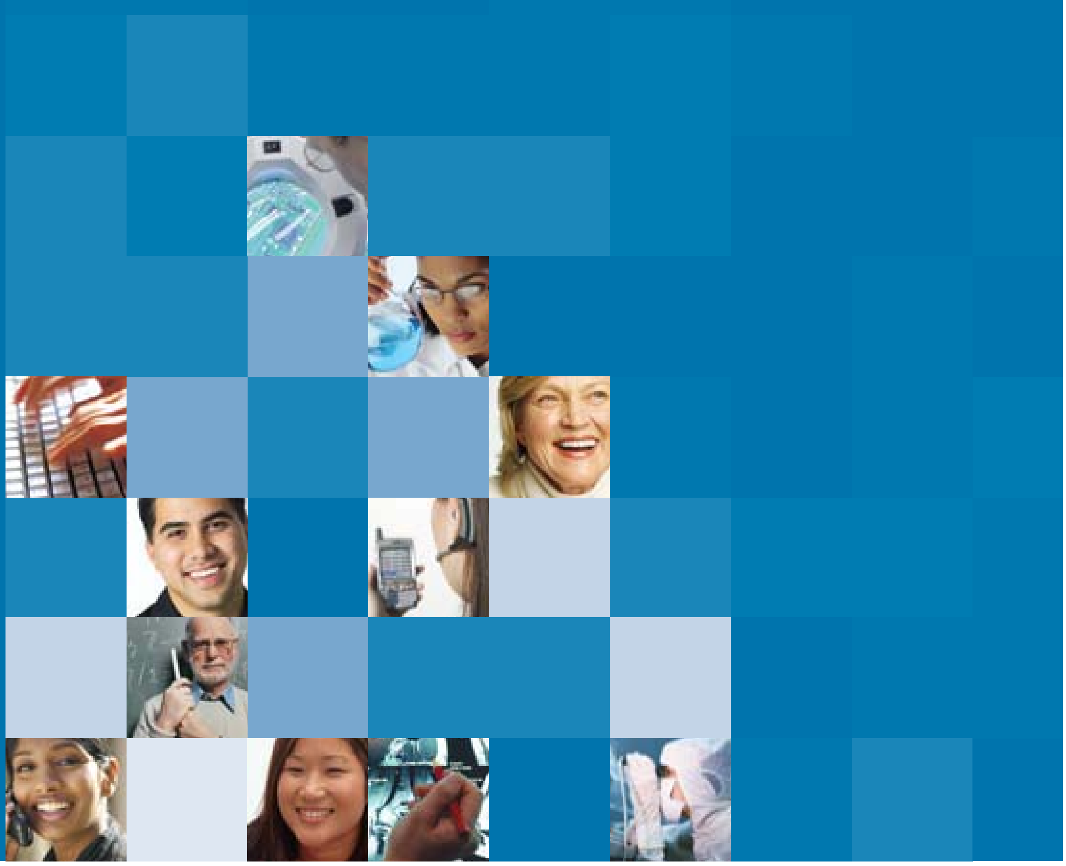

IEEE Membership

You are part of a community of 375,000 innovators, including 80,000 students, worldwide. IEEE membership brings you access to the latest technical information and research, global networking and career opportunities, and exclusive discounts on education and insurance products.

This brochure is your guide to some of the exclusive benefits and services available to IEEE members. Visit **www.ieee.org/benefits** for more information. Member benefits vary by location and grade.

Your IEEE Membership Makes a Difference

As an IEEE member, you are helping to make a difference in the world. Your membership enables efforts such as university program accreditation, advocacy for the profession, the introduction of technology careers to young people around the world, and more.

Get Involved

Participate through local Section, Chapter, and Branch activities. You automatically become a member of your local IEEE Section when you join. IEEE has over 300 local Sections, 1,700 Technical Chapters and 1,600 Student Branches. Participate in events and begin building relationships that will last through your career.

Both undergraduate and graduate students can get involved in the local Student Branch at their university, adding professional skills and experience to enhance formal education.

For more information, visit: www.ieee.org/sections

KNOWLEDGE

Stay Current with the Fast Changing World of Technology

IEEE membership offers exclusive access to the latest research and advances in technology, keeping you knowledgeable and current. As an IEEE member, you have access to the world's **best** collection of technical information.

IEEE Member Benefits include:

- Ask*IEEE
- Education Partners Program
- *IEEE Potentials* Magazine
- *IEEE Spectrum* Magazine
- IEEE.tv
- Online access through the IEEE *Xplore*® digital library
- What's New @ IEEE Newsletters

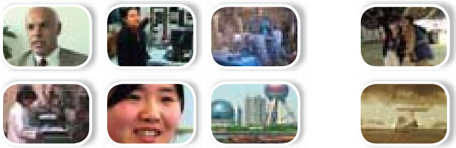

The IEEE.tv Internet-based television network offers conference highlights, author interviews, careers and technology features, and much more. Access via myIEEE.

www.ieee.org/**benefits**

SOCIETY MEMBERSHIPS

Customize Your Membership and Engage Your Technical Interests

Gain access to current information specific to your chosen technical interest, allowing you to network with peers who best understand your technical expertise. IEEE members often choose to add one or more of these listed below memberships to customize and add value to their IEEE membership.

IEEE Society & Special Interest Memberships:

- IEEE Aerospace and Electronic Systems Society
- IEEE Antennas and Propagation Society
- IEEE Broadcast Technology Society
- IEEE Circuits and Systems Society
- IEEE Communications Society
- IEEE Components, Packaging, and Manufacturing Technology Society
- IEEE Computational Intelligence Society
- IEEE Computer Society
- IEEE Consumer Electronics Society
- IEEE Control Systems Society
- IEEE Dielectrics and Electrical Insulation Society
- IEEE Education Society
- IEEE Electromagnetic Compatibility Society
- IEEE Electron Devices Society
- IEEE Engineering in Medicine and Biology Society
- IEEE Geoscience and Remote Sensing Society
- IEEE Industrial Electronics Society
- IEEE Industry Applications Society
- IEEE Information Theory Society
- IEEE Instrumentation and Measurement Society
- IEEE Intelligent Transportation Systems Society
- IEEE Lasers and Electro-Optics Society
- IEEE Magnetics Society
- IEEE Microwave Theory and Techniques Society
- IEEE Nuclear and Plasma Sciences Society
- IEEE Oceanic Engineering Society
- IEEE Power Electronics Society
- IEEE Power & Energy Society
- IEEE Product Safety Engineering Society
- IEEE Professional Communication Society
- IEEE Reliability Society
- IEEE Robotics and Automation Society
- IEEE Signal Processing Society
- IEEE Society on Social Implications of Technology
- IEEE Solid-State Circuits Society
- IEEE Standards Association
- IEEE Systems, Man, and Cybernetics Society
- IEEE Ultrasonics, Ferroelectrics, & Frequency Control Society
- IEEE Vehicular Technology Society
- IEEE Women in Engineering

Most IEEE Society memberships include a personal subscription to the core journal, magazine, or newsletter of that society, with many including full-text online archive access as well.

www.ieee.org/benefits

ADD-ON BENEFITS

Broaden Your Membership and Move Your Career Forward

IEEE Membership also provides you with additional resources to stay current, expand your community, and advance your career. These additional benefits require an additional fee.

Additional Benefits include:

- IEEE Education Partners Program
- IEEE Expert Now
- IEEE Member Digital Library
- IEEE Society Memberships (see full list inside)
- *Proceedings of the IEEE* Journal
- and other IEEE publications

IEEE Member Digital Library brings you access to nearly 2 million IEEE documents —the most essential information in technology. With an optional subscription to the IEEE Member Digital Library, you may access up to 25 articles a month from any of the more than 140 highly cited IEEE periodicals, 800 annual conferences, and other literature.

Exclusive Member Programs and Services are offered to IEEE members at competitive, members-only rates. This wide range of exclusive services and products are specially priced and customized for the needs of engineers.

www.ieee.org/benefits

There are also a number of financial and personal benefits available* to IEEE members, including:

- Group Term Life Insurance, 10-Year and 20-Year Level Term Life Insurance
- Professional Liability Insurance
- Group Disability Insurance
- High Limit Accident Insurance
- Long Term Care Insurance
- Group Dental Insurance
- Short Term Medical Insurance
- Financial Planning Services
- International Travel Accident & Medical Insurance
- The IEEE-Endorsed Credit Card—available in the U.S., Canada, United Kingdom, Puerto Rico, Hong Kong, and Singapore
- Express Shipping
- Moving & Storage Services
- and more!

*benefits vary by location

IEEE Expert Now offers members the availability to purchase the best of IEEE's conference tutorials, short courses, and workshops in one-hour online learning modules. IEEE Expert Now also provides Continuing Education Units (CEUs) or Professional Development Hours (PDHs).

www.ieee.org/benefits

How Can We Help You?

IEEE Member Services
445 Hoes Lane
Piscataway, NJ 08854 USA

Phone
+1 800 678 4333
(toll free, US & Canada)

+1 732 981 0060
(Worldwide)

Fax
+1 732 562 6380

Web
www.ieee.org/memberservices

610472
08-MEMB-0369 10/08

IEEE
Celebrating 125 Years
of Engineering the Future

ieee spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
12.08

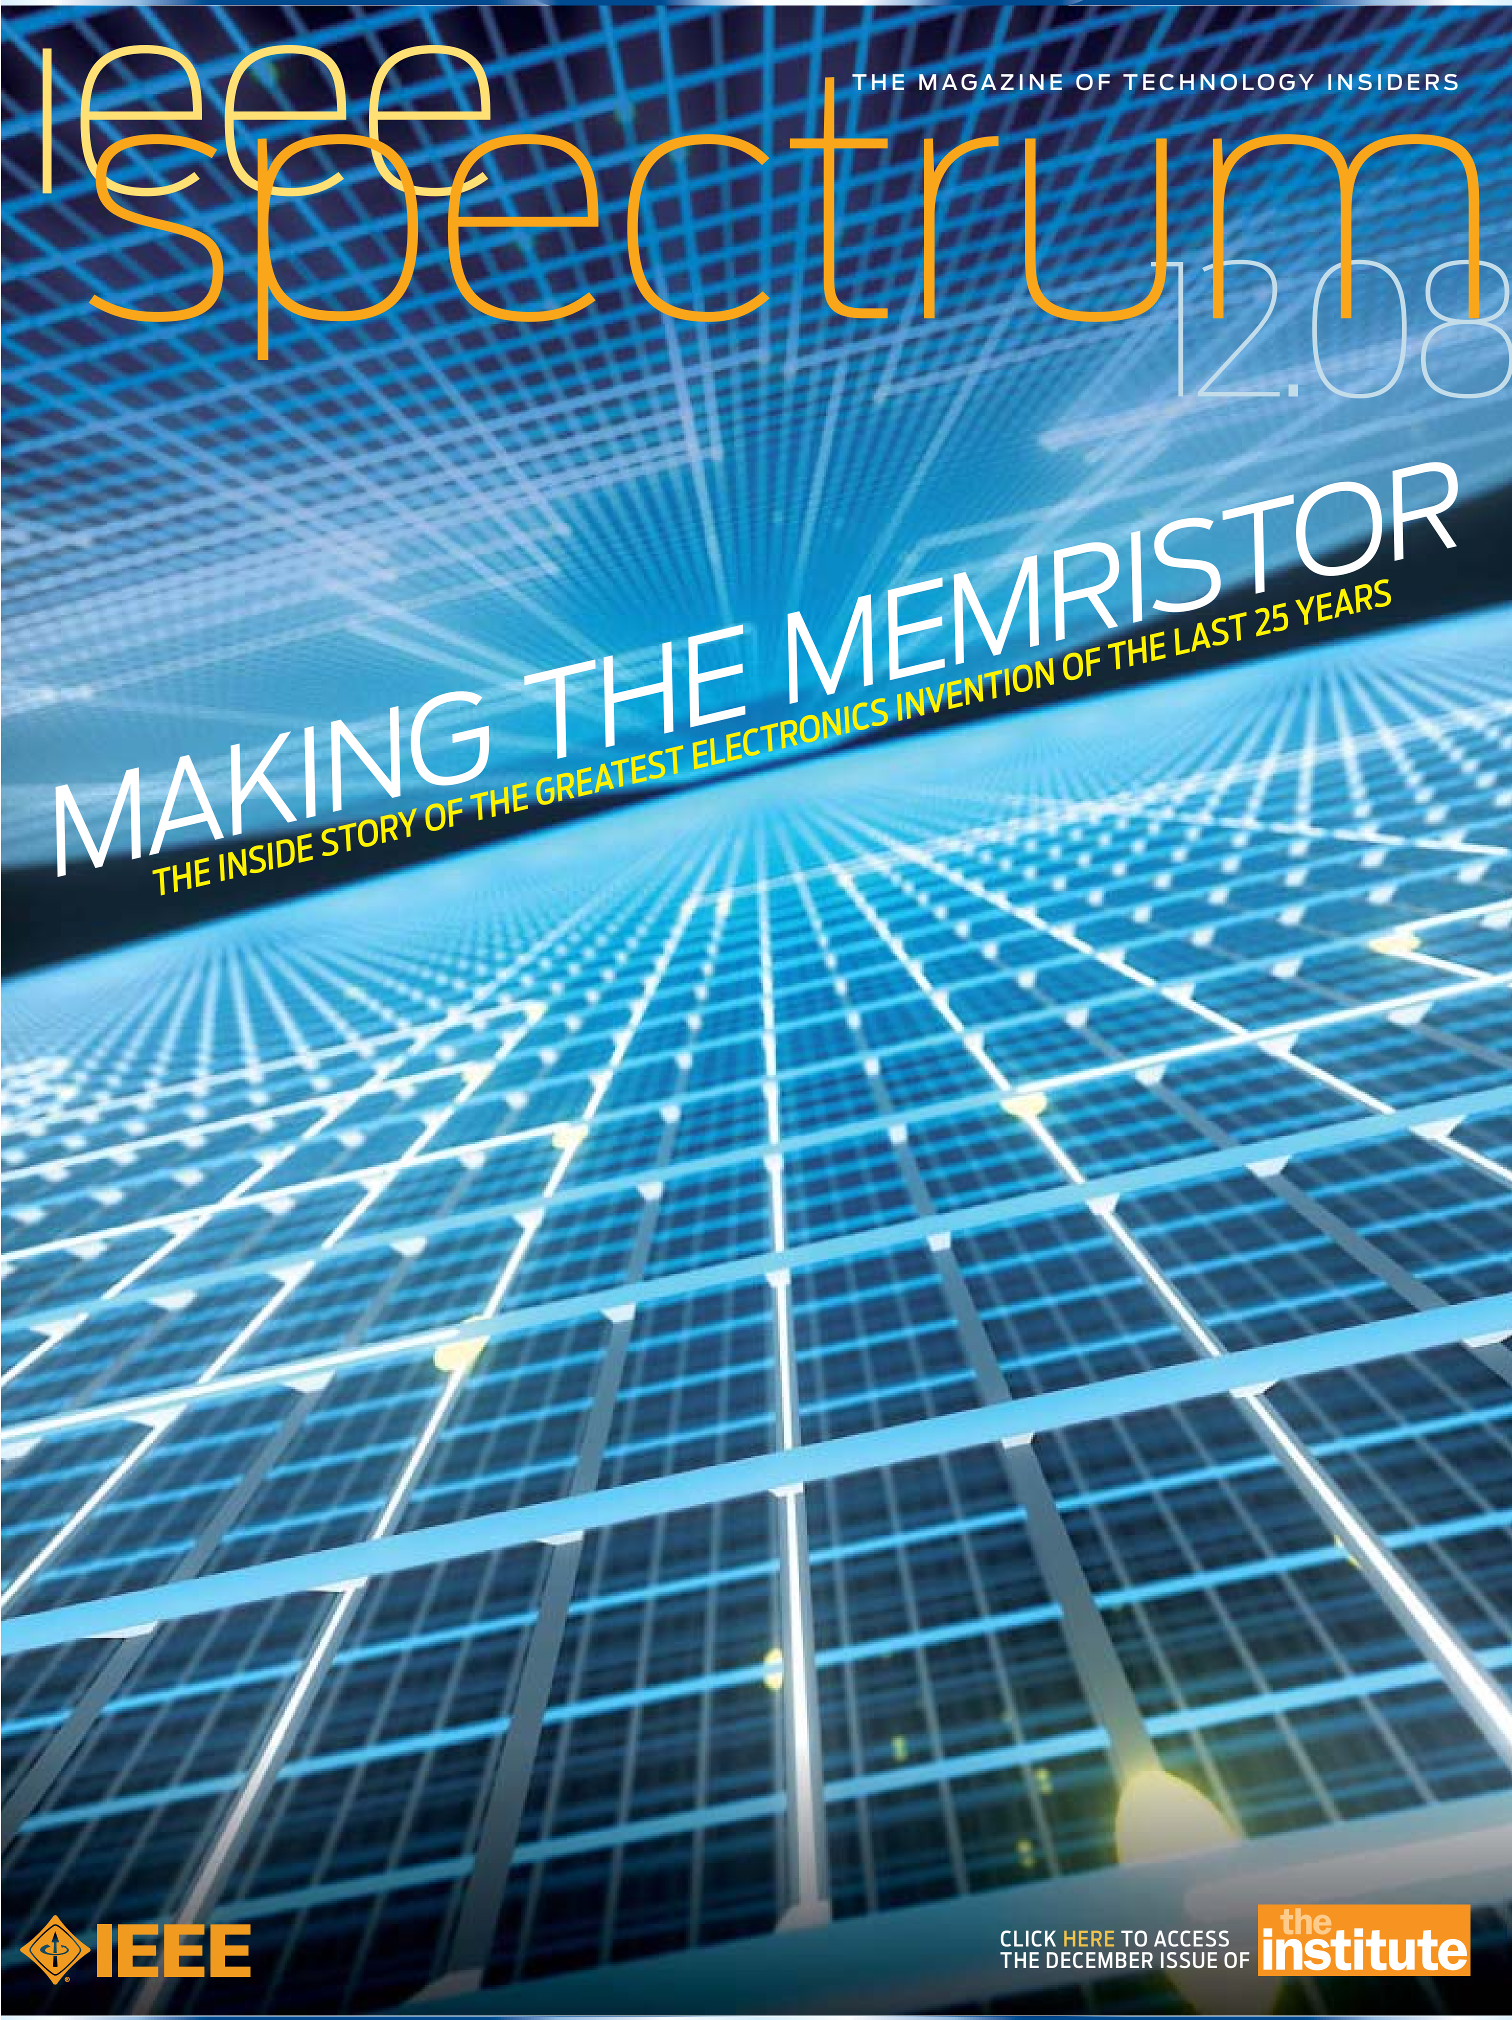

MAKING THE MEMRISTOR

THE INSIDE STORY OF THE GREATEST ELECTRONICS INVENTION OF THE LAST 25 YEARS

IEEE

CLICK HERE TO ACCESS THE DECEMBER ISSUE OF the institute

volume 45 number 12 international

ieee spectrum 12.08


32

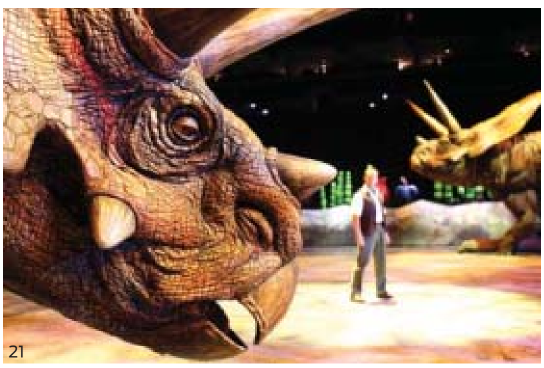
21

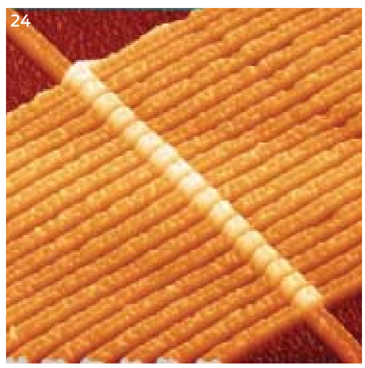
24

WHAT A CONCEPT: A novel engine technology propels the Mercedes F700 [top]; after millions of years, dinosaurs return to tour the United States [bottom left]; the memristor's architecture [bottom right] will enable a new generation of neural networks.

COVER: BRYAN CHRISTIE DESIGN

THIS PAGE, CLOCKWISE FROM TOP LEFT: STEFFEN JAHN; R. STANLEY WILLIAMS/ HP LABS; JOAN MARCUS/ GLOBAL CREATURES

COVER STORY

UPDATE

OPINION

DEPARTMENTS

ieee spectrum
volume 45 number 12 international
12.08

TOP: MICROSOFT; BOTTOM: JAMES ANDREW SMITH

SPECTRUM.IEEE.ORG

AVAILABLE ONLINE 1 DECEMBER

THE SURFACE IS THE COMPUTER

At a bar in Las Vegas, patrons can use the surfaces of their cocktail tables to flirt with other customers. In several AT&T stores, visitors can place two mobile phones side by side on a table to generate an instant comparison. And at Disneyland's Tomorrowland kitchen, the counter recognizes ingredients placed on it and suggests recipes for using them. This is surface computing, circa 2008. The technology, also called tabletop computing, has just hit the marketplace after brewing in research laboratories for the past 15 years. *IEEE Spectrum* takes you on a tour through the past—and the future—of surface computing.

ONLINE FEATURES:

PATENTS THAT POWER PROFITS: Which technology companies are filing the most patents, and which patents are having the biggest impact? Our latest patent-power rankings include companies providing e-ink technology for electronic books.

FOODIE ENGINEERS VIDEO: Meet the ex-engineers who now develop tasty treats instead of high-tech gadgets.

ALSO ONLINE:

- Webcasts
- Podcasts
- News
- Blogs
- Jobs
- Career Accelerator Forum
- IEEE Xplore® digital library
- White papers
- Opinions
- More!

WWW.IEEE.ORG/THEINSTITUTE

AVAILABLE 5 DECEMBER ON THE INSTITUTE ONLINE

THE ELECTION RESULTS ARE IN

Find out who won the race for 2009 IEEE President-Elect and read about other winners in the annual IEEE election.

A DECADE OF ONLINE PUBLISHING

IEEE's electronic library turned 10 this year. Learn about its history and how IEEE got into online publishing.

HANDS-ON PROJECTS FOR BUDDING ENGINEERS

Eight projects that won IEEE's Real-World Engineering Project, including such assignments as designing a prosthetic hand and developing error-correction codes for wireless communication systems, are now available online for free.

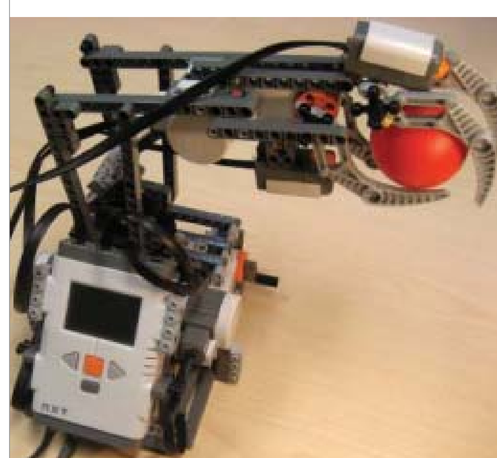

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2008 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine 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.

back story

Dance Into The Light

STARING AT the Mercedes-Benz F700 sitting in his Stuttgart, Germany, studio, photographer Steffen Jahn is ready for the shot. But this time he won't be behind the camera. Instead, he walks over to the car, grabs a fluorescent lamp, and tells an assistant to shut off all the lights. Then he does a little dance.

In a long, continuous movement, Jahn sweeps the lamp, wrapped with colored plastic strips, just above the Mercedes. With its shutter open, his camera captures not only his luminous brushstroke but also the intricate reflections off the car's shiny body [above].

By combining majestic effects like this with bright photos of the F700, Jahn created the vibrant, kinetic compositions that illustrate this issue's "The Soul of a New Mercedes," an in-depth look at how the German carmaker engineered this superluxurious and yet remarkably fuel-efficient concept car.

Jahn says he's seen many futuristic vehicles, but "nothing as impressive as the F700," with its compact and clean engine, active suspension, leather-and-cork interior, and gleaming metallic paint. "The trunk was filled with computers," he says. "Mercedes even sent a mechanic to help set up the car."

Concept cars have itineraries as tightly packed as those of U.S. presidential candidates, and the F700 is no exception. Jahn had only a single day with it. He concocted several different special light effects and made dozens of exposures, but he says, wistfully, that he'd have loved more time with the car. Alas, it had to go to the Stuttgart airport the next morning. It had a date with a Saudi sheik. ❑

STEFFEN JAHN

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. 12 (INT), December 2008, p. 9, or in *IEEE Spectrum*, Vol. 45, no. 12 (NA), December 2008, p. 11.

ieee spectrum

EDITORIAL

EDITOR IN CHIEF Susan Hassler, s.hassler@ieee.org

EXECUTIVE EDITOR Glenn Zorpette, g.zorpette@ieee.org

MANAGING EDITOR Elizabeth A. Bretz, e.bretz@ieee.org

SENIOR EDITORS Harry Goldstein (Online), h.goldstein@ieee.org; Jean Kumagai, j.kumagai@ieee.org; Samuel K. Moore (News), s.k.moore@ieee.org; Tekla S. Perry, t.perry@ieee.org; Philip E. Ross, p.ross@ieee.org; David Schneider, d.a.schneider@ieee.org; William Sweet, w.sweet@ieee.org

SENIOR ASSOCIATE EDITOR Steven Cherry (Resources), s.cherry@ieee.org

ASSOCIATE EDITORS Sally Adee, s.adee@ieee.org; Erico Guizzo, e.guizzo@ieee.org; Joshua J. Romero (Online), j.j.romero@ieee.org; Sandra Upson, s.upson@ieee.org

ASSISTANT EDITOR Willie D. Jones, w.jones@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org

COPY EDITOR Michele Kogon, m.kogon@ieee.org

EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

EXECUTIVE PRODUCER, SPECTRUM RADIO Sharon Basco

ASSISTANT PRODUCER, SPECTRUM RADIO Francesco Ferorelli, f.ferorelli@ieee.org

ADMINISTRATIVE ASSISTANTS Ramona Gordon, r.gordon@ieee.org; Nancy T. Hantman, n.hantman@ieee.org

INTERN Monica Heger, m.heger@ieee.org

CONTRIBUTING EDITORS John Blau, Robert N. Charette, Peter Fairley, Alexander Hellemans, David Kushner, Robert W. Lucky, Paul McFedries, Kieron B. Murphy, Carl Selinger, Seema Singh, John Voelcker

ART & PRODUCTION

SENIOR ART DIRECTOR Mark Montgomery

ASSISTANT ART DIRECTOR Brandon Palacio

PHOTO EDITOR Randi Silberman

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy

EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia

SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

WEB PRODUCTION COORDINATOR Jacqueline L. Parker

WEB PRODUCTION SPECIALIST Michael Spector

EDITORIAL CORRESPONDENCE

IEEE Spectrum, 3 Park Ave., 17th Floor, New York, NY 10016-5997
Attn: Editorial Dept. Tel: +1 212 419 7555 Fax: +1 212 419 7570
Bureau: Palo Alto, Calif.; Tekla S. Perry +1 650 328 7570
Responsibility for the substance of articles rests upon the authors, not the IEEE or its members. Articles published do not represent official positions of the IEEE. Letters to the editor may be excerpted for publication.

ADVERTISING CORRESPONDENCE

IEEE Spectrum, 3 Park Ave., 17th Floor, New York, NY 10016-5997
Attn: Advertising Dept. +1 212 419 7760
The publisher reserves the right to reject any advertising.

contributors

BOB BRISCOE describes how to ease Internet congestion by remaking the way we share bandwidth in "A Fairer, Faster Internet" [p. 38]. He says the problem isn't bandwidth hogs but the Internet's sharing protocol itself. Briscoe is chief researcher at BT's Networks Research Centre, in England. He is working with the Trilogy Project to fix the Internet's architecture.

VADIM BULITKO, JONATHAN SCHAEFFER, and **MICHAEL BURO** are all part of the GAMES group at the University of Alberta, in Canada. They describe how they are using artificial intelligence to develop the next generation of interactive video games in "Bots Get Smart" [p. 44]. As its acronym suggests, their research group creates software for games, with the goal of beating—or at least seriously challenging—the human competitor. In 1994, *Chinook*, the team's checkers program, became the first game software to win a championship against humans, earning it a place in *Guinness World Records*.

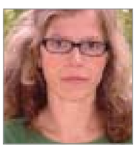

NANA RAUSCH, originally from Heidelberg, Germany, says that figuring out how to depict the bandwidth-sharing protocol for "A Fairer, Faster Internet" [p. 38] was challenging. Eventually she decided to go for a shopping analogy. "Everybody has been in a supermarket line with just one item, and there are people in front of them with their big carts," she says. QuickHoney, the design studio she owns with Peter Stemmler, has created illustrations for *Fast Company*, *GQ*, and *The New York Times*.

CARL SELINGER, an *IEEE Spectrum* contributing editor, shares some insights on navigating the workplace terrain in "Dealing With Difficult People" [p. 20]. A private consultant with 40 years of experience in business, government, and academia, Selinger gives seminars to engineers on nontechnical skills. His 2004 book, *Stuff You Don't Learn in Engineering School* (Wiley-IEEE Press), has now been published in China.

JOHN VOELCKER went to Daimler headquarters in Stuttgart, Germany, to report on the Mercedes-Benz F700, a prototype for future gas-sipping luxury cars [p. 32]. "They brought me into the top-secret design lab," he says, "and, unexpectedly, rolled the full-size styling model of the F700 out into the sunlight and onto a turntable surrounded by 8-meter-tall hedges. I got exactly the same viewing that the top brass at Daimler did!"

R. STANLEY WILLIAMS, a senior fellow at Hewlett-Packard Labs, wrote this month's cover story, "How We Found the Missing Memristor" [p. 24]. Earlier this year, he and his colleagues shook up the electrical engineering community by introducing a fourth fundamental circuit design element. The existence of this element, the memristor, was first predicted in 1971 by IEEE Fellow Leon Chua, of the University of California, Berkeley, but it took Williams 12 years to build an actual device.

