THE MAGAZINE OF TECHNOLOGY INSIDERS

ieee spectrum

04.09

DESERT RUNNERS

LIZARDS, CRABS, AND COCKROACHES ARE SHOWING RESEARCHERS HOW TO BUILD THE ROBOTS THAT WILL SOMEDAY SCURRY ON THE SANDS OF MARS

THE 10 BEST HIGH-TECH CARS

THE UNIVERSAL CELLPHONE HANDSET, SOFTWARE DEFINED

IEEE

Celebrating 125 Years
of Engineering the Future

volume 46 number 4 international

ieee spectrum 04.09

LEAN AND MEAN: The Ford Fusion features a four-cylinder gasoline engine and also a 93-kilowatt electric motor—powerful enough to carry the car at speed all by itself.

COVER: BRYAN CHRISTIE DESIGN
THIS PAGE: FORD MOTOR CO.

volume 46 number 4 international

ieee spectrum 04.09

PHOTOS: LEFT, FESTO; RIGHT, TOM MERTON/ GETTY IMAGES

WWW.SPECTRUM.IEEE.ORG

AVAILABLE 1 APRIL ON SPECTRUM ONLINE

ROBOTS GONE WILD

Mechatronic jellyfish, sensor-laden seals, and lobster look-alikes are just a few of the robots whose main design techniques derive from the animal kingdom. These artificial jellyfish [above] are blooming in the labs of engineers at Festo, a German automation company. Inside each AquaJelly's umbrella-shaped top are 11 LEDs that pulse infrared signals. The lights allow the robots to communicate with their mechanical brethren, enabling them to operate as a swarm and adjust their positions relative to each other. IEEE Spectrum Online shows you how today's robots are gaining new abilities to swim, skip, and nuzzle in a convincingly lifelike way.

ONLINE FEATURES:

COULD THE BIGGEST photovoltaic power plant in the eastern United States be justified without renewables mandates and generous subsidies? *IEEE Spectrum*'s Bill Sweet investigates.

KEEP PACE with the latest major software crashes and system failures on Robert N. Charette's Risk Factor blog.

FOLLOW SPECTRUM ON TWITTER for the latest tech news and views.

ALSO ONLINE:

- Webcasts
- Podcasts
- News
- Videos
- Blogs
- Jobs
- Career Accelerator
- IEEE Xplore® digital library
- Interviews
- Opinions
- More!

WWW.IEEE.ORG/ THEINSTITUTE

AVAILABLE 6 APRIL ON THE INSTITUTE ONLINE

WHY JOB HUNTERS SHOULD VOLUNTEER

If you're unemployed, the best thing you can do is volunteer for an organization like IEEE—not just because you'll pick up new skills but also because it's the best way to network. You'll make contacts and meet leaders in your field, any one of whom may lead you to a job offer.

GIVING BACK

IEEE Secretary Barry L. Shoop helped establish an engineering school in Afghanistan, which recently graduated its first class.

IEEE AND ETA KAPPA NU SIGN MERGER AGREEMENT

IEEE has signed a merger agreement with the engineering honor society Eta Kappa Nu, which has nearly 200 university chapters and more than 100 000 members who have made significant accomplishments in electrical and computer engineering.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 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.

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.

back story

YVONNE BOYD

Running With Robots

HOW HARD can it be to make a lizard run? Just ask the researchers in Daniel Goldman's laboratory at Georgia Tech. "It's very high tech," says Goldman [left]. "We startle it."

They wave bits of cardboard, flash lights, and gently pinch the lizards' tails. Then they do it again. And again. And again. "Sometimes I try to make a noise, but that doesn't seem to work," says Chen Li, a physics graduate student [right].

The 65 lizards, crabs, and cockroaches in Goldman's lab can be ornery companions to the one robot that beats a steady path down a track in the lab. In "March of the SandBots," Goldman and his colleagues explain how they invested in that robot the tricks they'd learned about sand scampering from the finest desert runners on the planet.

But building a bio-inspired robot can involve a few odd negotiations. To acquire some of the creatures in Goldman's menagerie, one student coaxed collectors in the Middle East into selling their prized Egyptian desert cockroaches (*Polyphaga aegyptiaca*) for about US $10 a bug.

Goldman, too, contributes a catch or two. He grew up chasing ghost crabs on the sandbar islands of North Carolina and five-lined skinks around his home in Richmond, Va. At the age of 10 he learned how to snag one of those lizards: He'd dangle a coil of rope on a pole and stand several steps from his target. Then he'd lower the noose around its neck and pull up, ensnaring the thrashing skink and delivering it unharmed to a cloth bag.

But unlike lizards in the wild, captive animals can be exhaustingly placid. "I can say, 'Animal, go!' or 'Animal, move your legs in a certain way,' and it won't do that," Goldman reports. Indeed, an annoyed ghost crab may instead go on the offensive, snapping its claws energetically at its appointed frightener.

Li, who grew up watching animal documentaries in China, hadn't anticipated how unpredictable the animals would be. So when it's the robot's turn to strut, the physicists breathe a sigh of relief. "The robot gives us data that's more repeatable," Li says. "It's very nice." ❑

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. 4 (INT), April 2009, p. 7, or in *IEEE Spectrum*, Vol. 45, no. 4 (NA), April 2009, p. 11.

contributors

BILL ARNOLD has been chief scientist at ASML since 1998, which means he's helping to shape the next generation of lithographic processes. His article, "Shrinking Possibilities" [p. 22], is about the future of chipmaking, the inevitability of extreme ultraviolet lithography, and how the continued realization of Moore's Law will affect the semiconductor industry.

TAVIS COBURN, based in Toronto, draws inspiration from printed materials of the 1920s to the 1960s, including the Russian avant-garde and classic comic books. He started out creating illustrations by hand with paint and silk screens. Now he achieves similar effects with digital tools. For the opening image of "The Universal Handset" [p. 32], Coburn says, "I went with old radio towers, which look a lot cooler than modern cell antennas."

DANIEL P. DERN wrote for us in December 2008 about his search for the perfect ultralight laptop. In January, he spent a good portion of this year's Consumer Electronics Show looking at new battery technologies that would add hours—but not ounces—to small gadgets. As he reports in Update [p. 12], a number of sophisticated battery-recharging systems are almost ready for widespread adoption.

DANIEL GOLDMAN, HALDUN KOMSUOGLU, and **DANIEL KODITSCHEK** describe how lizards, crabs, and cockroaches are teaching robots new tricks for conquering sandy terrain in "March of the SandBots" [p. 26]. Goldman, an assistant professor of physics at Georgia Tech, studies how the animals move their limbs, while his colleagues at the University of Pennsylvania—Koditschek, an electrical engineering professor, and Komsuoglu, a postdoctoral fellow in Koditschek's lab—refine their robots' abilities to perceive and respond to their environments.

PETER KOCH and **RAMJEE PRASAD,** who explain how software-defined radio will soon transform cellphones in "The Universal Handset" [p. 32], are professors at Aalborg University, in Denmark. Koch works at the university's Center for Software Defined Radio and also operates an amateur radio station for fun. He's shooting to reach other hams in all parts of the world. "I'm not there yet," he says. Prasad, an IEEE Fellow, heads the university's Center for TeleInfrastructure. He, too, enjoys making international contacts, but rather than doing so wirelessly, he regularly travels to the far corners of the world.

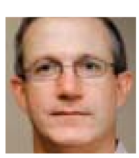

JOHN VOELCKER, *IEEE Spectrum*'s automotive editor, covers the part of the business that's slighted by most industry mavens: technology. In this year's iteration of the annual Top 10 Tech Cars report [p. 38], his sixth, he argues that technology will matter even more now that the auto industry is passing through its biggest contraction in a lifetime.

IEEE
Celebrating 125 Years
of Engineering the Future

IEEE MEDIA

STAFF DIRECTOR; PUBLISHER, *IEEE SPECTRUM* James A. Vick, j.vick@ieee.org
ASSOCIATE PUBLISHER, SALES & ADVERTISING DIRECTOR Marion Delaney, m.delaney@ieee.org
RECRUITMENT SALES DEVELOPMENT MANAGER Michael Buryk, m.buryk@ieee.org
BUSINESS MANAGER Robert T. Ross
MARKETING & PROMOTION MANAGER Blanche McGurr, b.mcgurr@ieee.org
INTERACTIVE MARKETING MANAGER Ruchika Anand, r.t.anand@ieee.org
LIST/RECRUITMENT MARKETING MANAGER Ilia Rodriguez, i.rodriguez@ieee.org
REPRINT SALES +1 212 221 9595, EXT. 319
DEPARTMENT ADMINISTRATOR Faith H. Jeanty, f.jeanty@ieee.org
ADVERTISING SALES +1 212 419 7760
TELEPHONE ADVERTISING/SALES REPRESENTATIVE John Restchack +1 212 419 7578
ADVERTISING PRODUCTION MANAGER Felicia Spagnoli
SENIOR ADVERTISING PRODUCTION COORDINATOR Nicole Evans
ADVERTISING PRODUCTION +1 732 562 6334
IEEE STAFF EXECUTIVE, PUBLICATIONS Anthony Durniak

IEEE BOARD OF DIRECTORS

PRESIDENT & CEO John R. Vig
+1 732 562 3928 FAX: +1 732 465 6444 president@ieee.org
PRESIDENT-ELECT Pedro A. Ray
TREASURER Peter W. Staecker
SECRETARY Barry L. Shoop
PAST PRESIDENT Lewis M. Terman

VICE PRESIDENTS

Teofilo Ramos, *Educational Activities;* Jon G. Rokne, *Publication Services & Products;* Joseph V. Lillie, *Member & Geographic Activities;* W. Charlton Adams, *President, Standards Association;* Harold L. Flescher, *Technical Activities;* Gordon W. Day, *President, IEEE-USA*

DIVISION DIRECTORS

Giovanni De Micheli (I); Robert E. Hebner Jr. (II); Curtis A. Siller Jr. (III); Roger W. Sudbury (IV); Deborah M. Cooper (V); Mark I. Montrose (VI); John D. McDonald (VII); Stephen L. Diamond (VIII); Frederick C. Mintzer (IX); Richard A. Volz (X)

REGION DIRECTORS

Howard E. Michel (1); William P. Walsh Jr. (2); William B. Ratcliff (3); Don C. Bramlett (4); David J. Pierce (5); Leonard J. Bond (6); Ferial El-Hawary (7); Jozef W. Modelski (8); Enrique E. Alvarez (9); Yong Jin Park (10)

DIRECTORS EMERITUS

Eric Herz, Theodore W. Hissey

IEEE STAFF

HUMAN RESOURCES Betsy Davis, SPHR
+1 732 465 6434, e.davis@ieee.org
PUBLICATIONS Anthony Durniak
+1 732 562 3998, a.durniak@ieee.org
EDUCATIONAL ACTIVITIES Douglas Gorham
+1 732 562 5483, d.g.gorham@ieee.org
STANDARDS ACTIVITIES Judith Gorman
+1 732 562 3820, j.gorman@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
CORPORATE STRATEGY & COMMUNICATIONS Matthew Loeb, CAE
+1 732 562 5320, m.loeb@ieee.org
BUSINESS ADMINISTRATION Richard D. Schwartz
+1 732 562 5311, r.schwartz@ieee.org
TECHNICAL ACTIVITIES Mary Ward-Callan
+1 732 562 3850, m.ward-callan@ieee.org
MANAGING DIRECTOR, IEEE-USA Chris Brantley
+1 202 530 8349, c.brantley@ieee.org

IEEE PUBLICATION SERVICES & PRODUCTS BOARD

Jon G. Rokne, *Chair;* Tayfun Akgul, Tariq S. Durrani, Mohamed E. El-Hawary, Gerald L. Engel, Gerard H. Gaynor, Marion O. Hagler, Jens Hannemann, Lajos Hanzo, Hirohisa Kawamoto, Michael R. Lightner, Lloyd A. Morley, William W. Moses, Adrian V. Pais, Saifur Rahman, Sorel Reisman, Edward A. Rezek, Barry L. Shoop, W. Ross Stone, Robert J. Trew, Leung Tsang, Karl R. Varian, Stephen Yurkovich

IEEE OPERATIONS CENTER

445 Hoes Lane, Box 1331, Piscataway, NJ 08854-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

CLOCKWISE FROM TOP LEFT: ASML; RANDI SILBERMAN; ERNEST LILLEY; TAVIS COBURN spectral lines

Cap and Trade or Cap and Tax?

WHENEVER OIL and gas prices plummet, renewable energy projects tend to hit the skids, because the economic incentive to develop these new technologies simply disappears. Low prices bring on the Hummers; high prices have consumers clamoring for Toyota Priuses.

The current global economic crisis has consumers wanting neither. So U.S. President Barack Obama is looking for ways to kick the economy, and renewable energy development, into gear.

In his address to the U.S. Congress in February, Obama called on lawmakers "to send me legislation that places a market-based cap on carbon pollution and drives the production of more renewable energy in America." The question is, Will a cap-and-trade system deliver the desired result—a robust renewable energy program—better than a straight carbon tax?

A discussion held last October at Columbia University, in New York City, between Yvo de Boer, the executive secretary of the United Nations Framework Convention on Climate Change, and Jeffrey Sachs, director of Columbia's Earth Institute, sheds some light on the subject.

De Boer argued the merits of the Kyoto Protocol, which sets limits on greenhouse-gas emissions for signatory industrial nations and led to the European Union's carbon-emission trading system. A central element of the Kyoto Framework is the Clean Development Mechanism. The CDM allows industrialized countries to finance clean technology projects in developing countries. In return for their investment, the industrialized countries receive emission-reduction credits they can apply toward their own carbon-reduction goals or sell to others.

De Boer reported that there are now 1170 registered CDM projects in 49 developing countries. Critics argued that these projects would have gone ahead even without outside investment and that the CDM is no more than a cheap way for industrialized countries to buy their way out of carbon compliance.

De Boer countered that carbon-cap-and-trade programs and the CDM, though imperfect, are worthy works in progress. And, he said, he did not oppose a complementary set of carbon taxes. "We need a set of tools to spur both private and public money flows," he said. "Climate change is a global problem, and we will need all the tools at our disposal."

Sachs, for his part, argued that market-based mechanisms like cap and trade have not demonstrated that they can turn the trajectory of carbon emissions as "sharply and dramatically as we need to do it." The CDM, said Sachs, "is unfortunately a very small, marginal tool that isn't going to really change the broad framework of how energy is produced and technology distributed." What's more, he argued, cap and trade encourages the continued deployment of questionable, risky financial instruments. "I am not so keen on sending our best and brightest off to do more financial market engineering" for the carbon markets, he said. "I think the meltdown shows how we took a generation of brilliant young people and put them to tasks that don't solve problems."

Sachs believes that a carbon tax is much easier to administer than a cap-and-trade scheme. "There are just a few places we get carbon from, and by taxing upstream—at the refinery or the wellhead—you reach a carbon price for the whole economy," he said. Cap-and-trade systems, on the other hand, require monitoring the compliance of hundreds of thousands of enterprises, which, Sachs suggested, keeps a lot of regulators, attorneys, and auditors busy but has a questionable impact on carbon emissions.

Unlike a cap-and-trade system that thrives on price volatility, a carbon tax would put a floor on the price of carbon, holding it to a fixed price. A more certain price for carbon would in turn encourage long-term investment in clean technologies, said Sachs. Rather than allowing Wall Street's financial engineers to make money from cap and trade, a carbon tax would let real engineers figure out ways to get carbon dioxide under control. But, Sachs said, he isn't holding his breath waiting for Congress to act on his suggestion.

—SUSAN ARTERIAN CHANG

Editor's note: At publication time, U.S. Representative John Larson (D-Conn.) introduced the bill known as America's Energy Security Trust Fund Act of 2009, which proposes a carbon tax to reduce carbon emissions.

A version of this column appeared in IEEE Spectrum Online's Energywise blog on 25 February. Susan Arterian Chang reports on the impacts of climate change from a policymaking and financial markets perspective.

JACQUELYN MARTIN/AP PHOTO

forum

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.

FIRST, TAKE A DEEP BREATH...

As the director of primary care diving medicine with Jefferson Internal Medicine Associates, I note that the article on the Poseidon Discovery rebreather [Tools & Toys, January] discounts the fact that manual or electronically driven rebreathers still have a fatality rate greater than that of open-circuit scuba, even among highly trained technical-level divers. Maybe the Discovery Mark VI's software can react faster than humans can, but if we still don't have software fully in control of airplanes or driving our cars in slow-moving traffic, how will a fully automated rebreather perform? Is the reliability-over-liability ratio so great that Poseidon would be willing to field enough of these devices to net a profit? Given the supposed minimal skill levels needed to use it, this rebreather is in full control and is likely fully liable.

Marvin E. Gozum
IEEE Member
Philadelphia

Executive Editor (and scuba diver) Glenn Zorpette responds:
I regret that space constraints did not allow us to more adequately address the safety issue. Rebreathers are high tech and amazing, but anecdotal evidence, at least, suggests that rebreather dives are still not as safe as open-circuit, recreational (no-decompression) dives.

CODELAND SECURITY

Thanks for the informative article "Cryptographers Take On Quantum Computers" [Update, January]. However, the state of the art in lattice-based cryptography is further advanced than the article implies. The NTRUEncrypt lattice-based system, first proposed in 1996, has been extensively researched and deployed. In fact, it was standardized in December 2008, by none other than IEEE (as IEEE Std 1363.1). NTRUEncrypt has been the subject of dozens of papers, has been licensed to companies such as Infospace, and is under consideration for standardization by the Accredited Standards Committee X9 for use in financial services. Therefore, it is the clear leading candidate for adoption if quantum computers force RSA and elliptic-curve cryptography to be withdrawn.

William Whyte
Acton, Mass.

