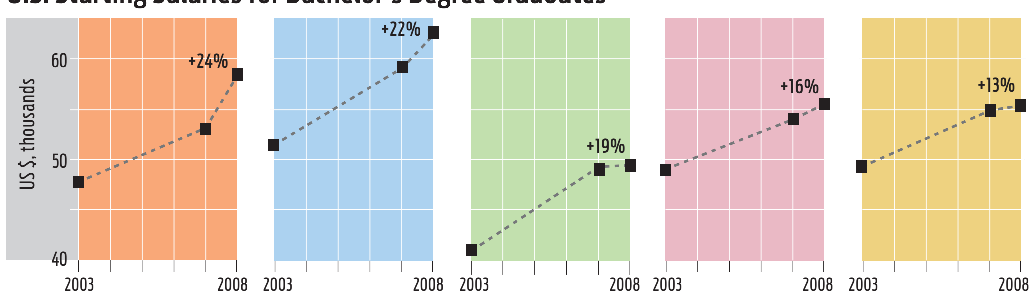

EE Salaries, BSEE vs. MSEE

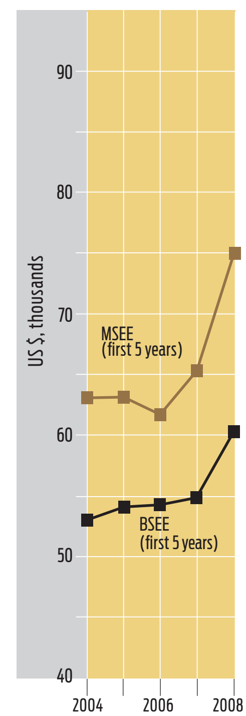

U.S. Salaries for EE Specialties

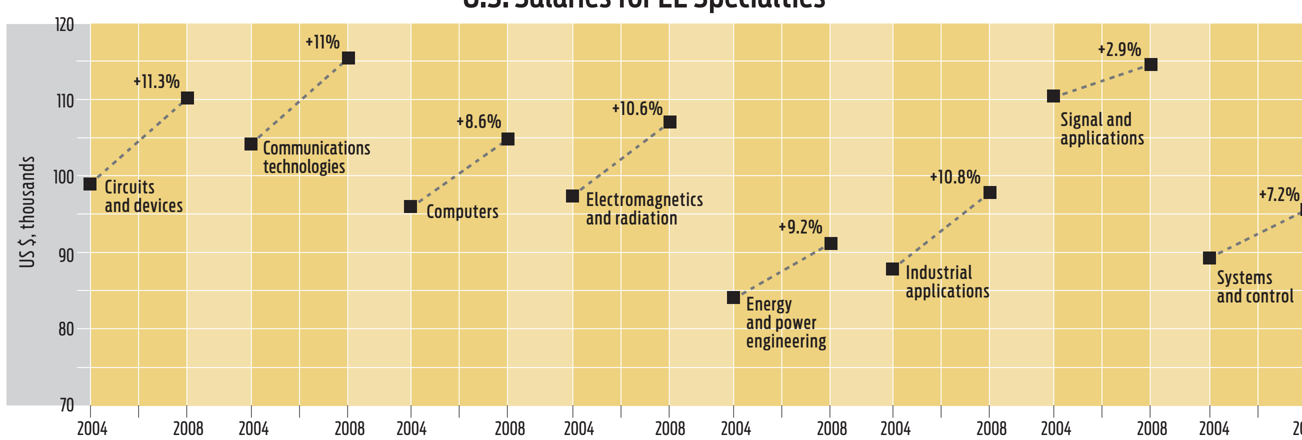

Sources: U. S. Department of Labor, Bureau of Labor Statistics; National Association of Colleges and Employers; IEEE-USA.

THREE AIRCRAFT, A SINGLE MODEL,
AND 80% COMMON CODE.

THAT'S MODEL-BASED DESIGN.

To develop the unprecedented three-version F-35, engineers at Lockheed Martin created a common system model to simulate the avionics, propulsion, and other systems and to automatically generate final flight code.
The result: reusable designs, rapid implementation, and global teamwork. To learn more, visit mathworks.com/mbd

MATLAB® & SIMULINK®

The MathWorks™
Accelerating the pace of engineering and science