CLOCKWISE FROM TOP LEFT: LYN BRISCOE; STEFANIE CASSIN; RANDI SILBERMAN; HEWLETT-PACKARD LABS; PETER STEMMLER spectral lines

The Audacity Of Technology

BARACK OBAMA believes in the power of technology. The cellphone-toting, BlackBerry-packing next president of the United States ran a brilliant Web 2.0 campaign, demonstrating how everything from social networking to GPS could be deployed to organize and mobilize volunteers and voters alike. He's not a technologist, but he has surrounded himself with tech-savvy strategists.

Obama clearly understood that technology could help him get where he wanted to be. Do he and his advisors have a plan for using it to get the United States where it needs to go? He's certainly saying all the right things. His Web site, BarackObama.com, makes the case for a formidable list of technology to-dos that includes:

- making broadband access available to all;
- overhauling the national electricity grid and other critical infrastructure;
- increasing funding for clean energy resources, such as biofuels and low-carbon coal technology;
- improving technology education and literacy; and
- building a bigger workforce of home-grown scientists and engineers and creating jobs for them.

There are two things Obama should do right now to signal his seriousness about science and technology. First, he should appoint his science advisor *before* his inauguration, to show he believes that the science advisor's job is as important as other senior-level positions. You'll recall that the current administration's science advisor, physicist John Marburger, wasn't appointed until near the end of President George W. Bush's first year in office, after decisions about divisive issues like stem cell research had already been made. Marburger also lost the title "Assistant to the President," which his predecessors had held. And while he's at it, Obama should fill the new cabinet-level position he's been talking about, that of chief technology officer, again to demonstrate serious intent.

Second, he should work with the U.S. Congress to get the America COMPETES Act (ACA) of 2007 fully funded, ideally in the first quarter of 2009. ACA is supposed to improve U.S. competitiveness and innovation by substantially increasing the amount of federal money spent on supporting national science agencies and institutes, R&D, science, technology, engineering, and math education, and an adjunct teachers corps. The bill also aims to reform the immigration system so that technologists from around the world can work in the United States with fewer restrictions.

This law was fast-tracked by Congress but fell victim to budgetary infighting. In 2005, Senators Lamar Alexander and Jeff Bingaman asked the U.S. National Academies to study the state of U.S. competitiveness in science and technology. The result was a report called *Rising Above the Gathering Storm: Energizing and Employing America for a Brighter Economic Future*, which painted a grim picture of the United States' eroding technological advantage but also offered a number of concrete proposals for turning the situation around. The report, by an elite committee of scientists and technologists led by former Lockheed Martin chairman and CEO Norman Augustine, created such a stir that a bipartisan congressional committee, which included Senator Obama, got the ACA into legislation in early 2007. By August, President Bush had signed it into law.

But then Congress and the president got deadlocked over the domestic spending budget—US $21 billion kept them apart—and the bill was never funded. Russell Lefevre, president of IEEE-USA—both Lefevre and IEEE-USA are longtime proponents of ACA—calls it "the train wreck of December '07." The hopes of the *Gathering Storm* report's authors were dashed. Some federal science agencies, like the U.S. National Science Foundation, were left with budget increases that were less than the rate of inflation.

An authorization bill without appropriations and $3 will get you a cup of coffee at Starbucks, as Lefevre says. But the 2009 budget is not yet finalized—Congress is waiting for the new administration to arrive—and so there is an opportunity for Obama and the new Congress to fully fund ACA early next year.

With so many other problems to deal with, things like rebuilding the American economy, ending the wars in Iraq and Afghanistan, fixing the health-care system, and restoring diplomatic relations with the rest of the world, why take this on?

Because if the president-elect really believes that technology can help realize our hopes and dreams, that technology drives social change and economic growth and stability, that technology is vital to health-care reform and the key to a transparent and engaged democracy, he, working together with the new Congress, should actively support science and technology development, starting now. Doing so will help seed solutions to some of the gravest problems facing all of us today. —SUSAN HASSLER

JAY L. CLENDENIN/AURORA PHOTOS

forum


spectral lines

Digital Dilemma

Converting to digital television is supposed to be simple. It's not

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.

THE BIG SWITCH

I ENJOYED Tekla S. Perry's "Digital Dilemma" on the trials of setting up digital television [Spectral Lines, October], but I hate to tell her that going to cable won't help. Here in northern Nevada, our cable provider offers not only standard analog television but also digital television, where you are promised four different PBS channels, a channel with old-time TV shows, and lots of other interesting stuff. The provider charges more for this service, even though, as we all know, the reason for going digital is not quality but to cram more channels into a given spectrum. For the digital service, the cable provider also demands that you rent a converter for each television. The problem is that the digital stuff comes in blocky or not at all, even though our analog reception has always been excellent. We took the converter back to the cable provider and noticed a long line of people doing the same thing. The representative did not question us when we said it simply didn't work. Now the local cable provider is seeking to move several popular channels to digital only. Guess we'll have to spend more time in the library.

ROBERT D. PARKER
IEEE Senior Member
Galena Forest, Nev.

MY CONCLUSION is that broadcast digital TV is getting a bum rap from people who don't really give it a chance and then declare it unsatisfactory. Receiving these broadcasts is not difficult, but as with any other technical project, a little quantitative design and configuration is in order.

Perry is certainly on the right track with replacing her old VHF rooftop antenna with a new UHF model. However, because digital TV will be broadcast on both the VHF high band (174 to 216 megahertz) and the UHF band (470 to 806 MHz), one of the newer antennas optimized for 174 to 806 MHz would have been a better choice.

I doubt that the aging cables in her house have anything to do with her problems, although cleaning or replacing the connectors might help. Distributing digital TV signals around her house probably requires preamplification at the antenna and possibly amplification at the head end of the cable runs. The higher-frequency signals involved in digital TV encounter noticeable attenuations in typical coaxial cable, and the splitters used to divide the signals into their multiple cable runs add their own losses. At 800 MHz, a 100-foot (30-meter) run of cable has a loss of about 10 to 12 decibels. Add that to the combined loss of about 8 to 10 dB for splitting the signal into three to five cable runs, and it's no wonder her installation isn't satisfactory. When adapting existing house cabling to digital TV, you must take care to assure that any splitters or amplifiers are rated for at least 174 to 806 MHz and that the old cabling has clean, uncorroded connector ends.

The picture-perfect quality of DTV is more than worth the nominal extra expense of building or adapting a home antenna and distribution system to accommodate it.

JAMES PALMER
IEEE Life Member
Santa Ana, Calif.

THE VOTES ARE IN

YOUR ARTICLE "Open-Source Voting" [Update, October] was interesting, but I'm surprised that no one has yet proposed a fairly obvious approach to clean elections that draws on both traditional practice and fault-tolerant computing designs.

In traditional elections, observers from each political party are part of the process, from voter check-in through vote counting and recounts of paper ballots. Optical character-recognition (OCR) voting machines, which read markings on paper ballots, allow election workers the option of manually recounting those ballots if a very close election warrants it. That manual recount option is typically lost with both mechanical and touch-screen machines. As an example of extreme fault tolerance, consider the space shuttle's five computers, which vote on mission-critical decisions. If the machines don't agree, special procedures are invoked.

It is easy to envision putting these models together: feeding the same stack of OCR ballots through multiple scanners from different vendors, one brought by each party, and perhaps even additional machines brought by watchdog agencies. If the tallies agree, all is copacetic. If not, then it's time for a manual recount.

MILES FIDELMAN
IEEE Member
Boston

more online at www.spectrum.ieee.org

update

The Price Is Wrong for Oil Shale and Tar Sand Tech

Falling energy prices could squeeze oil-extraction research

THE HUGE RUN-UP in oil prices over the last several years, reaching a peak of close to US $150 per barrel this past summer, has given energy companies a big incentive to find new ways of harvesting unconventional oil, especially in North America. Technology firms targeted oil from tar sands in Canada and from shale, a sedimentary rock abundant in the western United States. But in the fourth quarter of 2008, oil prices plummeted, and that could put the brakes on the development of new extraction technologies, say experts.

"It's a difficult time to come out with a new technology," says Jim Sledzik, an investment manager at Energy Ventures, a venture capital firm in Stavanger, Norway, that invests in oil and gas technology companies. He estimates that oil prices have to be above about $65 a barrel for heavy oil extraction—whether from oil sands or oil shale—to be viable. Sledzik predicts that fewer high-risk projects will receive funding.

According to a 2004 U.S. Department of Energy report, the United States is home to 2 trillion barrels of the world's estimated 2.6 trillion barrels of shale oil. Almost all of the U.S. shale is found in the Green River Basin in Colorado, Utah, and Wyoming.

These reserves have not been tapped, because getting oil from shale has been too costly, energy intensive, and dirty. Recovering shale oil also has the potential to contaminate ground water and produce toxic waste. Estonia, a country dependent on shale for energy, first mines the rock, then, without extracting the oil, burns it to run a generator [see "New Tech, Old Fuel" *IEEE Spectrum,* February 2007]. But burning the shale produces polluting ash, carbon dioxide, nitrogen oxides, airborne particulates containing heavy metals, and sulfur dioxide, which causes acid rain.

Newer technologies seek to mitigate these impacts by pumping the oil out of the ground without mining the shale.

Raytheon has developed a technique that uses radio-frequency energy to extract the oil. Oil producers would lower radio antennas into a well and then heat the shale

ABANDONED HOPE: Colony Oil gave up on this oil shale plant in Colorado 20 years ago. Technology that might revive shale's prospects is in development.

PHOTO: DAVID ZALUBOWSKI/ AP PHOTO

110.7 Number of scientific and technical articles published per million people in 2005, according to the *Organisation for Economic Co-Operation and Development Science, Technology and Industry Outlook 2008*

update

with radio waves, reducing the oil's viscosity enough to pump it to the surface. According to John Cogliandro, Raytheon senior principal engineer, the technique consumes the equivalent of one barrel of oil for every six produced.

Shell has developed a similar recovery process, except that instead of using RF energy, the company places electric heaters into wells to heat the shale. A technique recently developed by researchers at the University of Alberta, in Canada, improves on it by injecting an iron powder solvent, which in the lab decreased the viscosity of the oil and increased the amount recovered by between 20 and 40 percent.

Technology's impact on the economics of oil shale will go only so far, says Judson Jacobs, research director at Cambridge Energy Research Associates. "Regardless of how it is produced, oil shale [extraction] will be very carbon intensive," he says. If carbon is traded or taxed nationally, that too will be an impediment to oil shale development.

New technology is also being tested for oil sands extraction. Today the oil in Canada's tar sands is extracted either by mining, if the deposits are shallow, or by steam injection, for deeper sands. Both of these methods have a huge environmental cost.

Calgary-based E-T Energy is testing a technology to reduce the environmental impact by extracting the oil in situ with electricity. Current is sent down through wells to steel electrodes, which heat the oil, making it easier to pump to the surface. CEO Bruce McGee says E-T's technology will reduce emissions and water usage and will cost less than either mining or steam injection. McGee says the recent drop in oil prices has not affected the company's plans. "We are still moving forward," he says. "We're taking the view that these recent drops in oil prices are temporary."

Some analysts think that strategy is the right one. Cambridge Energy's Jacobs says it's too soon to predict the impacts of the recent drop in oil prices. He and Energy Ventures' Sledzik agree that most oil companies plan further into the future and aren't going to adjust their investments based on daily price fluctuations.

—Monica Heger

Exascale Computing by 2015?

The U.S. Defense Advanced Research Projects Agency (DARPA) asked a group of eminent engineers and computer scientists to figure out how to get a thousandfold increase in supercomputing ability by 2015. The group, led by University of Notre Dame professor and IEEE Fellow Peter M. Kogge, found that if technology trends continue into the future, exascale computing will still be just a dream in 2015. According to the report, getting to exascale by then—or ever—requires among other things an intensive focus on reducing the power devoted to moving data around in computers. However, some critics point out that a similar report in the mid-1990s—predicting the technologies needed for petascale computing—has now proved to be off the mark. And some say the DARPA group did not properly consider new types of applications supercomputers will be running in the future. Kogge explains the report's findings on Spectrum Radio at http://www.spectrum.ieee.org/dec08/scompod.

PHOTO-ILLUSTRATION: ARMAND VENEZIANO; ORIGINAL PHOTO: RYAN MCVAY/GETTY IMAGES

Google Searches for Ad Dollars in Social Networks

New patents aim to pry profits from patterns

THE U.S. PATENT and Trademark Office recently published a series of intriguing patent applications from Google. They raise questions about the search giant's significance for the profitability of social networks—and whether anyone has figured out how best to translate Web 2.0 hype into bankable income. Dozens of social-networking sites such as MySpace, Facebook, Bebo, and Friendster continue to flourish like Web start-ups in the dot-com heyday—consuming engineering talent, computing resources, and thousands of lines of code along the way. But no one has yet found the golden keys to profitability. The three Google patents, which rely on language processing and other technology to search for patterns in data, could ratchet up social networks' ad revenue by better targeting ads to individuals, experts say.

Most social network sites rely on ads for revenue, but according to New York City research firm eMarketer, those sites account for just one US $1.4 billion slice of the $50 billion online advertising pie. A good click-through rate for advertising on traditional media sites is 2 percent, but "on Facebook, you're lucky if you're going to get 0.3 percent," says Jaffer Ali, CEO of online advertising agency Vidsense, based in Mokena, Ill.

The patents—Open Profile Content Identification, Custodian Based Content Identification, and Related Entity Content Identification—and the algorithms behind them would let Google find patterns in users' profiles, pages, and friend lists in order to better target ads to them. Ideally, they would make the users more likely to click through.

Google's Related Entity patent, for one, involves prying information from a user's list of friends or user groups. But Google is not alone in this field. In June the social ad firm SocialMedia Networks said it had invented an algorithm called FriendRank that also scours a user's friendship lists for friends whose names might be dropped in a targeted ad.

The Open Profile and Custodian patents would mine data from, say, a MySpace user's profile and the profile of the MySpace page the user is visiting. The Open Profile patent, for instance, would consider a user profile like "I really enjoy hiking, especially long hikes when you can camp out for a few days. Indoor activities don't interest me at all, and I really don't like boring outdoor activities like gardening."

Using smart language-processing algorithms to detect the user's sentiments ("enjoy" or "don't like" near "hiking" or "gardening") and other linguistic cues, the system would then potentially serve up active outdoor sports-related ads to this user but avoid ads about more hobbyist-oriented activities.

While none of Google's proposed patents look into business strategies that social networking ad agencies haven't tried already, says Hussein Fazal, CEO of Toronto-based agency AdParlor, its language-processing and pattern-recognizing algorithms are probably key to the whole enterprise. However, Google did not disclose its particular pattern-searching algorithms, and a Google spokesperson declined an interview.

AdParlor's ad targeting, Fazal says, typically examines four or five factors, such as a social-network user's gender, age, and location. But Google's computing resources, he says, "might be able to analyze everything about that user."

Is there a patent fight brewing in all this algorithm activity? Jeremy Pinkham, chief technical officer of social-media advertising company Lotame Solutions, based in Elkridge, Md., doesn't think so.

"There's a lot of room for different folks to try different approaches," he says. Google's new patents help to "validate that this industry is worth people's attention." —MARK ANDERSON

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

HELPING HELICOPTERS: BBN Technologies is adapting a sniper-sensing system for military helicopters.
PHOTO: JOHN MOORE/GETTY IMAGES

Spotting Snipers With Sound

BBN's gunshot-detection system used for trucks in Iraq is being modified for helicopters

IN THE MOVIES, action heroes can dodge bullets or even catch them with their hands. In the real world, your average soldier is happy just to know where the shooter is.

For three years now, soldiers in Iraq and Afghanistan have had help locating snipers. An acoustics system known as Boomerang, developed by BBN Technologies, in Cambridge, Mass., determines the location of a shooter by sensing the blast at the rifle's muzzle and the air pressure of the bullet as it whizzes past. Now BBN is taking the technology further, building and testing a version for helicopters, which it hopes to ship next year.

The U.S. military first began to deploy the system in Iraq in 2005. Today more than 1000 of the units are in Iraq and in Afghanistan. BBN claims that the systems detect 95 percent of all bullets that fly anywhere within 30 meters of their sensors, never mistaking some other loud noise for a bullet.

BBN is best known for building the original Internet network in the 1960s. But it has also done a lot of work in acoustics.

In 1978, the U.S. Congress asked the company to analyze a motorcycle policeman's audiotape of the 1963 assassination of President John F. Kennedy. "That was the point at which we started thinking about gunshots and gunshot detection via acoustic signatures," says Dave Schmitt, BBN program manager for Boomerang. In the mid-1990s the Defense Advanced Research Projects Agency (DARPA), ran a small competition to pursue that line of inquiry, which BBN won. The United States wasn't involved in any hot wars at the time, "so they put that technology on the shelf," he says. DARPA's interest revived in the summer of 2003, when postinvasion resistance in Iraq led to frequent sniper attacks on U.S. troops. The agency asked BBN to build a system that would operate on the move. The company came up with a first-generation design in about two months, and 50 units were quickly assembled and sent over.

Using an array of seven microphones as its sensors, the system first detects the bullet's supersonic shock wave, from which it determines the angle from the sensor to the shooter. The system then listens for the muzzle blast. The time difference between the shock wave and the arrival of the blast—and some clever algorithms—yields complete information about the direction, elevation, and range, all in less than a second and a half. The shooter's direction is indicated as a clock position on a small console, the range and elevation are displayed on an LED screen, and all the information is spoken aloud by a recorded voice.

Once the microphone array and the signal processing had been worked out, the principal technical challenge was eliminating false positives—sounds that might be misidentified as gunshots. Early on, slamming a Humvee door would be enough to trigger the system, Schmitt says. "It would look enough like a [bullet's] shock wave that we would do the processing and report it falsely." After some tinkering, BBN got Boomerang into shape. The trade-off in reducing false positives was in diminishing the system's ability to figure the origin of bullets that do not come within 30 meters of a moving vehicle or 50 meters of a stationary one. But a soldier doesn't care about gunfire that misses by such a wide margin, says Schmitt. And Boomerang hasn't registered a single false alarm in the past two years in Iraq and Afghanistan, he claims.

At the same time, BBN began testing a bullet-detection system that can be mounted on helicopters. Such a system is much more complicated, in part because a helicopter can move much faster and in more directions. Bullets are so fast that Boomerang needn't factor in the movement of a Humvee, even when traveling at 100 kilometers per hour. But a helicopter's greater speed can't be ignored, so the new system uses accelerometers to factor it in. A single set of sensors arranged on a mast strapped to the top of a vehicle serves all types of cars and trucks, but this won't work on a helicopter, which can be shot at from below as well as above. The microphones have to be carefully positioned, and each model will require a custom configuration. Finally, the new system uses only the shock wave—and more sophisticated algorithms—not the muzzle blast, which cannot be reliably measured from a helicopter.

Last June and then again in August, a prototype was tested at Fort Rucker, in Alabama. Previous tests had collected data to be processed off-line back at BBN; these were the first to process data on board. Both BBN and the Department of Defense declined to comment on their outcomes, but Schmitt says that he believes a deployable version will be ready sometime in 2009.

—STEVEN CHERRY

A Battery-Capacitor Hybrid—for Hybrids

Engineers give lead-acid batteries a makeover by crossing them with ultracapacitors

BATTERY TEST: A hybrid-electric vehicle is put through its paces using a combination ultracapacitor/lead-acid battery.

PHOTO: ADVANCED LEAD-ACID BATTERY CONSORTIUM/CSIRO

LEAD-ACID BATTERIES are relics that haven't changed much since their invention nearly 150 years ago. Heavy and unable to withstand rapid charge-discharge cycles, they are unsuitable for the automotive world's killer app, hybrid-electric vehicles. Hybrids instead use expensive nickel-metal hydride (NiMH) batteries or, experimentally, lithium batteries. But a new, souped-up version of lead-acid batteries could change that, cutting the cost of hybrids and also improving the function of power grids and a range of other applications.

The new design combines lead-acid chemistry with ultracapacitors, energy-storage devices that can quickly absorb and release a lot of charge, which they store along the roughened surface of their electrodes. Unlike ordinary lead-acid batteries, which are slowed by the movement of chemicals within them, these could provide quick bursts of power for acceleration and then recharge during braking, a must for hybrid-electric and electric vehicles. A hybrid's rapid recharging cycles and high currents would destroy the lead electrodes in standard batteries, because lead sulfate would build up on them. The new batteries can go through at least four times as many charging cycles as lead-acid batteries, and, crucially, would cost about a quarter of NiMH batteries.

At least two lead-acid/ultracapacitor technologies are now poised for market release. Battery giant East Penn Manufacturing Co., in Lyon Station, Pa., licensed the technology for the UltraBattery in September from Furukawa Battery, in Yokohama, Japan, which has already begun manufacturing the devices. Researchers at Australia's Commonwealth Scientific and Industrial Research Organisation (CSIRO), who invented the UltraBattery, tested it early this year in a Honda Insight hybrid, which ran for 160 000 kilometers.

Meanwhile, Axion Power International, based in New Castle, Pa., has developed a slightly different design, which it will test for U.S. Marine Corps assault vehicles; the company got US $1.2 million from the Department of Homeland Security in October for the tests. A bank of 1000 of Axion's batteries will also soon be tested as a utility-grid buffer in upstate New York. Axion CEO Thomas Granville says that the new technology "lets us get into markets that have been in the past closed to lead-acid batteries."

The new batteries' advantage over standard lead-acid batteries comes from simple tweaks of the negative electrode. Instead of a lead plate, Axion makes the electrode from activated carbon, the highly porous, spongelike material used in ultracapacitor electrodes. When a regular battery discharges, the lead electrode reacts with sulfate ions, forming lead sulfate and creating protons and electrons. Axion's activated carbon electrode directly releases and adsorbs protons from the sulfuric acid electrolyte during discharging and charging. The batteries recharge four times as fast as conventional ones, Granville says.

The UltraBattery is slightly different, says Lan Lam, project manager of the battery work at CSIRO. The negative electrode is split into two, one half made of lead and the other half of activated carbon. The two halves are connected in parallel so that their currents combine. This split-electrode design gives the battery the best of both technologies, according to Lam. While activated carbon provides quick energy bursts, it cannot store as much energy as the lead-acid chemistry. The combination gives the UltraBattery an energy capacity closer to that of a lead-acid battery than an ultracapacitor could get alone, Lam says.

Both designs have a big cost advantage. "Nickel-metal hydride, depending on the application, is as much as $800 to $1200 per kilowatt-hour," Granville says. "Axion's battery costs $200 per kilowatt-hour."

These battery/ultracapacitor combinations will have to compete with lithium-ion batteries as the successor to NiMH for hybrid vehicles. Cost and safety, however, are still a concern for lithium. Lithium-ion batteries can overheat, ignite, and even explode if mistreated.

The lead-acid/ultracapacitor batteries have other advantages. They are easier to recycle than NiMH or lithium, according to East Penn. Lithium-ion batteries don't have much usable metal, so they are usually incinerated, while the nickel from NiMH batteries is consumed in the steel industry. The military, meanwhile, is interested in Axion's batteries, not so much for hybrids but because they work at temperatures as low as –50° C and weigh less than standard lead-acid batteries, Granville says.

—PRACHI PATEL-PREDD

update

Tiny Implants Combat Chronic Pain

Spinal pacemakers mask discomfort from nerve damage

BACK IN 2006, Adam Hammond, a U.S. Army skydiver, experienced every jumper's worst nightmare when his parachute failed to deploy. "I was basically in free fall," Hammond says of his accident. "When I hit the ground, my helmet shattered and my shoes flew off my feet." Hammond was completely immobilized by chronic pain for two years. Earlier this year he received an implanted device that electrically stimulates his spinal cord. Now, instead of feeling a stabbing pain in his tailbone, he experiences just a tingling sensation. "My life has done a complete 180," he says.

SIZE MATTERS: Shrinking pain implants could increase the size of their market. PHOTO: ST. JUDE MEDICAL'S ANS DIVISION

Smaller, longer-lasting implants are broadening the appeal of pain management devices for patients who have not been well served by conventional medications. With smaller sizes, surgeons have more flexibility with where to place the implant. And with better techniques for transferring and storing energy, the implants can last longer and be placed deeper in the body, which increases their cosmetic appeal.

The device that Hammond received, Advanced Neuromodulation Systems' Eon Mini, is the smallest neurostimulator on the market. According to ANS, of Plano, Texas, depending on the power output used to block pain, patients can go between one week and a few months before the implant needs to be recharged—wirelessly—and the battery is expected to last 10 years. The device delivers the stimulation through 16 electrodes, which a physician can adjust individually to produce pulses of different intensities and frequencies. A patient can control the stimuli with an inductively coupled programming wand.

"What's unique about pain is that it's entirely subjective," says Tom Hickman, the vice president of product management at ANS, a division of St. Jude Medical. "What feels good for one patient won't feel good for another. But the outcomes, the return of activity and quality of life, are objective."

Physicians estimate that each year as many as 40 000 people worldwide receive spinal-cord-stimulation implants. That's a small fraction of the roughly 5 million patients who might benefit from the treatment, say analysts at Medtech Insight, a market research firm in Irvine, Calif. But the mechanisms of why the therapy works are not well understood, and only about half of the implant recipients experience some relief from pain, says Linda Porter, a program director at the National Institute of Neurological Disorders and Stroke, in Bethesda, Md. The uncertainty, coupled with the approximately US $20 000 price tag, have made many physicians and hospitals wary.

Though surgeons have been using these stimulators for pain for about 30 years, technical innovation is fueling new research, and stimulators are currently the focus of a half dozen clinical trials in the United States.

Two factors have limited the efficacy of the device: the lifetime of the battery and the implant's overall size. Occasionally, patients report that the pulse generator can shift and erode through the skin. The electrodes, too, can move, and their leads can break under stress from the frequent movements of the back. Few medical centers are prepared to handle migrating neurostimulators, which means patients might need to travel long distances for surgery to adjust the implant.

"The size of the battery is getting smaller, and with that, the less painfully [the device] nicks away under the skin," says Krishna Kumar, a neurosurgeon at Regina General Hospital, in Canada. Improvements in the devices power supplies are helping, too. "Three years ago, you had to throw the battery away every two years or so," Kumar says.

Before a recent redesign, ANS's implant consisted of a titanium can encasing the electronics, with a charging coil attached to the outside of the can. Like most neurostimulators, it was bulky.

The key to shrinking the device was to use both a smaller battery and a smaller charging coil, which was tucked inside the casing. Then the new coil was tuned so that it could receive energy efficiently through the titanium-alloy can without slowing down the battery's recharging rate or sacrificing the depth at which the implant could be embedded. The new device is about half the thickness of its predecessor.

Though spinal cord stimulation has been approved in the United States to treat people with chronic limb pain and those who have had unsuccessful back surgery, the therapy is also used in Europe for patients with vascular disease and angina and is now being explored for other conditions, including migraines and obesity. —SANDRA UPSON

the big picture

A NEW SPIN ON ARCHITECTURE

It's your LegoLand building-block fantasy come to life. When the twisted tower in this rendering is completed in Dubai in 2010, each of its 80 floors will rotate independently around a central column that will house its elevator shafts, stairwells, and utilities. Sections of each floor, complete with cabinets, electrical wiring, and bathroom fixtures, will be preassembled off-site, hoisted up, and hung, intact, on the central column. In the gaps between each floor, wind turbines will spin; these and solar panels on the edges of each floor's partially exposed roof will generate enough power to make the 420-meter-high rotating skyscraper independent of the energy grid. The curvy, statuesque beauty will house offices, a luxury hotel, apartments, and a few full-floor villas with space enough for residents to park their cars safely inside.

IMAGE: ROTATING TOWER TECHNOLOGY INTERNATIONAL

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

technically speaking BY PAUL McFEDRIES

Trafficking in Words

Traffic is the expression of human purpose.
—**Tom Vanderbilt**, *Traffic*

ON RESIDENTIAL streets, engineers attempt to control car speeds by installing **traffic-calming devices**. You probably know all about the venerable **speed bump**, but these days you're more likely to drive over the lower and wider **speed hump** or the even wider **speed table**. Many people refer to these generally as **sleeping policemen**, a richly evocative, perhaps even poetic locution.

In 1861, when the few steam-powered automobiles around were known as road locomotives (or light locomotives), the British Parliament established the world's first speed limit of 10 miles per hour (16 kilometers per hour). In 1865 a revised law slashed it to 4 mph in the country and 2 mph in towns and villages. Apparently alarmed by even these modest speeds, Parliament also decreed the world's first traffic-calming devices: each vehicle had to be preceded by a person walking 60 yards (about 55 meters) ahead and waving a red flag to warn others of the vehicle's approach.

The world has seemed like a too-fast place ever since, and today's traffic engineers spend much of their time getting people to slow down. Along the way they've developed a remarkable lingo.

One traffic-calming technology popular of late is the **neckdown** (also called a **bulbout**), which is a curb that extends into the street on both sides of an intersection; narrowing the intersection forces cars to slow down. A similar idea is the **chicane**, an S-shaped curve that creates **curb extensions** drivers must negotiate.

EARL ELIASON/ISTOCKPHOTO

Speed bumps, humps, and tables fall under the category of **vertical deflection**, while neckdowns and chicanes are **horizontal deflections**.

For many people, anything that forces motorists to slow down is a good thing. Hence the popularity in some places of the **mobile speed bump**, which is a car that travels at the speed limit to force the cars behind to do the same. Tom Vanderbilt, in his terrific book *Traffic: Why We Drive the Way We Do (and What It Says About Us)*, mentions that the notoriously reckless drivers of New Delhi are often forced to slow down thanks to a natural traffic-calming device: the occasional cow lounging in the middle of the road.

Traffic-calming features aren't universally loved, however. They make driving a hassle; moreover, some drivers accelerate between humps in an apparent attempt to make up for lost time, a problem for pedestrians. First responders hate the way they are slowed in an emergency. (One novel solution to this last complaint is the **speed cushion**. Picture a single speed bump split into three smaller bumps; the spaces on either side of the middle bump are just far enough for the wheels of emergency vehicles to pass through, while narrower cars must still negotiate at least one of the bumps.)