CHIPPED OFF

"Multicore Made Simple" [Winners & Losers VI, January] contained some misconceptions about Intel's future Larrabee chip and Nvidia's parallel-computing architecture, Compute Unified Device Architecture (CUDA), which is now shipping in all Nvidia graphics-processing units (GPUs). Intel's *x*86 architecture is not the solution to the parallel-computing problem—namely, how to split your problem into many parts to speed up computation. Few applications are able to make use of a quad-core central-processing unit (CPU) today, so how will companies solve this problem when they move to multiple cores? Also, the fact that the chip is *x*86 based does not mean it will become immediately prevalent, nor will it necessarily be simple to program. The *x*86 architecture is somewhat like the engine under the hood of a car: Everyone knows how to drive, but almost no one knows how the engine actually works or how to tweak it for maximum performance. More than 110 million CUDA-enabled GPUs have been sold and deployed in data centers and on desktops, giving users 20 to 200 times as much computing speed as before. Intel graphics chips have lagged for years, and they have a lot to prove before they deserve the credibility they are being given.

Andrew Humber
Senior PR Manager
Computing Products
Nvidia
Santa Clara, Calif.

CORRECTION

The letter "Memristor Mechanics" [Forum, February] should have been attributed to H. Robert Schroeder, IEEE Member, Ewing, N.J.

update

more online at www.spectrum.ieee.org

SMART SPAN: The Rion-Antirion Bridge in Greece, studded with 300 sensors, is among the smartest structures in the world.
PHOTO: GRANT FAINT/GETTY IMAGES

Despite Stimulus Money, Most U.S. Bridges Might Stay Dumb

Sensors are starting to prove themselves in the biggest, most complex bridges, but the technology isn't ready for the hundreds of thousands of smaller ones

THE 2.9-KILOMETER Rion-Antirion Bridge in Greece, with its 300 sensors, is a testament to how smart a piece of infrastructure can be. It routinely tells operators when an earthquake, frequent in those parts, or high winds warrant shutting down traffic.

"The bridge tolls are meant to collect thousands of euros per day," says Alexandre Chaperon, an engineer at the company that designed the system, Advitam, in Vienna, Va. "Without the monitoring system, the bridge would be closed after every earthquake, more than three days in some cases, instead of 5 minutes."

Dozens of the largest and most complex bridges in the world are already studded with strain and displacement gauges, three-dimensional accelerometers, tiltmeters, temperature sensors, and other instruments. They are wired to central data-acquisition units—though some newer bridges have wireless systems—which collect and analyze the information and relay it to engineers, in hopes of catching signs of distress before human inspectors could. With the United States injecting US $27.5 billion into revamping the country's roadways and bridges as part of an $800 billion economic stimulus effort, it might seem like a perfect opportunity to add smarts to more bridges.

But monitoring system costs are too high and the benefits unproven for most of the nearly 600 000 bridges and overpasses

update

HIGH IQ: Accelerometers, tiltmeters, and strain gauges send their signals to remote engineers. *PHOTOS: ADVITAM*

in the United States, experts say. Sensor systems seem to make sense only for big, complex bridges and for those that are already known to be in trouble. And it will be many years before these systems can be used on other bridges.

An installation of sensors on an existing bridge that would help save the owners money and mitigate disasters has not been demonstrated yet, says Emin Aktan, director of the Intelligent Infrastructure and Transport Safety Institute at Drexel University, in Philadelphia. "There is no case where we can say sensors have saved a lot of lives," he says. For most bridges and overpasses "it will take 5 to 10 more years before we have sufficient fundamental knowledge to start deploying such systems and expect returns."

John DeWolf, a civil and environmental engineering professor at the University of Connecticut, in Storrs, who has equipped six bridges with up to 50 sensors each over the past decade, has some of the longest-term data in the country. At today's costs and level of knowledge, he says, a monitoring system might be of most use on "an older bridge that is nearing the end of its life cycle or a bridge with known problems." However, there are potentially an awful lot of those. In 2008, the U.S. Federal Highway Administration listed about 152 000 bridges as structurally deficient or obsolete.

The systems also make sense for bridges with new designs that need to be tested, such as the Rion-Antirion, or for structures like the long suspension and cable-stay bridges that span waterways and cost hundreds of millions of dollars to build. "With these very complex bridges there is a benefit, and there are enough resources available to apply that level of effort and technology," says Andrew Foden, a supervising engineer with infrastructure consulting firm Parsons Brinckerhoff.

The monitoring system for the €630 million Rion-Antirion should be able to tell bridge operators when and where maintenance is needed, potentially saving 15 to 25 percent of long-term maintenance costs, says Advitam's Chaperon. The system has already identified two abnormal vibrations in the nearly five-year-old bridge. One of them, if not treated, could have led to damage to the stay cables. So engineers have installed additional dampers on cables in certain locations to prevent the problem.

Foden says structure monitoring is costly mainly because there is no "one size fits all" system. Every bridge is unique and is exposed to different traffic, weather, and terrain conditions. The instruments must be tailored to each structure and placed at crucial spots. And the more sensors, the better.

When the I-35 bridge collapsed in Minneapolis in 2007, it was due to a gusset plate failure, Foden says, and a strain gauge at that particular gusset might have given a warning. "But there can be thousands of gusset plates on a bridge, and to put sensors everywhere becomes impractical," he adds. Analyzing data from thousands of sensors would only add to the cost.

The National Institute of Standards and Technology recently released new funding for research on low-cost wireless sensors and smart materials that can be built into a structure to measure changes, as well as other technologies that could bring down the cost of smart bridges. But results from that research are likely years away. —Prachi Patel

11.9% The amount that PC shipments are expected to decrease this year, according to market research firm Gartner, in Stamford, Conn. Remember how bad things seemed during the tech crash of 2001? PC shipments fell only 3.2 percent then.

Automated Auditors to Chase Down Cheats

Data mining and math tricks might catch a Madoff or an Enron earlier

In December, when Bernard Madoff's hedge fund turned out to be a gigantic Ponzi scheme that was US $50 billion in the red, everybody wondered how this had gone undetected for more than a decade. Amazingly, in separate audits, government auditors saw no evidence of wrongdoing.

Embarrassed into action, the U.S. Securities and Exchange Commission (SEC) has announced plans to implement an off-the-shelf electronic monitoring system that tracks transactions conducted by 8000 hedge funds. The system will allow SEC accountants to examine hedge funds using scenario analysis and dozens of different graphical representations of the data it collects. But the SEC is intended to be the backstop in the financial accounting system, with most crooked firms getting caught by internal or outside auditors. So technological sleuthing capability belongs in the average auditor's toolbox and not just the government's.

Firms such as PricewaterhouseCoopers, a Big Four auditor, have been working on data-mining systems—software that looks for fraud by finding hidden patterns in financial data. PwC's automated Ledger Analyzer (formerly known as Sherlock) detects anomalies such as a higher-than-normal number of credits to expense accounts or late-posted transactions. These are among the more than 100 indicators of a significant risk of error or fraud taken into account by the system, says David M. Steier, director of research at the company's Center for Advanced Research in San Jose, Calif. Researchers at the auditing firm are refining the software so that as more data points are added, it becomes better at spotting these anomalies, while reducing the number of false positives.

But there are big limits to data mining's fraud-finding potential, says Chao-Hsien Chu, a professor of information sciences and technology at Pennsylvania State University's Smeal College of Business, in University Park. He's conducted a review of the literature on the subject and found that although data mining alone has been used to detect illicit activity such as credit card fraud, insurance fraud, and unauthorized computer system access, it is not a reliable method for checking all financial statements.

One problem, says Chu, is that companies are less likely than ever to share the amount of general ledger data required to make data mining's prediction, classification, and clustering functions effective, citing privacy as a concern.

Mark J. Nigrini, a professor of accounting and information systems at the College of New Jersey's School of Business, in Ewing, is working on the problem using a method that might require less invasive access to data. Nigrini determines whether a given data set has been ginned up artificially by using Benford's Law. This statistical phenomenon describes the relative frequency with which the digits 1 through 9 should appear in a particular decimal place. For example, the digit 1 should appear roughly 30 percent of the time in the decimal place farthest to the left. Lists of numbers manufactured by people, regardless of the purpose, tend to violate this statistical principle, giving an auditor a clue that something is amiss.

Nigrini uses off-the-shelf software to create statistical plots showing where and how widely a company's numbers deviate. Asked if Benford analysis could be used to reverse engineer a passable list of fraudulent numbers, Nigrini says, "Most fraudsters aren't that sophisticated." However, Nigrini acknowledges that his method has limitations, making it clear that a suite of tools is necessary.

Though data mining and pattern recognition make it possible to search through a general ledger with millions of transactions, some of which have 50 000 lines of detail, "technology alone can't solve the fraud detection problem," says Steier. "No matter how good the system is, it's no substitute for an auditor's professional judgment."

Asked whether Benford analysis would have foiled the schemes of Enron and WorldCom executives, Nigrini says that the Enron gang would have been caught. But WorldCom's plot, which used real-world numbers that had been intentionally mislabeled, would have passed the Benford sniff test. As this issue went to press, Nigrini was trying it out on Madoff. —Willie D. Jones

CON MAN: Could software have stopped disgraced financier Bernie Madoff? PHOTO: LUCAS JACKSON/REUTERS

1 GIGATON How much less carbon dioxide will be emitted per year as a result of the global recession, according to New Energy Finance, in London. Unfortunately, the lost investment in renewables, the research and consulting group says, could postpone the peak of world CO_2 emissions by more than a decade.

update

The Silence of the Cellphones

Engineers are working on ways to disable contraband cellphones

You're enjoying a rare bit of peace and quiet, perhaps dozing on a train, when someone begins a loud cellphone conversation nearby. Electromagnetic countermeasures are available to reclaim the silence, but for the most part they are illegal. In the United States, for example, only federal agents can legally use signal jammers to block cellular communications by transmitting interfering signals. Even state prison officials, who want jammers to prevent inmates from talking on contraband cellphones, are barred from using them.

This aspect of U.S. law has created quite a stir in recent years. Things boiled over in January, when lawyers for the CTIA, a wireless-industry association, petitioned the courts to forbid even a 30-minute test of a cellphone-jamming system in a Washington, D.C., prison. On the flip side, legislation proposed in Congress in January would legalize cellphone jamming in prisons.

One person who believes he has a better solution for prisons, concert halls, or wherever cellphone communications are prohibited is J. David Derosier, a founder of Cell Block Technologies, in Fairfax, Va. "I went out to a nice dinner with my wife—it was supposed to be romantic—and the guy at the table next to us started yelling at his kid on his cellphone," Derosier says of the incident that inspired him to create the start-up. After briefly considering the jamming of all cellphone frequencies, he decided to find a more discerning way to prevent most cellphone conversations, while allowing emergency calls to get through.

LOCK DOWN: Technology that would squelch cellphone signals rankles regulators. *PHOTO: MARTIN MEISSNER/AP PHOTO*

Derosier's prototype convinces nearby handsets to think it's the cellular base station with the best signal. But instead of handling most calls, the unit fools phones into operating on a channel with no service. Derosier says his system could be used legally in many parts of the world, including the European Union, although not in the United States.

Aaron Dow, an R&D engineer for Alcatel-Lucent in Wellington, New Zealand, and his colleagues have tried a similar approach to block cellphones at one prison, using a full-blown cellular base station installed nearby—"a honey pot" as he calls it. The idea is to make sure the phones all connect with this station, which then doesn't service the calls.

But this approach is expensive—up to several hundred thousand U.S. dollars, according to Dow. Costs should plummet, though, with a technology that seems destined to become ubiquitous: femtocells. These are small cellular base stations that residential consumers and businesses will likely begin buying en masse in the next year or so to improve cellular coverage indoors. Any cellphone in the vicinity of a femtocell will connect to it wirelessly, and in turn the femtocell will hook up with the mobile carriers' networks through a broadband Internet connection. "If we could adapt the femto solution to do the blocking, it would be much more attractive," says Dow.

Indeed, someone could employ a femtocell to limit the use of cellphone handsets, not just in prisons but in any building. The gadget itself is just a wireless access point, notes Manish Singh, vice president for product-line management at Continuous Computing, a San Diego–based company that provides, among other things, software for femtocells. A suitably smart server attached to the femtocell could screen calls in all sorts of ways. "You can block the traffic right then and there," says Singh.

A business setting up such a system would have the technical ability to permit or deny calls to or from cellphones located within its walls—but would that practice be allowed? In contrast to Wi-Fi routers, femtocells operate in a licensed portion of the radio spectrum, limiting what sorts of functions vendors will be permitted to offer, which is why Singh regards such filtering not as a technological challenge but as "a regulatory issue." And it's a thorny one at that. —David Schneider

Engineers Map Volcanic Lightning

Lightning sensors could lead to better eruption warnings

FLASH, CRACKLE, POP! Lightning might warn of imminent eruptions.
PHOTO: CARLOS GUTIERREZ/UPI /LANDOV

VOLCANIC ERUPTIONS are often accompanied by spectacular bursts of lightning—Krakatoa, Mount St. Helens, and Vesuvius have provided some relatively recent examples—and yet these breathtaking bolts are not well understood. Obtaining insight into volcanic lightning, besides being of considerable scientific interest, could make it possible to get earlier warnings of eruptions and might even yield clues to the origins of life. With those ends in mind, electrical engineers at the New Mexico Institute of Mining and Technology, in Socorro, have installed compact sensing stations of their own design at Mount Redoubt in Alaska.

Ronald Thomas, professor of electrical engineering at the institute, and his colleagues plan to map the lightning from that volcano's eruption in three dimensions, hoping to illuminate what causes electrification during some eruptions and how volcanic lightning compares with thunderstorm lightning, which itself is not fully understood.

The sensors, boxed in modified picnic coolers, record the time and magnitude of the radio-frequency impulses that lightning creates. Correlating the time that the waves hit each receiver, the researchers triangulate the position of the radiation source in the sky to within 12 meters. They can then reconstruct the charge structure inside storm clouds, helping them understand what causes lightning and when and how it touches the ground.

To study volcanic lightning, the researchers pack the sensors—along with 160 gigabytes of memory, worth three months of recording time—into 20-kilogram boxes. Then the researchers must get the sensors to the right place at the right time. On the first two occasions they tried this, they didn't quite make it in time to get all the data they wanted.

During the January 2006 eruption of Alaska's Mount Augustine, the team arrived after the eruption had started and were able to set up only two sensors. The data was not enough to generate three-dimensional images, but it revealed a new type of lightning. Until then, lightning in volcano plumes was known to resemble thunderstorm lightning—highly branched flashes that last about half a second. But at Mount Augustine the researchers also found continuous, explosive sparks that lasted only a few milliseconds, which appeared at the mouth of the volcano just when it started erupting. This indicated that the eruption itself, not just the ejected ash and rock, had created a large amount of charge. Thomas is not sure how the charge is generated, something he hopes the Redoubt experiment will reveal.

When the Chaitén volcano erupted in Chile in May 2008, Thomas's team also arrived later than was ideal, but this time they were able to get four sensors in place, giving them their first 3-D maps. Preliminary analysis showed horizontal lightning up to 8 kilometers long.

At Redoubt, which is expected to erupt soon, they'll have a head start and will for the first time be recording data right at the first eruption. Thomas's hopes are high: "We'll get a lot better estimate of what's going on inside the volcanic cloud."

In the kind of storm clouds that generate conventional lightning, ice particles and soft hail collide, building up positive and negative charges, respectively. They separate into layers, and the charge builds up until the electric field is high enough to trigger lightning. The conventional wisdom has been that in volcanic eruptions, charged ash and rock debris produce lightning by analogous processes. From what Thomas and his team have already learned, volcanic lightning might be more complex than that.

If they are successful in developing a mapping system, it could provide useful warning that an eruption has actually begun. "Just because a volcano is rumbling and making lots of seismic noise, you can't tell whether it erupted," Thomas says. Tamsin Mather, a volcano researcher at the University of Oxford, in England, adds that the sensors could be a handy warning system especially for "remote volcanoes in Alaska or Kamchatka that don't have people watching them all the time but have plenty of planes that fly in the vicinity." Airplanes have unknowingly flown into ash, which has sometimes choked their engines.

Volcanic lightning could also yield clues about Earth's geological past, Mather says. And it could answer questions about the beginning of life on our planet. Scientists suspect that volcanoes on a primeval, sweltering Earth could have been the cradle of life. They had the right ingredients: water, hydrogen, ammonia, and methane. Lightning would have been the essential spark that converted these molecules into amino acids, the building blocks of protein. —PRACHI PATEL

update

First Affordable Fuel Cells for Mobile Gear

Medis Technologies has a cheap, disposable power pack. Others aren't far behind

According to the U.S. Fuel Cell Council, an industry association, there are at least 40 different fuel cells for portable, stationary, and transportation applications. But as recently as last August, industry analysts were saying that fuel cells that were *both* small and cheap enough to power cellphones, PDAs, and digital cameras were still "a year or so away."

Medis Technologies, in New York City, seems to have proved them wrong. That August the company started selling a recyclable fuel cell with a reusable power management cable for US $34.99, with fresh power packs for $20. This makes Medis the first to market. But competition is coming later this year from a number of companies.

The initial version of Medis's cell, the Power Pack Portable Charger, provides 1 watt and is intended for only the low-power end of the gadget market—no touch screens, for example—and only four or five full charges at that. Medis subsequently released its higher-powered Xtreme 24/7 Portable Power Generator, which can charge a wider range of products, including iPhones.

While far from perfect—Medis's power packs can't be refilled—the cell is a breakthrough, says Bob Wichert, technical director of the U.S. Fuel Cell Council. "The difficulties in bringing fuel cells to the product stage have been greater than many people anticipated," he says.

In fuel cells, hydrogen reacts with oxygen over a catalyst to make electricity and water, but because a canister of hydrogen is difficult to come by, many fuel cells start with the hydrogen bound up in a hydrogen-rich fuel such as methanol and produce carbon dioxide as a by-product. Medis's technology is unusual in that it combines a hydrogen-rich liquid, borohydride, with oxygen to generate electricity, water, and the mineral borax. "Borohydride is expensive compared to methanol, but it's probably the least technologically demanding to bring to market," says Louis Stuhl, an inorganic-chemistry consultant and founder of ChemMotif, in Concord, Mass. Borohydride chemistry prevents carbon dioxide from poisoning the fuel cell catalyst, which is a big problem for many other technologies, such as methanol fuel cells.