To solve these problems, traffic engineers are coming up with novel ideas that create **psychological traffic calming**. For example, large trees near the side of the road, lowered curb heights, and even a child's bicycle parked by the roadside create uncertainty in the mind of the driver, and that uncertainly causes the driver to slow down. In other words, it creates a kind of **mental speed bump** in the mind of the driver.

Some engineers have taken this idea to extremes and created **naked streets**, which have no signs, road markings, or traffic lights. This sounds like a recipe for car chaos, but by forcing drivers to pay attention to their surroundings, the resulting streets are safer than traditional streets festooned with lights, signs, and other **street furniture**.

A similar idea is the **complete street**, a street designed to accommodate cars, public transit, bicycles, and pedestrians. A common feature is the **pedestrian scramble**, a traffic-light-and-crosswalk system that stops cars in all four directions.

In 2007 more than 41 000 people died in traffic-related accidents in the United States alone, and that's the *lowest* number in recent years. We need more sleeping policemen, more naked streets, and, perhaps, armies of volunteers walking 55 meters in front of each car, waving red flags. ❑

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

careers

FLY FERNANDEZ/ZEFA/CORBIS

THE EE GENDER GAP IS WIDENING

Electrical engineering faces an age-old question: What do women want?

WALK INTO a classroom of environmental engineering students and, odds are, nearly half of them will be women. Now head next door to an electrical engineering class: you'll likely find eight men for every woman.

The failure to recruit and retain more women in electrical and computer engineering—large fields with lots of students—is a major reason the representation of women in U.S. engineering as a whole has remained so low for so long. Last year, only 18.1 percent of all bachelor's degrees in engineering awarded by U.S. schools went to women. And things are getting worse: that's the lowest level in more than a decade.

Electrical engineering has one of the lowest proportions of bachelor's degrees awarded to women, a meager 12.4 percent last year, down from an already low 14 percent the year before, according to the American Society for Engineering Education. It's not the lowest—computer engineering, at 9.2 percent, is at the bottom—but its proportion of women is less than a third that of biomedical or environmental engineering.

Experts see in these differences a larger pattern: "Women seek areas where the societal benefits are very apparent," says Diane Matt, executive director of the Women in Engineering Proactive Network, in Denver. "They want careers that have a positive impact on the world."

Not that there aren't social benefits to all areas of engineering—the problem is one of perception. Eleanor Baum, dean of engineering at Cooper Union, in New York City, and an IEEE Fellow, says electrical engineers are not sending a compelling message about their profession. "Instead of trying to explain what's a power engineer, what's a computer engineer, and so on, we should focus on a single coordinated message: electrical engineers do this, and this affects people's lives."

Karen Panetta, an electrical engineering professor at Tufts University, in Medford, Mass., and editor in chief of *IEEE Women in Engineering Magazine*, has found a way of doing just that. She created the organization Nerd Girls, which helps women students develop projects involving energy and the environment. "It works because the young women can immediately understand the importance of what they're doing," she says. "Bringing renewable energy to places like an island or an emerging country are achievements that everyone can be proud of."

"While many areas within EE make clearly positive contributions to society, we don't market that aspect of our field well enough," says Mary Baker, an electrical and computer engineering professor at Texas Tech University, in Lubbock. "It is up to us to show that electrical engineering does indeed have a more human side." Baker and her colleagues have organized EE-specific summer camps for female high school juniors as well as summer internships for community college and high school students that focus on biomedical and other topics known to appeal to women.

MIT has a similar program. Its Women's Technology Program in Electrical Engineering & Computer Science holds summer courses for high school girls with team-based projects and hands-on classes taught by female MIT students.

Other programs try to establish a more direct and durable connection between young women and EE professionals. MentorNet and MAGIC (More Active Girls in Computing) help connect high school girls with women who have careers in technology. And STAR (Student-Teacher and Research Engineer/Scientist), created by the IEEE Women in Engineering program, promotes involvement of IEEE members with high school teachers and students.

Will female participation in electrical engineering ever equal that of disciplines like chemical and environmental engineering? One thing is certain: the first step is to remove that glaring gender disparity in the EE classroom. —ERICO GUIZZO

For more statistics and resources on women in engineering, see http://www.spectrum.ieee.org/dec08/women.

DEALING WITH DIFFICULT PEOPLE

You can't get rid of them, so you'd better learn to get along with them

A ROCK NAMED JUDY

The ancient Greeks named the stars after heroes. Now the Intel International Science and Engineering Fair is doing it too. That's how a lump of rock numbered 23155 became Asteroid Judithblack, in honor of a 19-year-old student who won a US $1500 prize for her scheme to improve the efficiency of an automated box-cutting machine in a factory near her home in Warrenpoint,

Northern Ireland.

Black, who is now an EE major at Queen's University, in Belfast, says getting an asteroid named for her "was better than the money!"

The next Intel ISEF will be held 10 to 15 May 2009 in Reno, Nev. (see http://societyforscience.org).

—*Susan Karlin*

THE MERE mention of "difficult people" brings knowing looks to our faces. These people won't go away, so you need to deal with them effectively before they cause any real damage. Easier said than done!

First, learn to recognize difficult people. Some key characteristics: they're never satisfied, they don't communicate easily with others, they can be argumentative and unreasonable, and they often use their power to obstruct.

You can't change them, but you can change the way you deal with them.

Back when I worked at an airport, I needed approval from Fred—who was always hard to pin down—on a pricing strategy. He wouldn't get back to me, so I drafted a memo to force him to decide and placed it on his chair before he arrived at work. When I phoned him for his reaction and he said he hadn't received the memo, I asked him to get up off his chair—look at what you're sitting on! A potentially irritating situation turned into laughter. The lesson here is *be overly nice and cooperative with difficult people. Let them bother others, not you.*

I once got a call from an airport manager asking me how I'd gotten a fellow named Cliff, famous for avoiding new tasks, to volunteer for one. Cliff represented his unit on a task force I led. He said that a certain off-airport baggage check-in procedure "couldn't be done," so I told him that an airline was actually doing it and suggested that he might call the airline to find out how. Cliff agreed to do it—and sure enough, he did. My answer to his manager's question was to *assume that even a difficult person is really trying to do the right thing. Appeal to the better angels of his nature.*

Another difficult person was Dave, a business rep who oversaw my projects. When I first met him to discuss my responsibilities, he bluntly pushed back: "What the heck do you know about business development?" (Dave actually used stronger language.) I needed to keep my eye on Dave, so I asked his secretary, whom I knew, to give me a heads up if Dave seemed mad at me.

Sometime later she called: "Dave is on the warpath" about one of my phone-card vending machines, she said. I first called our phone-card contractor to see if he could meet at the airport that day; then I called Dave on the pretext of setting up a meeting. Sure enough, I got an earful about the locations of newly installed machines. I asked him if it would be okay to visit him that day along with the contractor, who could fix the problem. We did, and Dave was very pleased. The solution here was to *keep your eyes and ears open to head off situations with difficult people before they can develop into crises.*

Often, a person you may not work well with gets along just fine with others. I had a strained relationship with one airport executive until I took a seminar that analyzed personality types according to a test called the Myers-Briggs Type Indicator. It helped me understand what motivates different people, including myself. The executive was a numbers-oriented person, and so I started giving him more facts and figures. Our relationship improved immediately. My manager glowed and said, "That stuff really works!" *If you understand yourself, it becomes easier to understand—and accommodate—others.*

That brings us to our last tip: *look in the mirror.* In the middle of my career, my new manager told me she learned I had a bad reputation—people in other departments had said I was too hard on them. This was news to me. She counseled me to work more closely with each department. "And Carl—be nice," she said. I needed to remember the lesson of Pogo, the cartoon character who once observed, "We have met the enemy, and he is us." *Don't be a difficult person yourself.* —CARL SELINGER

TOP: ADRI BERGER/GETTY IMAGES; BOTTOM: GARY CORNHOUSE/GETTY IMAGES

media

THE JURASSIC, LIVE AND IN PERSON

A traveling troupe of life-size dinosaurs tour the United States

BABY STEPS: The show's only human character, a paleontologist-cum-narrator, looks on as a mother *T. rex* protects her young in a scene from *Walking With Dinosaurs: The Live Experience*.

PHOTO: JOAN MARCUS/GLOBAL CREATURES

AS FAR as my two *Jurassic Park* fanatics are concerned, dinosaurs are a current fact. It's just a matter of time until Coby, 6, pulls his own lab together to create his own raptor. Sam, 3, pleads for him to stick to "plant eaters."

So from the kid point of view, seeing 15 life-size dinosaurs fill New York City's Madison Square Garden wasn't very different from watching the New York Knicks come running out of the locker room. It was Mom whose breath was taken away by *Walking With Dinosaurs: The Live Experience*, a show that began touring the United States this past spring and is slated to travel throughout North America for the rest of 2009.

Billed as the world's largest puppet show, *Walking With Dinosaurs* is an electrical and mechanical engineering marvel. True, some of the smaller dinosaurs are powered by a puppeteer wearing his creature's suit, and an *Ornithocheirus* with a wingspan of 12 meters flies overhead. But two brachiosaurs, an adult *Tyrannosaurus rex*, and seven other enormous creatures, as heavy as 1400 kilograms, are driven by car-size sleds. An essential sense of scale—the sole reminder of just how massive these creatures are (excuse me, were)—is provided by a lone human actor playing a paleontologist master of ceremonies.

A team of humans—some onstage, some off—control the largest dinosaurs. An engineer drives each sled and is also responsible for watching all dino systems, including a tangle of hydraulics run on 10- and 12-volt dc engines powered by up to 12 truck batteries. Once the dinosaurs are onstage, their legs are controlled by microprocessors that run preset walking sequences and also control how the dinosaurs bob up and down as they walk. Facial expressions and smaller movements are controlled remotely by puppeteers radio-linked to their characters from a platform set up in the audience. Unbeknownst to the audience, these dinosaurs, framed in steel and covered in latex, are also inflated by continuously running fans.

Even jaded New Yorkers collectively *aww*-ed when the daddy brachiosaur gave his baby a neck snuggle. We cheered when the mom *T. rex* saved her baby from a hungry group of raptors. And much to Sam's relief (and Coby's outright dismay), nobody actually got eaten.

This is mechanical life imitating computer life imitating biological life, says Tim Haines, producer of the television documentary *Walking With Dinosaurs,* which inspired the live show. Haines came from Australia, where the show originated in 2007, to Madison Square Garden with his wife and five children to gape and gasp with the rest of us. He says that *Jurassic Park*—the movie that convinced my kids that dinosaurs exist now—did indeed change the state of the art in dinosaur representations. Television, Haines says, had to follow, leading eventually to the "live" dinosaurs that pounded the arena floor.

One mechanical decision in particular contributed to the sense that these creatures were indeed alive. Because large hydraulic mechanisms tend to bounce back with force when they are stopped—making evident the steel structures inside—the show's designers let the uncontrollable be, well, uncontrolled. For example, a head may be turned to the right with an initial thrust, but then the head is left to drift in that direction. The final stopping point may be inexact, but there is no shake.

The entire show was designed and built in less than a year, in the process giving birth to the Creature Technology Company, in Melbourne. These dinosaurs were like nothing else its designers had ever built, says chief engineer Trevor Tighe. The state of the art in animatronics is usually focused on a few bits of an alien or monster for close-ups. The rest is usually generated by computer. When a whole creature is built, it is rarely larger than your average dog.

Working on such a spectacular scale has led to a few spectacular failures. "We had a *T. rex* head collapse in the first few weeks in Sydney," Tighe says. Although there have been other embarrassing moments on stage, "nobody asks for their money back," he says. —SHERRY SONTAG

New performances are booked regularly; check http://www.dinosaurlive.com.

tools & toys

MacBook Air

SMALL IS BIG IN NOTEBOOKS—BUT NOT TOO SMALL

Ultralight computers add back a few more ounces and a lot more usability

A LOT OF road warriors have learned, much to their chagrin, that a notebook computer can indeed be too light or too thin. Shrink one much below 1.8 kilograms (4 pounds)—the class known as ultralight—and it gets maddeningly hard to use for much more than e-mail or Web browsing.

That's why manufacturers have recently begun reemphasizing a slightly larger breed—call them not-quite-ultralights. These machines sport displays of at least 13 inches; fuller, if not full-size, keyboards; built-in optical drives; Ethernet and other ports; and more processing power than their rail-thin predecessors.

"You can use a 13.3-inch screen full time," says John Jacobs, a research director at DisplaySearch, a display market research and consulting company based in Austin, Texas. "I don't see people with smaller ones at their desk unless they have an external monitor."

But an ample monitor isn't enough. Apple's elegant MacBook Air, with its 13.3-inch screen and full-size keyboard, might win "best biggish screen and full-size keyboard in its class"—if you don't mind the lack of a removable battery, Ethernet jack, or optical drive.

The belle of the ultraportable ball is, arguably, Lenovo's ThinkPad X300/X301 line, with a weight starting at just under 1.4 kg, a built-in optical drive, a full array of ports, and the same 13.3-inch screen that larger ThinkPads have. The 140-gram (5-ounce) increase in size and bulk over a comparable ThinkPad X61s is more than made up for by the increased usability, which translates to greater productivity.

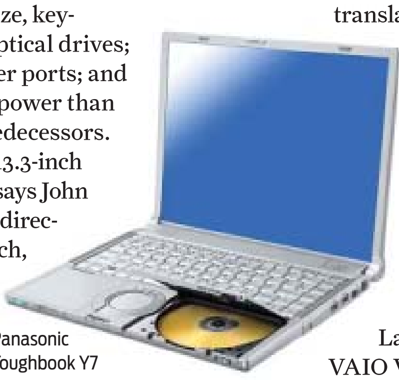

Panasonic Toughbook Y7

Ditto for the Panasonic Toughbook Y7, with its 14.1-inch screen, versus its smaller sibling, the Toughbook W7. Other noteworthy contenders in this weight and screen-size class include Fujitsu's LifeBook S6510/6520, Dell's Latitude E4300, and Sony's VAIO VGN-SZ650N/C.

The new category of ultralights stems partly from some new technologies and partly from some existing ones that have matured—plus some hard design work to minimize the productivity compromises. Howard Locker, director of new technology, desktop and mobile systems at Lenovo, says, "There's no perfect answer. It's what's the best compromise between thinness and lightness, battery life and functionality that gives the best experience."

Here are some of the new technologies:

Displays now get their backlighting from LEDs rather than thicker fluorescent tubes. Jacobs says that slimming a display can reduce power consumption by up to 20 percent, extending a typical 4-hour battery's life to almost 5 hours.

Dell Latitude E4300

Solid-state drives weigh less and use less power than hard disks, adding, to be sure, several hundred dollars to the purchase price. "The affordable SSD is still a few years away," says Jacobs. "But if performance is more important than cost, customers will pay for it today."

Lower-power microprocessors, which incorporate a panoply of power-management features and sometimes integrate graphics, also bring size down and battery life up.

Wi-Fi based on the new IEEE 802.11n standard uses power more efficiently than 802.11a/b/g. And manufacturers can now embed 3G and other broadband chips and antennas, yielding better

Lenovo ThinkPad X300

performance and consuming less power.

Since ultralights are meant for travel, durability counts. New "drip channels" protect key components from spilled liquids. Some makers are replacing plastic chassis with magnesium alloys, while Lenovo's X300 includes a carbon fiber "roll cage" case.

Even power supplies have gotten smaller and lighter, such as Lenovo's card-deck-thin Slim AC/DC Combo Adapter, half the size of its predecessor. Also, many manufacturers offer feather-weight three-cell batteries as an alternative to longer-lasting six-cell ones.

So what's next?

Expect more storage capacity and more solid-state drives as memory prices continue to fall, predicts Jacobs. On the other hand, don't get your hopes up yet for those nice, bright, thin organic LED screens. "It will be a year and a half to two years before we see an OLED display in a notebook platform," says Jacobs. And don't expect a notebook in this class with flexible roll-out screens. "They're great for mobile phones or e-books," he says. "The problem is that the response time is slow." Lastly, look for further design improvements that go toward reliability and usability—for example, how the display opens when you're in a cramped airplane seat in coach.

So consider buying a slightly bigger machine. Its extra ounces won't send you screaming to the chiropractor, and it sure will be easier on your eyes, fingertips, and temper. —Daniel P. Dern

hands on

A Poor Man's Tesla

MIT students electrify a classic sports car

IF YOU'RE itching to own a sexy electric roadster but can't afford the US $109 000 price tag for Tesla Motors' sleek new model, there's another option you might want to consider: building your own.

Two years ago, Yang Shao-Horn, a professor in MIT's department of mechanical engineering and head of its electrochemical energy laboratory, bought a 1976 Porsche 914 roadster on eBay for about $5000. She then donated it to her students so they could electrify the car, which is a popular model among hobbyists for electric conversion.

The budding engineers put together a vehicle that, like Tesla's acclaimed sports car, uses lithium-ion batteries (about $36 000 worth, donated by Valence Technology, of Austin, Texas) as well as an ac motor and controller. And like the Tesla roadster, which is based on the Lotus Elise, the MIT vehicle should inherit the superb handling of a midengine sports car. We haven't seen its top speed yet—the students are only now preparing to carry out high-speed highway tests—but they have calculated the expected performance, and it's very respectable for a DIY conversion.

If you can spare about $60 000 and are willing to put in some long hours in the garage, you too could have such a set of wheels. You'd still be smoked coming out of a stoplight by any Tesla driver, but you could probably turn just as many heads with your retro electric roadster. And the car's 210-kilometer range is surely enough to take you, laughing, all the way to the bank. —*Steven Cherry & David Schneider*

A Tesla on the Cheap

	Tesla Roadster	MIT Electric Porsche 914
Cost	US $109 000	$60 000 plus labor
Seats	2	2
Motor	185-kilowatt ac induction	43-kW ac induction
Weight	1220 kilograms	2500 kg
Battery pack	53 kilowatt-hour lithium-ion	24 kWh lithium-ion
Acceleration	4 seconds (0 to 100 km/h)	23 seconds (0 to 100 km/h)
Range	350 kilometers	210 km

PORSCHE: DONNA COVENEY/MIT

BATTERIES *ARE* INCLUDED: Students at MIT take their electric Porsche 914 for a spin on campus; team leader Irene Berry walks alongside. ***Inset***: A state-of-the art battery array nearly fills the Porsche's trunk.

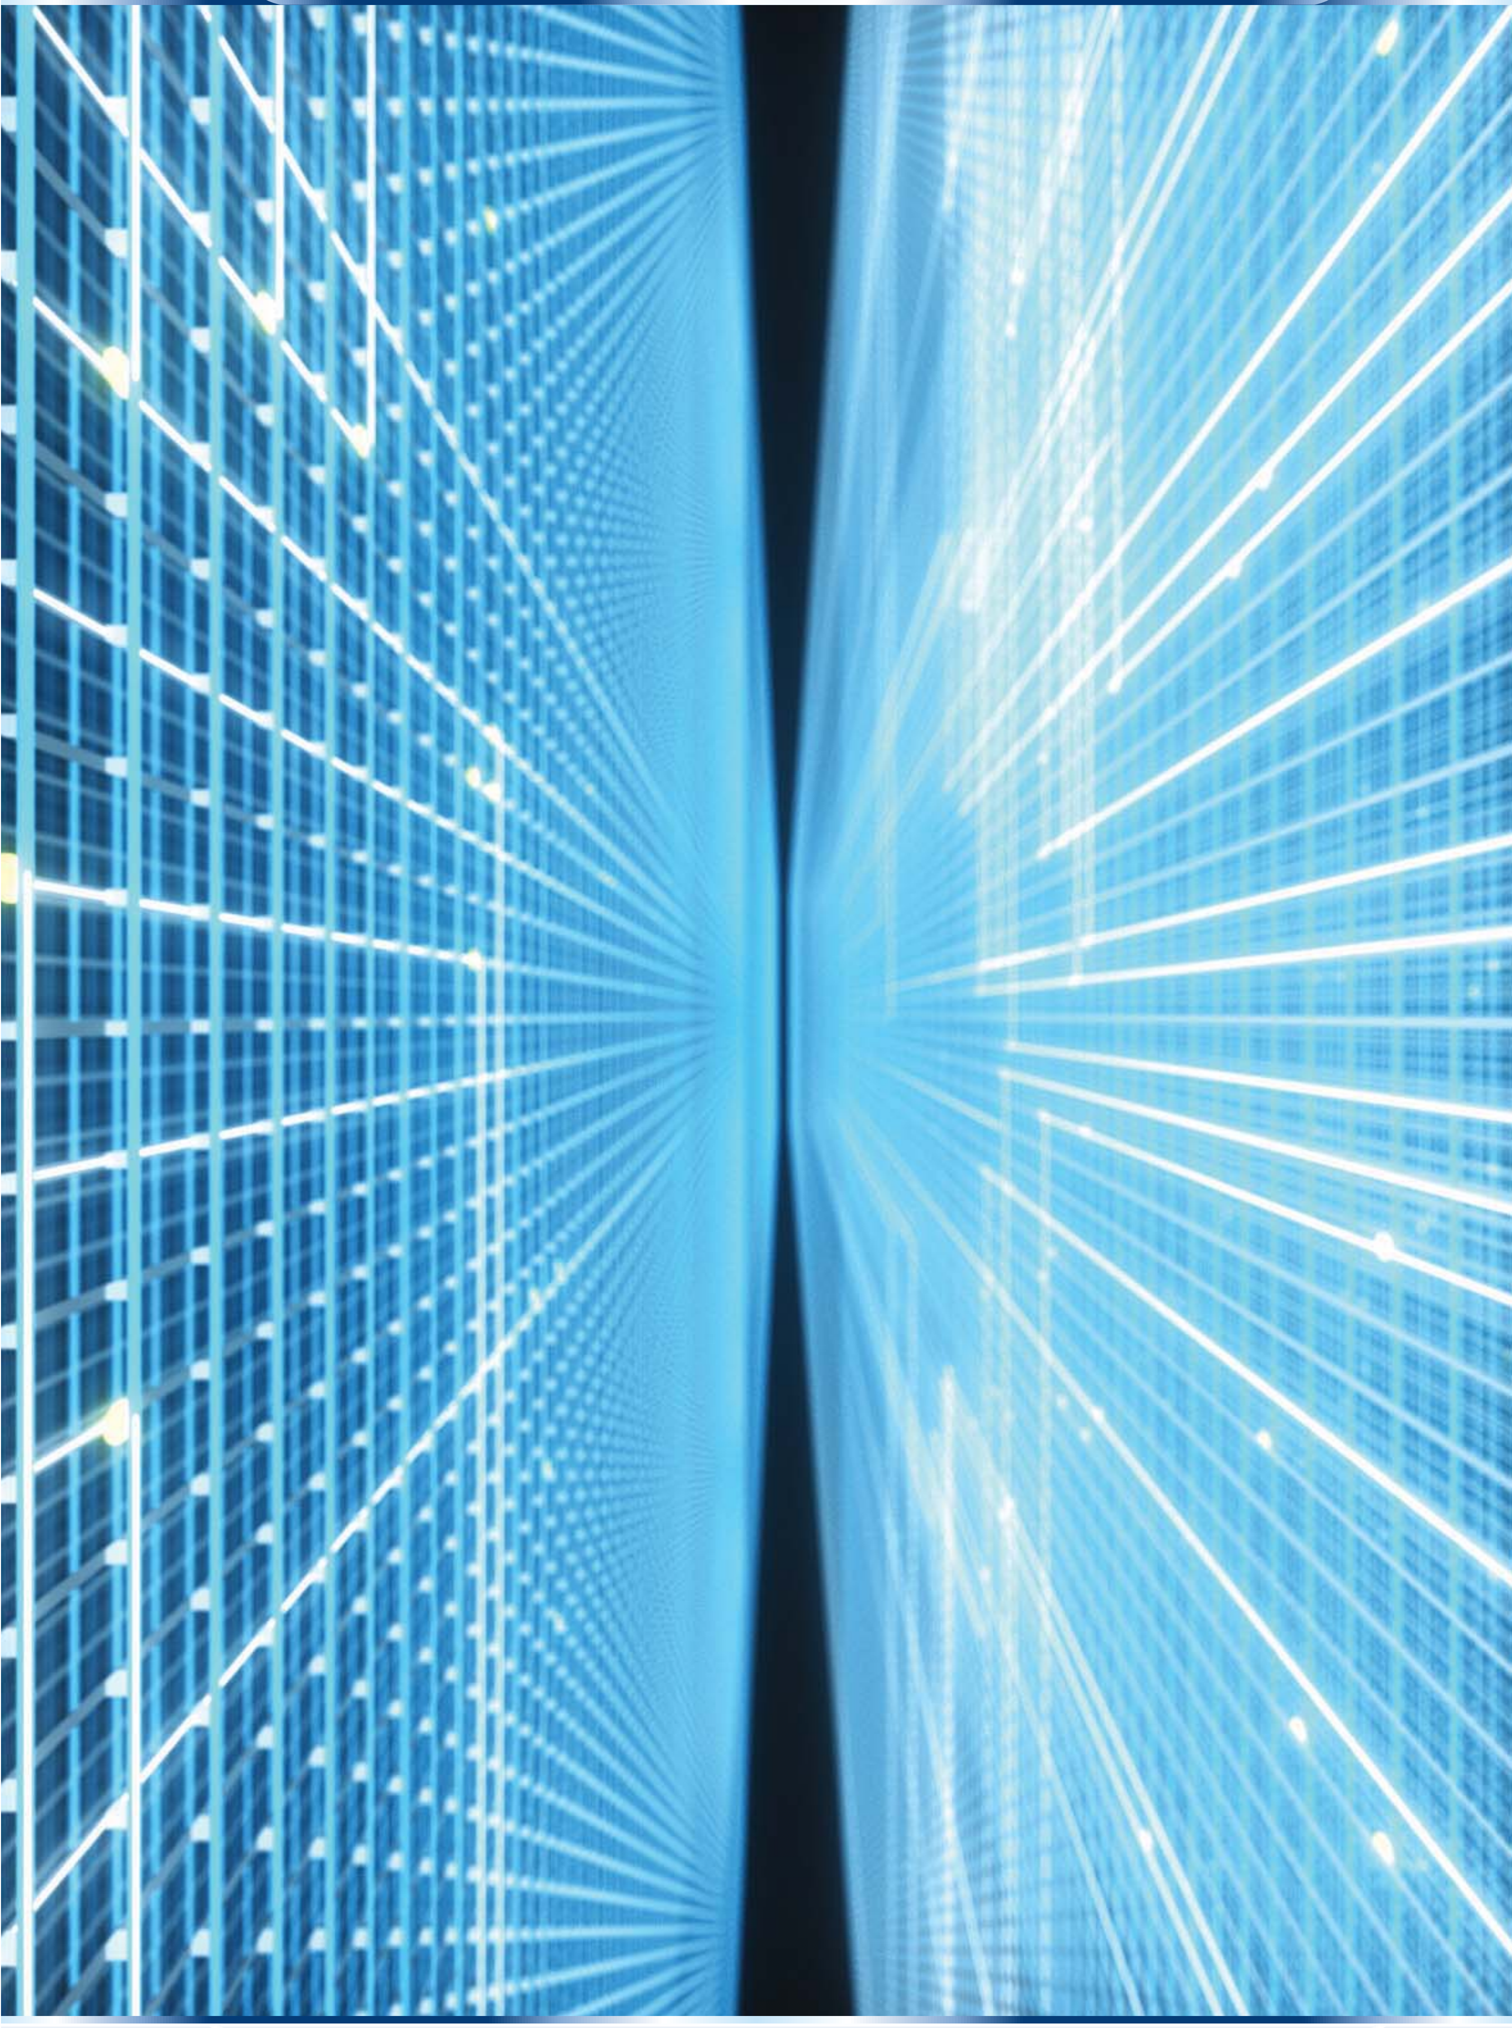

HOW WE FOUND THE MISSING MEMRISTOR

The memristor—the functional equivalent of a synapse—could revolutionize circuit design *By R. Stanley Williams*

It's time to stop shrinking. Moore's Law, the semiconductor industry's obsession with the shrinking of transistors and their commensurate steady doubling on a chip about every two years, has been the source of a 50-year technical and economic revolution. Whether this scaling paradigm lasts for five more years or 15, it will eventually come to an end. The emphasis in electronics design will have to shift to devices that are not just increasingly infinitesimal but increasingly capable.

Earlier this year, I and my colleagues at Hewlett-Packard Labs, in Palo Alto, Calif., surprised the electronics community with a fascinating candidate for such a device: the memristor. It had been theorized nearly 40 years ago, but because no one had managed to build one, it had long since become an esoteric curiosity. That all changed on 1 May, when my group published the details of the memristor in *Nature*.

Combined with transistors in a hybrid chip, memristors could radically improve the performance of digital circuits without shrinking transistors. Using transistors more efficiently could in turn give us another decade, at least, of Moore's Law performance improvement, without requiring the costly and increasingly difficult doublings of transistor density on chips. In the end, memristors might even become the cornerstone of new analog circuits that compute using an architecture much like that of the brain.

For nearly 150 years, the known fundamental passive circuit elements were limited to the capacitor (discovered in 1745), the resistor (1827), and the inductor (1831). Then, in a brilliant but underappreciated 1971 paper, Leon Chua, a professor of electrical engineering at the University of California, Berkeley, predicted the existence of a fourth fundamental device, which he called a memristor. He proved that memristor behavior could not be duplicated by any circuit built using only the other three elements, which is why the memristor is truly fundamental.

Memristor is a contraction of "memory resistor," because that is exactly its function: to remember its history. A memristor is a two-terminal device whose resistance depends on the magnitude and polarity of the voltage applied to it and the length of time that voltage has been applied. When you turn off the voltage, the memristor remembers its most recent resistance until the next time you turn it on, whether that happens a day later or a year later.

Think of a resistor as a pipe through which water flows. The water is electric charge. The resistor's obstruction of the flow of charge is comparable to the diameter of the pipe: the narrower the pipe, the greater the resistance. For the history of circuit design, resistors have had a fixed pipe diameter. But a memristor is a pipe that changes diameter with the amount and direction of water that flows through it. If water flows through this

THINKING MACHINE: This artist's conception of a memristor shows a stack of multiple crossbar arrays, the fundamental structure of R. Stanley Williams's device. Because memristors behave functionally like synapses, replacing a few transistors in a circuit with memristors could lead to analog circuits that can think like a human brain. IMAGE: BRYAN CHRISTIE DESIGN

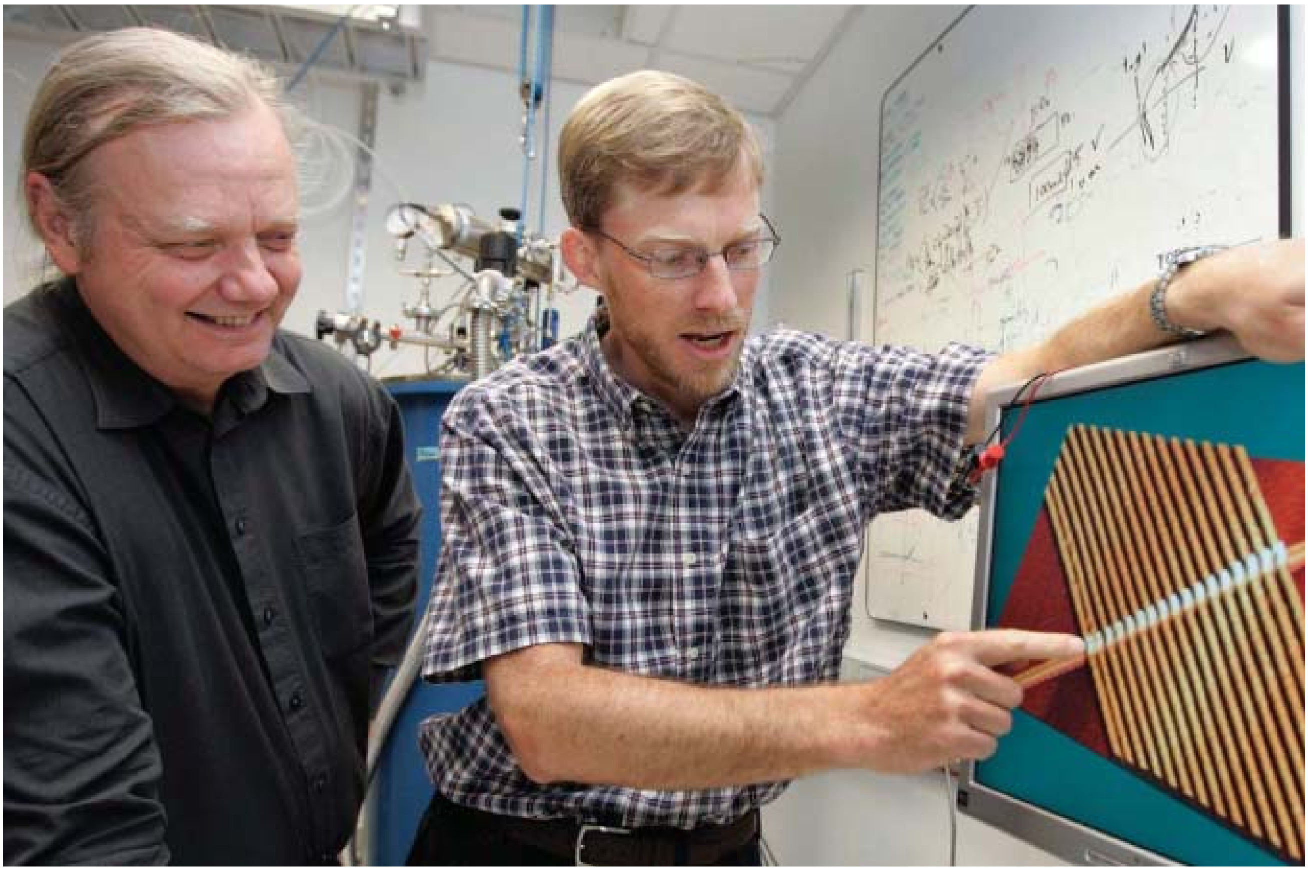

PICTURING MEMRISTANCE: HP Labs senior fellow R. Stanley Williams [left] and research physicist Duncan Stewart [right] explain the fourth fundamental circuit element. Williams worked with nearly 100 scientists and engineers to find the memristor. PHOTO: PAUL SAKUMA/AP PHOTO

pipe in one direction, it expands (becoming less resistive). But send the water in the opposite direction and the pipe shrinks (becoming more resistive). Further, the memristor remembers its diameter when water last went through. Turn off the flow and the diameter of the pipe "freezes" until the water is turned back on.

That freezing property suits memristors brilliantly for computer memory. The ability to indefinitely store resistance values means that a memristor can be used as a nonvolatile memory. That might not sound like very much, but go ahead and pop the battery out of your laptop, right now—no saving, no quitting, nothing. You'd lose your work, of course. But if your laptop were built using a memory based on memristors, when you popped the battery back in, your screen would return to life with everything exactly as you left it: no lengthy reboot, no half-dozen auto-recovered files.

But the memristor's potential goes far beyond instant-on computers to embrace one of the grandest technology challenges: mimicking the functions of a brain. Within a decade, memristors could let us emulate, instead of merely simulate, networks of neurons and synapses. Many research groups have been working toward a brain in silico: IBM's Blue Brain project, Howard Hughes Medical Institute's Janelia Farm, and Harvard's Center for Brain Science are just three. However, even a mouse brain simulation in real time involves solving an astronomical number of coupled partial differential equations. A digital computer capable of coping with this staggering workload would need to be the size of a small city, and powering it would require several dedicated nuclear power plants.

Memristors can be made extremely small, and they function like synapses. Using them, we will be able to build analog electronic circuits that could fit in a shoebox and function according to the same physical principles as a brain.

A hybrid circuit—containing many connected memristors and transistors—could help us research actual brain function and disorders. Such a circuit might even lead to machines that can recognize patterns the way humans can, in those critical ways computers can't—for example, picking a particular face out of a crowd even if it has changed significantly since our last memory of it.

THE STORY OF THE MEMRISTOR is truly one for the history books. When Leon Chua, now an IEEE Fellow, wrote his seminal paper predicting the memristor, he was a newly minted and rapidly rising professor at UC Berkeley. Chua had been fighting for years against what he considered the arbitrary restriction of electronic circuit theory to linear systems. He was convinced that nonlinear electronics had much more potential than the linear circuits that dominate electronics technology to this day.

Chua discovered a missing link in the pairwise mathematical equations that relate the four circuit quantities—charge, current, voltage, and magnetic flux—to one another. These can be related in six ways. Two are connected through the basic physical laws of electricity and magnetism, and three are related by the known circuit elements: resistors connect voltage and current, inductors connect flux and current, and capacitors connect voltage and charge. But one equation is missing from this group: the relationship between charge moving through a circuit and the magnetic flux surrounded by that circuit—or more subtly, a mathematical doppelgänger defined by Faraday's Law as the time integral of the voltage across the circuit. This distinction is the crux of a raging Internet debate about the legitimacy of our memristor [see sidebar, "Resistance to Memristance"].

Chua's memristor was a purely mathematical construct that had more than one physical realization. What does that mean? Consider a battery and a transformer. Both provide identical voltages—for example, 12 volts of direct current—but they do so by entirely different mechanisms: the battery by a chemical reaction going on inside the cell and the transformer by taking a 110-V ac input, stepping that down to 12 V ac, and then transforming that into 12 V dc. The end result is mathematically identical—both will run an electric shaver or a cellphone, but the physical source of that 12 V is completely different.

Conceptually, it was easy to grasp how electric charge could couple to magnetic flux, but there was no obvious physical interaction between charge and the integral over the voltage.

Chua demonstrated mathematically that his hypothetical device would provide a relationship between flux and charge similar to what a nonlinear resistor provides between voltage and current. In practice, that would mean the device's resistance would vary according to the amount of charge that passed through it. And it would remember that resistance value even after the current was turned off.

He also noticed something else—that this behavior reminded him of the way synapses function in a brain.

Even before Chua had his eureka moment, however, many researchers were reporting what they called "anomalous" current-voltage behavior in the micrometer-scale devices they had built out of unconventional materials, like polymers and metal oxides. But the idiosyncrasies were usually ascribed to some mystery electrochemical reaction, electrical breakdown, or other spurious phenomenon attributed to the high voltages that researchers were applying to their devices.

As it turns out, a great many of these reports were unrecognized examples of memristance. After Chua theorized the memristor out of the mathematical ether, it took another 35 years for us to intentionally build the device at HP Labs, and we only really understood the device about two years ago. So what took us so long?

It's all about scale. We now know that memristance is an intrinsic property of any electronic circuit. Its existence could have been deduced by Gustav Kirchhoff or by James Clerk Maxwell, if either had considered nonlinear circuits in the 1800s. But the scales at which electronic devices have been built for most of the past two centuries have prevented experimental observation of the effect. It turns out that the influence of memristance obeys an inverse square law: memristance is a million times as important at the nanometer scale as it is at the micrometer scale, and it's essentially unobservable at the millimeter scale and larger. As we build smaller and smaller devices, memristance is becoming more noticeable and in some cases dominant. That's what accounts for all those strange results researchers have described. Memristance has been hidden in plain sight all along. But in spite of all the clues, our finding the memristor was completely serendipitous.

CROSSBAR ARCHITECTURE: A memristor's structure, shown here in a scanning tunneling microscope image, will enable dense, stable computer memories. *IMAGE: R. STANLEY WILLIAMS/HP LABS*

In 1995, I was recruited to HP Labs to start up a fundamental research group that had been proposed by David Packard. He decided that the company had become large enough to dedicate a research group to long-term projects that would be protected from the immediate needs of the business units. Packard had an altruistic vision that HP should "return knowledge to the well of fundamental science from which HP had been withdrawing for so long." At the same time, he understood that long-term research could be the strategic basis for technologies and inventions that would directly benefit HP in the future. HP gave me a budget and four researchers. But beyond the comment that "molecular-scale electronics" would be interesting and that we should try to have something useful in about 10 years, I was given carte blanche to pursue any topic we wanted. We decided to take on Moore's Law.

At the time, the dot-com bubble was still rapidly inflating its way toward a resounding pop, and the existing semiconductor road map didn't extend past 2010. The critical feature size for the transistors on an integrated circuit was 350 nanometers; we had a long way to go before atomic sizes would become a limitation. And yet, the eventual end of Moore's Law was obvious. Someday semiconductor researchers would have to confront physics-based limits to their relentless descent into the infinitesimal, if for no other reason than that a transistor cannot be smaller than an atom. (Today the smallest components of transistors on integrated circuits are roughly 45 nm wide, or about 220 silicon atoms.)

That's when we started to hang out with Phil Kuekes, the creative force behind the Teramac (tera-operation-per-second multiarchitecture computer)—an experimental supercomputer built at HP Labs primarily from defective parts, just to show it could be done. He gave us the idea to build an architecture that would work even if a substantial number of the individual devices in the circuit were dead on arrival. We didn't know what those devices would be, but our goal was electronics that would keep improving even after the devices got so small that defective ones would become common. We ate a lot of pizza washed down with appropriate amounts of beer and speculated about what this mystery nanodevice would be.

We were designing something that wouldn't even be relevant for another 10 to 15 years. It was possible that by then devices would have shrunk down to the molecular scale envisioned by David Packard or perhaps even *be* molecules. We could think of no better way to anticipate this than by mimicking the Teramac at the nanoscale. We decided that the simplest abstraction of the Teramac architecture was the crossbar, which has since become the de facto standard for nanoscale circuits because of its simplicity, adaptability, and redundancy.

The crossbar is an array of perpendicular wires. Anywhere two wires cross, they are connected by a switch. To connect a horizontal wire to a vertical wire at any point on the grid, you must close the switch between them. Our idea was to open and close these switches by applying voltages to the ends of the wires. Note that a crossbar array is basically a storage

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

How Memristance Works

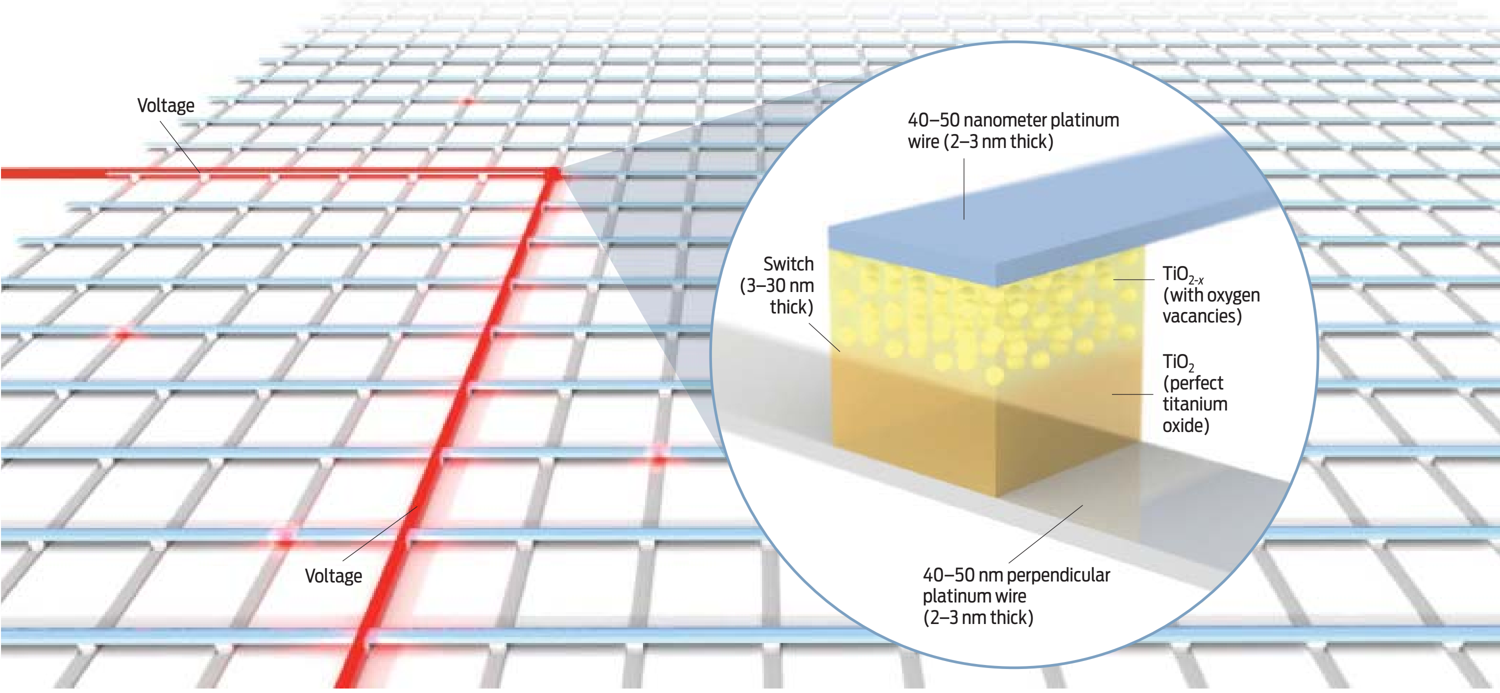

THE CROSSBAR ARCHITECTURE: The crossbar architecture is a fully connected mesh of perpendicular wires. Any two crossing wires are connected by a switch. To close the switch, a positive voltage is applied across the two wires to be connected. To open the switch, the voltage is reversed.

THE SWITCH: A switch is a 40-nanometer cube of titanium dioxide (TiO_2) in two layers: The lower TiO_2 layer has a perfect 2:1 oxygen-to-titanium ratio, making it an insulator. By contrast, the upper TiO_2 layer is missing 0.5 percent of its oxygen (TiO_{2-x}), so x is about 0.05. The vacancies make the TiO_{2-x} material metallic and conductive.

APPLIED MEMRISTANCE: The oxygen deficiencies in the TiO_{2-x} manifest as "bubbles" of oxygen vacancies scattered throughout the upper layer. A positive voltage on the switch repels the (positive) oxygen deficiencies in the metallic upper TiO_{2-x} layer, sending them into the insulating TiO_2 layer below. That causes the boundary between the two materials to move down, increasing the percentage of conducting TiO_{2-x} and thus the conductivity of the entire switch. The more positive voltage is applied, the more conductive the cube becomes.

A negative voltage on the switch attracts the positively charged oxygen bubbles, pulling them out of the TiO_2. The amount of insulating, resistive TiO_2 increases, thereby making the switch as a whole resistive. The more negative voltage is applied, the less conductive the cube becomes.

What makes this switch special—memristive—is that when the voltage is turned off, positive or negative, the oxygen bubbles do not migrate. They stay where they are, which means that the boundary between the two titanium dioxide layers is frozen. That is how the memristor "remembers" how much voltage was last applied.

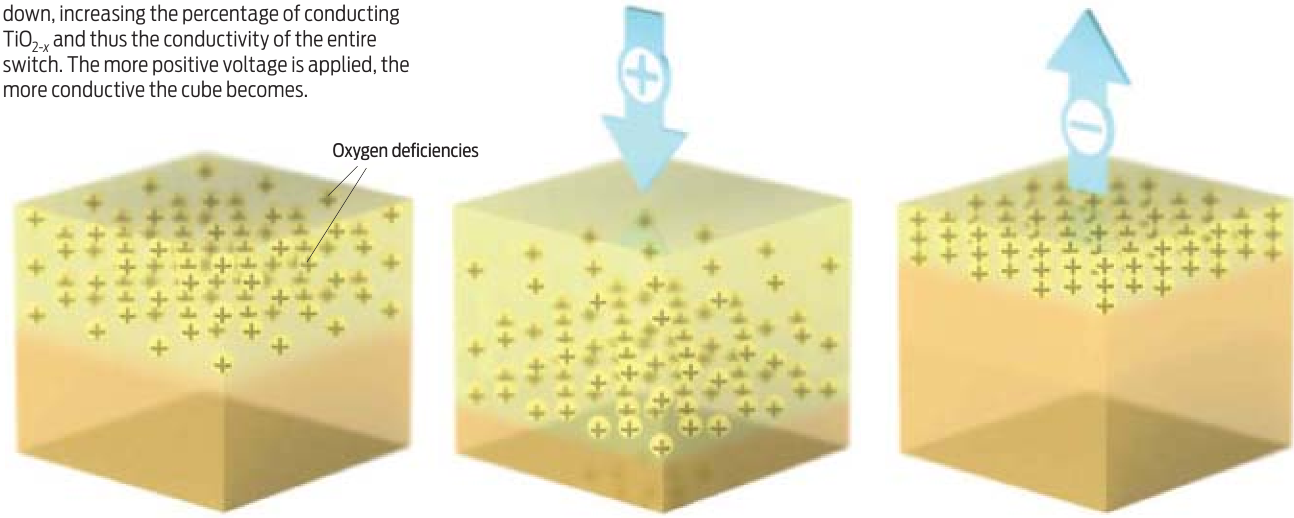

BRYAN CHRISTIE DESIGN

system, with an open switch representing a zero and a closed switch representing a one. You read the data by probing the switch with a small voltage.

Like everything else at the nanoscale, the switches and wires of a crossbar are bound to be plagued by at least some nonfunctional components. These components will be only a few atoms wide, and the second law of thermodynamics ensures that we will not be able to completely specify the position of every atom. However, a crossbar architecture builds in redundancy by allowing you to route around any parts of the circuit that don't work. Because of their simplicity, crossbar arrays have a much higher density of switches than a comparable integrated circuit based on transistors.

But implementing such a storage system was easier said than done. Many research groups were working on such a cross-point memory—and had been since the 1950s. Even after 40 years of research, they had no product on the market. Still, that didn't stop them from trying. That's because the potential for a truly nanoscale crossbar memory is staggering; picture carrying around the entire Library of Congress on a thumb drive.

One of the major impediments for prior crossbar memory research was the small off-to-on resistance ratio of the switches (40 years of research had never produced anything surpassing a factor of 2 or 3). By comparison, modern transistors have an off-to-on resistance ratio of 10 000 to 1. We calculated that to get a high-performance memory, we had to make switches with a resistance ratio of at least 1000 to 1. In other words, in its off state, a switch had to be 1000 times as resistive to the flow of current as it was in its on state. What mechanism could possibly give a nanometer-scale device a three-orders-of-magnitude resistance ratio?

We found the answer in scanning tunneling microscopy (STM), an area of research I had been pursuing for a decade. A tunneling microscope generates atomic-resolution images by scanning a very sharp needle across a surface and measuring the electric current that flows between the atoms at the tip of the needle and the surface the needle is probing. The general rule of thumb in STM is that moving that tip 0.1 nm closer to a surface increases the tunneling current by one order of magnitude.

We needed some similar mechanism by which we could change the effective spacing between two wires in our crossbar by 0.3 nm. If we could do that, we would have the 1000:1 electrical switching ratio we needed.

Our constraints were getting ridiculous. Where would we find a material that could change its physical dimensions like that? That is how we found ourselves in the realm of molecular electronics.

CONCEPTUALLY, OUR DEVICE was like a tiny sandwich. Two platinum electrodes (the intersecting wires of the crossbar junction) functioned as the "bread" on either end of the device. We oxidized the surface of the bottom platinum wire to make an extremely thin layer of platinum dioxide, which is highly conducting. Next, we assembled a dense film, only one molecule thick, of specially designed switching molecules. Over this "monolayer" we deposited a 2- to 3-nm layer of titanium metal, which bonds strongly to the molecules and was intended to glue them together. The final layer was the top platinum electrode.

The molecules were supposed to be the actual switches. We built an enormous number of these devices, experimenting with a wide variety of exotic molecules and configurations, including rotaxanes, special switching molecules designed by James Heath and Fraser Stoddart at the University of California, Los Angeles. The rotaxane is like a bead on a string, and with the right voltage, the bead slides from one end of the string to the other, causing the electrical resistance of the molecule to rise or fall, depending on the direction it moves. Heath and Stoddart's devices used silicon electrodes, and they worked, but not well enough for technological applications: the off-to-on resistance ratio was only a factor of 10, the switching was slow, and the devices tended to switch themselves off after 15 minutes.

Our platinum devices yielded results that were nothing less than frustrating. When a switch worked, it was spectacular: our off-to-on resistance ratios shot past the 1000 mark, the devices switched too fast for us to even measure, and having switched, the device's resistance state remained stable for years (we still have some early devices we test every now and then, and we have never seen a significant change in resistance). But our fantastic results were inconsistent. Worse yet, the success or failure of a device never seemed to depend on the same thing.

We had no physical model for how these devices worked. Instead of rational engineering, we were reduced to performing huge numbers of Edisonian experiments, varying one parameter at a time and attempting to hold all the rest constant. Even our switching molecules were betraying us; it seemed like we could use anything at all. In our desperation, we even turned to long-chain fatty acids—essentially soap—as the molecules in our devices. There's nothing in soap that should switch, and yet some of the soap devices switched phenomenally. We also made control devices with no molecule monolayers at all. None of them switched.

We were frustrated and burned out. Here we were, in late 2002, six years into our research. We had something that worked, but we couldn't figure out *why*, we couldn't model it, and we sure couldn't engineer it. That's when Greg Snider, who had worked with Kuekes on the Teramac, brought me the Chua memristor paper. "I don't know what you guys are building," he told me, "but this is what I want."

To this day, I have no idea how Greg happened to come across that paper. Few people had read it, fewer had understood it, and fewer still had cited it. At that point, the paper was 31 years old and apparently headed for the proverbial dustbin of history. I wish I could say I took one look and yelled, "Eureka!" But in fact, the paper sat on my desk for months before I even tried to read it. When I did study it, I found the concepts and the equations unfamiliar and hard to follow. But I kept at it because something had caught my eye, as it had Greg's: Chua had included a graph that looked suspiciously similar to the experimental data we were collecting.

The graph described the current-voltage (I-V) characteristics that Chua had plotted for his memristor. Chua had called them "pinched-hysteresis loops"; we called our I-V characteristics "bow ties." A pinched hysteresis loop looks like a diagonal infinity symbol with the center at the zero axis, when plotted on a graph of current against voltage. The voltage is first increased from zero to a positive maximum value, then decreased to a minimum negative value and finally returned to zero. The bow ties on our graphs were nearly identical [see graphic, "Bow Ties"].

That's not all. The total change in the resistance we had measured in our devices also depended on how *long* we applied the voltage: the longer we applied a positive voltage, the lower the resistance until it reached a minimum value. And the longer we applied a negative voltage, the higher the resistance became until it reached a maximum limiting value. When

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

Resistance To Memristance

INTRODUCING A NEW fundamental circuit element earned R. Stanley Williams some grief along with his newfound fame. After the *Nature* article appeared in May, online comments pages boiled over with skepticism. "Is this a hoax?" someone asked on the Wikipedia memristor page on 30 April 2008, the day the news broke, in one of the milder statements of disbelief. Seven months later, the debate continues.

Skeptics argue that the memristor is not a fourth fundamental circuit element but an example of bad science. The crux of their argument rests on two fundamental misunderstandings: first, skeptics overlook the expanded design space that arises from working with nonlinear circuit elements. The second and more profound misunderstanding concerns Leon Chua's mathematical definition of a memristor.

At first, most people—including Williams—assumed that Chua defined memristance strictly as the relationship between electric charge and magnetic flux. However, the actual definition of memristance is more general. Linking electric charge and magnetic flux is one way to satisfy the definition, but it's not the only one. In fact, it turns out you can bypass magnetic interaction altogether.

Chua's general memristance definition has two parts. The first equation defines how the memristor's voltage depends on current and a "state variable"—that is, a quantity that measures some physical property of a device, like the length of a column of mercury in a thermometer. The column's length correlates to the thermometer's temperature, and adding or removing heat makes the column longer or shorter. In Williams's memristor, the state variable is the thickness of the stoichiometric titanium dioxide in the switch; increasing or decreasing that thickness causes the device's resistance to increase or decrease.

The second equation expresses how the changing state variable (the TiO_2's thickness) depends on the amount of charge flowing through the device. In Williams's memristor, the TiO_2's thickness depends on the distribution of the oxygen vacancies throughout the material.

Here is what you need to remember: one, a magnetic interaction is not necessary for memristance. Two, in nonlinear circuit elements, memristance is not the same thing as nonlinear resistance. Three, because no combination of passive devices can reproduce the properties of a memristor, memristance is a fundamental circuit quantity.

Williams himself is sanguine about the memristor's reputation. "A hundred years after Einstein proposed his theory of relativity," he says, shrugging, "you still have some people arguing against it." *—Sally Adee*

we stopped applying the voltage, whatever resistance characterized the device was frozen in place, until we reset it by once again applying a voltage. The loop in the I-V curve is called hysteresis, and this behavior is startlingly similar to how synapses operate: synaptic connections between neurons can be made stronger or weaker depending on the polarity, strength, and length of a chemical or electrical signal. That's not the kind of behavior you find in today's circuits.

Looking at Chua's graphs was maddening. We now had a big clue that memristance had something to do with our switches. But how? Why should our molecular junctions have anything to do with the relationship between charge and magnetic flux? I couldn't make the connection.

Two years went by. Every once in a while I would idly pick up Chua's paper, read it, and each time I understood the concepts a little more. But our experiments were still pretty much trial and error. The best we could do was to make a lot of devices and find the ones that worked.

But our frustration wasn't for nothing: by 2004, we had figured out how to do a little surgery on our little sandwiches. We built a gadget that ripped the tiny devices open so that we could peer inside them and do some forensics. When we pried them apart, the little sandwiches separated at their weakest point: the molecule layer. For the first time, we could get a good look at what was going on inside. We were in for a shock.

WHAT WE HAD was not what we had built. Recall that we had built a sandwich with two platinum electrodes as the bread and filled with three layers: the platinum dioxide, the monolayer film of switching molecules, and the film of titanium.

But that's not what we found. Under the molecular layer, instead of platinum dioxide, there was only pure platinum. Above the molecular layer, instead of titanium, we found an unexpected and unusual layer of titanium dioxide. The titanium had sucked the oxygen right out of the platinum dioxide! The oxygen atoms had somehow migrated through the molecules and been consumed by the titanium. This was especially surprising because the switching molecules had not been significantly perturbed by this event—they were intact and well ordered, which convinced us that they must be doing something important in the device.

The chemical structure of our devices was not at all what we had thought it was. The titanium dioxide—a stable compound found in sunscreen and white paint—was not just regular titanium dioxide. It had split itself up into two chemically different layers. Adjacent to the molecules, the oxide was stoichiometric TiO_2, meaning the ratio of oxygen to titanium was perfect, exactly 2 to 1. But closer to the top platinum electrode, the titanium dioxide was missing a tiny amount of its oxygen, between 2 and 3 percent. We called this oxygen-deficient titanium dioxide TiO_{2-x}, where x is about 0.05.

Because of this misunderstanding, we had been performing the experiment backward. Every time I had tried to create a switching model, I had reversed the switching polarity. In other words, I had predicted that a positive voltage would switch the device off and a negative voltage would switch it on. In fact, exactly the opposite was true.

It was time to get to know titanium dioxide a lot better. They say three weeks in the lab will save you a day in the library every time. In August of 2006 I did a literature search and found about 300 relevant papers on titanium dioxide. I saw that each of the many different communities researching titanium dioxide had its own way of describing the compound. By the end of the month, the pieces had fallen into place. I finally knew how our device worked. I knew why we had a memristor.

The exotic molecule monolayer in the middle of our sandwich had nothing to do with the actual switching. Instead, what it did was control the flow of oxygen from the platinum dioxide into the titanium to produce the fairly uniform layers of TiO_2 and TiO_{2-x}. The key to the switching was this bilayer of the two different titanium dioxide species [see diagram, "How Memristance Works"]. The TiO_2 is electrically insulating (actually a semiconductor), but the TiO_{2-x} is conductive, because its oxygen vacancies are donors of electrons, which makes the vacancies themselves positively charged. The vacancies can be thought of like bubbles in a glass of beer, except that they don't pop—they can be pushed up and down at will in the titanium dioxide material because they are electrically charged.