In methanol-based fuel cells, the fuel is diluted with more than nine parts water to one part fuel in order to release the methanol's hydrogen. "We can have a much higher fuel concentration," says Mark Kinkelaar, Medis's senior vice president for technology and operations. Because no dilution is needed, "our fuel goes right to the anode, with no valves or pumps," he says. Pumps reduce efficiency and add complexity—and cost. Moreover, methanol fuel cells can exceed 400 °C, whereas the Medis unit operates close to the ambient temperature—another reason it's easier and cheaper to manufacture.

HYDRO-CHARGERS: Fuel cells for portables have arrived.
PHOTOS: FROM TOP, HORIZON FUEL CELL TECHNOLOGIES; MEDIS TECHNOLOGIES; JADOO POWER

So the Medis Power Pack is cheaper, but is it cheap enough? The obvious comparison is to a sealed-unit AA battery pack. Medis's start-up cost is on a par with, say, the Duracell Instant Power Charger, which can be had for $30. Then there's the $20 you spend for every four or five charges. The Medis fuel cell may make sense only when you don't expect to have access to a recharge outlet or if the added weight of a recharging unit is a concern—all-day air travel, for example, or a camping trip.

While Medis can claim to be the first with a consumer-priced fuel cell for mobile devices, it may not be the only player for long. At the January Consumer Electronics Show in Las Vegas, Horizon Fuel Cell Technologies showed its MiniPak fuel cell, which it hopes to market sometime in 2009 for an expected price of $50. The charger will use drop-in solid-state metal-hydride cartridges of different sizes that will cost about $1 per watt-hour. And Jadoo Power, which makes a 12-watt-hour recharger, unfortunately requiring a 2-kilogram hydrogen canister, has a whole new line of fuel cell products coming out this year.

Medis isn't standing still either. It's working on a cartridge-based version of its power pack and plans to work on a refuelable version; each would cut the total cost. "We know 20 places to make it better, but we need to start to launch," says Kinkelaar.

—Daniel P. Dern

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

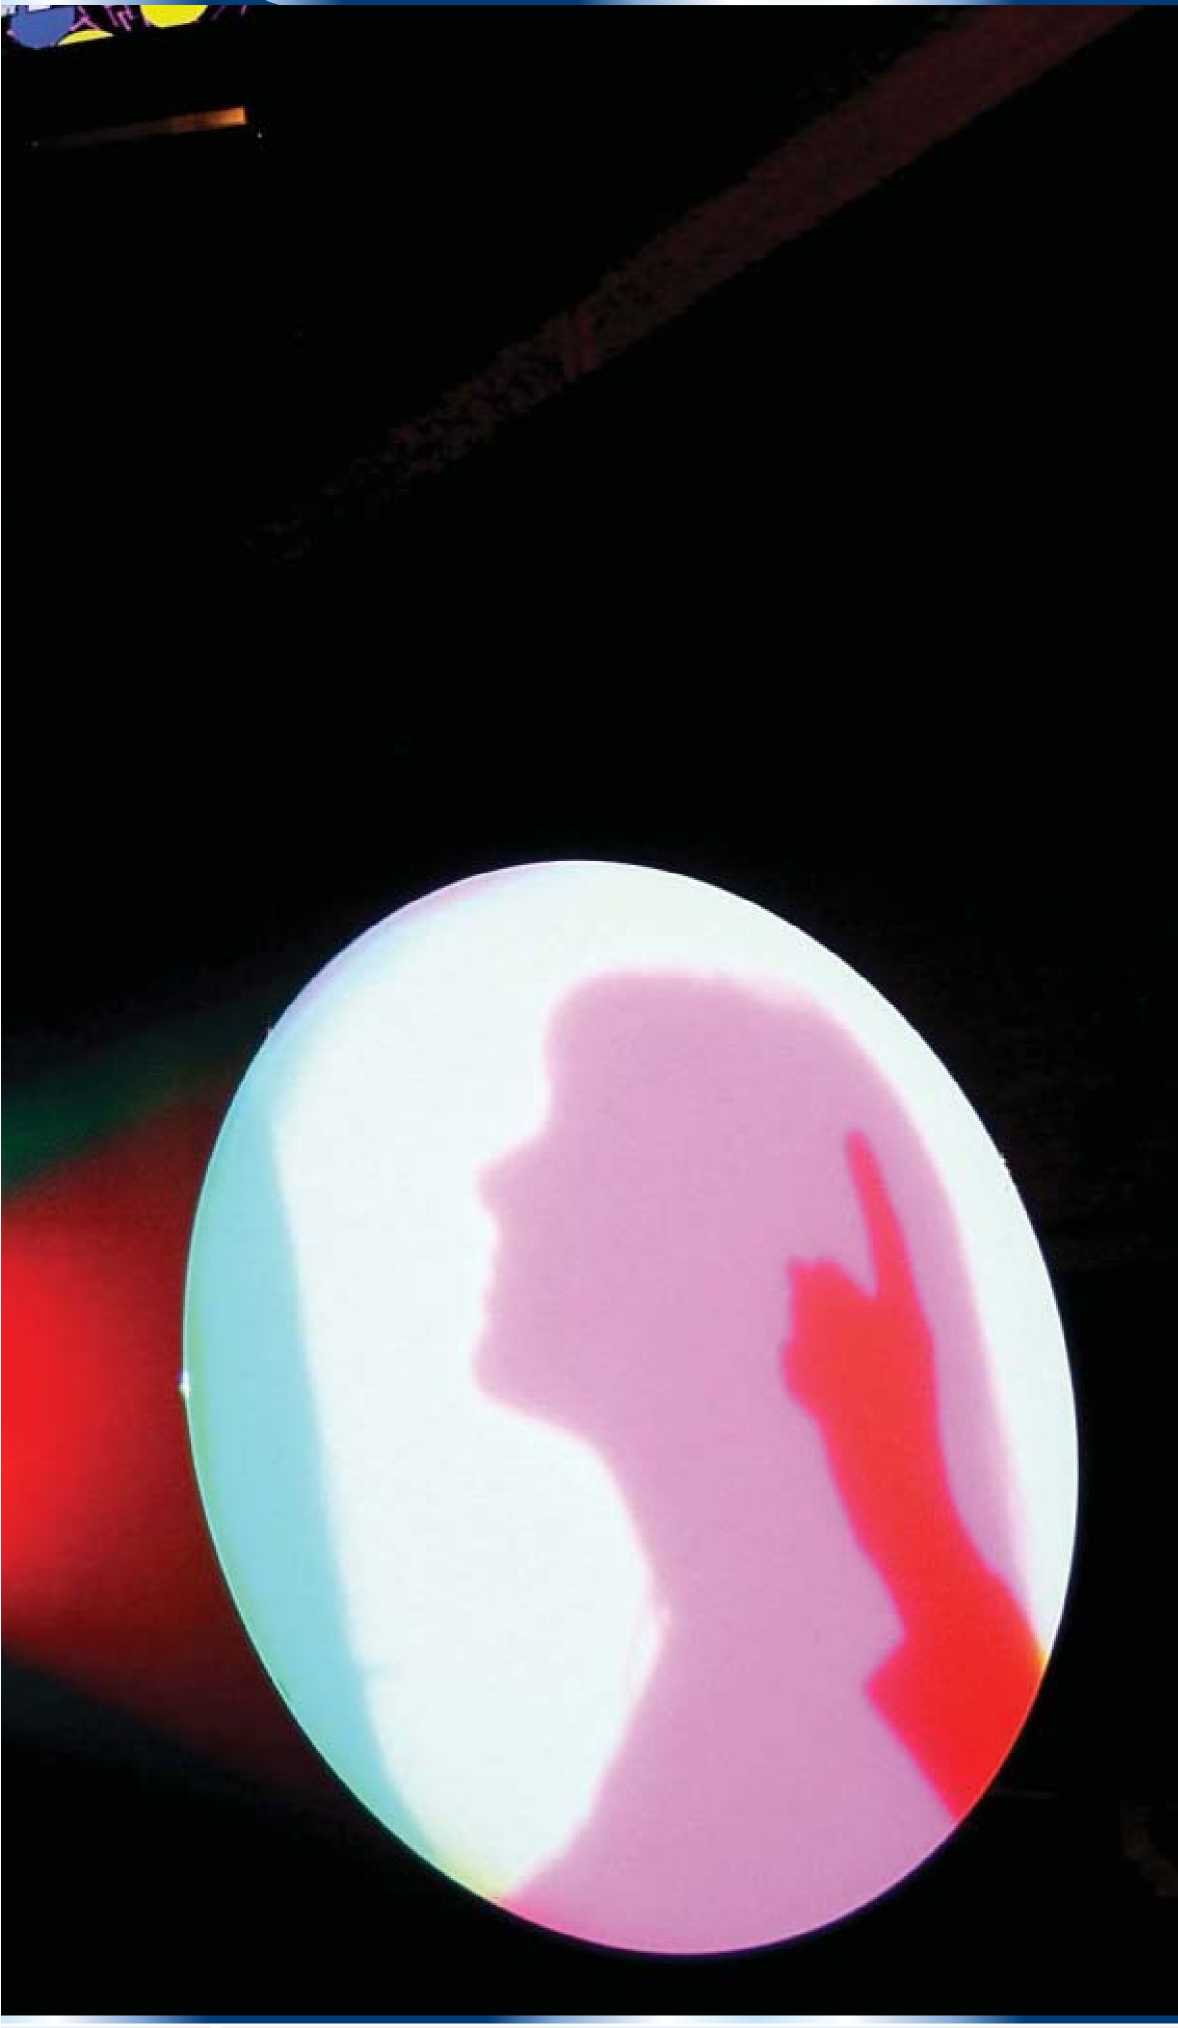

the big picture

PRIMARY COLORS

What is television but the deft arrangement of red, green, and blue? At a late January event sponsored by Sony, British artist Paul Cocksedge made the point that no matter how fancy and feature-filled the TV set, red, green, and blue lights remain encoded in its DNA. His installation at the Tramshed gallery in London juxtaposed a simulation of the basic elements of a TV pixel—in the form of colored spotlights converging on a single white spot—with some advanced new Sony sets, including one with an organic LED screen and another as thin as a CD case.

IMAGE: PETER MACDIARMID/GETTY IMAGES

hands on

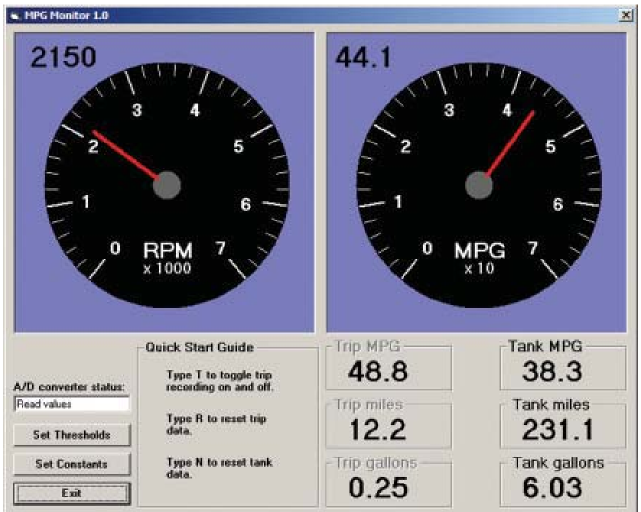

A FUEL-ECONOMY GAUGE FOR THE REST OF US

Most cars—even older, cheaper ones—control their fuel electronically. Here's how to tap into that signal

THE PRICE of gasoline is down from its stratospheric highs of last July, but there are still plenty of reasons to conserve fuel and plenty of ways to do so. Drivers can make changes in their cars and better yet, in themselves; unfortunately, the human element isn't easy to tweak. That's partly because we lack detailed real-time data about which driving behaviors are wasteful.

The information is easy to come by if you pay the price. A fuel-use monitor that plugs into the vehicle's OBD-II jack—the onboard diagnostic interface that mechanics use to diagnose problems—will set you back US $160 (for the ScanGauge II, for example) to about $280 (for the PLX Kiwi). That's assuming you have an OBD-II port—cars made before 1996 typically don't.

If you're willing to do a little hacking to your ride, there are cheaper ways. Do you have an old laptop gathering dust? Does it have a stereo line-in jack? If so, and your car's engine is fuel injected, you may be able to track your fuel economy with just a few scrounged parts. I'll first describe doing that and then an even better way—one that uses a new $40 gadget instead of an aging laptop.

The trick in either case is to tap two wires at your engine-control unit (ECU)—the one from the vehicle-speed sensor (VSS) and the one that controls the opening of the engine's fuel injectors. Finding which wires to splice into may be a bit of a challenge, but it's easy if you have a copy of your car's factory service manual, as I did. A resource that may help others in identifying ECU pin-outs is available on Innovate Motorsports' Web site. (My vehicle isn't listed, so I can't attest to this site's reliability. For this and other online resources, see the list of URLs below.)

Gaining access to these connections on my 1997 Geo Metro was straightforward, because the ECU is conveniently located directly behind the glove compartment. For each signal, I used two resistors wired to provide a 10-to-1 voltage divider, and two back-to-back germanium diodes to clamp the divided-down signals to plus or minus 0.3 volt. I placed these components close to the ECU to minimize the risk that a shorted wire would cause the injector to spray fuel. I then routed the VSS and injector signals to my laptop's line-in jack, confident that they wouldn't overload the sound-input circuitry. If this computer were doing anything besides gathering dust, I might worry about voltage transients getting past these diodes, but that really wasn't a concern for my ancient Pentium machine.

Having resurrected a 1990s-era laptop to monitor a 1990s-era automobile, it was only fitting to use software of a similar vintage—Visual Basic 6.0—to write the code that transforms the two signals into distance traveled, gasoline used, and their quotient, miles per gallon (others might prefer liters per 100 kilometers). Using the FreeView Sound driver (available at Technical-

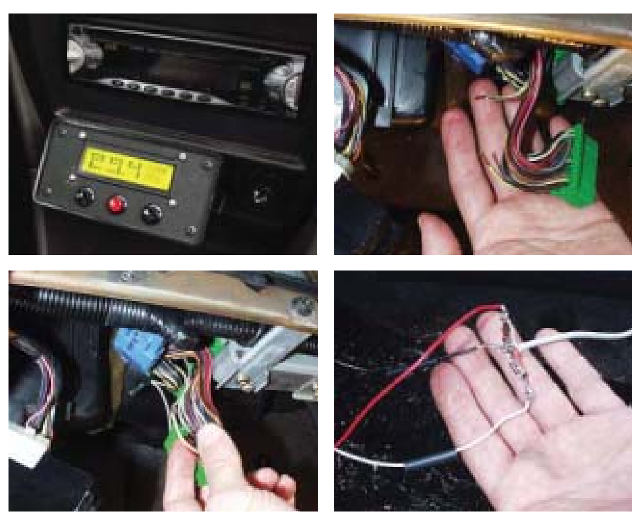

MPGUINO ON BOARD: Adding this open-source fuel-economy monitor [top left] requires tapping two signals at the car's engine control unit. Cutting the appropriate wires [top right] and soldering the ends together with an additional lead [bottom left] is the easiest approach. A few diodes and resistors [bottom right] limit the voltage swing.
PHOTOS: DAVID SCHNEIDER

computing.com) made it a snap to access the computer's audio-recording hardware. Still, I spent a long weekend getting the rest of the code written—including such niceties as a screen for viewing the raw signals and two simulated dashboard gauges for instantaneous mpg and revolutions per minute [opposite page].

It worked great, except for a couple of complications I hadn't counted on. The computer I used took forever to boot—so much for hopping in the car and darting off. Also, there was no good place to position the bulky laptop. So while the performance was good and the price was right, this strategy ultimately proved impractical.

I thus decided to replace the laptop with an open-source fuel-economy computer called MPGuino. This gadget is based on the Arduino microcontroller and is available in kit form for less than $40 from Spiff's Electronics Notebook.

Assembling the MPGuino was a quick job, and having already located the VSS and injector signals, the biggest challenge (as with many such do-it-yourself projects) was fashioning a suitable enclosure. I ended up with a RadioShack project box, which I mounted in front on the car's used-only-for-loose-screws ashtray. The calibration values listed for my model on the MPGuino wiki (at EcoModder.com) turned out to be spot-on, so setup wasn't difficult.

This little "carputer" has opened my eyes to how my car uses gas and what I can do to limit the burn rate. I've learned, for example, that idling while stopped really does waste a significant amount of fuel—it's outright painful to watch your trip-averaged mpg tick downward while waiting for a red light to change. On the other hand, sprightly acceleration burns less fuel than I'd expected, and it's great to watch the instantaneous-mpg numbers shoot up as you disengage the clutch, lift your foot off the gas pedal, and coast. But leaving the car in gear to decelerate is often a good conservation strategy, as it can put the engine in fuel-cutoff mode, which the MPGuino shows.

I plan to hone my ecodriving skills with this electronic aid. Perhaps by this July, my 12-year-old, three-cylinder Geo will be challenging my town's many Priuses for bragging rights. In any event, I'll be ready if gasoline prices again climb to levels that make your ears pop.

—David Schneider

ECU pin-out diagrams: *http://innovatemotorsports.com/resources/ecu_pinout.php*
FreeView Sound Software: *http://technical-computing.com/freeviewsound.htm*
Arduino microcontroller: *http://arduino.cc*
Spiff's Electronics Notebook: *http://spiffie.org/electronics*
MPGuino wiki: *http://ecomodder.com/wiki/index.php/MPGuino*

Maplesoft
Maple 12, MapleSim

tools & toys

NEW MATH

Major updates from Maplesoft, Wolfram, Design Science, and Tera Analysis present engineers with a calculatory cornucopia

The venerable math programs Maple and Mathematica long ago outgrew pure math to include numerical and symbolic math, graphics, document preparation, and more. Maplesoft and Wolfram have each recently released a major update—Maple 12 and Mathematica 7—and Maple has a new simulation package as well. The company is increasingly pitching to engineers—an approach that has long been successful for another big math program, MathWorks' Matlab.

For example, Maple's new Dynamical Systems package helps users study the behavior of discrete and continuous time systems, central to many signal-processing and control applications. Maple can now communicate with the two leading computer-aided design packages, Autodesk Inventor and SolidWorks, allowing an engineer to optimize a design in Maple and then send data back to the CAD program.