Now I was able to predict the switching polarity of the device. If a positive voltage is applied to the top electrode of the device, it will repel the (also positive) oxygen vacancies in the TiO_{2-x} layer down into the pure TiO_2 layer. That turns the TiO_2 layer into TiO_{2-x} and makes it conductive, thus turning the device on. A negative voltage has the opposite effect: the vacancies are attracted upward and back out of the TiO_2, and thus the thickness of the TiO_2 layer increases and the device turns off. This switching polarity is what we had been seeing for years but had been unable to explain.

On 20 August 2006, I solved the two most important equations of my career—one equation detailing the relationship between current and voltage for this equivalent circuit, and another equation describing how the application of the voltage causes the vacancies to move—thereby writing down, for the first time, an equation for memristance in terms of the physical properties of a material. This provided a unique insight. Memristance arises in a semiconductor when both electrons and charged dopants are forced to move simultaneously by applying a voltage to the system. The memristance did not actually involve magnetism in this case; the integral over the voltage reflected how far the dopants had moved and thus how much the resistance of the device had changed.

We finally had a model we could use to engineer our switches, which we had by now positively identified as memristors. Now we could use all the theoretical machinery Chua had created to help us design new circuits with our devices.

Triumphantly, I showed the group my results and immediately declared that we had to take the molecule monolayers out of our devices. Skeptical after years of false starts and failed hypotheses, my team reminded me that we had run control samples without molecule layers for every device we had ever made and that those devices had never switched. And getting the recipe right turned out to be tricky indeed. We needed to find the exact amounts of titanium and oxygen to get the two layers to do their respective jobs. By that point we were all getting impatient. In fact, it took so long to get the first working device that in my discouragement I nearly decided to put the molecule layers back in.

A month later, it worked. We not only had working devices, but we were also able to improve and change their characteristics at will.

But here is the real triumph. The resistance of these devices stayed constant whether we turned off the voltage or just read their states (interrogating them with a voltage so small it left the resistance unchanged). The oxygen vacancies didn't roam around; they remained absolutely immobile until we again applied a positive or negative voltage. That's memristance: the devices remembered their current history. We had coaxed Chua's mythical memristor off the page and into being.

EMULATING THE BEHAVIOR of a single memristor, Chua showed, requires a circuit with at least 15 transistors and other passive elements. The implications are extraordinary: just imagine how many kinds of circuits could be supercharged by replacing a handful of transistors with one single memristor.

The most obvious benefit is to memories. In its initial state, a crossbar memory has only open switches, and no information is stored. But once you start closing switches, you can store vast amounts of information compactly and efficiently. Because memristors remember their state, they can store data indefinitely, using energy only when you toggle or read the state of a switch, unlike the capacitors in conventional DRAM, which will lose their stored charge if the power to the chip is turned off. Furthermore, the wires and switches can be made very small: we should eventually get down to a width of around 4 nm, and then multiple crossbars could be stacked on top of each other to create a ridiculously high density of stored bits.

Greg Snider and I published a paper last year showing that memristors could vastly improve one type of processing circuit, called a field-programmable gate array, or FPGA. By replacing several specific transistors with a crossbar of memristors, we showed that the circuit could be shrunk by nearly a factor of 10 in area and improved in terms of its speed relative to power-consumption performance. Right now, we are testing a prototype of this circuit in our lab.

And memristors are by no means hard to fabricate. The titanium dioxide structure can be made in any semiconductor fab currently in existence. (In fact, our hybrid circuit was built in an HP fab used for making inkjet cartridges.) The primary limitation to manufacturing hybrid chips with memristors is that today only a small number of people on Earth have any idea of how to design circuits containing memristors. I must emphasize here that memristors will never eliminate the need for transistors: passive devices and circuits require active devices like transistors to supply energy.

The potential of the memristor goes far beyond juicing a few FPGAs. I have referred several times to the similarity of memristor behavior to that of synapses. Right now, Greg is designing new circuits that mimic aspects of the brain. The neurons are implemented with transistors, the axons are the nanowires in the crossbar, and the synapses are the memristors at the cross points. A circuit like this could perform real-time data analysis for multiple sensors. Think about it: an intelligent physical infrastructure that could provide structural assessment monitoring for bridges. How much money—and how many lives—could be saved?

I'm convinced that eventually the memristor will change circuit design in the 21st century as radically as the transistor changed it in the 20th. Don't forget that the transistor was lounging around as a mainly academic curiosity for a decade until 1956, when a killer app—the hearing aid—brought it into the marketplace. My guess is that the real killer app for memristors will be invented by a curious student who is now just deciding what EE courses to take next year. ❑

Bow Ties

LEON CHUA'S original graph of the hypothetical memristor's behavior is shown at top right; the graph of R. Stanley Williams's experimental results in the *Nature* paper is shown below. The loops map the switching behavior of the device: it begins with a high resistance, and as the voltage increases, the current slowly increases. As charge flows through the device, the resistance drops, and the current increases more rapidly with increasing voltage until the maximum is reached. Then, as the voltage decreases, the current decreases but more slowly, because charge is flowing through the device and the resistance is still dropping. The result is an on-switching loop. When the voltage turns negative, the resistance of the device increases, resulting in an off-switching loop. *—R.S.W.*

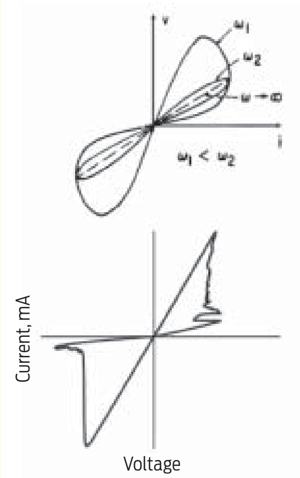

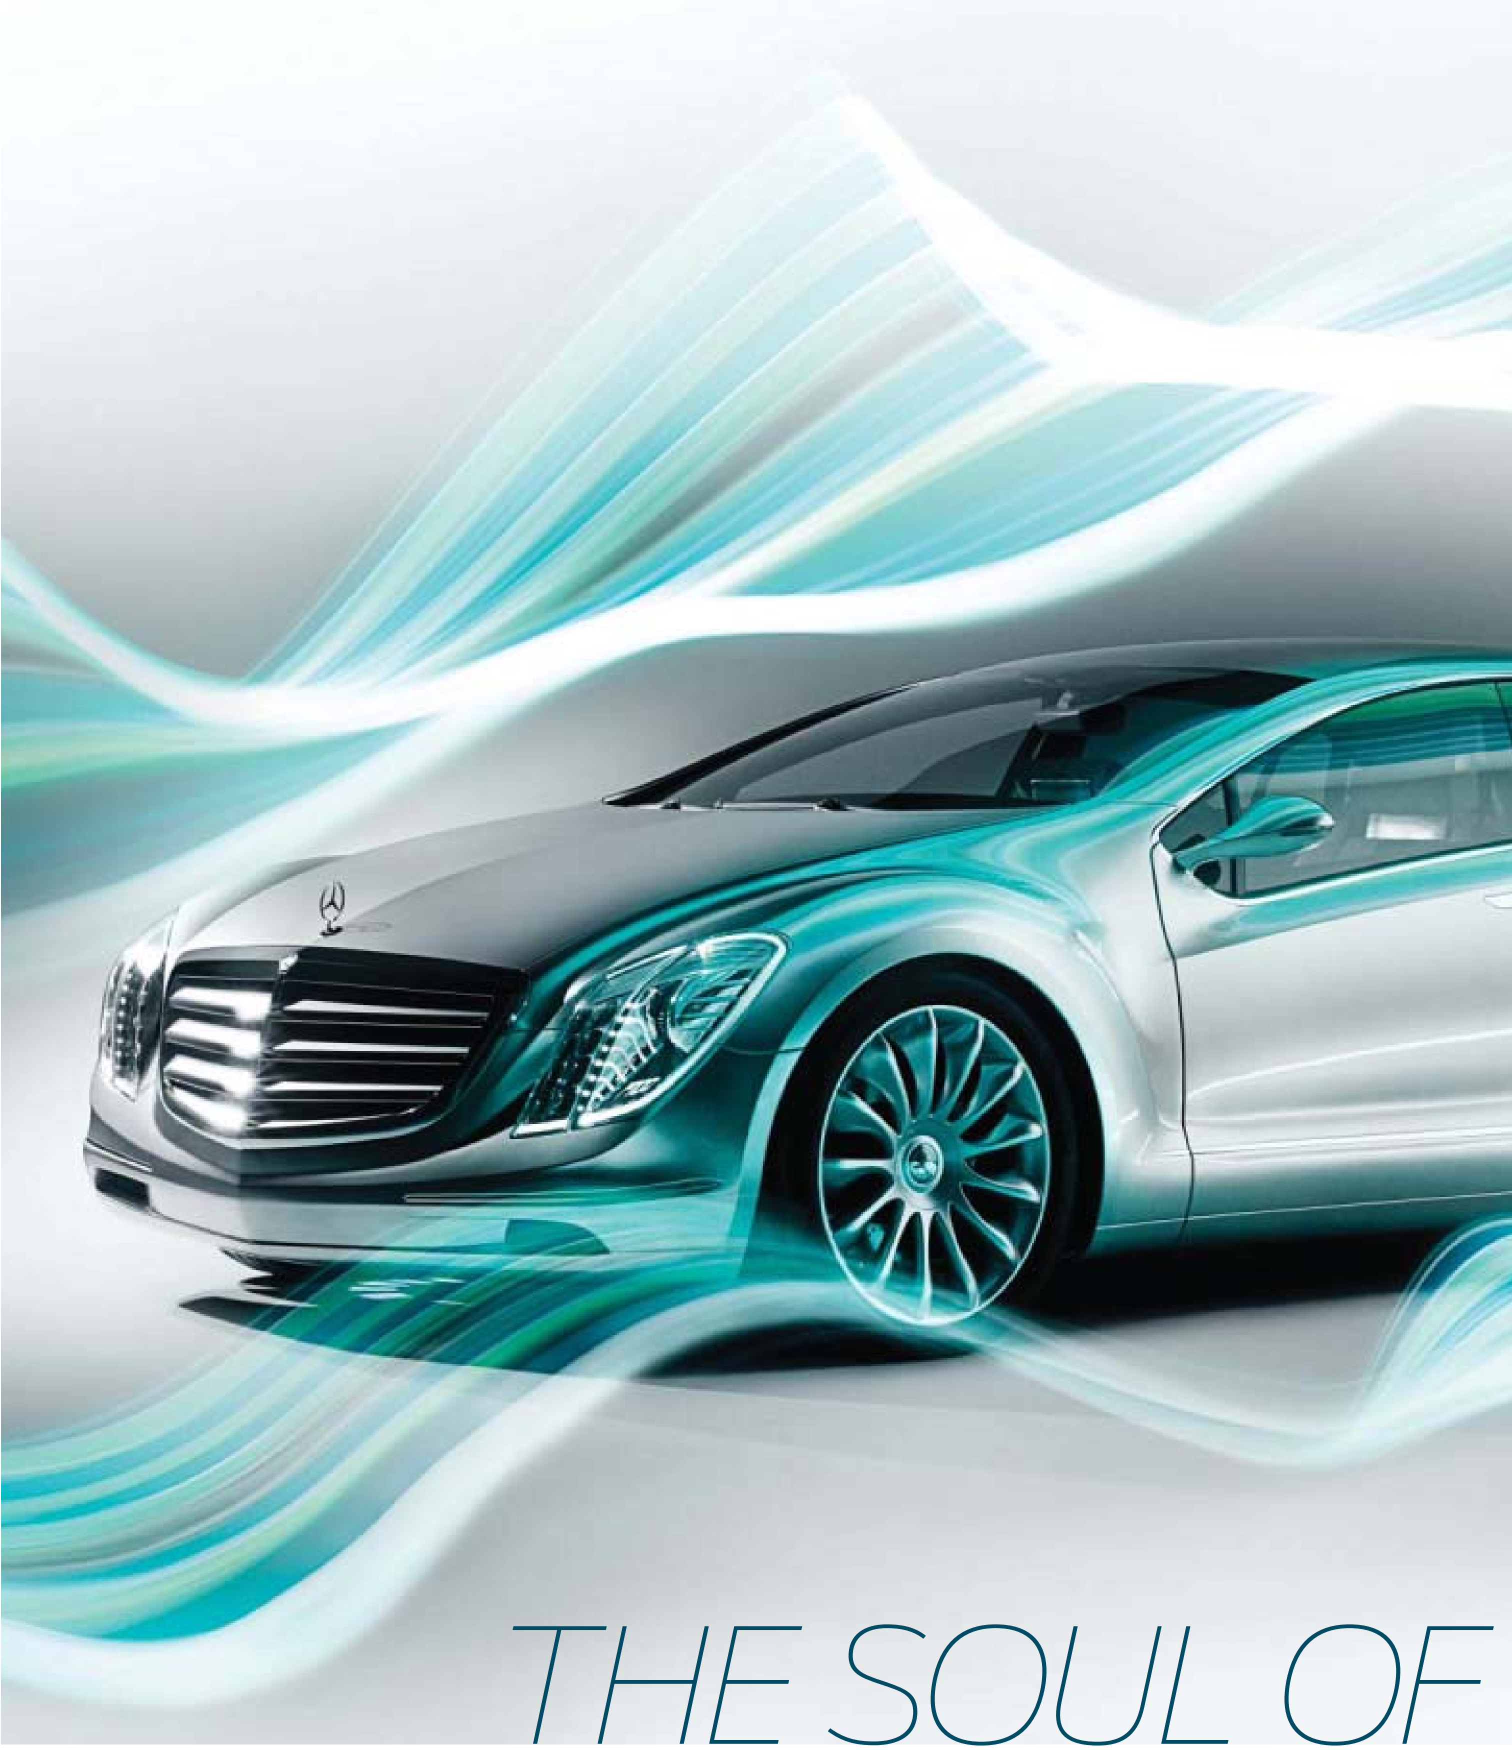

THE SOUL OF

How the F700 concept vehicle offers luxury-car performance, comfort, and

STEFFEN JAHN

OBJECT OF DESIRE: The F700 is a research car, but that didn't stop some wealthy customers from making offers.

A NEW MERCEDES

econo-car fuel efficiency *By John Voelcker*

CAR BUYERS everywhere are spurning gas guzzlers, and some carmakers feel the pain more than others. Take Mercedes-Benz, long known for its luxury cars—its flagship S-Class is a full-size sedan that starts at around US $90 000. Without sacrificing performance or comfort, the company must meet the expectations of its increasingly environmentally conscious customers.

Mercedes, a division of Daimler, has ingeniously applied technology to come up with a stunning solution. It's the F700 research car, which debuted last year. The F700 was created to give customers a preview of what they might

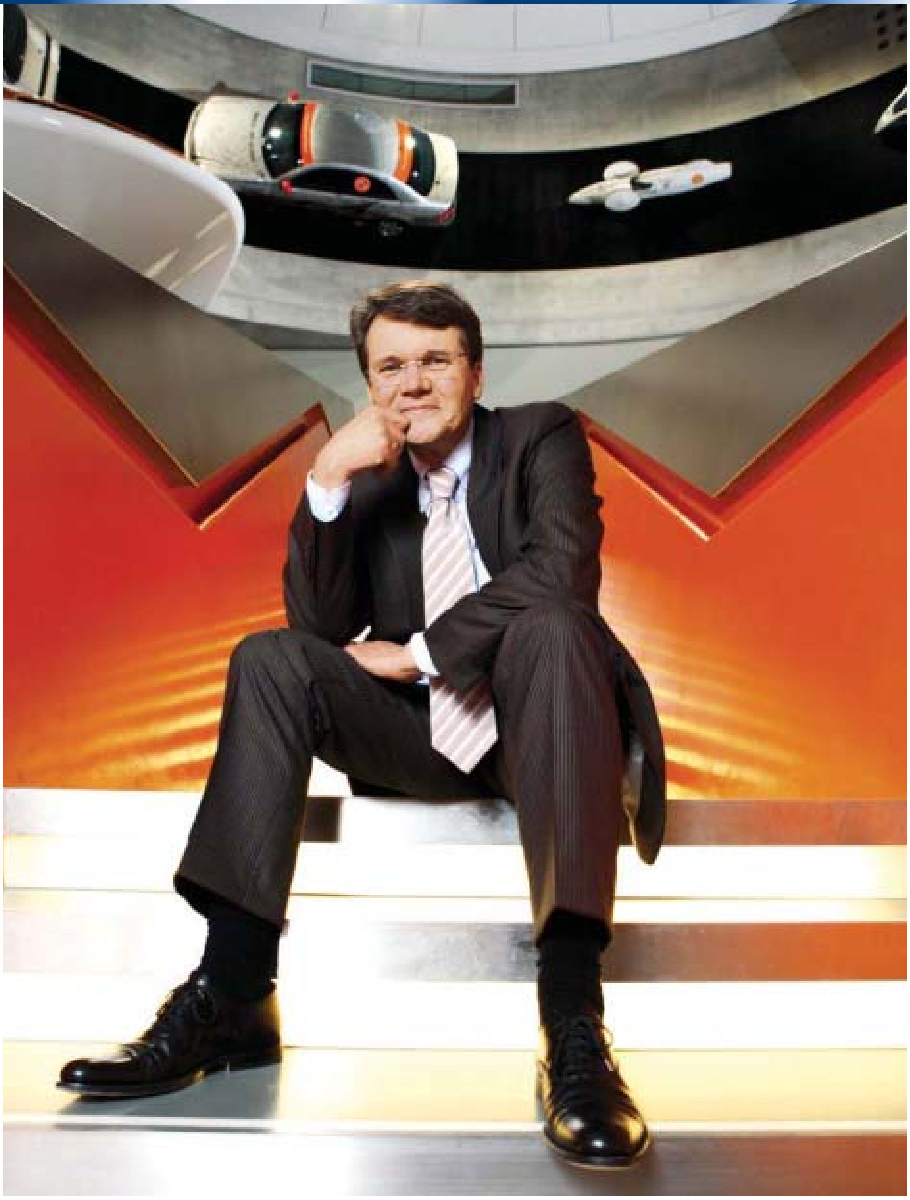

VEHICLE VISIONARY: Herbert Kohler mobilized Mercedes's advanced engineering groups to create the F700.

DAIMLER

expect from an S-Class sedan 10 years from now. The interior is hugely spacious, modern, and filled with electronics; the exterior is shiny and sleek, a dramatic evolution from Mercedes's current upright look. The car is also fast: it goes from 0 to 100 kilometers per hour in 7.5 seconds.

And here's the best part: the massive F700 sips no more fuel than a Toyota Prius.

The critical advance is the power plant. The F700 is a gasoline-electric hybrid, but it wrings most of its efficiencies out of good old-fashioned internal combustion. Mercedes is betting that such engines will be a mainstay of the industry for a very long time to come. But in the face of climate change and soaring global demand for oil, the company decided to show how much room for improvement there was in the 130-year-old internal combustion engine.

What makes this feat possible is a futuristic technology known as HCCI, or homogeneous charge-compression ignition. HCCI combines the low emissions of gasoline engines and the fuel efficiency of diesel engines. Automakers have long been working to develop HCCI engines and the electronic controllers needed to tame their combustion. GM and Volkswagen, among others, have running prototypes. But Mercedes's small power plant—a 1.8-liter four-cylinder twin-turbo HCCI engine that burns regular gasoline—shocked the experts.

"'A lot more from a lot less' is a fundamentally exciting vision for the future of an industry that's under so many clouds right now," says Nigel Griffiths, a managing director at the industry analysis firm Global Insight, in London.

Mercedes says the engine delivers the same performance as the 3.5-L V6 offered in European S-Class models while consuming just half the fuel.

"We had to demonstrate that good fuel consumption and premium luxury cars are not a contradiction," says Herbert Kohler, who as head of group research and advanced engineering for vehicle and power train is responsible for all of Mercedes's concept cars.

The company added the motor-alternator hybrid system to provide punch during acceleration. It also suspended the F700 with hydraulics that use infrared laser sensors to read and adapt to road conditions—not just in real time but ahead of time. In other words, the suspension isn't active; it's proactive.

Thanks to the compact engine, designers were able to extend the wheelbase, making the F700 no longer than an S-Class but far roomier. Finally, to convey the linked concepts of elegance and efficiency, Mercedes stylists found inspiration in...dolphins.

"The astounding thing about the F700 is that it showcases a vision of future technologies that doesn't assume continued downsizing by consumers," Griffiths says. "If they can really get those technologies working, we're looking at a virtual quantum leap."

KOHLER, a measured and precise man, can pinpoint the exact steps that led to the conception of the F700. Sitting in his book-lined office in Daimler's headquarters in Untertürkheim, on the outskirts of Stuttgart, Germany, he says the project began to take shape in the fall of 2005. That year, Mercedes unveiled a concept known as the Bionic Car, an exercise in structural and aesthetic design. It was natural to look next for improvements that could be made under the hood.

When cruising, the F700's engine switches from conventional spark ignition to a more fuel-efficient compression-ignition mode

Kohler considered a conventional hybrid-electric drive but concluded that it wouldn't give radical fuel reductions, especially during high-speed autobahn cruising. Could the car be fully electric? The problem was that even the best lithium-ion cells projected for 2018 would be too heavy to provide the 500 kilometers of range expected from luxury cars. Hydrogen fuel cells also appeared too far-fetched, as the necessary hydrogen fueling infrastructure was likely even further away.

Digging deeper into the company's research, Kohler became convinced that the novel HCCI technology was ready to make the leap from lab bench to demonstration vehicle—in concert with a handful of additional engine modifications. HCCI could offer the advantages of gasoline and diesel engines while avoiding the disadvantages peculiar to each.

A gasoline engine achieves a well-controlled, efficient combustion with relatively low emissions by using a spark plug to ignite a vapor of fuel and air, compressed typically at a 10:1 ratio inside the cylinders. Diesel engines, on the other hand, have no spark plugs; combustion happens spontaneously when an injector sprays atomized diesel fuel into the cylinders, filled with extremely hot air compressed at ratios as high as 18:1. The higher compression ratio makes the engine more energy efficient but at a cost: higher temperatures generate nitrogen oxides, a major cause of air pollution.

The idea behind HCCI is to take advantage of both the cleaner burn of spark ignition and the higher efficiency of compression ignition. In a gasoline engine, the burn starts at the spark plug and propagates outward from there. In a diesel, combustion begins at the edges and moves inward. But in an HCCI engine, the fuel-air mixture is so thoroughly combined that it ignites at many points simultaneously. The benefits are twofold: the engine can use leaner fuel-air mixtures without suffering a loss in power, and the combustion temperatures are too low to form nitrogen oxides.

To take HCCI from the lab to the test track, Kohler turned to Gunther Ellenrieder, an old colleague who was director of group research and advanced engineering for vehicle concepts and human factors at the Mercedes-Benz Technology Center, at the company's largest plant, in Sindelfingen, Germany. Ellenrieder, who is tall and prone to gestures to sketch out his explanations, nodded emphatically as he recalled their first conversation about HCCI. "The calculations looked pretty good," he said, "but that didn't mean it was mature enough"—meaning that getting HCCI working in a drivable car would be quite the challenge.

In early 2006, Ellenrieder mobilized a team of development engineers and gave them a tight deadline: September 2007, when Mercedes was to unveil the F700. To keep the fuel consumption close to the 50 percent mark, they settled on a four-cylinder HCCI engine of 1.8 L—the same size as those used in midsize European sedans. The team added two turbochargers, pumps driven by exhaust gases that pack additional air into each cylinder, allowing the injection of more fuel. As the engine starts to rev from idle, the pump with the smaller output nozzle boosts pressure immediately; the pump with higher output, from a larger nozzle, takes over at higher speeds.

The one thing the new engine couldn't manage was the torque of a larger power plant. So the engineers added a 15-kilowatt electric motor, which provides an electric assist when maximum power is needed, and an integrated starter-generator, tucked inside the seven-speed automatic transmission, to both start the engine and recharge a battery pack on overrun.

Under regular loads, the engine acts conventionally, using spark ignition. Under heavy loads or for better acceleration—passing on a highway or going up a hill—the electric motor may add power as well. But under a partial load—level highway cruising, for instance—the engine switches over to work as an HCCI, cutting fuel consumption. And if there is no load at all, the engine simply switches off until it's needed again.

MERCEDES engine wizards hunkered down in the Sindelfingen lab for weeks to assemble the F700's HCCI heart. Some parts they borrowed and adapted from production vehicles, including air-intake valves that open at different times and fuel injectors that squirt gasoline into the cylinders at different pressures.

But the engineers also perfected two systems not normally found in other engines. The first retained exhaust gas left over from the previous combustion cycle, using it to prewarm the gasoline-air mixture to reach ignition temperature. This required pricey pressure transducers inside the combustion chamber; they would feed data to the engine controller, which would determine in real time how much to open the exhaust and intake valves for each cylinder and when to inject the fuel.

The engineers next devised a mechanism to change the position of the piston at the top of its stroke. This complex piece of engineering raises the compression ratio to create the optimal environment for self-ignition under light loads and lowers the ratio for greater power in spark-ignition mode. It comes as no surprise that Mercedes won't reveal the specifics, declining to answer questions about the range of compression ratios or the mechanical methods used.

The assembled engine block is a compact cube of metal and plastic. The company's marketers christened it DiesOtto, combining the names of the two 19th-century German pioneers of diesel and gasoline combustion engines, Rudolph Diesel and Nikolaus Otto, respectively. Mercedes says that while several other makers are researching HCCI concepts, the DiesOtto engine is unique in its combination of spark-ignition and prewarmed HCCI modes, its variable compression ratio, and the application of turbocharging.

And then there's the engine control unit. Conventional

MORPHING MOTOR: The F700's HCCI engine can vary its compression ratio on the fly.

STEFFEN JAHN

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

The F700's "magic carpet" active suspension maps the road so accurately that a driver going over a pothole at high speed barely feels a jolt

engine controllers keep combustion efficient and emissions low by constantly adjusting the amount of fuel injected and the timing of the ignition. They make adjustments based on data received from dozens of sensors measuring parameters as varied as engine rotation, atmospheric pressure, and tailpipe oxygen levels.

Designing such controllers is fiendishly complicated, even for standard cars. Designing the F700 controller—which had to manage such a complex range of combustion modes—was "an order of magnitude" worse, as Ellenrieder puts it. The switching between spark and compression-ignition modes proved especially hard. In the end, the engineers realized that the controller can figure out which mode to use from the data coming from the pressure transducers inside each cylinder. The controller then automatically adjusts the compression ratio, the amount of fuel injected, the valve timing, and other parameters. It's the main brain behind the DiesOtto.

SUSPENSION was next on Kohler's list. Naturally, the car would have the most advanced electronic systems for safety and comfort used in Mercedes cars—traction control that varies power and braking on each wheel and adaptive cruise control that maintains a safe distance behind the car ahead at any road speed. But as a vision statement, the F700 needed more.

Once again, Kohler turned to Ellenrieder, who suggested what he called the "magic carpet." Like today's active suspension systems, it would take in data like road speed, acceleration rate, brake application, and force from each suspension component. But on top of that, it would monitor the conditions of the road just ahead of the vehicle.

If the system worked, the F700 would whiz over the asphalt almost as if floating in midair, already prepared to smooth out the bumps and irregularities it "knew" were coming—something no other vehicle could do. There were doubts inside Mercedes on whether the system—aptly named Pre-Scan—was far enough along to be exposed to the global glare of publicity. But Ellenrieder prevailed, after what he diplomatically termed a "push-pull process," and the system went into full-on development.

The key innovation was the use of lidar, or light detection and ranging technology. Originally developed for military remote-sensing applications, lidar is now used by geoscientists and construction crews to map the elevations of a terrain. The device emits pulsed beams of infrared laser light, detects their reflections, and computes the distance to the target. The idea was to use lidar sensors in the headlights to scan the road so that the car's suspension could firm up or soften its damping and direct the hydraulic shocks to absorb or counter the loads on each wheel.

Programming the suspension's controller, fitted with multiple processors, brought its own challenges. Among other things, the controller had to keep track of the motion of the car body—where the sensors were mounted—relative to that of the wheels and suspension. And because the lidar sensors mapped the road so accurately—they could "measure" the thickness of the painted lines on the asphalt—the controller also had to cleverly smooth out the input data to make sense of different road surfaces. But the result paid off: sitting behind the wheel of the F700, a driver going over a pothole at high speed barely feels a jolt.

FINALLY, it was time to consider what the car would look like. Kohler discussed his early ideas with Peter Pfeiffer, head of the company's design department, and they settled on rough dimensions. The length would be about 5 meters, the size of the current S-Class. With the small engine up front, the wheelbase could be a full 3.5 meters—about as long as that of ultraluxury vehicles like the Maybach 57 or a Rolls-Royce Phantom—giving interior space that would be impossible in regular cars.

Kohler asked Pfeiffer what to do with this roominess. Pfeiffer told him that the company's advanced design studio for interiors in Como, Italy, had experimented with repositioning seats to let passengers have more relaxing conversations or do business more effectively. One idea that emerged: if the rear seats rotated 180 degrees, passengers would be able to talk face to face, limousine style.

With those and other specs in hand, Kohler and Pfeiffer contacted the company's design studios in Germany, Japan, and California and asked them for proposals. Above all, the design had to denote a sense of inherent fuel efficiency and environmental concerns—and, of course, still convey the instant emotional impression of a large, prestigious Mercedes-Benz.

In July 2006, Mercedes's board of directors reviewed several proposals and chose a design submitted by the California studio, the Advanced Design North America division, in Irvine. Its chief designer is Chris Rhoades, who, like several of his colleagues, is a surfer. "Design is inspired *by* nature," he says, "to give back *to* nature." After Rhoades sketched, pondered, surfed, and sketched some more, what emerged in his flowing shapes and design elements was the image of a dolphin breaching the waves. Sleek, smart, efficient, it seemed to Rhoades the perfect metaphor for the F700. The board agreed, and the dolphin got the nod.

The result was a flowing roof and gently rounded shoulder lines, highest in the center, and a short snout up front. A dorsal fin down the car's roof added aerodynamic stability—and undeniable marine overtones. Over the next five months, the design team worked frantically to build a full-size clay study, refine every detail and surface, and produce a final hard-foam model. From this model, the template for the actual F700 metal body finally emerged.