The new Exploration Assistant lets a user vary a parameter in an equation or graph and see the results. And inexperienced Maple users now have new templates that take common commands. Mathematica has had many of these features for some time. It's great to see that competition between the two companies is easing the learning curve users faced on earlier releases.

The new MapleSim is an elegant product that runs in parallel with Maple 12. With a few mouse clicks, you can create, say, a model of a circuit, or a thermal or mechanical system, and the program will transparently call on Maple 12 to show how it performs. Models

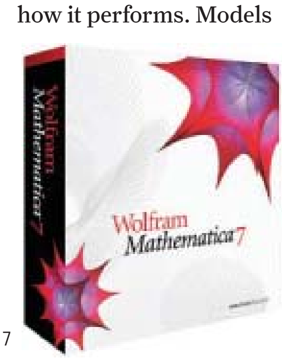

Wolfram
Mathematica 7

can even handle several domains, including motors, for example, as well as mechanical elements and the attendant friction and backlash.

MapleSim has a lot in common with MathWorks' Simulink, but, interestingly, it works by finding the differential equations for a system, which Maple simplifies analytically before solving. This process is potentially faster and more accurate than Simulink's purely numerical approach.

I found MapleSim easy and even fun to use. But automating it has some limits. Many of the things an engineer would want to do—for example, fitting data to a model to obtain a best set of parameters or tweaking a model's parameters to meet design constraints—are likely to force you to delve deeply into Maple and possibly require extra-cost software as well. There is also a key restriction to lumped parameter models (electrical circuits perhaps, but no electromagnetics).

Mathematica tries even harder than Maple to be all things mathematical to all people. Its latest version can locate information about subatomic particles, the human genome, historical weather data, the geometry of polyhedrons, and on and on. The program also includes many new graphical and mathematical functions and supports parallel computing on computers with multicore processors.

Users fortunate enough to work in institutions with maintenance contracts will probably have received these latest versions by now. The rest of us should carefully examine the vendor literature to decide

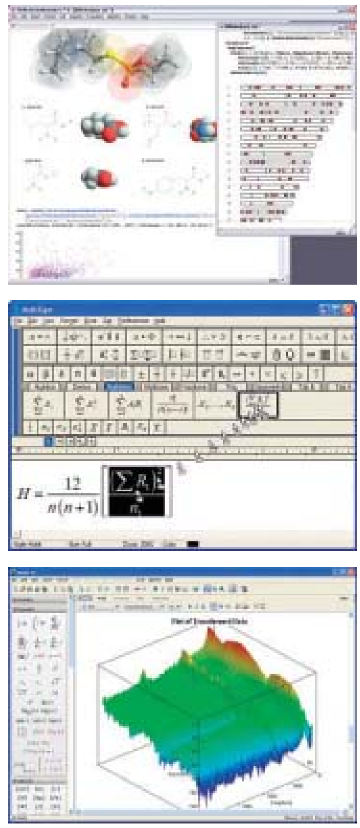

MATH PATH: Mathematica [top] makes genetic imaging easy—well, easier. MathType [center] lets you drag frequently used expressions to the toolbar for easy pasting. Maple 12's three-dimensional plots [bottom] let you visualize complex data.

whether the additional features of either program are worth a pricey upgrade.

Other programs for the mathematically inclined continue to evolve as well. Design Science has a new version of its stand-alone equation editor, MathType 6.5. Engineers acquainted only with the stripped-down MathType that's embedded in Microsoft Word will be pleasantly surprised to find a product that supports various mathematics markup languages and a wide variety of application software.

Do you want to copy an equation from Wikipedia, Mathematica, or Maple and produce a nicely formatted and editable equation in Microsoft Word? Or copy an equation from a Word document and paste it as an executable equation into Maple? Developers have tried for years to build this kind of capability into applications using the language MathML [see "MathType 5 With MathML for the WWW," *IEEE Spectrum*, December 2001] only to be stymied by the way different applications use different flavors of MathML. But now MathType has translators built into it for many different math packages, which should make copy-and-paste operations simple for the user. Neat.

Finally, there's QuickField, which I first reviewed back in 1993 as a program that numerically solved problems in a variety of physical domains—thermal, mechanical, electrical, mixed—using the finite element method. It featured a nice graphical interface and blazingly fast calculations. Now, numerous versions later, QuickField has acquired many advanced features, including optimization, publication-quality graphics, and so on. There's even a useful textbook, James Claycomb's *Applied Electromagnetics Using QuickField & Matlab* (Infinity Science Press, 2008), to go with a free student version that you can download from QuickField's Web site.

—Kenneth R. Foster

Maple 12: *http://maplesoft.com*
Mathematica 7: *http://wolfram.com*
MathType 6.5: *http://dessci.com*
QuickField 5.6: *http://quickfield.com*

media

Money for Nothing

Artist-programmer Aaron Koblin scored a 2005 viral hit with his *Flight Patterns*, a time-lapsed loop based on a day's worth of airplane flight data. More recently, he and fellow San Francisco artist Takashi Kawashima tapped Amazon.com's Mechanical Turk network to create a collage titled *Ten Thousand Cents*.

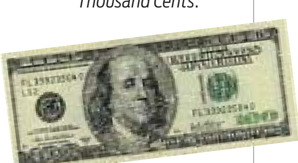

They scanned a US $100 bill, cut it into 10 000 pieces, and paid anonymous contributors a penny each to copy a piece with a software drawing tool, without knowing the overall design they were contributing to. The work appears through 12 April at the Pasadena Museum of California Art. See http://www.aaronkoblin.com.

—*Susan Karlin*

AARON KOBLIN AND TAKASHI KAWASHIMA/TEN THOUSAND CENTS

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

JOHN LAMB/GETTY IMAGES

careers

'TIL LAWSUITS DO US PART

Noncompete clauses may hinder you less than you fear

LAST NOVEMBER, Apple happily announced that it had hired Mark Papermaster, a 26-year IBM veteran, to head up its iPod and iPhone hardware engineering teams. IBM immediately sued Papermaster, claiming that his employment contract with Big Blue included a noncompete agreement prohibiting just such a move. A U.S. federal district court quickly ordered Papermaster out of Apple until further notice. The two companies agreed to a settlement in late January; Papermaster will finally sit down at an Apple desk later this month.

Noncompete clauses are essential to protect a company's intellectual property, especially its trade secrets—after all, they're not good secrets if employees can leave and tell them to the company's competitors. But courts are wary of these clauses and balance the protection of trade secrets against an employee's basic rights.

In its complaint, IBM asserted that Papermaster was "in possession of significant and highly confidential IBM trade secrets and know-how, as well as highly sensitive information regarding business strategy and long-term opportunities." Papermaster had been IBM's top blade-server guru, and the company offered him a substantial raise as an enticement to stay—it even went so far as to promise a year's salary in exchange for not joining a competitor. In his counterclaim, Papermaster said his work for Apple would not impinge upon the agreement because he wouldn't be working on blade servers there and because Apple's focus is on consumer electronics, not large-enterprise applications.

Predicting how a trial court will rule in a dispute over the validity of a noncompete agreement is a tricky business, says Ross Dannenberg, a partner in the Washington, D.C., law offices of Banner & Witcoff. "There are three types of limitations found in a noncompete agreement—geographical, temporal, or line of business," he says. "Courts base their rulings on whether any of them creates an unreasonable restriction on a worker's ability to earn a living." Thus, a court may void the contract clause if it finds that forcing you to remain unemployed for more than, say, a year does your career irreparable harm by rendering you a high-tech dinosaur.

Another consideration, Dannenberg says, is whether the agreement was the result of a true arms-length negotiation: "Prospective employees at the bottom of the food chain have little bargaining power and are in the position where they either sign or don't get the job." Dannenberg notes that was probably not the case with Papermaster, who was a company vice president when he signed his contract with IBM in 2006.

What a U.S. court finds also depends on where it's located. California state courts are unusual in that they generally refuse to enforce such agreements because of a statute outlawing them. This has permitted rapid job switching in Silicon Valley, which in turn has allowed the most talented people to move to the most innovative firms. An interesting wrinkle in the Papermaster case was that although Apple is headquartered in California—which would ordinarily leave IBM little recourse—the agreement he signed made it clear that any dispute would be decided according to New York state law.

This variability in enforcement can lead to a certain amount of schizophrenia for those who cross state lines when they jump from one job to the next. Take the curious case of Matt Marx, a software engineer who is now a Harvard Business School doctoral student. In 2001, still bound by a noncompete agreement he signed when he joined a Boston-based technology firm specializing in speech recognition, Marx was recruited by another speech-recognition firm, in Silicon Valley. He let the noncompete clause in his Boston contract hold him fast until he found out about California's ban on such agreements. On the other hand, when Marx later returned to Harvard, he declined an offer of part-time consulting for one of the founders of the Boston company because he knew Massachusetts courts would enforce the noncompete clause in his California contract.

No one can say how common noncompete agreements are because, as private contracts, they needn't be registered with a state authority (unlike a deed or marriage license). We hear only about the cases, like Papermaster's, that go to court. And even in the rare instance of a David boldly and publicly taking on a Goliath, an out-of-court settlement is sure to follow, to avoid a pretrial discovery process that would require both sides to reveal how the employee in question fits into their business plans and other zealously guarded corporate secrets. —WILLIE D. JONES

technically speaking

BY PAUL McFEDRIES

Information at Your Fingertips

At the center of this will be the idea of digital convergence. That is, taking all the information—books, catalogs, shopping approaches, professional advice, art, movies—and taking those things in their digital form, ones and zeroes, and being able to provide them on demand on a device looking like a TV, a small device you carry around, or what the PC will evolve into.
—Bill Gates, "Information at Your Fingertips—2005," (COMDEX keynote speech, 1994)

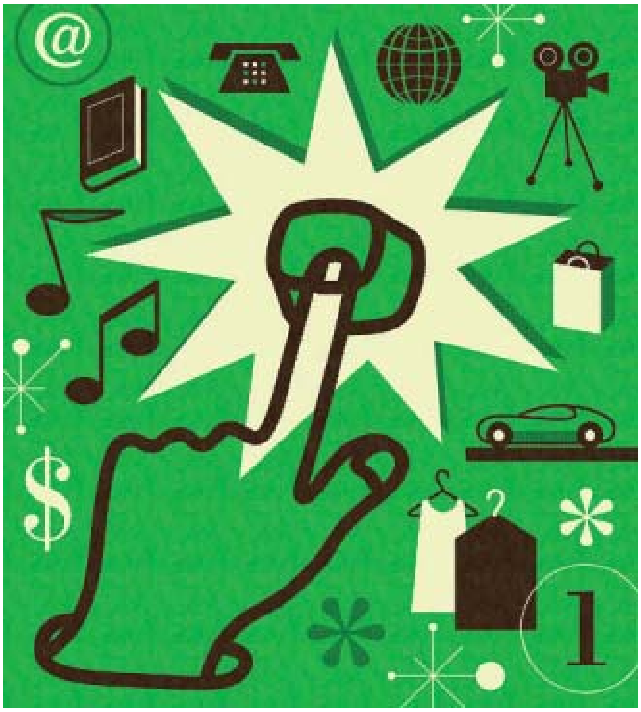

MICROSOFT FOUNDER Bill Gates is famous for many things, including being the world's richest person (until recently, anyway) and being the world's unofficial Alpha Geek. However, in tech circles he's also known for his long-standing belief that, as he once put it, "any piece of information you want should be available to you." The idea wasn't new—as far back as the 1970s, the motto of the Information Industry Association was "Putting Information at Your Fingertips"—but Gates championed it as early as 1989, and he was in a position to do something about it. It remained his overriding goal for the next two decades.

In fact, you could argue that **IAYF** (as the cooler geeks now call it) has been the goal for the entire tech sector for the past 20 years, particularly since the Internet broke out of its academic cloister and started cavorting in the mainstream. But a funny thing happened between then and now: Quietly and without much fuss, this seemingly futuristic goal has pretty much become a reality. Wondering if that restaurant you see from your car window is any good? Ask your car's GPS system. Somebody at dinner claims that Dustin Hoffman was in *Star Wars*? Whip out your iPhone and look it up in the Internet Movie Database—that information is **iPhoneable**.

Information at Your Fingertips went from pie-in-the-sky to data-in-your-smartphone seemingly overnight, and we're just starting to realize how much **findability** (also called **Googleability**) the world's data has. In fact, it's turning into what technologist John Seely Brown calls **calm technology**, remaining in the background until needed, thus letting you interact with it in a calm, engaged manner. IAYF will lead to an age of **ambient findability**, the ability to find anyone or anything from anywhere at any time.

The key is the rise of **small tech**, particularly **thin client** devices that are **Web enabled**. The slender iPhone is the archetype here, but **netbook** computers—tiny PCs optimized for Web viewing and other online activities—also enable folks, wherever they are, to Google away and access **cloudware**. Mobile hardware such as digital audio players, GPS devices, and even digital cameras are becoming increasingly **Web aware** (or, at least **Wi-Fi aware**). Radio-frequency identification (RFID) tags, real-world sensors, and other **ambient interfaces** are connecting (**deep networking**) all such devices with the world and its staggering wealth of data. Devices thereby become **environment aware**, giving rise to so-called **ecological computing** and **information surfacing**: revealing previously hidden data that we or our devices can use to better make our way in this brave new world.

The ultimate expression of Information at Your Fingertips will be **ubicomp** (ubiquitous computing), where we can harness data using everyday objects and interfaces (**ambient informatics** in the lingo), which some wag dubbed **everyware**. Like the atmosphere itself, if something is everywhere it becomes nowhere; we cease to notice it. Thus, **pervasive computing** (also the name of an IEEE magazine, by the way) becomes **transparent computing**.

Expect to see a privacy backlash against all this when people realize that once the world's information is at their fingertips, information about *them* is also at everyone else's fingertips. The new elite will be the **unGoogleables**, who've never posted anything online using a real name—no comments, no blog posts, no newsgroup rants, no Web pages—and so remain blissfully beyond the reach of Google and its ilk. Here's a can't-miss business idea for you, gratis: Start a company that specializes in **Google scrubbing**—removing all traces of a person from Google's servers. Let me know when you're a going concern and I'll look you up on my iPhone. ❑

GREG MABLY

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

Reach the next level

Pull ahead with CST MICROWAVE STUDIO. Explore your world with the No.1 technology for 3D EM simulation.

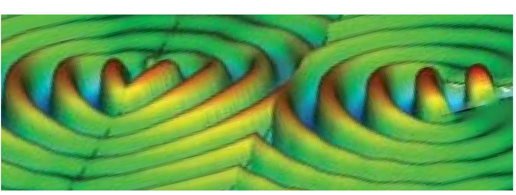

→ Get equipped for the job. CST's tools enable you to characterize, design and optimize electromagnetic devices all before going into the lab or measurement chamber. This can help save substantial costs especially for new or cutting edge products, and also reduce design risk and improve overall performance and profitability.

CST's flagship product, CST MICROWAVE STUDIO® [CST MWS], is the market leading time domain tool for 3D EM simulation, but its achievements don't stop there. CST MWS is the first commercial HF 3D EM code to offer the advantages of time and frequency domain, hexahedral, tetrahedral, and surface meshing, united in one interface. This gives you the flexibility to choose the technology best suited to solving your application. Embedded in an advanced design environment, CST MWS can be coupled with all CST STUDIO SUITE™ solver technology including circuit and thermal simulation.

Want to learn more? Take a look at why our customers have chosen CST technology: http://www.cst.com/Testimonials

→ Go further; draw your advantage from CST technology.

CST

CHANGING THE STANDARDS

CST – COMPUTER SIMULATION TECHNOLOGY | www.cst.com | info@cst.com

Shrinking Possibilities

Lithography will need multiple strategies to keep up with the evolution of memory and logic By Bill Arnold

SAY GOOD-BYE TO THE NODE. For 39 years, the node endured as the elusive and yet universally accepted metric that semiconductor specialists used to indicate how small their transistors were. Like depth readings on a wild descent into the infinitesimal, node figures were plotted out for the near future in a "road map" released annually by the semiconductor industry associations of Europe, Japan, Korea, Taiwan, and the United States. That map was, and is, a collection of the global semiconductor industry's best ideas about how it was going to fulfill the Moore's Law prophecy of a 30 percent shrink in transistor size—and consequent doubling in density—on chips every two years.

But now, in the first tremor of what promises to be a tectonic shift in the semiconductor industry, the node is no more. For decades, makers of logic chips used the concept of the "node" not only to measure their transistors but also to indicate how advanced their chip-fabrication lines were. Memory chipmakers, meanwhile, used a different measure, half-pitch, for the same purpose. Now lithography, the printing process at the heart of chipmaking, is being pushed to extremes to get to the end of that road map. These extremes will affect different devices on different timescales, but the end of the road looks the same for every device.

We're pulling out all the stops for the current generation of chips. And if that sounds like a platitude you've heard before, consider this fact: Nothing significant that we're using now will work to create the chips we plan to produce commercially just five or six years from now—least of all the current method of lithography. The next generation of chips won't be possible without the next generation of lithography. And that, in turn, means that the next generation of lithography will depend critically on what happens to the different chipmakers. For example, memory technology, an industry that sees prices falling at the staggering rate of about 40 to 50 percent per year, faces significant pressure to scale up faster than logic devices do.

Industry observers will not be surprised by the death of the node, as the node and half-pitch, once synonymous, have been diverging for some time. This minor change foreshadows a big change in the way the lithography business will deal with memory versus logic. For the first time, lithography will apparently have to adjust to follow both microprocessors and the different memories—including NAND flash, DRAM, and SRAM—down their respective paths, which have been diverging for decades.

Optical lithography, the most important and technologically demanding aspect of chipmaking, is a pillar that won't be easily toppled. But the technology is at a critical point. The technique, which uses radiation with about half the wavelength of purple light, is fast approaching steep, if not insurmountable, obstacles. Unfortunately, none of the various technologies proposed over the years to replace it has inspired confidence that it will be ready when the time comes.

Nevertheless, one thing is clear. From now on, the relationship between chips and lithography will be two-way. Not only will the fate of chips depend on the future of lithography, but also the reverse will be true.