But challenges still loomed. As it turned out, the paint job almost ruined the car. The process was multilayered, with many cycles of painting,

PROACTIVE SUSPENSION: Lidar sensors in the headlights scan the road ahead of the car.

STEFFEN JAHN

MAKING IT REAL: The F700 designers sought nature-inspired shapes to convey efficiency and elegance [top left]. Chassis, engine, and other parts converged at a top-secret studio for the final assembly [top right]. Lined with leather and cork, the interior [bottom] is roomier than that of a flagship S-Class. The iridescent silver paint job took a week to complete [center left].

TOP: DAIMLER (3); BOTTOM: STEFFEN JAHN

buffing, and sanding to leave fine metal flakes suspended in a deep iridescent silver surface. After a week in the paint shop, the car's surface emerged flawlessly smooth—but the wrong shade! It wasn't until a week later that the last coat of hardener finished curing in the sunlight and the slight green tint vanished. The paint shop foreman could only repeat, over and over, "It's a miracle!"

THE FULLY assembled F700 emerged for final review at the top-secret Mercedes-Benz design studio in the Sindelfingen complex. A circular building designed by world-renowned architect Renzo Piano, the studio has several roll-up garage doors that allow cars to be taken outside, where fences and shrubbery shield them from prying eyes. Here the F700 was positioned on a turntable, where it shone in the sunlight.

It looked just as good on the test track. The F700 has a rated power of 175 kW (238 horsepower) at maximum load. At cruising speeds, it consumes only 5.3 liters per 100 kilometers (44 miles per gallon), with CO_2 emissions of 127 grams per kilometer—and this for a full-size vehicle that weighs 1700 kilograms. By comparison, the S350's 3.5-L gasoline V6 produces 200 kW (268 hp) and takes 7.3 seconds to go from 0 to 100 km/h—just a bit better than the F700—but its fuel consumption is 10.3 L/100 km (23 mpg) and its CO_2 emissions are 242 g/km, almost twice those of the concept car.

The F700 debuted at the Frankfurt Motor Show on 11 September 2007. It was like no Mercedes the world had ever seen, and the reception was rapturous. Some wealthy customers who were shown the car made offers on the spot or asked if Mercedes would build one for them. Several German commentators used the phrase *aus einem Guss*, meaning "from a single cast," or seamless, all of a piece. The DiesOtto, in particular, received much acclaim, some even calling it a "green engine." Industry analysts and critics praised the car's other innovations as well—most of them, at least. The control interface of the navigation system was too far from the driver, some critics said, and talking to an avatar to input settings didn't work very well. And then there was the dolphin-like body: some observers simply found it too radical for an S-Class.

But the question that customers, analysts, and other carmakers have all been asking is when the F700 features—especially the innovative DiesOtto engine—will appear in production models. Kohler is coy. He says Mercedes has no exact timetable, although he can point to dozens of features in today's cars that first appeared in earlier concept cars.

Some industry analysts predict that HCCI engines will become practical for mass production no sooner than 2015. But with a new generation of S-Class models slated for 2012, some of the F700's innovations may trickle into the showroom over the next few years. One possibility is a direct-injection four-cylinder engine, which would bring some of the fuel-saving advances of the F700 project to the S-Class models. For a line of sedans that Mercedes has always fitted with big engines, this alone would be a radical move.

Just don't expect a dolphin-shaped body anytime soon. ❑

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

A Fairer, Faster Internet

TCP—the way we share bandwidth—needs a makeover *By Bob Briscoe*

The Internet is founded on a very simple premise: shared communications links are more efficient than dedicated channels that lie idle much of the time. ¶ And so we share. We share local area networks at work and neighborhood links from home. And then we share again—at any given time, a terabit backbone cable is shared among thousands of folks surfing the Web, downloading videos, and talking on Internet phones. ¶ But there's a profound flaw in the protocol that governs how people share the Internet's capacity. The protocol allows you to seem to be polite, even as you elbow others aside, taking far more resources than they do.

Network providers like Verizon and BT either throw capacity at the problem or improvise formulas that attempt to penalize so-called bandwidth hogs. Let me speak up for this much-maligned beast right away: bandwidth hogs are not the problem. There is no need to prevent customers from downloading huge amounts of material, so long as they aren't starving others.

Rather than patching over the problem, my colleagues and I at BT (formerly British Telecom) have worked out how to fix the root cause: the Internet's sharing protocol itself. It turns out that this solution will make the Internet not just simpler but much faster too.

QUICKHONEY

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

YOU MIGHT BE SHOCKED to learn that the designers of the Internet intended that your share of Internet capacity would be determined by what your own software considered fair. They gave network operators no mediating role between the conflicting demands of the Internet's hosts—now over a billion personal computers, mobile devices, and servers.

The Internet's primary sharing algorithm is built into the Transmission Control Protocol, a routine on your own computer that most programs run—although they don't have to. TCP is one of the twin pillars of the Internet, the other being the Internet Protocol, which delivers packets of data to particular addresses. The two together are often called TCP/IP.

Your TCP routine constantly increases your transmission rate until packets fail to get through some pipe up ahead—a tell-tale sign of congestion. Then TCP very politely halves your bit rate. The billions of other TCP routines around the Internet behave in just the same way, in a cycle of taking, then giving, that fills the pipes while sharing them equally. It's an amazing global outpouring of self-denial, like the "after you" protocol two people use when they approach a door at the same time—but paradoxically, the Internet version happens between complete strangers, even fierce commercial rivals, billions of times every second.

The commercial stakes could hardly be higher. Services like YouTube, eBay, Skype, and iTunes are all judged by how much Internet capacity they can grab for you, as are the Internet phone and TV services provided by the carriers themselves. Some of these companies are opting out of TCP's sharing regime, but most still allow TCP to control how much they get—about 90 percent of the 200 000 terabytes that cross the Internet each second.

THIS EXTRAORDINARY SPIRIT of global cooperation stems from the Internet's early history. In October 1986, Internet traffic persistently overran available capacity—the first of a series of what were called congestion collapses. The TCP software of the day continued to try to retransmit, aggravating the problem and causing everyone's throughput to plummet for hours on end. By mid-1987 Van Jacobson, then a researcher at Lawrence Berkeley National Laboratory, had coded a set of elegant algorithms in a patch to TCP. (For this he received the IEEE's prestigious Koji Kobayashi Computers and Communications Award in 2002.)

Jacobson's congestion control accorded well with the defining design principle of the Internet: traffic control is consigned to the computers around the edges of the Internet (using TCP), while network equipment only routes and forwards packets of data (using IP).

The combination of near-universal usage and academic endorsement has gradually elevated TCP's way of sharing capacity to the moral high ground, altering the very language engineers use. From the beginning, equal rates were not just "equal," they were "fair." Even if you don't use TCP, your protocol is considered suspect if it's not "TCP-friendly"—a cozy-sounding idea meaning it consumes about the same bit rate as TCP would.

SADLY, AN EQUAL BIT RATE for each data flow is likely to be extremely unfair, by any realistic definition. It's like insisting that boxes of food rations must all be the same size, no matter how often each person returns for more or how many boxes are taken each time.

Consider a neighborhood network with 100 customers, each of whom has a 2-megabit-per-second access line connected to a single shared 10 Mb/s regional link. The network provider can get away with such a thin shared pipe because most of the customers—let's say 80 of the 100—don't use it continuously, even over the peak period. These people might think they are constantly clicking at their browsers and getting new e-mail, but their data transfers might be active perhaps only 5 percent of the time.

However, there are also 20 heavy users who download continuously, perhaps using file-sharing programs that run unattended. So at any one moment, data is flowing to about 24 users—all 20 heavy users, and 4 of the 80 light ones. TCP gives 20 shares of the bottleneck capacity to the heavy users and only 4 to the light ones. In a few moments, the 4 light users will have stepped aside and another 4 will take over their shares. However, the 20 heavy users will still be there to claim their next 20 shares. They might as well have dedicated circuits!

It gets even worse. Any programmer can just run the TCP routine multiple times to get multiple shares. It's much like getting around a food-rationing system by duplicating ration coupons.

This trick has always been recognized as a way to sidestep TCP's rules—the first Web browsers opened four TCP connections. Therefore, it would have been

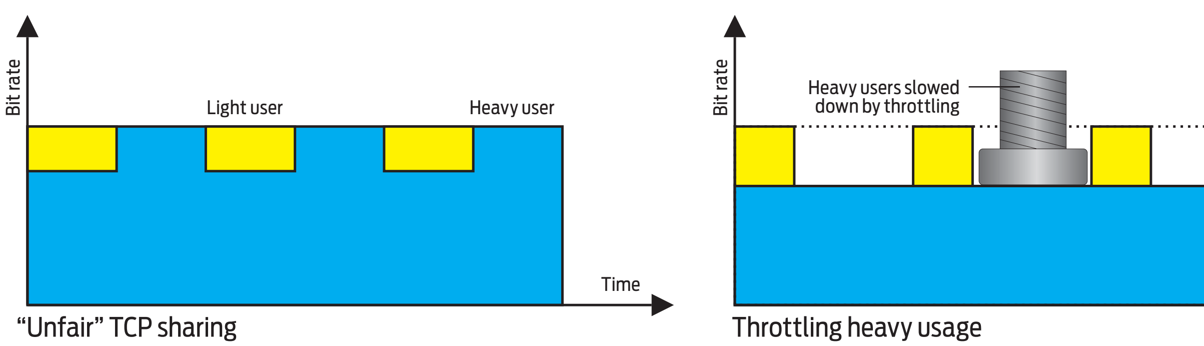

THROTTLE THIS: Throttling tries to correct today's TCP system [left] by clamping down on heavy users [center], but the technique misses a trick. With weighted TCP sharing [right], light users can go superfast, so they finish sooner, while heavy users slow only fleetingly, then catch up. All this can be done without any prioritization in the network.

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

remarkable if this ploy had not become more common.

A number of such strategies evolved through innocent experimentation. Take peer-to-peer file sharing—a common way to exchange movies over the Internet, one that accounts for a large portion of all traffic. It involves downloading a file from several peers at once. This parallel scheme, sometimes known as swarming, had become routine by 2001, built into such protocols as BitTorrent.

The networking community didn't immediately view connecting with many machines as a circumvention of the TCP-friendliness rule. After all, each transfer used TCP, so each data flow "correctly" got one share of any bottleneck it encountered. But using parallel connections to multiple machines was a new degree of freedom that hadn't been thought of when the rules were first written. Fairness should be defined as a relation between people, not data flows.

Peer-to-peer file sharing exposed both of TCP's failings. First, a file-sharing program might be active 20 times as often as your Web browser, and second, it uses many more TCP connections, typically 5 or even 50 times as many. Peer-to-peer thus takes 100 or 1000 times as many shares of Internet bottlenecks as a browser does.

Returning to our 100 broadband customers: if they were just browsing the Web and exchanging e-mail, each would get nearly the full benefit of a 2 Mb/s access pipe—if 5 customers were active at a time, they'd just squeeze into the 10Mb/s shared pipe. But if even 20 users started continuous parallel downloading, the TCP algorithm would send everyone else's bit rate plummeting to an anemic 20 kilobits per second—worse than dial-up! The problem isn't the peer-to-peer protocols; it's TCP's sharing rules.

WHY CAN'T THE SERVICE provider simply upgrade that stingy 10 Mb/s shared pipe? Of course, some upgrades are necessary from time to time. But as a general approach to the problem of sharing, adding capacity is like throwing water uphill.

Imagine two competing Internet service providers, both with this 80:20 mix of light and heavy users. One provider quadruples its capacity; the other doesn't. But TCP still doles out the upgrader's capacity in the same way. So the light users, who used to have a measly 20 kb/s share, now get a measly 80 kb/s—still barely better than dial-up. But now the 80 light users must pay substantially more for four times the long-distance capacity, which they hardly get to use. No rational network operator would upgrade under these conditions—it would lose most of its customers.

But there is plenty of evidence that Internet service providers are continuing to add capacity. This is partly explained by government subsidies, particularly in the Far East. Equivalently, weak competition, typical in the United States, allows providers to fund continued investment through higher fees without the risk of losing customers. But in competitive markets, common in Europe, service providers have had to attack the root cause: the way their capacity is shared.

Network providers often don't allow TCP to give all the new capacity straight to the heavy users. Instead they impose their own sharing regimes on their customers, thus overriding the worst effects of TCP's broken regime. Some limit, or "throttle," the peak-time bit rate of the peer-to-peer customers. Others partition the pipe to prevent heavy users encroaching on lighter ones. Increasingly, the share of Internet capacity you actually get is the result of this tussle between TCP and the service providers' allocation schemes.

THERE'S A FAR BETTER SOLUTION than fighting. It would allow light browsing to go blisteringly fast but hardly prolong heavy downloads at all. The solution comes in two parts. Ironically, it begins by making it easier for programmers to run TCP multiple times—a deliberate break from TCP-friendliness.

Programmers who use this new protocol to transfer data will be able to say "behave like 12 TCP flows" or "behave like 0.25 of a TCP flow." They set a new parameter—a weight—so that whenever your data comes up against others all trying to get through the same bottleneck, you'll get, say, 12 shares, or a quarter of a share. Remember, the network did not set these priorities. It's the new TCP routine in your own computer that uses these weights to control the number of shares it takes from the network.

At this point in my argument, people generally ask why everyone won't just declare that they each deserve a huge weight. The answer to the question involves a trick that gives everyone good reason to use the weights sparingly—a trick I'll get to in a minute. But first, let's check how this scheme ensures the lightning-fast browsing rates I just promised.

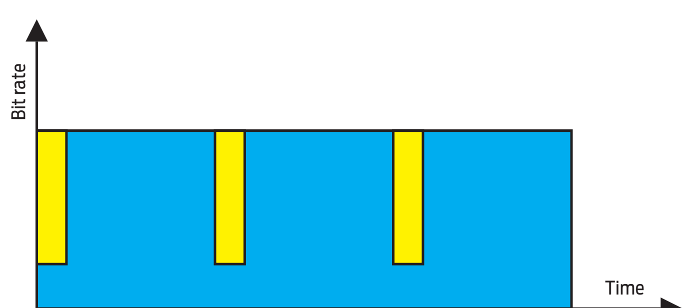

Weighted TCP sharing

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

The key is to set the weights high for light interactive usage, like surfing the Web, and low for heavy usage, such as movie downloading. Whenever these uses conflict, flows with the higher weighting—those from the light users—will go much faster, which means they will also finish much sooner. Then the heavy flows can expand back to a higher bit rate sooner than otherwise. This is why the heavy flows will hardly take any longer to complete. The weighting scheme uses the same strategy as a restaurant manager who says, "Get those individual orders out right away, then come serve this party of 12." But today's Internet has the balance of the weights exactly the other way around.

That brings us to the second part of the problem: how can we encourage everyone to flip the weights? This task means grappling with something that is often called "the tragedy of the commons." A familiar example is global warming, where everyone happily pursues what's best for them—leaving lights on, driving a big car—despite the effect this may have on everyone else through the buildup of carbon dioxide and other greenhouse gases.

On the Internet, what matters isn't how many gigabytes you download but how many you download when everyone else is trying to do the same. Or, more precisely, it's the volume you download weighted by the prevailing level of congestion. Let's call this your congestion volume, measured in bytes. Think of it as your carbon footprint for the Internet.

As with CO_2, the way to cut back is to set limits. Imagine a world where some Internet service providers offer a deal for a flat price but with a monthly congestion-volume allowance. Note that this allowance doesn't limit downloads as such; it limits only those that persist during congestion. If you used a peer-to-peer program like BitTorrent to download 10 videos continuously, you wouldn't bust your allowance so long as your TCP weight was set low enough. Your downloads would draw back during the brief moments when flows came along with higher weights. But in the end, your video downloads would finish hardly later than they do today.

On the other hand, your Web browser would set the weights high for all its browsing because most browsing comes in intense flurries, and so it wouldn't use up much of your allowance. Of course, server farms or heavy users could buy bigger congestion quotas, and light users might get Internet access with a tiny congestion allowance—for a lower flat fee.

But there's a snag. Today Internet service providers can't set congestion limits, because congestion can easily be hidden from them. As we've said, Internet congestion was intended to be detected and managed solely by the computers at the edge—not by Internet service providers in the middle.

Minimizing Congestion: A Refeedback Primer

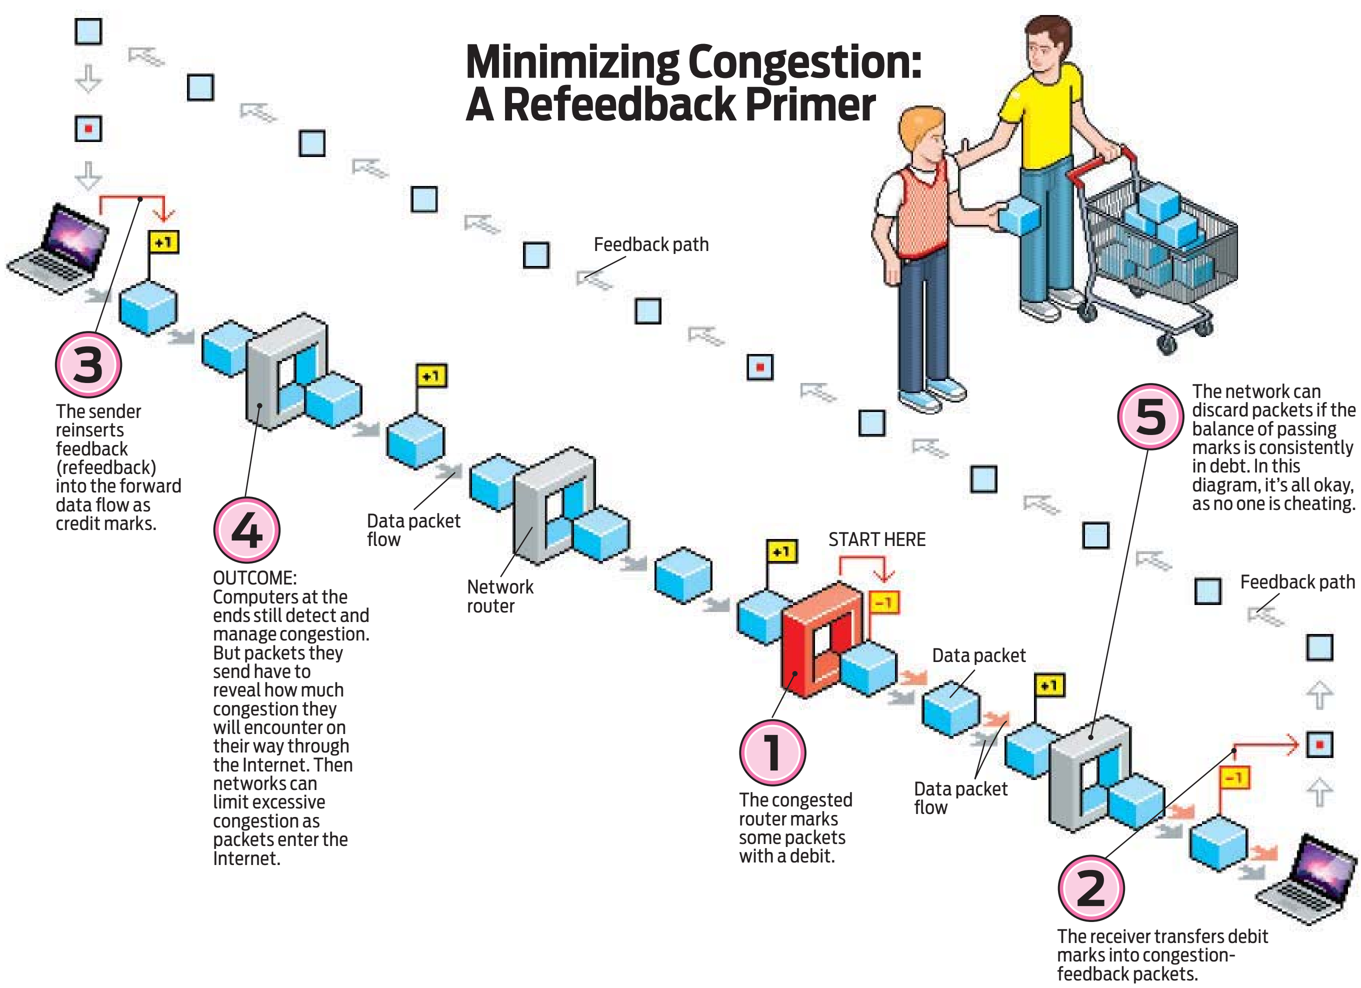

QUICKHONEY

Certainly, the receiver does send feedback messages about congestion back to the sender, which the network could intercept. But that would just encourage the receiver to lie or to hide the feedback—you don't have to reveal anything that may be used as evidence against you.

Of course a network provider does know about packets it has had to drop itself. But once the evidence is destroyed, it becomes somewhat tricky to hold anyone responsible. Worse, most Internet traffic passes through multiple network providers, and one network cannot reliably detect when another network drops a packet.

Because Internet service providers can't see congestion volume, some limit the straight volume, in gigabytes, that each customer can transfer in a month. Limiting total volume indeed helps to balance things a little, but limiting congestion volume does much better, providing extremely fast connections for light users at no real cost to the heavy users.

My colleagues and I have figured out a way to reveal congestion so that limits can be enforced. We call it "refeedback." Here's how it works. Recall that today the computers at each end of an exchange of packets see congestion, but the networks between them can't. So we built on a technique called Explicit Congestion Notification—the most recent change to the TCP/IP standard, made in 2001. Equipment that implements that change marks packets during impending congestion rather than doing nothing until forced to drop them. The marks—just a change in a single bit—let the network see congestion directly, rather than inferring it from gaps in the packet stream. It's also particularly neat to be able to limit congestion before anyone suffers any real impairment.

Although the 2001 reform reveals congestion, it is only visible downstream of any bottleneck as packets leave the network. Our scheme of refeedback makes congestion visible to the upstream network before it enters the Internet, where it can be limited.

Refeedback introduces a second type of packet marking—think of these as credits and the original congestion markings as debits. The sender must add sufficient credits to packets entering the network to cover the debit marks that are introduced as packets squeeze through congested Internet pipes. If any subsequent network node detects insufficient credits relative to debits, it can discard packets from the offending stream.

To keep out of such trouble, every time the receiver gets a congestion (debit) mark, it returns feedback to the sender. Then the sender marks the next packet with a credit. This reinserted feedback, or refeedback, can then be used at the entrance to the Internet to limit congestion—you *do* have to reveal everything that may be used as evidence against you.

Refeedback sticks to the Internet principle that the computers on the edge of the network detect and manage congestion. But it enables the middle of the network to punish them for providing misinformation.

The limits and checks on congestion at the borders of the Internet are trivial for a network operator to add. Otherwise, the refeedback scheme does not require that any new code be added to the network's equipment; all it needs is that standard congestion notification be turned on. But packets need somewhere to carry the second mark in the "IP" part of the TCP/IP formula. Fortuitously, this mark can be made, because there is one last unused bit in the header of every IP packet.

In 2005, we prepared a proposal documenting all the technical details and presented it to the Internet Engineering Task Force (IETF), the body that oversees Internet standards.

At this point, the story gets personal. Because I had set myself the task of challenging the entrenched principle of TCP-friendliness—equality of flow rates for all TCP connections—I decided to talk only about the change to IP, omitting any mention of weighted TCP. Instead I played up some other motivations for adding refeedback to IP. I even showed how refeedback could enforce equal flow rates—pandering to my audience's faith while denying my own. But I just looked like yet another mad researcher pushing a solution without a problem.

After a year of banging my head against a wall, I wrote an angry but—I trust—precise attack on the dogma that equal flow rates were "fair." My colleagues got me to tone it down before I posted it to the IETF; evidently I'd softened it enough at least to be invited to present my ideas at a plenary session in San Diego late in 2006. The next day, a nonbinding straw poll of the large audience showed widespread doubt about using TCP-friendliness as a definition of fairness. Elwyn Davies of the Internet Architecture Board e-mailed me, saying, "You have identified a real piece of myopia in the IETF."

I was hardly the first to challenge these myths. In 1997 Frank P. Kelly, a professor at the University of Cambridge, put together some awe-inspiringly elegant and concise mathematical arguments to prove that the same weighted sharing would maximize the value that users get from their Internet throughput. However, to create the right incentives, he proposed varying the prices charged for the packets as they were received, and everyone balked. People like to control, in advance, what they will pay.

Objections to Kelly's pricing scheme blinded the Internet community to all the other insights in his work—particularly the message that equalizing flow rates was not a desirable goal. That's why my team built the refeedback mechanism around his earlier ideas—to limit congestion within flat fees, without dynamic pricing.

Everyone's subsequent obsession with bandwidth hogs, and thus with volume, is also misdirected. What matters is congestion volume—the CO_2 of the Internet.

Meanwhile, our immediate task is to win support in the Internet community for limiting congestion and for a standards working group at the IETF to reveal the Internet's hidden congestion. The chosen mechanism may be refeedback, but I won't be miffed if something better emerges, so long as it makes the Internet as simple and as fast. ❑

TO PROBE FURTHER Details on Internet fairness and the refeedback project can be found at http://www.cs.ucl.ac.uk/staff/bbriscoe/projects/refb/.

Bots Get Smart

Can video games breathe new life into AI research?

By Jonathan Schaeffer, Vadim Bulitko & Michael Buro

You're following a gloomy corridor into a large boiler room, dimly lit by a flickering fluorescent lamp and echoing with the rhythms of unseen machinery. Three enemy soldiers suddenly appear on a catwalk high above the floor. They split up, one of them laying down suppressive fire, which forces you to take cover. Although you shoot back, the attackers still manage to creep forward behind a curtain of smoke and flying debris.

Moments later, a machine gun rings out, and you are cut down in a shower of bullets. Then, as you lie dying, you glimpse the soldier who flanked you from behind while his two buddies drew your attention.

Thankfully, it was only a video game, so in fact you're not mortally wounded. Still, your ego might well be bruised, because you were not only outgunned but also outsmarted by artificial intelligence (AI).

The game is called *F.E.A.R.*, short for First Encounter Assault Recon, and its use of AI, along with its impressive graphics, are its prime attractions. The developer, Monolith Productions of Kirkland, Wash., released it in 2005 to rave reviews, including the GameSpot Web site's Best Artificial Intelligence award. Such recognition means a lot to the game's creators, who face stiff competition in what has become a multibillion-dollar industry.

The game is a far cry from the traditional diversions that AI researchers like ourselves have long studied, such as chess and checkers. Whereas the goal in the past was to write computer programs capable of beating expert players at such board games, now the metric of success for AI is whether it makes video games more entertaining.

Because a high fun factor is what sells, the video-game industry has become increasingly keen to make use of developments in AI research—and computer scientists have taken notice. A watershed came in 2000, when John E. Laird, a professor of engineering at the University of Michigan, and Michael van Lent, now chief scientist at Soar Technology, in Ann Arbor, Mich., published a call to arms that described commercial video games as "AI's killer application." Their point was that research to improve AI for such games would create spin-offs in many other spheres.