Let's start by defining our terms. Today's most advanced microprocessors use a 32-nanometer process, and thus are said to be at the 32-nm node. To get a sense of how infinitesimal 32 nm is, consider that to span the width of the lowercase letter *l* on this page, you would need to bunch together more than 9500 32-nm objects. *Node* in this context has historically been used to refer to the size of the smallest parts of the transistors on the chips. Until the late 1990s, that was typically a feature called a gate. But there is a very fuzzy relationship between the technology node's number and the actual dimensions of the gate it purports to signify. In fact, the International Technology Roadmap for Semiconductors, the industry's guide star, abandoned the term in 2005, but its usage has persisted.

In both logic and memory chips, each of the vast profusion of transistors acts like a switch that allows electrons to flow through the device. A metal-oxide semiconductor field-effect transistor (MOSFET), the kind found on virtually all chips, has three main parts: a source, a drain, and a gate. A voltage applied to that gate lets the electrons flow from source to drain. Physically, the gate sits between the source and the drain.

On a chip, that translates to the distance between the parallel metal lines, called interconnects, that carry the electrons through the chip. These interconnects are stacked today in multiple levels, and as many as 10 can populate a chip (a cutting-edge chip could have 10 kilometers of interconnects). The distance between these metal lines at the first level is called the pitch, and logically, the half-pitch is half that distance.

The half-pitch of the metal lines on the first and densest level is special, because that distance was what once defined not only the half-pitch but also the gate and, consequently, the node. But by 2000, it was a dicey relationship. The half-pitch was becoming bigger than the node.

So, for example, in 2005 the gate width on an Intel microprocessor was 32 nm. The node was called "65 nm," but the half-pitch for the first level of wires was 105 nm. Confused yet?

The trouble with the terminology started in the early 1990s, when these gate widths fell down a steep slope [see table, "Pitch Counts"]. For logic devices, the gate length became the smallest feature, but for memory the half-pitch remained the smallest feature. Those were simpler, happier times. The industry sold microprocessors based on how fast the chips could process instructions, and that rate was pretty much directly related to how small the gate width was. We in the industry (I was working for Advanced Micro Devices, in Sunnyvale, Calif., at the time) called it the time of the megahertz wars. These wars drove the shrinking, with the result that the gate width got much smaller than the half-pitch.

So it was that, in the second half of the 1990s, the market's endless appetite for better-performing logic devices drove microprocessors—which used to lag two to three years behind memory in half-pitch—to start closing the gap. Because microprocessor speed was largely determined by the dimensions of the gate, by 2000 the gate had become the smallest feature produced in the semiconductor industry.

While speed is also a key parameter for memories, there was no similar war between memory manufacturers seeking to drive up the frequency with which transistors execute their instructions, known as clock frequencies. Memory makers focused on reducing the size of each memory cell on their chips so that they could squeeze ever more bits into less and less real estate.

So let's recap. Today's cutting-edge 32-nm-node logic chips are actually at a 50- or 56-nm *half-pitch*. Today's cutting-edge memory chips, if they were described in logic terms, would be at the 22-nm *node*, but they are in fact at about a 34-nm half-pitch, putting them ahead of logic development by a full chip generation.

The unifying factor is that both memory and logic have always been made with optical lithography. But because memories and logic have different shrink rates, memory makers will have to be the first to make the transition from optical lithography to the lithography technology that will succeed it, called extreme ultraviolet lithography (EUV). People who dread the transition to EUV often claim that it's not optical lithography, putting it in the same zoo with the other next-generation lithographies, such as electron beam lithography and imprint lithography. I prefer not to distinguish between optical and EUV: EUV *is* optical, inasmuch as it is electromagnetic radiation, only with a wavelength that's about one-fifteenth as long.

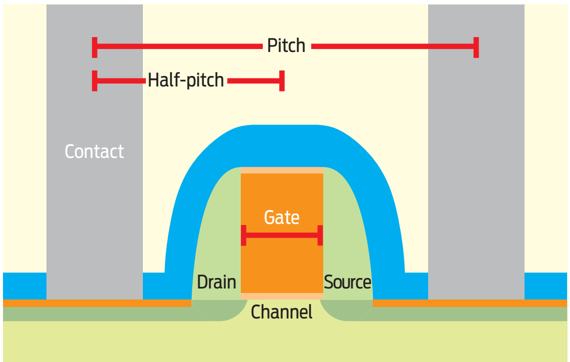

THE NODE'S DILEMMA: The half-pitch of the first wiring layer is the defining feature for memory chips, while the gate length is the gauge for logic manufacturers. Neither is entirely representative of the node.

The death of optical lithography, chipmakers have been known to say, is always seven years away. That joke was new 30 years ago, when chips were being fabricated at the 10-micrometer node. Today's chips squeeze 4 billion transistors into a space smaller than a postage stamp. The technology that brought about that stunning advancement, and the key driver of the semiconductor industry (which this year is expected to have revenues of US $200 billion), is this method of tiny writing called lithography.

Lithography is why Moore's Law endures after 44 years. Nevertheless, for the last 20 years, experts have been uneasy about lithography, which projects the fabulously complex patterns of a modern chip onto a semiconductor wafer using electromagnetic radiation with wavelengths shorter than those of visible light.

Fundamentally, optical lithography hasn't changed much in almost 50 years. It has become more sophisticated, but its kinship to old-time film-and-chemistry photography is still discernible. Microchips start out as small blank patches on a silicon wafer about the size of a dinner plate. The virgin wafer is shuttled through a series of machines in a chip-manufacturing plant the size of a couple of American football fields. At the end of its journey through hundreds of tools, the wafer emerges inscribed with the patterns of hundreds of identical microchips. The wafer is then broken up into these constituent microchips, which are sent out into the world to populate laptops, thumb drives, cellphones, and the GPS in your car. The patterns that distinguish all these different kinds of chips are created with optical lithography.

First, the wafer is covered with a thin insulating layer and then with a light-sensitive material called photoresist. Light streams onto the resist through an opaque mask with holes that let light through to form a pattern. This system projects postage-stamp-size patterns onto the wafer until the entire wafer is covered with identical microchip patterns. The exposed areas of the photoresist are weakened when the light hits them, and then a corrosive plasma etches the pattern into the silicon. The leftover photoresist is washed away, leaving the photoresist pattern engraved into the semiconductor wafer.

The process has gotten more complicated with each generation of shrinking features. Lithography toolmakers have had to reduce the wavelengths of light they use to project chip patterns through the masks. They've also had to find heroic optical tricks to finesse the light into depositing patterns far smaller than the wavelengths themselves.

The shorter the wavelength, the finer the resolution of the features you can print on the chips and the more transistors you can squeeze onto the chip. The history of semiconductor lithography is essentially the history of the search for stronger and shorter-wavelength sources of light. The first commercial lithography tools were manufactured in the early 1980s. They started skirting the edge of the visible spectrum with light at a wavelength of 436 nm, somewhere between violet and indigo. In 1987, steppers graduated to the ultraviolet 365 nm of mercury lamps and then to 248 nm in 1993. Finally, in 2001 the industry arrived at the 193-nm light, derived from an argon-fluoride laser. This laser, still used today to create patterns with feature sizes down to 38 nm, is projected through massive lenses that weigh nearly half a metric ton and cost several million dollars.

When semiconductor lithography began in the 1960s, the feature sizes of transistors were much larger than the wavelength of the exposure light. To print his original transistors, Gordon Moore actually cut patterns into Rubylith and projected them onto chrome-covered glass plates, or masks, using 16-mm movie-camera lenses that Robert Noyce had bought in a northern California camera shop for a few hundred dollars. Moore's transistors had a minimum feature size of around 100 µm, big enough to see with the naked eye.

Lithography had to keep up with feature sizes as they shrank to the size of the wavelengths of light itself—in the hundreds of nanometers—and then, more recently, vanished into mere fractions of the exposure wavelength.

As each image is exposed onto the semiconductor wafer, the lenses reduce the images 75 percent. Such sophisticated systems can expose more than 140 wafers per hour. Advances of this sort have improved resolution by a factor of about a thousand from the days of Moore and Noyce and the movie-camera lenses.

Impressive though it is to be printing today's 38-nm features with 193-nm light, what the industry would like more than anything else is to get back to printing features that aren't any smaller than the wavelength being used to print them. Here's why. Light diffracts when it shines through the mask, spreading out and blurring the edges. However, you want to capture as much of the light as possible to produce a good image. That means you need some pretty good tricks to corral all the *Continued on page 46*

Seeing the Light

The creation of 13.5-nanometer photons requires some creative contortions

A STRONG BEAM of EUV light is difficult to achieve. No traditional light source is capable of emitting 13.5-nanometer radiation. Therefore, only a small percentage of the radiation created at the 13.5-nm wavelength reaches the photoresist to make the pattern. We make EUV photons with a laser-produced plasma created by Cymer, in San Diego.

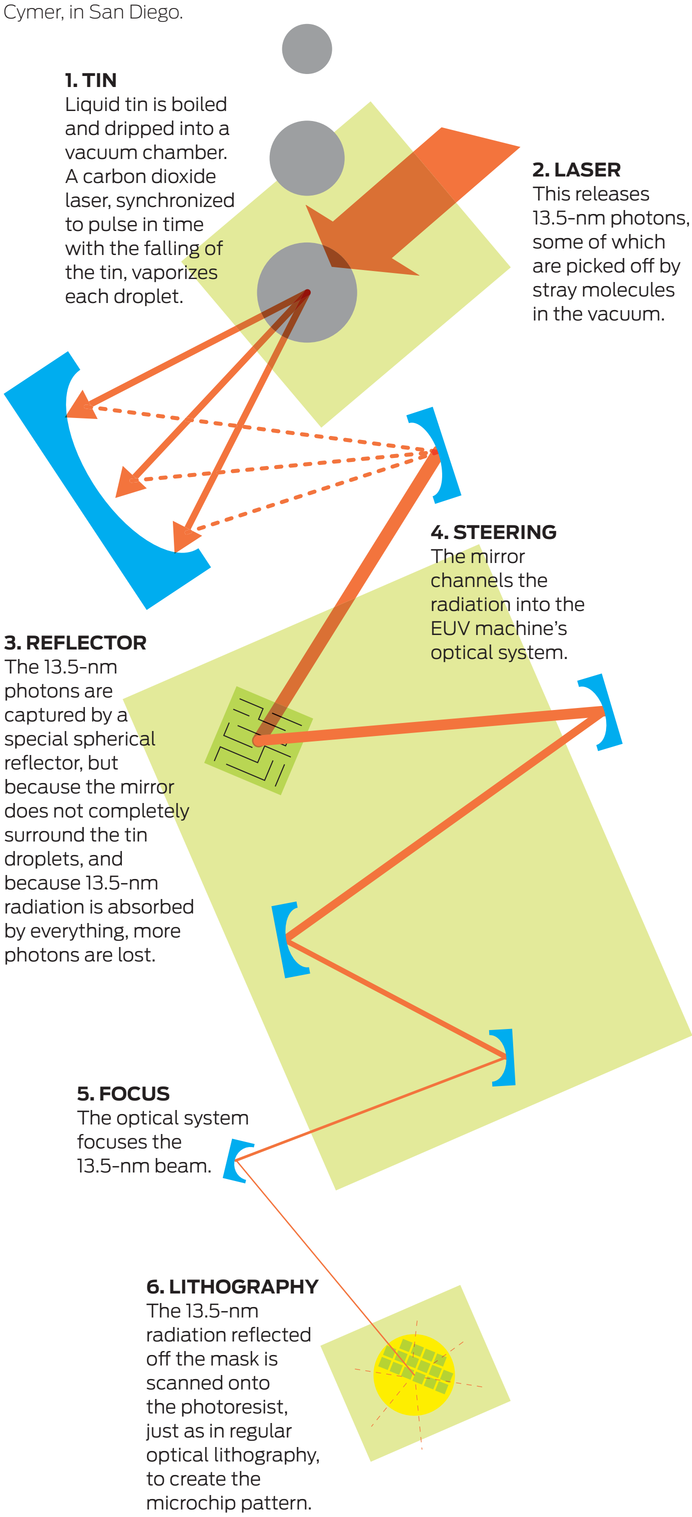

MARCH OF THE Sand

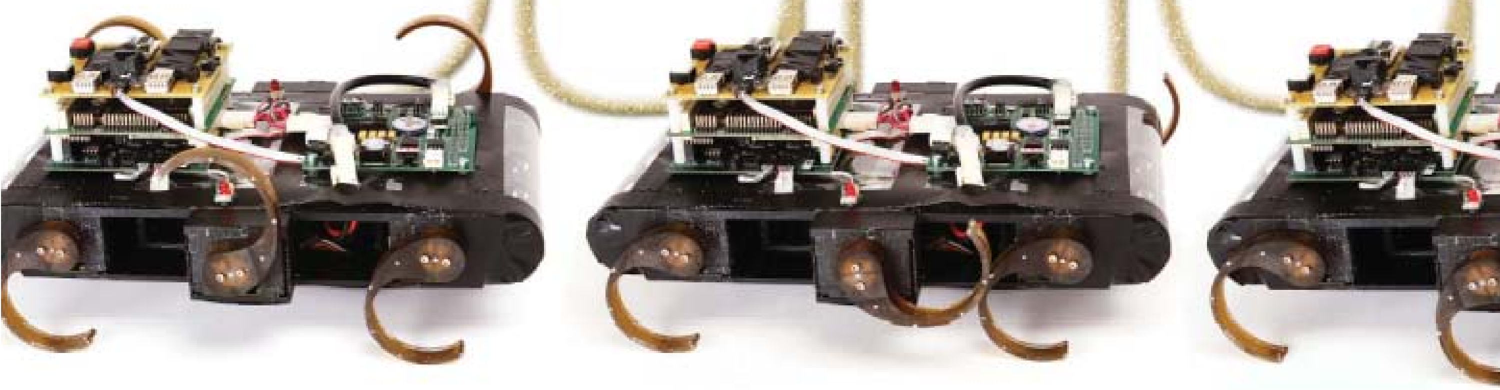

A new generation of legged robots will navigate the world's trickiest

A ZEBRA-TAILED lizard stands on a bed of tiny glass beads and shifts its weight. The beads slip underfoot, and the mottled beige creature stretches its spindly toes to get a better purchase. Suddenly it breaks into a run, blazing across the granular surface with stupendous agility, its toes stretching out flat as they hit the beads, its feet whipping back and forth in a blur. Each side of the lizard's body stretches and then coils in turn as the reptile darts ahead at several meters per second.

Scooped up a year ago in California's Mojave Desert and transplanted to a lab at Georgia Tech, the lizard holds our interest because of its truly peculiar feet. Those long, bony toes allow the reptile to navigate over sand, rocks, and the many other types of terrain it may face in the desert. In the lab, the bed of glass beads stands in for desert sand, and by blowing air through it or packing it down, we can make the ground looser or more solid. We then study how the lizard copes with the changes.

Our interest isn't purely biological. We—Goldman at Georgia Tech, Koditschek and Komsuoglu at the University of Pennsylvania, in Philadelphia, and our other collaborators—are hoping that by studying the zebra-tailed lizard and a menagerie of other desert-dwelling creatures, we can create more agile versions of our six-legged robot, SandBot. When traversing solid ground, the robot runs at a steady clip of two body lengths per second. (For comparison, a trotting dog covers four body lengths per second.) But on its first outing across the glass beads, SandBot dug holes fruitlessly with its crescent-shaped feet and got stuck after just a few steps.

Sand, it turns out, is one of the most difficult terrains for a robot to conquer. Sand is slippery, for one thing, and it is also inherently unstable: Its properties can easily flip between solid and fluid behavior in the course of a single footstep. Physicists still don't have a complete picture of the mechanics of sand, which is why we've turned our attention to the lizard and the clever strategies it has evolved to cope with sandy terrain. For example, we have noticed that the lizard's long toes sink deep into the sand at each step. It appears that this allows it to push off from sand that's deeper and more solid than the less stable surface layer. The

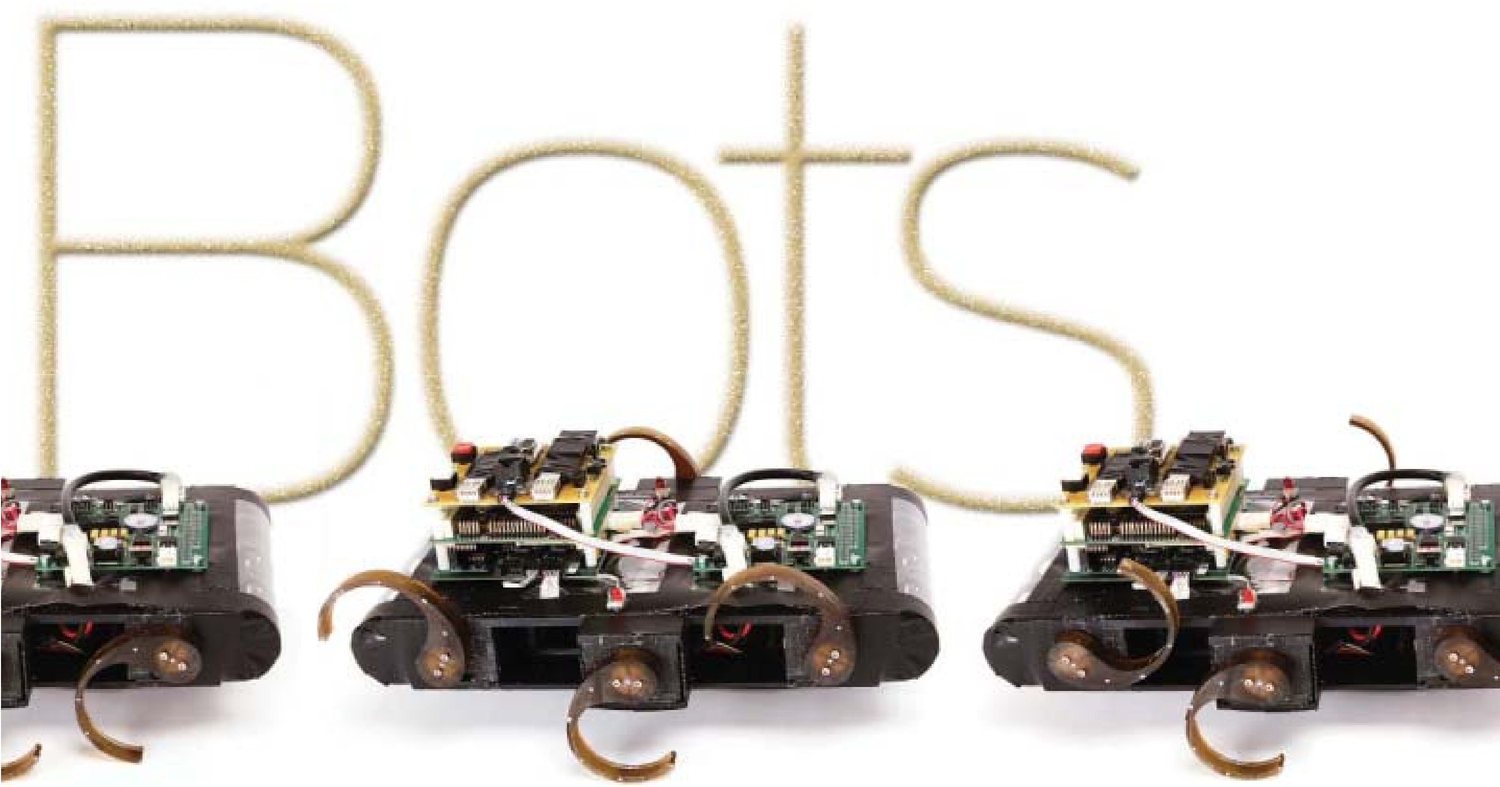

terrain By Daniel Goldman, Haldun Komsuoglu & Daniel Koditschek

effect, preliminary evidence suggests, is that the sinking enables the lizard to run as if on hard ground, allowing it to maintain speeds up to 75 percent of its pace on solid ground. Desert animals deal with sand with different levels of success, and their techniques provide valuable clues for refining SandBot.

Ultimately, we would like to build robots that can traverse any kind of terrain—bounding across hard ground like a gazelle, scaling tall trees and buildings like a squirrel, or maneuvering over slippery piles of leaves or mud like a snake. At least for short periods, a few robots already have managed to scale vertical walls, leaf-covered slopes, and even ice. Eventually, highly mobile robots could make a big difference in search-and-rescue missions and could explore all kinds of tricky terrain, not just on Earth but on the moon, Mars, and beyond.

First, though, our machines need to conquer sand. Had we been designing a wing for flying or a flipper for swimming, we would have been guided by the well-established rules for fluid flow, the Navier-Stokes equations. But for a complex material like sand, the equivalent models do not yet exist. So we had to start at the very beginning, by investigating the physical properties of granular materials. After about two years of study and experimentation, we in our small consortium of physicists, roboticists, and biologists think we have identified some basic rules describing movement across granular surfaces. Applying that knowledge to designing sandworthy robots, though, is not at all straightforward.

Consider how humans transport themselves over land. In places where massive investments have been made in roads and tracks, it's relatively simple to move about by car or train. In fact, our vehicles require all of that engineered smoothness—without it, they can't go far. But much of the Earth's surface is largely inaccessible to vehicles, including robots. About 30 percent of the land area is desert, and one-fifth of that is covered by some kind of sand.

Sand isn't the only issue. Disaster sites and battlefields—precisely the places where mobile robots are expected to be most useful—are full of unpredictable, impassable rubble. In 2001, for example, robots were sent in after the World Trade Center towers collapsed, but debris quickly clogged their tracks or caused

ALL PHOTOS: YVONNE BOYD

the robots to flip over. Likewise, when a coal mine collapsed in Sago, W.Va., in 2006, a rescue robot made it about 700 meters past the mine's entrance before getting stuck in mud. Even benign stuff like gravel and fallen leaves can stop a robot cold.

In short, robots that navigate on wheels and tracks are nearing their performance limits. Legged robots that mimic the movements of insects or animals offer a promising alternative, but figuring out the mechanics of walking hasn't been easy. Because not much is known about how the forces between a foot and the ground interact to create movement, the prevailing method for designing these robots has been essentially trial and error: Build the machine and hope for the best.

But we've come a long way. The first computer-controlled legged robot dates back to the 1960s, when Robert McGhee's Phony Pony took its first halting steps at the University of Southern California, in Los Angeles. McGhee then followed up on that project at Ohio State University, in Columbus, creating the first autonomous legged robot in 1976. This machine, known as Hexapod, could make its way slowly across some wooden blocks indoors.

A decade later, McGhee and his colleagues' 5-meter-long Adaptive Suspension Vehicle was the first autonomous legged machine to tackle the great outdoors. Moving ponderously at a fraction of a body length per second, the robot carefully placed each leg and then torqued its joints to generate the necessary ground-reaction forces to push its body forward.

The next phase in legged robots was ushered in with the dynamically dexterous machines built by Marc Raibert at Carnegie Mellon University, in Pittsburgh, and later at MIT. Dynamic dexterity is the ability to exchange potential energy and kinetic energy in a controlled manner—or the difference between a hopping kangaroo and a car. A kangaroo's bent legs store potential energy, which allows it to bound effortlessly over obstacles. The ability to direct its body's flow of mechanical energy is critical for a robot to navigate unpredictable terrain. Raibert's creations were essentially self-excited pogo sticks that used springs to balance, hop, and when yoked together, trot and bound. These robots still hold the ground speed record of 21 kilometers per hour, but they were strictly designed for controlled laboratory environments.

The RHex robot, designed by the roboticist Martin Buehler (then a professor at McGill University, in Montreal) and Koditschek's group in 1999, took running robots to the next level. This autonomous machine, inspired in part by integrative biologist Robert Full, of the University of California, Berkeley, has six legs that are attached outside its center of mass. This sprawled configuration grants the robot greater stability as it bounces over natural terrain. Faster runners have since appeared, but RHex remains, to our knowledge, the only legged machine that can traverse rugged, broken ground rapidly—at or above the pace of one body length per second.

RHex in turn became the model for a family of robots whose appendages are each driven by a motor located at the hip. Its progeny include, among others, the Aqua robot, which is basically RHex with flippers for swimming; a two-armed, wall-climbing robot named Dynoclimber; and SandBot.

In early 2007, Komsuoglu designed and built SandBot in less than a month, using the RHex model and a modular infrastructure of his own creation [see "Seeing Inside SandBot"]. At 2 kilograms, it is less than a quarter of the weight of RHex. Like RHex, SandBot has six compliant, independently controlled legs, each of which is a semicircular strip of plastic. Also like RHex, it walks with an "alternating tripod" gait, inspired by insects. The legs move in threes, with the front and rear leg on one

REGOLITH RUNNER: SandBot trundles down a track filled with poppy seeds, in preparation for the many kinds of dust, sand, and loose soil it will eventually encounter outside the lab.

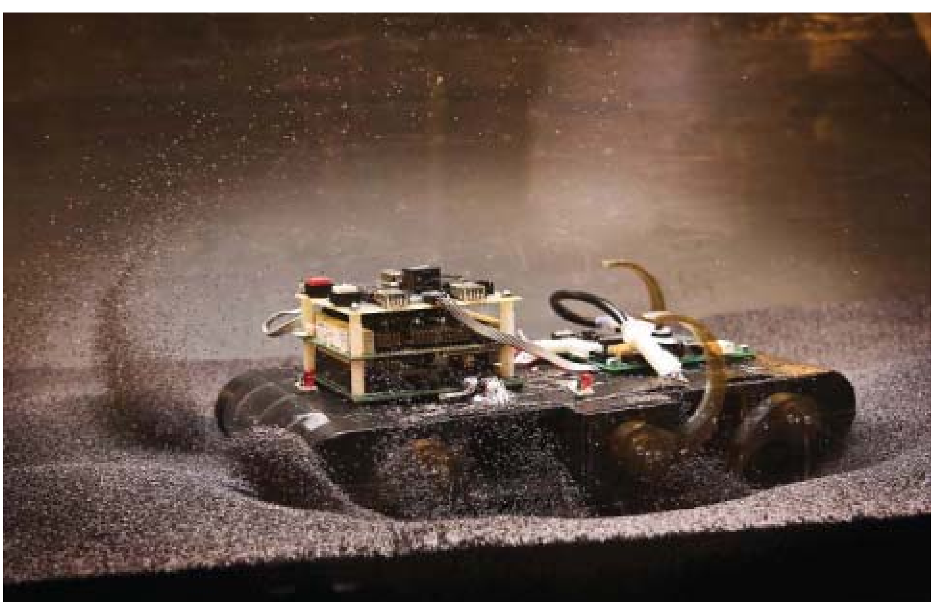

spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

Seeing Inside SandBot

SEVERAL BODY parts interact to allow SandBot to scurry over slippery grains. At the center is a computer that acts as the brain. In a typical jog, the legs support and propel the body by coordinating as two tripods—the two outer legs on one side and the middle leg on the opposite side. Relative to the ground, SandBot bounces like a two-legged pogo stick; the robot's legs, like those of a cockroach, are sprawled to give it more stability.

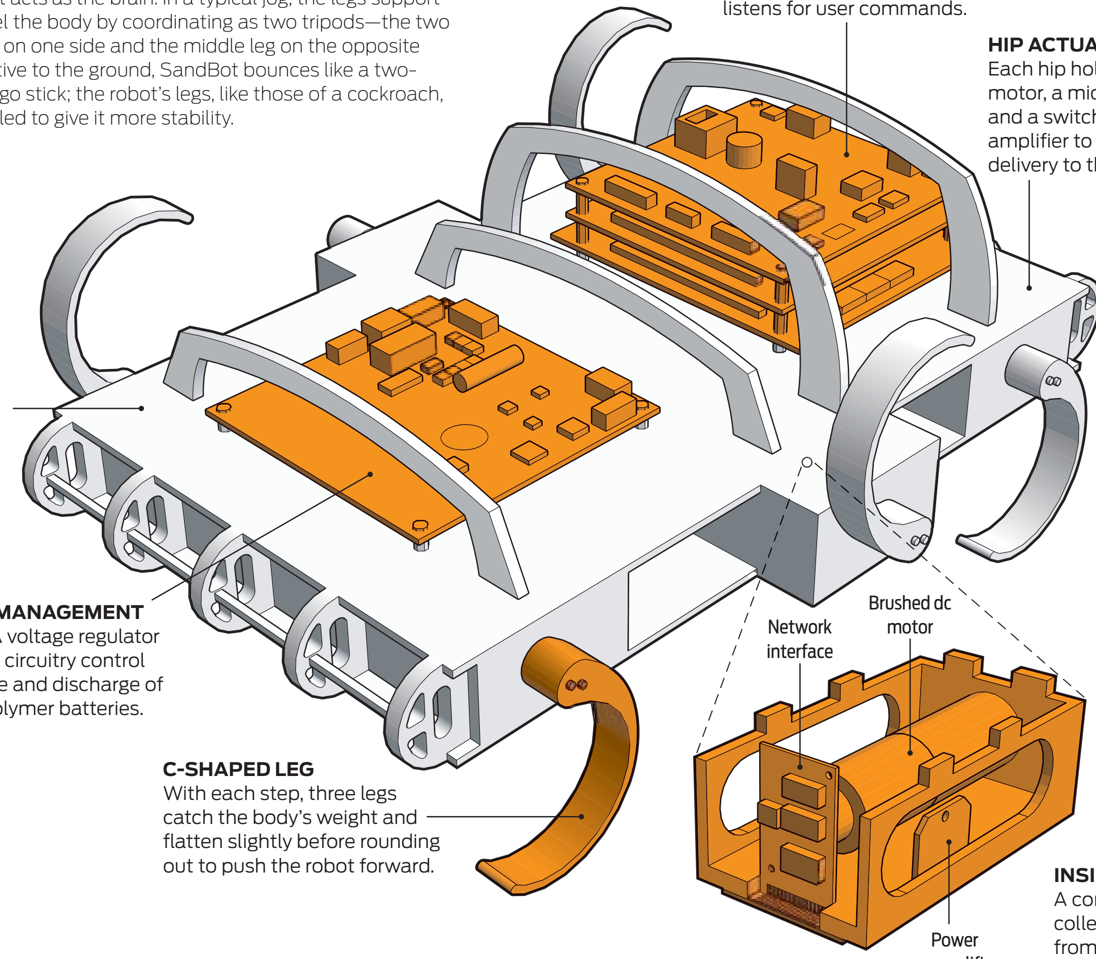

JASON LEE

side moving in sync with the middle leg on the opposite side. The two tripods alternate supporting and propelling the body, then circle around after each step.

On the inside, SandBot is composed of modular nodes that communicate through a real-time network called RiSEBus, inherited from an early version of its climbing sibling. At the hip joint of each leg sits an 11-watt brushed dc motor driven by a custom-designed motor controller board with a quadrature encoder, which senses the position of the motor's shaft and therefore the angular position of the leg. The six motor controllers link to a central computer, which functions as SandBot's brain and focuses on high-level behavioral decision making. Commands from the central processor instruct the motor controllers to bring the legs to a desired position and speed. A position-tracking controller determines the discrepancy between a leg's actual state and the desired state. The controller then computes the voltage needed to correct the error and applies it to the motor using a class-D power amplifier. This action gets the leg into position at the right speed.

To economize on the robot's computational power, the computer issues commands at the comparatively lazy rate of about 100 times a second, which frees up its cycles for other tasks. The central processor might, for example, tell one of the microcontrollers that its leg should move at a particular speed starting from a certain position. From then on, all tracking of that leg's position is carried out by the microcontroller, which can interrogate its sensors at the much higher frequency of 1 kilohertz.

This design allows the central processor to communicate with the legs using extremely compact data packets that require minimal computing power to decode. The separation of the control tasks frees up the central processor to perform longer-range planning. The central processor might use a camera to assess the difference between its relationship to a visual landmark and what it ought to be

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

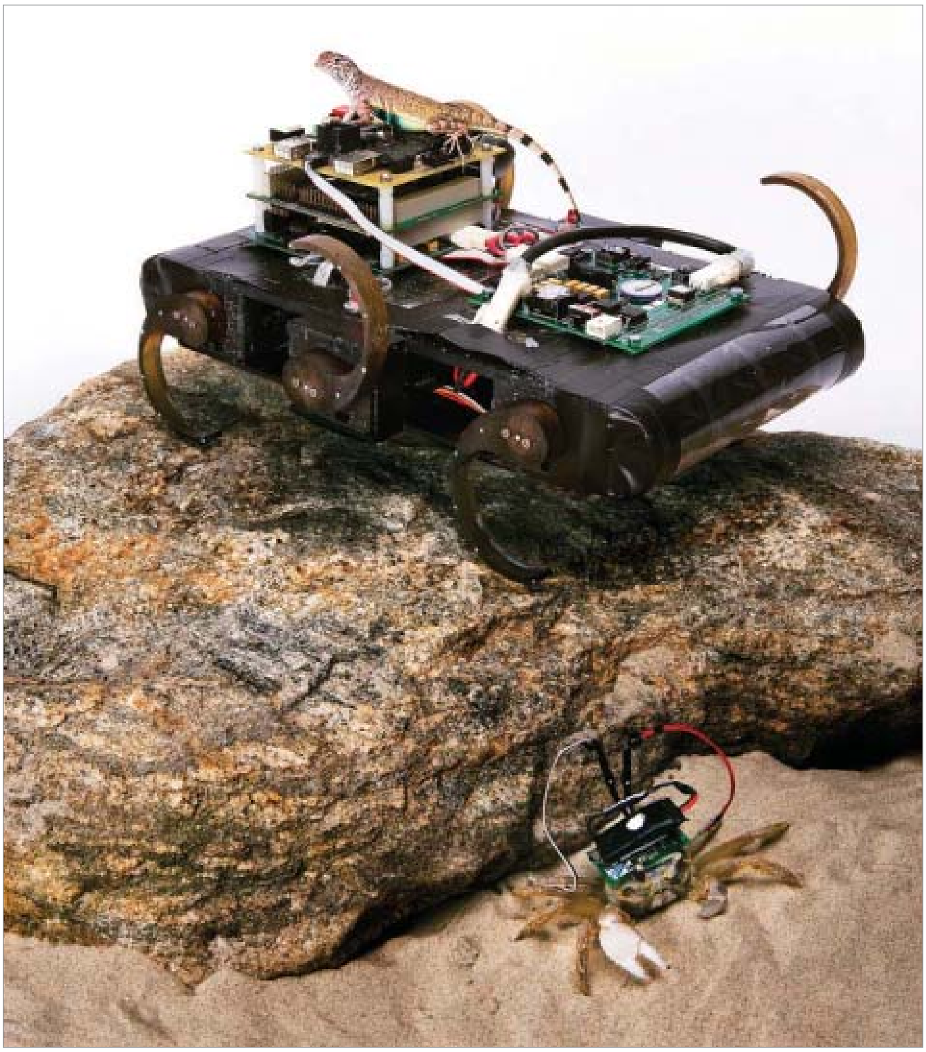

LEG ENVY: A zebra-tailed lizard and a ghost crab hold many secrets to fast running—such as how much a limb flexes on sand. The crab's backpack wirelessly sends measurements of the strain exerted by its legs.

or investigate how treacherous a surface is based on tactile feedback it retrieves from sensors in the legs.

SANDBOT'S DESIGN builds on experiments in Goldman's lab with real sand creatures. The zebra-tailed lizard, for example, can maintain a high speed over sand of almost any kind. The ghost crab, by contrast, is less versatile; on packed ground, its limbs and feet extend out from its 4-centimeter shell, and it scuttles along at a rapid 1 meter per second. But on looser soil the crab gets bogged down. The wind scorpion, for its part, can cover several body lengths per second even on granular slopes, where every step could trigger an avalanche.

Our observations of the lizard, crab, and scorpion under different conditions have helped shape our theory of sand locomotion. We believe this project represents the first attempt to combine direct measurements of a flowing physical substrate with observations of a runner's impact on the ground and its body movements. Broadly speaking, an animal's weight, foot shape, and gait all work together to apply a specific amount of stress to the sand. Under that model, the lizard is accessing the solid features of sand rather than slipping through the material and paddling, which is what the ghost crab ends up doing on softer terrain.