The main challenge is to make computer-generated characters—dubbed bots—act realistically. They must, of course, look good and move naturally. But, ideally, they should also be able to engage in believable conver-

BRYAN CHRISTIE DESIGN

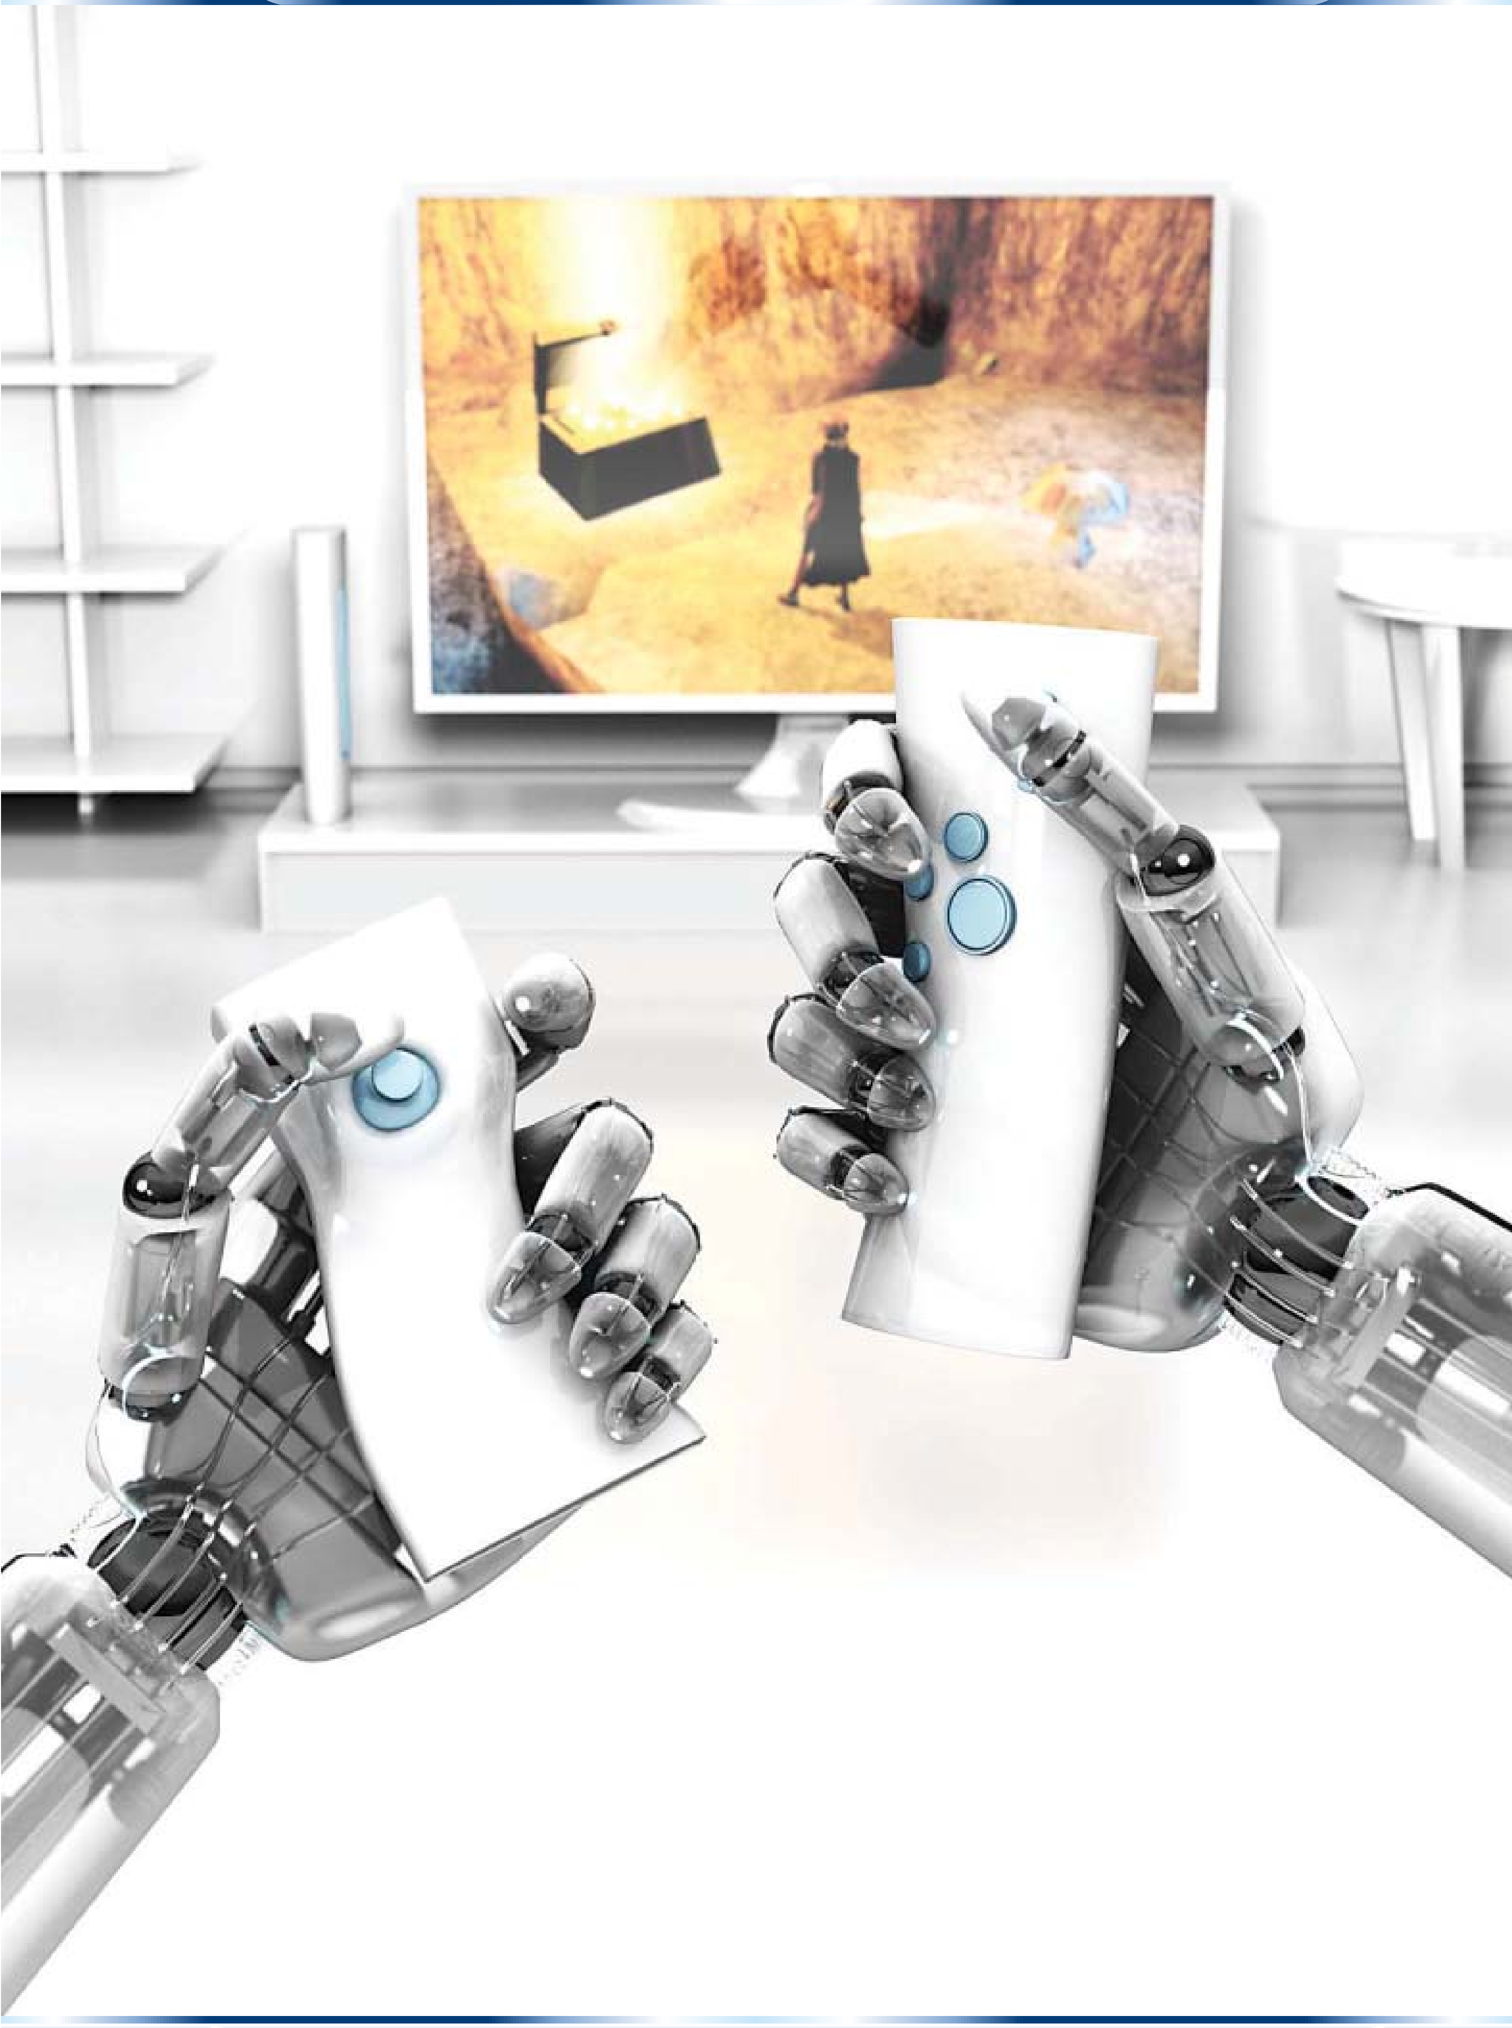

In the not-so-distant future, game bots will routinely use sophisticated AI techniques to shape their behavior

sations, plan their actions, find their way around virtual worlds, and learn from their mistakes. That is, they need to be smart.

Today many video games create only an illusion of intelligence, using a few programming tricks. But in the not-so-distant future, game bots will routinely use sophisticated AI techniques to shape their behavior. We and our colleagues in the University of Alberta GAMES (Game-playing, Analytical methods, Minimax search and Empirical Studies) research group, in Edmonton, Canada, have been working to help bring about such a revolution.

The AI of *F.E.A.R.* is based loosely on an automated planner called STRIPS (for STanford Research Institute Problem Solver), which Richard E. Fikes and Nils J. Nilsson, both now of Stanford University, developed way back in 1971. The general idea of STRIPS was to establish one or more goals along with a set of possible actions, each of which could be carried out only when its particular preconditions were satisfied. The planning system kept track of the physical environment and determined which actions were allowed. Carrying out one of them in turn modified the state of the environment, which therefore made other actions possible.

The designers of *F.E.A.R.* gave its soldiers such goals as patrolling, killing the player's character, and taking cover to protect their own virtual lives. The makers of the game also gave each kind of bot a set of possible actions with which to accomplish each of its goals. One advantage of this approach is that it saves the developers the burden of trying to specify a response to every situation that might arise. Further, it allows seemingly intelligent behaviors to appear almost magically—such as the maneuver described above.

In that instance, the three attackers were carrying out two types of basic actions. One is to move to covered positions that are as close as possible to the player's character. The other is simply to move around obstacles. The combination creates something that was not explicitly programmed into the game at all: a devastating flanking maneuver.

The spontaneous emergence of such complex behaviors is important because it provides a sense of deeper intelligence. That's really what gets your heart pounding when you play the game. But you'd also like your adversaries to become more cunning over time, and *F.E.A.R.* has no mechanism for accomplishing that.

Why do bots need to get smarter? Imagine a game of badminton in which your opponent always reacts to your serves in the same way, always falls for your drops, and never attempts to anticipate your smashes. It would be a boring match. Up until recently, AI had been able to offer video gamers no better: the imps of *Doom*, released in 1993, never shoot their fireballs preemptively, and the civil-protection officers in *Half-Life 2* (2004) always take the nearest cover while reloading their weapons—to mention just a couple of things players experience with two well-known releases.

The standard solution is to add an element of randomness to the code that controls a bot's decision making. Doing so varies a player's experience, but the result does not necessarily come across as being intelligent.

A better approach is for the computer to learn about the player and to adapt a bot's tactics and strategy appropriately. Of course, you don't want the bot to become so good that it will win all the time; you just want it to give the human player a good run for the money. This capability, known as machine learning, is found in very few commercial games: *Creatures*, from the now-defunct Creature Labs, employed machine learning as early as 1997, as did *Black & White*, developed by the UK-based Lionhead Studios a few years later. But most video games are not able to "learn" on the fly or otherwise adapt to the person playing. Our group is hoping to push things forward in this regard using a system we've created for research purposes called PaSSAGE, which stands for Player-Specific Stories via Automatically Generated Events.

PaSSAGE, as its name implies, is all about storytelling, which has long been a staple of various role-playing games. But video games of all types rely to some extent on engaging storytelling. You can categorize such games by the way they vary their repertoire to appeal to different people.

Some games—*Half-Life* (2001), for example—are immensely popular even though they feature just a single linear story. So good scriptwriting can clearly go a long way. Other games, such as *Star Wars: Knights of the Old Republic* (2003), offer several alternatives to the main plot. This gives you the impression that you can shape your virtual fate—what psychologists call a sense of agency. That feeling of being in control is usually limited, however, because the branching plot lines often merge later on.

Titles like *The Elder Scrolls IV: Oblivion* (2006) and *S.T.A.L.K.E.R.: Shadow of Chernobyl* (2007) work similarly, taking one main story and complementing it with episodes drawn from a library of side quests. Other games, such as *The Sims 2* (2005), go a step further by dispensing with a scripted plot altogether and creating an open-ended world in which players can effectively design their own happenings.

Although each of these techniques has enjoyed success, they all force the designer to make a trade-off between scriptwriter expressiveness and player agency. The approach we've taken with PaSSAGE avoids that conundrum by having the computer learn players' interests and preferences and mold the story to suit them as the game progresses.

PaSSAGE uses the same game engine as *Neverwinter Nights*, a fantasy adventure set in medieval times, produced by BioWare of Edmonton. With PaSSAGE, scriptwriters determine only the most general arc to the story and provide a library of possible encounters the player's character may have. The computer studies the player as he or she progresses and cues in the kinds of experiences that are most desired. For instance, if you like fighting, the game will provide ample opportunities for combat. If you prefer to amass riches, the game will conjure up ways for you to be rewarded for your actions. The software is able to make the sequence of events globally consistent by maintaining a history of the virtual world's chang-

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

ing state and modifying the player's future encounters appropriately. The game will therefore always appear to make sense, even though it unfolds quite differently for different people—or even for the same person as his moods and tastes change.

MACHINE LEARNING can also be used to formulate the tactics that bots use, a job that now must be handcrafted by a game's designers. Pieter Spronck and his colleagues, of the University of Tilburg, in the Netherlands, demonstrated this ability in 2005 using *Neverwinter Nights*. Spronck had one computer play against computerized opponents, programming it to get better over time by choosing the combat tactics that most often led to victory.

Members of our research group have been following through on Spronck's work with *Neverwinter Nights*, using a different learning algorithm. Other colleagues of ours at the University of Alberta aim to do something similar with a multiplayer online game called *Counter-Strike* (2003), which pits a group of terrorists against a squad of antiterrorist commandos. Each character can be controlled either by a person or by the computer. As with *F.E.A.R.*, players view the virtual world from the perspective of the characters they manipulate, making *Counter-Strike* an example of what's known as a first-person-shooter game.

This project has so far produced a formal system for analyzing and classifying a team's opening moves. That may not sound like much, but this task proved immensely challenging, because positions and actions are not

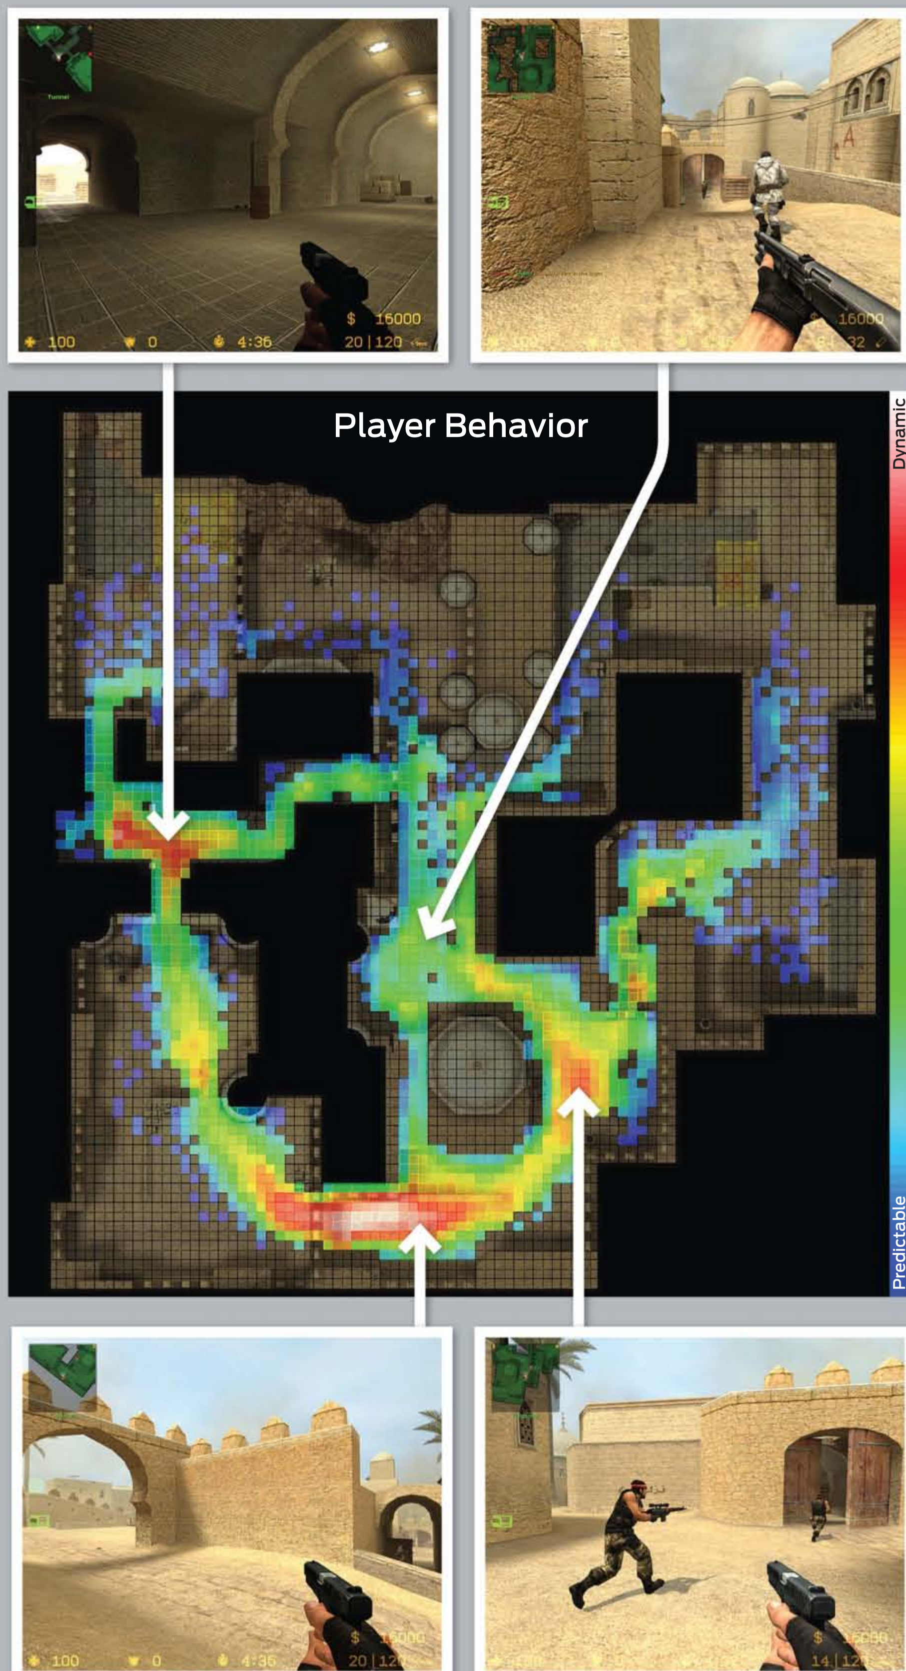

MAPPING OUT STRATEGIES: Tournament-play logs reveal the strategies experts use in the game *Counter-Strike*. But figuring out how to assess the state of the play in this first-person-shooter game is tricky, because there are so many possibilities for where the characters can be and for what they might be doing. This game map ("Dust2" scenario) shows where dynamic play [hot colors] is concentrated.

MAP: CHRIS RAYNER/THE UNIVERSITY OF ALBERTA GAMES GROUP; COUNTER-STRIKE: SOURCE IMAGES USED BY PERMISSION OF VALVE CORPORATION

The computer may need to move thousands of bots around their virtual stomping grounds without the action grinding to a crawl

nearly as constrained as they are in a game like chess. Researchers in our group have used this formalism to analyze computer logs of more than 50 hours of tournament-level play between seasoned *Counter-Strike* teams. Soon, we expect, computer bots programmed to learn tactics from such logs will play reasonably well—doing things a person might do. It'll be a long time before these bots will be able to beat expert human players, though. But that's not the objective, after all—they just need to make for entertaining adversaries.

Jeff Orkin and Deb Roy of MIT are undertaking a similar effort with something they call *The Restaurant Game*, for which they are applying machine learning to the task of making bots speak and act believably in social settings. In this case, the bots' behaviors are based on observations gleaned from more than 10 000 sessions of human play.

Machine learning can also pay off for poker, which has become an especially hot game in recent years with the explosion of opportunities for playing it online. The strongest programs for two-player fixed-bet-size poker attempt to calculate the mathematically optimal solution for winning each hand. It turns out that finding such solutions is computationally infeasible, at least right now—there are just too many possible combinations of cards and betting sequences. But members of our research group have devised ways to calculate near-optimal strategies using certain simplifying assumptions. For example, instead of allowing four rounds of betting—which is permitted in competition poker—the program sets the limit at three. By further reducing the complexity of the game in clever ways, the computational burden can be reduced to a reasonable level. BioTools, a commercial spin-off of our research group in Edmonton, has incorporated some of our group's work in this area in its Poker Academy software.

Although this program plays poker pretty well, it can't yet do what is most required—spot and exploit the other player's weaknesses. Figuring out how to program a computer to do that is extraordinarily hard. Why so? Studying an opponent should be easy, after all—and it is, but only if you have thousands of poker hands to analyze. What do you do if you have only a few? To make matters worse, human poker players make a point of changing their style so as to be hard to predict.

Right now, the best poker-playing programs to come out of our research group will make money off your average human player, and they are beginning to beat even some of the best in the world in organized competitions. This suggests that poker is just now joining the ranks of chess and checkers—games at which computers have trounced even world champions.

ONE LESSON that computer scientists learned from working on chess and checkers is that programs must strike a balance in how they decide what move to make next. At one extreme, the computer can look all the way to the end of a game, examine every possible final position, and evaluate whether each one constitutes a win, a draw, or a loss. Then it can work backward from those possibilities, assuming best play by both sides at every stage, to select the optimal move. But searching that far ahead would take a great deal of time—for chess, enough for the sun to burn out.

The alternative is to use an evaluation function that incorporates knowledge of the game, enough to go beyond just recognizing an outright win to sense, rather, the slightest inkling of an advantage. In the ideal case, such a program would play perfectly while looking only a single move ahead. Of course, such a sophisticated evaluation would also require a lot of computational power.

In actuality, chess-playing programs operate somewhere between these two extremes. The computer typically examines all the possibilities several moves ahead and evaluates each, say, by tallying points, having assigned a different number of points to a pawn, a knight, a rook, and so forth. The computer then works backward to the current board position. The result is a ranking of all the available next moves, making it easy to pick the best one.

The trade-off between blind searching and employing specialized knowledge is a central topic in AI research. In video games, searching can be problematic because there are often vast sets of possible game states to consider and not much time and memory available to make the required calculations. One way to get around these hurdles is to work not on the actual game at hand but on a much-simplified version. Abstractions of this kind often make it practical to search far ahead through the many possible game states while assessing each of them according to some straightforward formula. If that can be done, a computer-operated character will appear as intelligent as a chess-playing program—although the bot's seemingly deft actions will, in fact, be guided by simple brute-force calculations.

Take, for example, the problem of moving around intelligently in a virtual world—such as finding the shortest path to take from one spot to another. That's easy enough to figure out if you can fly like a crow. But what if you're earthbound and there are obstacles to contend with along the way?

A general algorithm for determining the best route between two points on a map has been around since the late 1960s. The problem with this scheme—known as A*—is that the amount of time the solution takes to compute scales with the size of the territory, and the domains of video games are normally quite large. So there isn't time to calculate the optimal path in this way. In some games, the computer needs to move hundreds—or even thousands—of bots around their virtual stomping grounds without the action grinding to a crawl, which means that computation times must often be kept to just a few milliseconds per bot.

To address this issue, our research group has developed a series of pathfinding algorithms that simplify the problem. Rather than considering each of the vast number of possible positions each bot can take, these algorithms seek good paths by using coarser versions of the game map. Some of

CHACUN À SON GOÛT: Adapting games to suit the tastes of individual players is possible with PaSSAGE (Player-Specific Stories via Automatically Generated Events), a system the authors have created for research purposes, based on the game engine of BioWare's *Neverwinter Nights*, a fantasy adventure set in medieval times. If the person playing PaSSAGE likes fighting, the game will cue in combat opportunities, including perhaps a battle with a giant spider [left]. If, however, the player prefers to amass wealth, the game can offer such possibilities, too [right].

THE UNIVERSITY OF ALBERTA GAMES GROUP/ATARI

these algorithms can use a set amount of time for planning each move, no matter how vast the playing field, so they can be applied to game worlds of any size and complexity. They are also suitable for environments that change frequently, for instance when paths are blocked, bridges destroyed, doors closed, and so forth. BioWare will be using some of our group's pathfinding algorithms in its forthcoming *Dragon Age: Origins*.

This same general approach can help computers master real-time strategy games, such as the *Warcraft* series, introduced in 1994, which was developed by Blizzard Entertainment of Irvine, Calif. In this popular genre, players control armies of game characters that work together to gather resources and battle enemies on uncharted terrain. The fast pace and large numbers of bots make these games too complex for today's AI systems to handle, at least at a level that would challenge good human players.

Our research tries to address this problem by considering only the relatively small set of high-level strategies each player can follow, such as having your army of characters rush the opponent or expand aggressively so as to take over more territory. The computer simulates what the outcome would be, given the current state of play, if each side picked one of these strategies and kept to it for the duration of the game. By taking into account whether its human opponent is using all or just a few particular strategies, the computer can choose the counterstrategy that is most likely to succeed. This approach works better than the scripted maneuvers computers now employ in real-time strategy games when pitted against a human player.

THE NEED for better AI in commercial video games is readily apparent—especially to the people playing them. And their thirst for more computer-generated intelligence will only continue to grow. Yet game makers rarely have the time or resources to conduct the research required to solve the many thorny problems involved, which is why they have come to recognize the value of engaging the scholarly community—a community that is hard at work in such places as Georgia Tech; Simon Fraser University, in Burnaby, B.C., Canada; the University of Teesside, in the UK; and the Technical University of Lisbon, to name but a few of the many research centers around the world involved in this kind of work.

With the increased participation of academics in game-related AI research, it will not be long before major improvements are apparent in the quality of the games entering the market. But there is a more significant reason to applaud the growing interest of AI researchers in the video-game industry—something Laird and van Lent pointed out to us and other computer scientists nearly a decade ago. The work we must do to make games feel more realistic will also take us a long way toward our ultimate goal of developing general-purpose machine intelligence. Now *that* sounds like a smart move. ❑

TO PROBE FURTHER
The full range of interests of the authors' research group is available at http://games.cs.ualberta.ca. The inner workings of F.E.A.R. *are detailed in "Three States and a Plan: The A.I. of* F.E.A.R.*," by Jeff Orkin, available at http://web.media.mit.edu/~jorkin/gdc2006_orkin_jeff_fear.pdf. Laird and van Lent's call to action can be found in "Human-level AI's Killer Application: Interactive Computer Games,"* AI Magazine, *Summer 2001.*

BAYLOR
UNIVERSITY

Baylor is seeking candidates from all areas of electrical and computer engineering for two positions:

(1) a tenure-track position at the assistant/ associate professor level, and (2) a full time lecturer.

Candidates must have excellent verbal and written communication skills, a commitment to teaching excellence and professional activities, and is a follower of Christ. Tenure-track applicants must have an earned doctorate, and demonstrated research achievement with potential to acquire external funding. Further, the successful tenure track candidate will also be expected to contribute to the department's progress toward increased research activity and the development of a Ph.D. program. Lecturer applicants will have an earned doctorate; although, outstanding applicants with a master's degree will also be considered. The successful candidate for lecturer will have demonstrated laboratory experience and effective classroom skills. Both positions will contribute to the current state of teaching excellence.

Chartered by the Republic of Texas in 1845, Baylor University is the oldest university in Texas, has an enrollment of 14,000 students, and is a member of the Big XII Conference. The Baylor campus is located in Waco along Interstate 35, a major route for international commerce. The community offers an array of educational, cultural, recreational, and residential opportunities and is within easy driving distance of other major cities and universities in Texas.

Baylor offers ABET/EAC-accredited B.S. programs in Electrical and Computer Engineering, Mechanical Engineering, and Engineering. The University is in mid-course to achieve a 10-year vision called Baylor 2012 which will strengthen its academic and Christian character. As part of that effort, the ECE department has established new master's programs (Master of Science in Electrical and Computer Engineering, Master of Science in Biomedical Engineering, Master of Engineering, joint Master of Engineering/Master of Business Administration), with accompanying growth in its faculty, student body and facilities. Hallmarks of Baylor engineering include a commitment to engineering education in a supportive Christian environment, faculty collegiality, small class sizes, a near-100% pass rate on the FE examination, and a high U.S. News and World Report peer-group ranking.

Applications will be reviewed beginning immediately and will be accepted until the positions are filled. To ensure full consideration, however, applications should be received by Jan. 9, 2009. Applications must include: 1) a letter of interest that identifies the position sought; 2) complete current vita, 3) statement of teaching and, for the tenure-track candidate, research interests, 4) statement of Christian faith and service, and 5) the names and contact information for at least three professional references.

Additional information is available at

http://www.ecs.baylor.edu.

Send materials to Dr. Robert J. Marks II, Baylor University, One Bear Place #97356, Waco, TX 76798-7356, or by email to Robert_Marks@baylor.edu.

Baylor is a Baptist university affiliated with the Baptist General Convention of Texas. As an Affirmative Action/Equal Employment Opportunity employer, Baylor encourages minorities, women, veterans, and persons with disabilities to apply.

Dakota Power

Chief Technology Officer and Research Engineer

Dakota Power, LLC, is seeking a Chief Technology Officer and a Research Engineer to join the team developing a revolutionary lightweight electric drive system. The Dakota Power drive system employs unique patented innovations to respond to the urgent needs for lightweight and fuel efficient power systems for both military and commercial applications.

The Chief Technology Officer will lead the design, testing, and evaluation program to optimize the efficiency and power density of the Dakota Power technologies for scalable electric drive systems. Ideally, the Chief Technology Officer will bring successful experiences in electric machine and switching circuit design using advances in modeling and evaluation techniques. An electrical engineering degree is required and a graduate degree is preferred.

The Research Engineer will lead the development of computer finite element and scalable design models to evaluate the efficiency and power density of the electric drive systems. Additionally, the Research Engineer will have responsibilities for the design, integration, and evaluation of systems to comply with military operational requirements. A degree in electrical engineering or computer science is required and a graduate degree is preferred.

Dakota Power, LLC, is located in Rapid City, South Dakota. Compensation is commensurate with experience and includes a competitive salary package, insurance, retirement, vacation and sick leave.

Additional information is available at **www.dplwed.com**

Please email your application detailing your experiences and three references to **hr@dplwed.com**. Dakota Power, LLC, is committed to diversity and equality in employment.

香港城市大學
City University of Hong Kong

The University invites applications and nominations for the following posts. Candidates with applied research achievements will receive very positive consideration. Relevant experience in business and industry will be a definite asset.

Professor (Chair)/Professor/ Associate Professor/Assistant Professor Department of Electronic Engineering *[Ref. A/551/36]*

Duties : Teach and supervise student projects at both undergraduate and postgraduate levels, conduct frontline research, and develop new research directions and courses when necessary.

Requirements : A PhD degree in Electronic Engineering or related disciplines, preferably in bio-engineering, nano-technology, system-on-chip and energy-related areas. Excellent academic records and English communication skills are essential. Ability to teach electronic engineering courses at both undergraduate and postgraduate levels, conduct original research, publish research results in high-impact journals and attract external research funding required. For the posts of Professor (Chair) and Professor, candidates should have demonstrated leadership in a specialized area, and outstanding records of teaching and research.