Much can be learned even from a single footstep. With each stride, the drag forces generated when a foot moves through sand can display both solid and fluid properties. If the stresses generated by the foot exceed a certain threshold, the material will flow. But it can also suddenly solidify if the stress drops sufficiently. That can happen, for example, if the downward forces produced by the limb and the weight of the robot are balanced by the amount of pressure within the sand, which is a function of its depth.

Another facet is that the behavior of sand depends on what's happened to it in the past. A section of solid sand disturbed by a footstep may be more loosely packed when the next foot hits the material, for example. The forces generated by a foot stepping into these different conditions can vary dramatically—the penetration resistance varies by a factor of 1.6 between a tightly packed material and a loosely packed one. That complicates the task of predicting how far a limb will penetrate in different granular states.

To learn how SandBot can best maneuver in sand, we have been subjecting it to a variety of precisely controlled granular environments. We control the environment using a 2.5-meter track built by Chen Li, a graduate student in Goldman's laboratory in Atlanta. The track looks sort of like a long bathtub, and it's filled with 90 kg of poppy seeds. There are tiny holes in the bottom through which we can blow air, causing the poppy seeds to lift off and dance before settling into a loosely packed state. (Why poppy seeds and not actual sand? We've found that each seed is large enough to keep us from worrying about it getting into the motors and yet light enough to be lofted by our air puffs. From separate experiments, we know that the exact material doesn't matter, as long as it is made up of granules.)

With sand and other granular media, we can describe the "strength" of the ground in terms of its solid volume fraction—that is, the fraction of the total volume occupied by the granules. Typically, the solid volume fraction falls between 58 and 64 percent for materials like sand or piles of seeds. A lower fraction means that, on average, there are fewer points

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

of contact between the grains and that the material is less solid. In our test track, an exact sequence of hundreds of air pulses carefully packs the poppy seeds to the desired volume fraction.

Because RHex had been so successful at walking on a variety of surfaces, we assumed that the smaller but relatively more powerful SandBot would perform well on sand. We were wrong. In an early experiment, we packed the material to a solid volume fraction of 63 percent, placed SandBot on the surface, and set the frequency of the alternating tripod gait to 5 revolutions per second. Earlier, the robot had bounced flawlessly across hard ground using those same parameters.

This time, though, it got stuck after just a few steps. Like a car's tires spinning in mud, the robot's rapidly rotating legs produced absolutely no forward motion on the poppy-seed-filled track. Discouraged, our first assumption was that SandBot was simply too heavy to walk on sand and that we would need to completely redesign the robot.

But we decided to play around with it a bit more. Komsuoglu, conferring by phone from his office at Penn, suggested that we modify the gait slightly to make the legs swing faster in parts of the cycle and slower in others. He knew from previous studies he'd done that some robots perform better with such a varied gait, at least on hard surfaces. It seemed worth a shot. As Komsuoglu told us over the phone which values to change, we entered them into the control program and, like magic, the robot started to move! The robot was still cycling its legs five times per second, but now it was scurrying down the track at one body length per second. Further study showed that each limb penetrated the poppy seeds until it supported the robot's weight, providing enough stability for the machine to thrust up and forward.

With Paul Umbanhowar, a mechanical engineer at Northwestern University, we subsequently developed a kinematic model explaining the relationship between the volume fraction, the limb rotation frequency, and the depth of the limb's penetration at each step. As both the model and empirical evidence show, if we increase the frequency with which the robot rotates its limbs, the robot sinks further into the material and the size of each step decreases, triggering a catastrophic loss of speed—quite the opposite of what happens on hard ground.

ANOTHER IMPROVEMENT we're working on is building SandBot a better foot, to give it the ability to grip sand just as the zebra-tailed lizard does. To that end, we've been measuring the forces on the foot during impact with and penetration of materials of different volume fractions. The tests look deceptively simple: We embed accelerometers into simple disc-shaped objects and then drop them on piles of sand. The results show that the forces produced when a foot hits the ground have different qualities in high- and low-volume-fraction materials. When the sample foot falls into a low-volume-fraction material, the force on it increases until the object comes to rest. When the object falls into a closely packed material, the force decreases during penetration. To also investigate the drag and lift forces that arise during the other parts of each step, we use a robotic arm to maneuver model feet and toes along granular paths.

To fully model the behavior of individual granules, we must resort to simulation. Yang Ding, a graduate student of Goldman's, has developed a computer simulation that models collisions of objects with sand, beads, and other granular media. We hope that eventually these foot experiments and simulations will feed into the development of a new sensing and control system for SandBot, to enable it to sense the shifting terrain ahead and swiftly adjust its gait to match. Sand isn't the only morphing environment that the robot could eventually tackle: Mud and loose leaf litter also display the solid and fluidizing features of granular media.

Indeed, with physics models built into their feet and brains, robots should one day be able to scramble across a rocky or sandy environment and learn, on their own, how to handle the changes in terrain from footstep to footstep. We can imagine thousands of SandBots scouring the surface of another world, stepping from a pile of rubble to a sandy patch with ease. That's still a big challenge for today's machines, but it's something even a hatchling sea turtle can handle. Despite having appendages that are better suited for swimming, these remarkable animals must climb out of a deep hole in the ground, clamber over grass and debris, and move across sand to reach the water, where they will spend much of the rest of their lives.

OUR OBSERVATIONS OF THE LIZARD, THE CRAB, AND THE SCORPION HAVE HELPED SHAPE OUR THEORY OF SAND LOCOMOTION

We're also looking below ground for inspiration. Using high-speed X-rays, we are now studying lizards called sandfish that can burrow into sand in the blink of an eye and then "swim" through the material underground. We're hoping these creatures will provide clues as to how robots could scramble through an unpredictable disaster area after an earthquake or flood or dig down to detect land mines. With nature as our guide, we expect that robots will soon master some incredible new abilities. ❑

TO PROBE FURTHER
For more about SandBot and its robotic relatives, see http://www.spectrum.ieee.org/apr09/moresandbot.

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

THE UNIVERSAL HANDSET

Software-defined radio will let cellphones speak Wi-Fi, 3G, WiMax, and more

by PETER KOCH & RAMJEE PRASAD

Time was when most radio sets had no software at all, and those that had any didn't do much with it. But Joseph Mitola III, an engineer working for a company called E-Systems (now part of Raytheon), envisioned something very different—a mostly digital radio that could be reconfigured in fundamental ways just by changing the code running on it. In a remarkably prescient article he wrote in 1992 for the IEEE National Telesystems Conference, he dubbed it software-defined radio (SDR).

A few short years later, Mitola's vision became reality. The mid-1990s saw the advent of military radio systems in which software controlled most of the signal processing digitally, enabling one set of electronics to work on many different frequencies and communications protocols. The first example was the U.S. military's Speakeasy radio, which allowed units from different branches of the armed forces to communicate effectively for the first time. But the technology was costly and rather unwieldy—the first design took up racks that only a large vehicle could carry around.

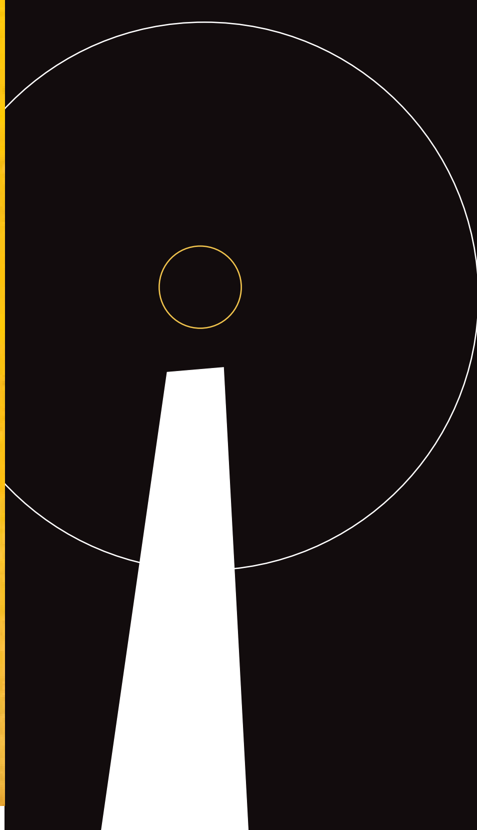

TAVIS COBURN

In the new millennium, SDR has spread from the battlefield to the commercial arena. Wireless service providers, in particular, have begun using it in the transceivers in cellphone base stations, allowing the same hardware to handle different cellular protocols. Next, SDR will spread to sets that fit in the palm of your hand.

That will come none too soon. Today's wireless mix is an often-turbulent alphabet soup of communication schemes: BGAN (Broadband Global Area Network), BT (Bluetooth), DECT (Digital Enhanced Cordless Telecommunication), EDGE (Enhanced Data Rates for GSM Evolution), GPRS (General Packet Radio Service), GSM (Global System for Mobile communication), IMT-A (International Mobile Telecommunications–Advanced), UMTS (Universal Mobile Telecommunications System), WiBro (Wireless Broadband), Wi-Fi, WiMax (Worldwide Interoperability for Microwave Access), and more. A mobile software radio that could communicate in all of these ways would, of course, be invaluable.

Up until now, SDR technology worked only in applications that didn't need to be small in size or frugal in power consumption. New technology should, however, overcome these constraints. Indeed, within the next year or so, you can expect to see people mothballing their old cellphones in favor of new software-defined handsets. By 2015, the transition should be nearly complete.

The first software-enabled sets to crawl out of the primordial ooze of traditional analog radio were modest affairs [see "Evolutionary Developments"]. They used embedded computers only to change the output level of the RF amplifier or to shift between individual RF front ends so that one unit could cover multiple bands.

In some of today's radios, software—often with the aid of digital hardware accelerators—does far more: It determines everything that happens to the signal after it's converted from RF to lower frequencies and before it's put in a form that's suitable for your ears. In these radios, only the RF front end and the amplifier that powers the speaker still use analog components.

The next era in SDR evolution will see what some call the ideal software radio or true software radio, in which the filtering and conversion from RF to lower frequencies that's normally accomplished in the radio's front end will be done digitally using the appropriate software—a strategy that requires moving the analog-to-digital converter (ADC) much closer to the antenna. These radios will still require a low-noise RF amplifier, though, because it's hard to imagine any ADC being able to pick up the micro- or even nanovolt signals generated in the antenna.

In the progression to that ideal radio, you'll probably notice that cellphones, mobile TV sets, GPS equipment, satellite phones, PDAs, digital music players, game consoles, and their kin will begin to look less and less distinct. As with the beaks of the duck and the platypus, the evolution of such gadgetry will converge toward the most functional form—in this case a small wireless unit that allows its user always to stay connected, from anywhere and for any type of content or use.

Designing such universal gizmos will be tough, of course. Perhaps the highest hurdle will be engineering the antenna, the size of which normally depends on the frequency of operation. Indeed, it's very difficult to make a radio with an antenna that is not a significant fraction of a wavelength in size. This dictate of physics introduces a fundamental problem, because you'd ideally like a single compact antenna to cover everything from FM reception, at roughly 100 megahertz, to satellite- and personal-network communications, which operate in the few-gigahertz range.

To cover such a large chunk of spectrum, you'd probably need a combination of something quite short, likely built into the unit's printed-circuit board, and something relatively long, such as the wire that connects with the user's earphones. But even if the frequency span isn't so great, designers probably won't be satisfied with just one antenna: RF engineers are quickly moving toward using multiple antennas, even for single-frequency operation. This strategy—known as multiple-input, multiple-output, or MIMO—allows for more reliable links and higher data rates. For example, IEEE 802.11n networking gear uses multiple antennas to communicate at about five times the speed of previous versions of Wi-Fi.

You can understand how MIMO works, at least in broad terms, with a simple thought experiment. Suppose you set up a transmitter with a single antenna and then move a receiver, also with a single antenna, far enough away for the reception to fade in and out once in a while. Such problems arise because the transmitted signal takes multiple routes to the receiver—some of it perhaps bouncing off a passing car, other parts reflecting off the steel beams of the building where the receiver is located. When the difference in length between two paths is half a wavelength (or three halves, or five halves, and so forth), the two waves will interfere destructively, clobbering the signal.

MIMO sidesteps that pitfall by multiplying the number of possible paths between transmitter and receiver. If the signal passed from one transmitting antenna to one receiving antenna fades, the signal from a different pair should still come in loud and clear, taking advantage of a phenomenon known to radio designers as transmit diversity.

Throw in some serious number crunching to process the digitized signals and you can achieve extraordinarily high data rates. Researchers at NTT DoCoMo, in Japan, which is developing such systems for 4G mobile communications, have managed 5 gigabits per second. And this wasn't just in a controlled laboratory setting; they achieved this rate outdoors, albeit with the receiver moving no faster than a swift walking pace (doing the same while traveling down the highway would be much more difficult). Impressive results with MIMO and other advanced antenna systems are also coming out of Stanford's Information Systems Laboratory, MIT's Lincoln Laboratory, and the Center for TeleInfrastructure at Aalborg University, in Denmark.

Another tricky issue for the makers of SDR handsets is designing the transmitter's power amplifier so that it can operate over a broad range of frequencies without mangling the signal. The challenge is not so great for FM transmission, but for communication schemes that require the amplitude of the wave to be manipulated, things can rapidly go awry.

Avoiding problems in such cases typically requires some kind of feedback mechanism. You can, for example, sample the output of the power amplifier and convert this RF signal to lower frequencies, which you can then compare with the signals used to modulate the amplifier. You can then compensate for any error you find by digitally adding the reverse distortion to

Evolutionary Developments

Radio steps into the digital age

HARDWARE RADIO

Only modification through physical intervention

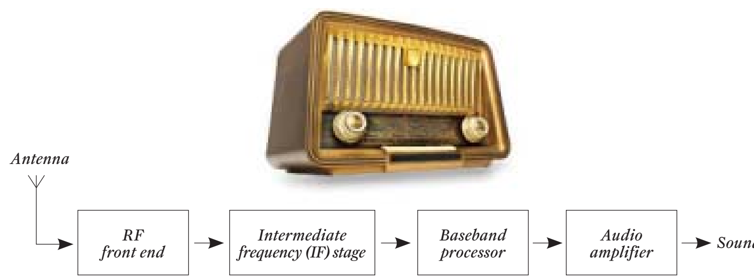

SOFTWARE-CONTROLLED RADIO

Computer selects circuitry to use

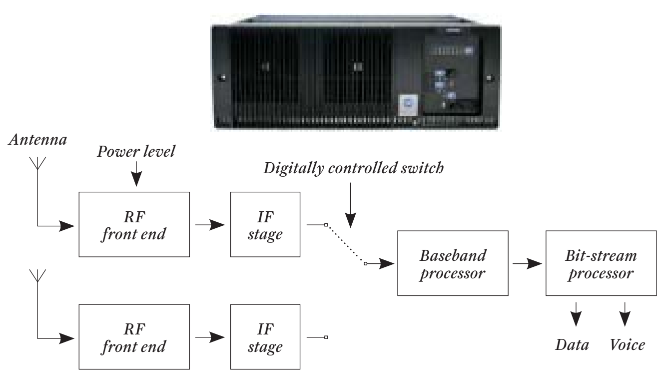

SOFTWARE-DEFINED RADIO

Software handles (de)modulation, frequency selection, security functions

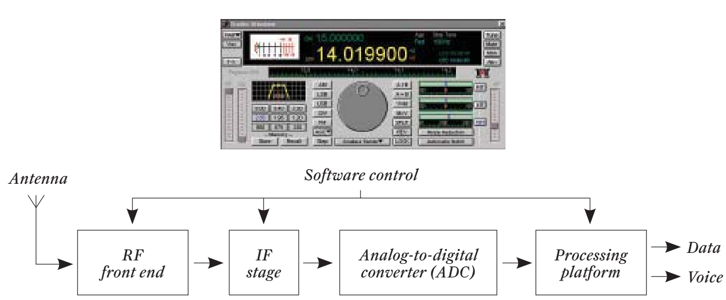

IDEAL SOFTWARE RADIO

All but initial amplification is done in software

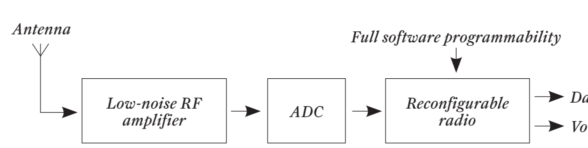

FROM TOP: ROLF BODMER/ISTOCKPHOTO; TAIT ELECTRONICS; TEN-TEC the input signal. Among the leading manufacturers of such designs are RF Micro Devices, in Greensboro, N.C.; Acco Semiconductor, in Saint-Germain-en-Laye, France; and Axiom MicroDevices, in Irvine, Calif.

Another challenge for SDR designers is making much faster ADCs. To avoid "aliasing"—the effect that makes rapidly spinning wagon wheels in old Westerns look as though they're turning slowly, or even backward—the ADC must sample the signal at a rate at least twice that of the highest-frequency component, and this may be quite high. The upcoming 4G technologies, for example, are expected to operate in the vicinity of 3.5 gigahertz, which means you'd need to take 7 billion samples per second—more than 10 times as fast as what today's best ADCs of sufficient resolution can manage.

Many SDR researchers consider this to be among the toughest obstacles ahead—not only because they must up the sampling rate so much but also because they'll simultaneously need to make significant improvements in the signal-to-noise ratio, power consumption, and physical size of this circuitry. Typically, you can better one of these parameters only by making trade-offs with the others. So achieving gains on all fronts at once is going to be extremely difficult.

There is, however, a strategy that might allow direct conversion of RF in the not-so-distant future: purposeful subsampling. The trick here is to arrange the sampling frequency of the ADC so that the inevitable aliasing that occurs works to your advantage. In one step, the operation both digitizes the RF signal and converts it to a lower frequency. This may seem a bit magical, but it's not so hard to understand. Just imagine the RF signal as one of those rapidly spinning wagon wheels. Adjust the frame rate of the motion-picture camera appropriately and your captured version of this wheel will turn at whatever lower frequency you want [see "Aliasing Harnessed"].