Salary and Conditions of Service
Remuneration package will be highly competitive, commensurate with qualifications and experience. Initial appointment will be made on a fixed-term gratuity-bearing contract. Fringe benefits include leave, medical and dental schemes, housing benefits and passage allowance where applicable. (The remuneration package is currently under review.)

Application and Information
Information about the posts and the University is available at http://www.cityu.edu.hk or from the Human Resources Office, City University of Hong Kong, Tat Chee Avenue, Kowloon, Hong Kong [Fax : (852) 2788 1154 or (852) 2788 9334/email : hrojob@cityu.edu.hk]. Please send the nomination or application enclosing a current C.V. with evidence of teaching ability in English and a concise (up to one page) statement of research interests and teaching philosophy to the Human Resources Office. **Applications will be considered until the positions are filled**. Please quote the reference of the post in the application and on the envelope. The University reserves the right to consider late applications and nominations, and not to fill the positions. Personal data provided by applicants will be used for recruitment and other employment-related purposes.

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

The Department of Electrical and Computer Engineering, University of Utah, Salt Lake City, seeks applications to fill at least two tenure-track positions at the assistant, associate or full professor level for an interdisciplinary research cluster in ***Micro and Nanosystem Integration and Packaging***.

We are particularly interested in candidates with background in electronic micro/nano-system integration and packaging, biocompatible materials and packaging, solid state devices, reliability, testing, and micro/nano system modeling and simulation. Information on department research activities and curricula may be found on the web at **www.ece.utah.edu**. The web site also has information on five more positions available in the department, including a Chair search. Information on the College of Engineering can be found at **www.coe.utah.edu**. Successful candidates will conduct research with tenure track appointments in the Department of Electrical and Computer Engineering, but may also be appointed in other departments such as Materials Science, Bioengineering or Mechanical Engineering. Suitable candidates may be considered for joint appointments with the College of Science or the Medical School at the University of Utah.

These positions are part of the Utah Science, Technology and Research Initiative (USTAR), which was funded by the Utah State Legislature to attract focused teams of outstanding researchers who have the potential to help build major research programs and create new technology that can ultimately lead to commercial products and/or new industries for Utah. The USTAR initiative is also supporting a new interdisciplinary building which will house a new nanofabrication laboratory and characterization facilities that will cater to solid state devices, MEMS, sensor and packaging research and development, as well as the handling of biomedical samples. The building will facilitate communication for researchers such as the ones hired under this solicitation, from engineering, sciences and the medical school, as well as offering lab access for selected industrial stake holders. Information on the USTAR initiative can be found under **www.ustar.utah.gov**. Candidates for this initiative should have a demonstrated track record of successful, funded projects and an interest or track record in technology commercialization, entrepreneurial or industrial experience.

The positions are also associated with and partially supported by the **Fraunhofer Institute for Reliability and Microintegration IZM**, and leverage a strong collaborative and international research program with a Fraunhofer IZM branch laboratory in Utah. Fraunhofer support includes in-house access to Fraunhofer infrastructure, know-how, and resources. Selected positions may be associated with joint Fraunhofer appointments, possibly at a center director's or co-director's level.

Résumés with names, contact information for at least three references, and statements for research and teaching goals should be sent to Ms. Debbie Sparks, USTAR Faculty Search Committee, University of Utah, Electrical and Computer Engineering Department, 50 South Central Campus Drive, Room 3280, Salt Lake City, UT 84112-9206. Email applications are accepted at **dsparks@ece.utah.edu**. Applications will be reviewed starting September 1, 2008, and will be accepted until the positions are filled.

Faculty responsibilities include developing and maintaining an internationally recognized research program, effective classroom teaching at the undergraduate and graduate levels, and professional service. Applicants must hold a Ph.D. by the time of appointment. The University of Utah values candidates who have experience working in settings with students from diverse backgrounds and possess a strong commitment to improving access to higher education for historically underrepresented students. The University is an AA/EO employer, encourages applications from women and minorities, and provides reasonable accommodations for known disabilities of applicants and employees.

TEXAS A&M UNIVERSITY
The Department of Electrical and Computer Engineering

has openings for two tenure-track faculty at the assistant professor rank in the areas of power systems and power electronics. General areas of interest in the power systems area include systems aspects, such as stability, state estimation, steady state and dynamic modeling, deregulation and distributed generation. For the power electronics position, areas of interest include electronic power conversion and control, power supplies, inverters, power conditioning systems for Solar/Wind/Fuel-cells, electric power quality and harmonic mitigation. Candidates with expertise in the sustainable energy area will be given preference for both positions.

Applicants must have a Ph.D. or equivalent degree, or completion of all requirements by date of hire as well as demonstrated potential for quality teaching and research.

The Department of Electrical & Computer Engineering currently has 68 faculty members and its graduate program has been ranked in the top 20 in recent years. Further information about the department may be obtained by visiting

http://ece.tamu.edu.

Applicants should send a complete resume, including names and addresses of three references, to:
Dr. Costas N. Georghiades, Department Head
Texas A&M University
Department of Electrical and Computer Engineering
TAMU 3128 College Station, TX, 77843-3128.

The deadline for applications is February 1, 2009 but positions will remain open until filled.

Texas A&M University is an equal opportunity/affirmative action employer and actively seeks the candidacy of women and minorities.

GIFT POSTECH

POHANG UNIVERSITY OF SCIENCE AND TECHNOLOGY (POSTECH)

GRADUATE INSTITUTE OF FERROUS TECHNOLOGY (GIFT)

Tenure-track Professorship Positions The Graduate Institute of Ferrous Technology (GIFT), located in South Korea, is an English language academic initiative initiated by Pohang University of Science and Technology (POSTECH) to address the need of the basic commodity materials industry for outstanding professionals specialized in metallurgical technologies.

The Institute has the ambition of leading the world in all aspects of the science of advanced steel products and processing. We have openings for candidates with an excellent and dedicated record of research in the field of iron and steel. Tenure-track professors in assistant or associate levels are sought in the following field:

Control and Automation (CA)
Position CA-1 Candidates are expected to have knowledge in control theory as well as its applications.

Position CA-2 Candidates are expected to have deep knowledge in sensors, industrial measurements and automation technology.

Candidates must hold a doctoral degree in the relevant area. Entry positions may be at the Assistant and Associate level, depending on the merits of the individual applicants. Further information about the Graduate Institute of Ferrous Technology is available from ***http://www.postech.ac.kr/gift***.

Interested candidates should respond with a cover letter, full resume with a photograph and the list of publication, and names, addresses and telephone numbers of six references to:
Professor Hae-Geon Lee
Dean, Graduate Institute of Ferrous Technology
Pohang University of Science and Technology
San31, Hyoja-Dong, Nam-Ku
Pohang, 790-784, South Korea
gift-recruit@postech.ac.kr

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Faculty Position in Mechanical Engineering at the Ecole polytechnique fédérale de Lausanne (EPFL)

The School of Engineering of EPFL invites applications for a **tenure-track faculty position** in Mechanical Engineering. The main focus of this search is for a junior position, however, exceptionally well-qualified candidates may be considered at a more senior level. We encourage applications in the broad area of **advanced mechanical engineering design.** Particular areas of interest include, but are not limited to: fundamentals in design of smart mechanical systems/mechatronics; instrumentation and sensors; actuators, monitoring and control; and design at micro- and nano-scales with advanced engineering materials. Evidence of strong research and teaching capabilities are expected.

The successful candidate is expected to initiate independent, creative research programs and participate in undergraduate and graduate teaching. Significant start-up resources and research infrastructure will be available. Internationally competitive salaries and benefits are offered.

Applications should include a curriculum vitae with a list of publications, a concise statement of research and teaching interests, and the names and addresses (including e-mail) of at least five referees. Applications should be uploaded to **http://igmeca-2search.epfl.ch.** The deadline for applications is **16 January 2009.**

Enquiries may be addressed to:
Prof. John Botsis,
e-mail: **position.igm@epfl.ch**

For additional information on EPFL and the School of Engineering, please consult the web sites: **http://www.epfl.ch** and **http://sti.epfl.ch**

EPFL aims to increase the presence of women amongst its faculty, and qualified female candidates are strongly encouraged to apply.

ROGERS DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING UNIVERSITY of TORONTO

The Edward S. Rogers Sr. Department of Electrical and Computer Engineering at the University of Toronto invites applications for a tenure-stream Assistant or Associate Professor position in the area of **electronics**, beginning July 1, 2009. Research areas of interest include, but are not limited to: **analog or mixed-signal circuit design, high-performance ADCs and/or DACs, or RF circuits.** Candidates must have (or are about to receive) a Ph.D. in the relevant area.

The department ranks among the top 10 ECE departments in North America. It attracts outstanding students, has excellent facilities, and is ideally located in the middle of a vibrant, artistic, and diverse cosmopolitan city. The department offers highly competitive salaries and start-up funding, and faculty have access to significant Canadian research operational and infrastructure grants. Additional information can be found at: ***www.ece.utoronto.ca***

The successful candidate is expected to pursue excellence in research and teaching at both the graduate and undergraduate levels. The successful candidate will join a highly active research group in electronics.

Applicants must submit their application by electronic mail to Professor Glenn Gulak, Electronics Search Committee Chair, The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto using the following address:

ElectronicsSearch@ece.utoronto.ca

Please submit only Adobe Acrobat PDF documents. Applicants will receive an email acknowledgement.

All applications should include: a curriculum vitae, a summary of previous research and proposed new directions, and a statement of teaching philosophy and interests.

In addition, applicants must arrange to have three confidential letters of recommendation sent directly (by the referee) by email to:

ElectronicsSearch@ece.utoronto.ca

Applications and referee-sent references should be received by January 15, 2009.

The University of Toronto is strongly committed to diversity within its community and especially welcomes applications from visible minority group members, women, Aboriginal persons, persons with disabilities, members of sexual minority groups, and others who may contribute to the further diversification of ideas.

All qualified candidates are encouraged to apply; however, Canadians and permanent residents will be given priority. Salary will commensurate with qualifications and experience.

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Faculty Position in BioMicroengineering

at the Ecole polytechnique fédérale de Lausanne (EPFL)

The Institutes of Bioengineering and Microengineering at EPFL are seeking a **tenure track assistant professor** in the field of **biomicro/nanosystems.** Exceptionally qualified candidates may also be considered at a more senior level. Bioengineering at EPFL is well integrated between the Schools of Engineering and Life Sciences. In this search, an appointment is sought within the School of Engineering, in which Microengineering is also situated. The open faculty position is offered in an environment of both theoretical and experimental research, rich for the development of novel basic technologies as well as for seeking deeper understanding of integrative (patho) physiological mechanisms and developing novel technological and biotherapeutic approaches at the levels of genes, biomolecules, cells and tissues.

The Institutes seek to grow at the interface of engineering with biology in the domain of nano- and microtechnologies and integrated systems. We particularly encourage candidates with strong expertise in the areas of bio-MEMS/NEMS, sensing and actuation for in vitro use, for example in basic biological investigation, systems biology, and diagnostics, or in vivo use, for example in diagnostic and sensing systems, integrated systems, and drug delivery. The interface with robotic, surgical and imaging systems is of specific interest. EPFL has strong research facilities, in particular in micro/nanofabrication, imaging, and cytometry.

Successful candidates are expected to initiate independent, creative research programs and participate in undergraduate and graduate teaching. We offer internationally competitive salaries, start-up resources and benefits.

Applications should include a curriculum vitae with a list of publications, a concise statement of research and teaching interests, and the names and addresses (including e-mail) of at least five referees. Applications should be uploaded to **http://biomems-rec.epfl.ch.** The deadline for applications is **15 January 2009.**

Enquiries may be addressed to:
Prof. Jeffrey A. Hubbell,
e-mail: **biomems-rec@epfl.ch**

For additional information on EPFL, School of Engineering, Institute of Bioengineering, and Institute of Microengineering, please consult the web sites: **http://www.epfl.ch, http://sti.epfl.ch, http://ibi.epfl.ch, http://imt.epfl.ch**

EPFL aims to increase the presence of women amongst its faculty, and qualified female candidates are strongly encouraged to apply.

ALBERT-LUDWIGS-UNIVERSITÄT FREIBURG

The Faculty of Applied Sciences at the University of Freiburg, with its Departments of Computer Science and Microsystems Engineering, invites applications for a tenure-track **Professor (W1) of Microsystems Engineering**

The successful candidate will establish a comprehensive research and teaching program in the area of systems and control theory in the Department of Microsystems Engineering.

The candidate should have a extensive experience in Microsystems Engineering or related fields. The appointment will be tenure track. The initial appointment is limited to a maximum term of six years – four years plus a possible extension of two years; within this time, the tenure process may be initiated.

The University of Freiburg aims to increase the representation of women in research and teaching, and therefore expressly encourages women with suitable qualifications to apply for the post.

Information about the Department of Microsystems Engineering can be obtained from **www.imtek.uni-freiburg.de** and the Faculty of Applied Sciences at **www. faw.unifreiburg.de**.

Applications, including a curriculum vitae, publication list and statement of research interests, should be sent by January 31, 2009 to the Dean of the Faculty of Applied Sciences, University of Freiburg, Georges-Köhler-Allee 101, 79110 Freiburg, Germany

Applicants should request an application form from the Dean's office by emailing to: **dekanat@faw.uni-freiburg.de**

TU VIENNA
Technische Universität Wien
Vienna University of Technology

Vienna University of Technology, Austria invites applications for the faculty position of a **full professor in Industrial Automation**. Candidates should have an excellent scientific background and strong expertise in at least one of the following research areas: process and production control engineering, autonomous sensor-driven systems, distributed control, in-process sensor and actuator systems, network controlled systems, industrial robotics for automation, autonomous und mobile robots, visual servoing, and process monitoring and diagnosis. Applications are expected until January 31st, 2009.

More information concerning the requirements for this position is available at

http://www.acin.tuwien.ac.at/fileadmin/IndustrialAutomation.pdf

UNIVERSITÄT DES SAARLANDES

Saarland University is seeking to establish several **Junior Research Groups** (W1/W2) within the recently established **Cluster of Excellence** *"Multimodal Computing and Interaction"* which was established by the German Research Foundation (DFG) within the framework of the German Excellence Initiative.

M²CI CLUSTER OF EXCELLENCE

The term "multimodal" describes the different types of digital information such as text, speech, images, video, graphics, and high-dimensional data, and the way it is perceived and communicated, particularly through vision, hearing, and human expression. The challenge is now to organize, understand, and search this multimodal information in a robust, efficient and intelligent way, and to create dependable systems that allow natural and intuitive multimodal interaction. We are looking for highly motivated young researchers with a background in the research areas of the cluster, including algorithmic foundations, secure and autonomous networked systems, open science web, information processing in the life sciences, visual computing, large-scale virtual environments, synthetic virtual characters, text and speech processing and multimodal dialog systems. Additional information on the Cluster of Excellence is available on **http://www.mmci.uni-saarland.de**. Group leaders will receive junior faculty status at Saarland University, including the right to supervise Bachelor, Master and PhD students. Positions are limited to five years.

Applicants for W1 positions (phase I of the program) must have completed an outstanding PhD. Upon successful evaluation after two years, W1 group leaders are eligible for promotion to W2. Direct applicants for W2 positions (phase II of the program) must have completed a postdoc stay and must have demonstrated outstanding research potential and the ability to successfully lead their own research group. Junior research groups are equipped with a budget of 80k to 100k Euros per year to cover research personnel and other costs.

Saarland University has leading departments in computer science and computational linguistics, with more than 200 PhD students working on topics related to the cluster (see **http://www.informatik-saarland.de** for additional information). The German Excellence Initiative recently awarded multi-million grants to the Cluster of Excellence "Multimodal Computing and Interaction" as well as to the "Saarbrücken Graduate School of Computer Science". An important factor to this success were the close ties to the Max Planck Institute for Computer Science, the German Research Center for Artificial Intelligence (DFKI), and the Max Planck Institute for Software Systems, which are co-located on the same campus.

Candidates should submit their application (curriculum vitae, photograph, list of publications, short research plan, copies of degree certificates, copies of the five most important publications, list of five references) to the coordinator of the cluster, Prof. Hans-Peter Seidel, MPI for Computer Science, Campus E1 4, 66123 Saarbrücken, Germany. Please, also send your application as a single PDF file to **applications@mmci.uni-saarland.de**.

The review of applications will begin on January 15, 2009, and applicants are strongly encouraged to submit applications by that date; however, applications will continue to be accepted until January 31, 2009. Final decisions will be made following a candidate symposium that will be held during March 9 – 13, 2009.

Saarland University is an equal opportunity employer. In accordance with its policy of increasing the proportion of women in this type of employment, the University actively encourages applications from women. For candidates with equal qualification, preference will be given to people with physical disabilities.

THE UNIVERSITY OF UTAH

The Department of Electrical and Computer Engineering, University of Utah, Salt Lake City, seeks applications to fill four tenure-track positions at the assistant professor level. Outstanding applicants with significant experience may also be considered at the associate or full professor level. We are particularly interested in candidates with **expertise in electromagnetics, power systems & power electronics, solid-state, and communications**. Information on department research activities and curricula may be found on the web at **www.ece.utah.edu**.

The web site also has information on additional positions available for a department chair and for the Utah Science, Technology and Research Initiative. Faculty responsibilities include developing and maintaining an internationally recognized research program, effective classroom teaching at the undergraduate and graduate levels, and professional service. Résumés with names and contact information for at least three references should be sent to Ms. Debbie Sparks, Faculty Search Committee, University of Utah, Electrical and Computer Engineering Department, 50 South Central Campus Drive, Room 3280, Salt Lake City, UT 84112-9206. Email applications are accepted at **dsparks@ece.utah.edu**.

Applications will be reviewed starting September 1, 2008, and will be accepted until the positions are filled. Applicants must hold a Ph.D. by the time of appointment. The University of Utah values candidates who have experience working in settings with students from diverse backgrounds and possess a strong commitment to improving access to higher education for historically underrepresented students. The University is an AA/EO employer, encourages applications from women and minorities, and provides reasonable accommodations for known disabilities of applicants and employees.

THE CHINESE UNIVERSITY OF HONG KONG

Applications are invited for:-

Department of Information Engineering

Professor(s) / Associate Professor(s) / Assistant Professor(s)

(Ref. 08/216(370)/2) (Closing date: February 16, 2009)

Applicants should have (i) a PhD degree; and (ii) a strong research record in security-related disciplines (e.g. systems security and cryptography). Apart from teaching, the appointees will actively undertake related research projects. Appointments will normally be made on contract basis for up to three years initially commencing August 2009, leading to longer-term appointment or substantiation later subject to mutual agreement. Further information about the Department is available at *http://www.ie.cuhk.edu.hk*.

Salary and Fringe Benefits
Salary will be highly competitive, commensurate with qualifications and experience. The University offers a comprehensive fringe benefit package, including medical care, plus a contract-end gratuity for appointments of two years or longer, and housing benefits for eligible appointees.

Further information about the University and the general terms of service for teaching appointees is available at *http://www.cuhk.edu.hk/personnel*. The terms mentioned herein are for reference only and are subject to revision by the University.

Application Procedure
Please send full resume, copies of academic credentials, a publication list and/or abstracts of selected published papers, together with names, addresses and fax numbers/e-mail addresses of three referees to whom applicants' consent has been given for their providing references (unless otherwise specified), to *recruit@ie.cuhk.edu.hk* by the closing date.

The Personal Information Collection Statement will be provided upon request. Please quote the reference number and mark 'Application - Confidential' on cover.

Texas A&M University at Qatar

(TAMUQ) is a branch campus of Texas A&M University (TAMU) at College Station, Texas. TAMUQ began teaching undergraduate students in chemical, electrical, mechanical, and petroleum engineering in the Fall of 2003 and started conferring degrees in December 2007. The coursework undertaken by the students is materially identical to the programs offered at the TAMU main campus. The TAMUQ campus is situated within a brand new building and is part of Education City, Doha, Qatar, a consortium of educational and research institutions hosted by the Qatar Foundation (QF) for Education, Science and Community Development.

The electrical and computer engineering (ECE) department at TAMUQ currently offers BS degrees in Electrical Engineering but a graduate program in electrical engineering is currently awaiting approval. All formal instruction is given in English. More information about the ECE program at TAMUQ can be found at

http://ecen.qatar.tamu.edu/.

The ECE department at TAMUQ invites applications for faculty positions at all ranks with research specializations in the following and related areas:

Power Systems: Electric power generation, transport and distribution, electrical energy conversion, deregulation, forecasting, electrical installations, photovoltaic applications for buildings, electrical renewable energy technologies.

Electronics & Electromagnetics: Analog and mixed signal circuits and systems, design, implementation and application of CMOS wireless transceivers; CMOS sensors and circuitry, solid state devices (electronic/opto-electronic). Propagation, Antennas and RF Systems.

Applicants with interdisciplinary research and teaching interests are particularly encouraged. Applicants must have a Ph.D. or equivalent degree, or completion of all requirements by date of hire. For senior positions, applicants should have a proven record of scholarly contributions and a proven ability to attract research funding. For junior positions, candidates should have demonstrated potential for quality teaching and research.

Starting rank and salary will depend on qualifications and experience. The appointment also includes the following benefits: Fully furnished housing; coverage of local tuition fees for school-age dependent children; annual home leave allowance for family members; air tickets to Doha on appointment; and local transportation allowance. Fringe benefits include health and medical insurance as well as an enrollment in a retirement plan. Initial appointment will normally be on a two-year contract, but a local tenure process is currently under consideration for TAMUQ. Re-appointment will be subject to mutual agreement.

Applications, including full curriculum vitae with list of publications, statement of teaching, statement of research as well as the names, addresses (regular mail and E-mail), fax, and phone numbers of three references to should be sent to:

Dr. Costas N. Georghiades, Department Head
c/o Ms. Debbie Hanson
Department of Electrical and Computer Engineering
Texas A&M University
College Station, TX, 77843-3128.

Texas A&M University at Qatar is an equal opportunity/affirmative action employer and actively seeks the candidacy of women and minorities. The deadline for applications is January 15, 2009 but applicants will be considered until the positions are filled.

the data

The Rise of the Machines

THERE ARE now 1 million industrial robots toiling around the world, and Japan is where they're the thickest on the ground. It has 295 of these electromechanical marvels for every 10 000 manufacturing workers—a robot density almost 10 times the world average and nearly twice that of Singapore (169), South Korea (164), and Germany (163).

Although the top three countries are in Asia, Europe gets the regional title as the epicenter of global automation; it has a robot density of 50, compared to 31 in the Americas and 27 in the Asia/Pacific region.

IEEE Spectrum computed the robot density for 67 nations in all, using data from the International Federation of Robotics and the International Labour Organization [see http://www.spectrum.ieee.org/dec08/robodata for the complete list].

By 2011, the world's industrial robot population is expected to rise to 1.2 million. Which countries these new robots will call home—and how the machine-to-human balance will change—remains to be seen.

—ERICO GUIZZO

TOP 10 COUNTRIES BY ROBOT DENSITY

(Industrial robots per 10 000 manufacturing workers)

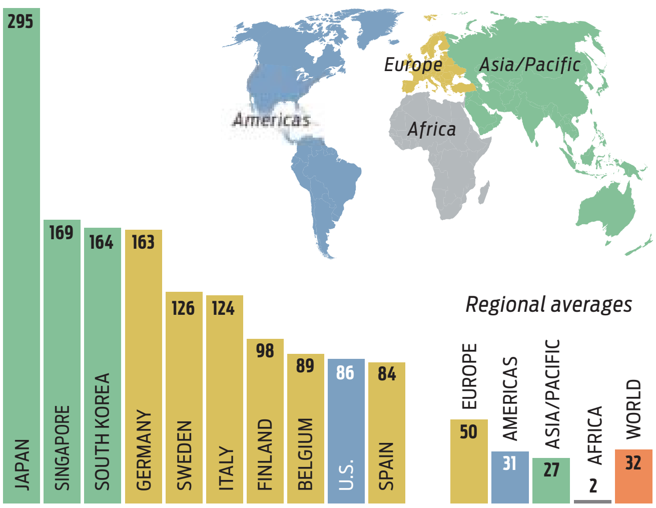

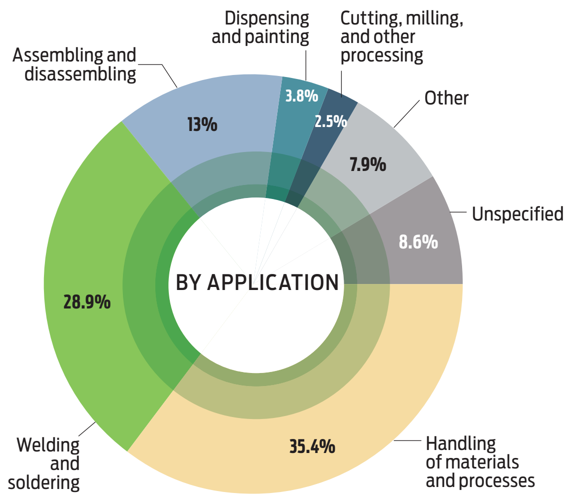

INDUSTRIAL ROBOTS

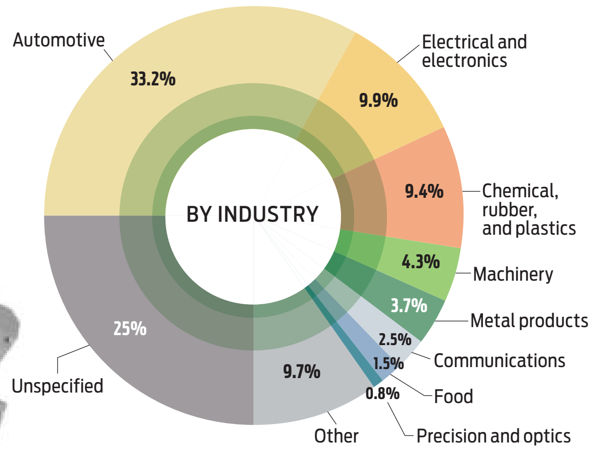

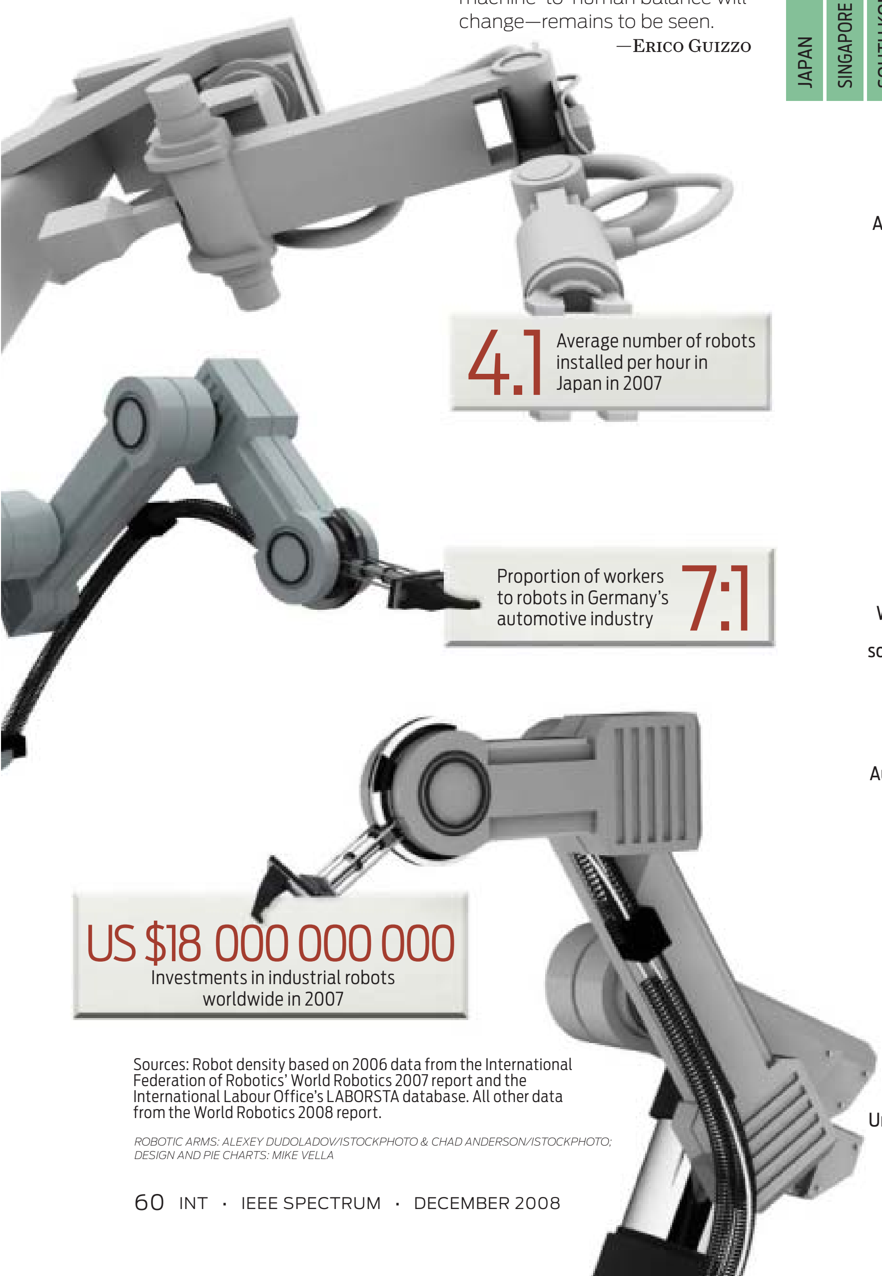

4.1 Average number of robots installed per hour in Japan in 2007

Proportion of workers to robots in Germany's automotive industry **7:1**

US $18 000 000 000 Investments in industrial robots worldwide in 2007

Sources: Robot density based on 2006 data from the International Federation of Robotics' World Robotics 2007 report and the International Labour Office's LABORSTA database. All other data from the World Robotics 2008 report.

ROBOTIC ARMS: ALEXEY DUDOLADOV/ISTOCKPHOTO & CHAD ANDERSON/ISTOCKPHOTO; DESIGN AND PIE CHARTS: MIKE VELLA


From Imagination to Market