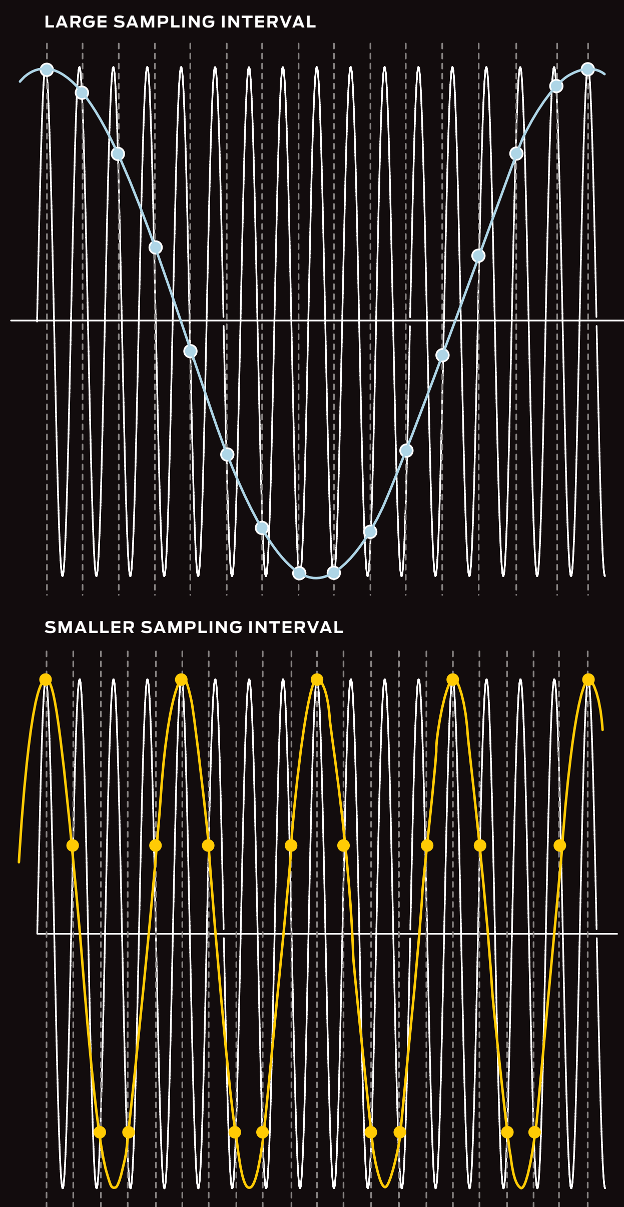

Aliasing Harnessed

Subsampling can work to your advantage

If the RF signal [white] is not sampled at least twice per cycle, aliasing will occur. But by properly adjusting the sampling interval [indicated by vertical lines], you can down-convert the RF to whatever lower frequency is desired [blue and yellow].

Astute readers might notice that the difficulties we've outlined so far all involve hardware. Software-defined handsets will have some challenging software, too. It'll manage the modulation, demodulation, encoding, decoding, encryption, and decryption, as well as the packing and unpacking of the data needed for the communications protocol employed—all computationally intensive tasks. What's the best kind of microprocessor for such heavy lifting?

Most SDR designers struggling with that question instinctively fixate on the MIPS rate—how many million instructions per second the processor can execute. That's because it must carry out a huge number of arithmetic operations—largely multiplications and additions—to massage the digitized signal. Specialized digital signal processors (DSPs) are usually the best chips for such things, but they may not be the only solution for SDR handsets. The reason is that these radios must do other kinds of signal processing, too.

In particular, SDR handsets need to detect and correct errors in the received digital bit stream, and the algorithms for that consist less of multiplications and additions than of "if-then-else" statements. Those branching operations are better done by a general-purpose processor, which would normally also be assigned the tasks of running the unit's real-time operating system, keyboard, and display. So a software-defined handset will need to have such a chip around anyway.

A general-purpose processor will also be required to host the software interfaces that connect different applications with the underlying hardware. In the near term at least, that "middleware" is likely to conform to the rules laid out in the U.S. military's Software Communications Architecture, an object-oriented computing framework that has become the de facto standard for software radios intended for combat use.

Designers of future SDR handsets will revel in the flexibility afforded by having the software control so much of the signal processing. And designers will no doubt want as much of the set's hardware as possible to be reconfigurable—that is, they'll want software to do not only signal processing but also be able to switch between different antennas and RF front ends. Such capability would allow you to turn a cellphone into a satellite-radio receiver, say, at the touch of a button.

While the technology for accomplishing this has been around for years, until now it's been too bulky and power hungry to be used in handsets. But consumers are now on the verge of enjoying the fruits of this approach, implemented with modest amounts of power and in very small packages. In February 2008, BitWave Semiconductor, of Lowell, Mass., announced its BW 1102 Softransceiver RFIC, a chip intended to bring SDR to both cellphones and femtocells (small wireless base stations that can be set up in a home or business to improve cellular coverage indoors). The BW 1102 is a single complementary metal-oxide-semiconductor integrated circuit containing a transceiver that supports a variety of wireless protocols and can operate anywhere on the spectrum from 700 MHz to 3.8 GHz.

Suppose, however, that you are a radio designer and want more than BitWave's chip can handle, such as the ability to receive FM broadcasts—and maybe even transmit on FM, too, so that you can play your favorite MP3 files on your car radio. How hard would it be to create the perfect IC for that? Hard indeed, it turns out, and that's why BitWave still has essentially no competitors.

But let's say you're keen to try. You might start by estimating the allowable execution time for each of the radio's intended functions and its power consumption, physical size, and other properties, including the frequency bands to be covered. Based on that assessment, you'd decide how to divide the overall system into hardware and software. Although this exercise isn't trivial, tools for hardware-software codesign are available.

Now comes the more difficult job: You've got to come up with detailed designs for each piece. Fortunately, you won't have to do that from scratch. Suitable designs for at least some of the larger building blocks—a DSP here, a general-purpose processor there—should be possible to find and license. After the hardware has been pinned down, you'll need to pull together the software to run it, which itself should keep you and your team busy for a large number of programmer-years.

The next challenge is to verify that your radio works correctly. Unfortunately, even state-of-the-art simulation tools aren't guaranteed to show system performance properly—and subtle errors here might be lethal for your product. Worse, many of the expected mobile services may be safety critical, so a tiny slip-up could be a literal killer, too.

One way to address this uncertainty is to go a step further than simulation: You can prototype the digital portion of your newly designed SDR system using one or more field-programmable gate arrays (FPGAs), integrated circuits that contain a vast number of logic blocks and potential interconnections. These devices can be configured after their manufacture to serve almost any purpose, constituting entire systems on a chip.

The problem with FPGAs for production is that they are the energy hogs of the semiconductor world, lacking the power-management features of their hardwired counterparts. Moreover, FPGAs suffer from the integrated-circuit equivalent of suburban sprawl, taking up a relatively large area on a silicon wafer. They are also expensive, which helps to explain why we haven't seen FPGAs being used to manufacture SDR handsets—at least not yet. A few researchers are exploring low-power FPGA technologies, so it's not out of the question that they could one day serve for high-volume production of handsets.

In the meantime, FPGAs remain a convenient way to build and test SDR prototypes. Among the most interesting examples of this is the Berkeley Emulation Engine 2 (BEE2) project at the University of California, Berkeley. This test-bed setup consists of five high-performance FPGAs, which with proper programming can be turned into various next-generation SDR systems. Another example of this approach is the SDR-based design effort at San Diego State University, which became widely known through a 2007 article in *DSP Magazine* titled "How to Pack a Room of Analog FM Modulators Into a Xilinx FPGA."

No doubt, many people are waiting for the day when they'll carry just one handheld gadget they can instantly switch from cellphone mode to that of a satellite radio receiver, or from a wireless Web browser to a mobile TV set; indeed, their handset might carry out all of these functions at once. Others, including the world's many technophobes, might be less enthusiastic about such a prospect. But SDR technology offers something for them, too—the possibility that their wireless equipment will eventually become smart enough to adapt to its communications environment all by itself.

A radio intelligent enough to reconfigure itself—perhaps by detecting free spectrum and switching its frequency of operation to claim it—would make wireless services cheaper and more reliable for their users, most of whom will not even be aware that such marvelous things are going on under the hood. As with SDR, this is a concept that Mitola promoted early on, in a 1999 article he wrote with Gerald Maguire Jr., of the Royal Institute of Technology, in Stockholm. They called it cognitive radio.

Ah, to have a radio that not only switches function on demand but also configures itself into the most effective form possible without its user even knowing it. Now *that* will be a truly universal handset. ❑

CARMAGEDDON

TOP 10 TECH CARS

ELECTRIC-DRIVE CARS TAKE SECOND PLACE TO SHEER SURVIVAL

BY JOHN VOELCKER

IF NECESSITY REALLY IS THE MOTHER OF INVENTION, THEN surely the auto industry is on the verge of an era of blinding brilliance. Times are that bad.

In July, when the price of a barrel of oil shot up to US $147, buyers in the United States dumped their gas guzzlers and lined up for that iconic hybrid, the Toyota Prius. And, astonishingly, by year's end they had bought 25 000 Smart ForTwo cars, two-seaters so small they make the Mini Cooper, with its lavish complement of four seats, look like a land yacht.

Then, just six months later, oil prices dropped back down, hitting a low of $37 a barrel. Good for drivers, bad for automakers: Such seesawing makes it impossible for them to plan—as they must—what cars they will be selling in three to five years. Then came the capper to that grisly year—the financial meltdown and the ensuing worldwide recession. It humbled what was left of America's Detroit Three (General Motors, Ford, and Chrysler, which barely counts as a bona fide automaker anymore). It even gave mighty Toyota its first operating loss in 70 years. Auto writers fretted about "carmageddon" and, more linguistically vexing, heralded "the carpocalypse." By the end of the year, the world's auto manufacturers had just a single imperative: survival.

Consolidation is in the air. Soon there will be fewer car companies, and they will be making fewer kinds of cars.

ALL PHOTOS PROVIDED BY MANUFACTURERS

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

2010 FORD FUSION HYBRID (P. 43)

MERCEDES-BENZ CONCEPT BLUEZERO (P. 41)

BMW MINI E (P. 40)

2010 TOYOTA PRIUS (P. 43)

For a glimpse of the future, look at Volkswagen. It builds more than a dozen separate models, totaling well over 1 million units a year, on the basic components of its Golf/Jetta/Rabbit line. They are sold as Volkswagens, Škodas, even Audis.

For another view of the future, look at upstart Chinese automaker BYD Automotive, in Shenzhen. Owned by a huge battery-cell manufacturer, it began making cars only five years ago, yet in December it stunned observers by launching the world's first production plug-in hybrid (see the 2009 BYD F3DM). Early reviewers called the car crude. But the Model T was crude, too. And by getting a plug-in hybrid to market before the end of the calendar year, BYD beat Toyota and GM, both of which are also working on plug-in hybrids, by almost two years.

If oil prices stay low, it may dampen consumers' interest in the next big thing: electric-drive cars. Big financial incentives from federal and state governments could speed things along, but those governments are now struggling themselves. Nevertheless, lots more partially electric and electric-drive vehicles are about to roll off assembly lines. Here's the complete list of major world automakers that *aren't* experimenting with batteries and electric motors: Mazda. Even the Germans have begun supplementing their beloved high-performance diesels with a handful of electric cars.

Europe's diesel diehards, meanwhile, despair that their hard-won improvements in fuel usage have been undermined by the surge in diesel fuel prices in the United States. And in more bad news for diesels, the Chinese government appears to have shifted its development priority from diesel to electric drive. Meanwhile, researchers all over the world are creating elaborate models to test whether running vehicles on grid power can really cut greenhouse-gas emissions overall. (The answer is usually yes, depending on how "clean" the local generating mix is.)

In the passenger compartment, technology continues to advance smartly. Telematics—the delivery of data by cellphone technology—is pumping traffic, weather, and gas-price information into car navigation systems in real time.

Digital entertainment also keeps getting better. This summer, the Mercedes-Benz S-Class sedan will offer an in-dash video screen that filters light onto adjacent pixels to display different images to different viewers. The driver can see routing maps, for instance, while the passenger happily watches a movie.

But most of these goodies seem to come from a time that suddenly seems very distant, the time before the meltdowns of the automotive industry and the world economy. Just as contractors will stop building McMansions and investment advisers will stop pushing mortgage-backed securities, so automakers may have to rethink personal transportation.

Usually the hottest new car technology tells us quite a lot about what to expect from the industry, but today that rule may not apply. Desperate times call for radical changes, and that means that even today's most outrageous concept cars may seem strangely outmoded and quaint in a few years' time—just as electric cars did in the 1920s. ❑

CONCEPT: MINI E

0 TO 100 KM/H IN 8.5 SECONDS ■ 240-KM RANGE

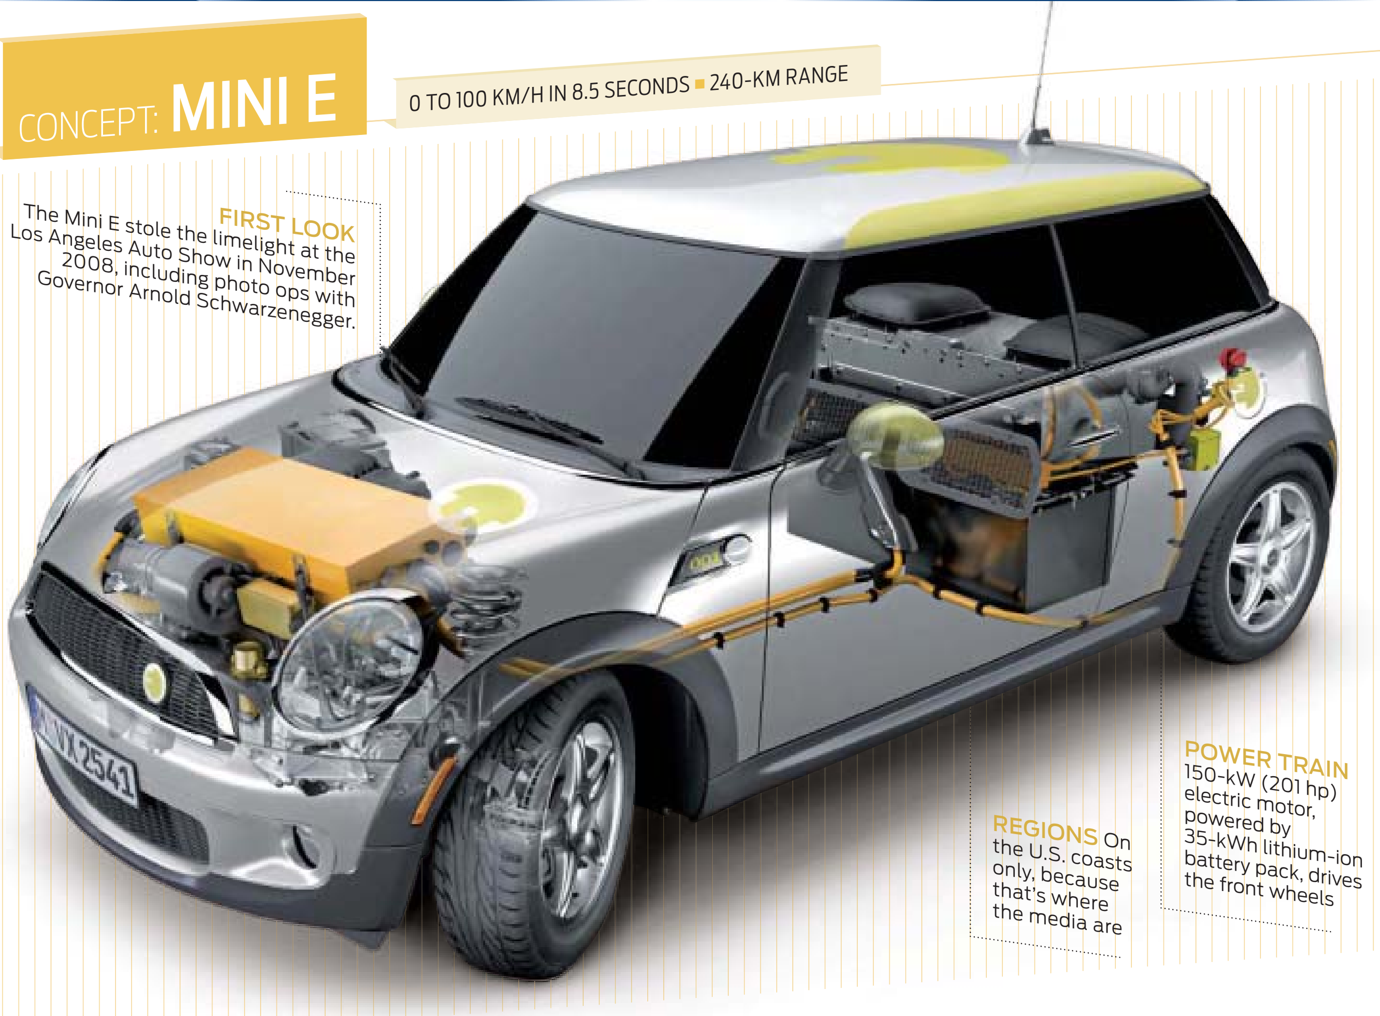

FIRST LOOK The Mini E stole the limelight at the Los Angeles Auto Show in November 2008, including photo ops with Governor Arnold Schwarzenegger.

REGIONS On the U.S. coasts only, because that's where the media are

POWER TRAIN 150-kW (201 hp) electric motor, powered by 35-kWh lithium-ion battery pack, drives the front wheels

BMW CATCHES UP ON ELECTRIC CARS—IN A HURRY

MODEL: BMW MINI E
AVAILABILITY: CALIF., N.Y., N.J.
RELEASE (BETA): APRIL 2009

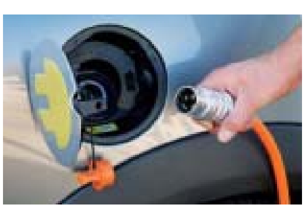

CHARGE IT: The driver plugs the car into a box that hangs on the garage wall and provides enough amperage to recharge the batteries in just 2.5 hours.

ALMOST ALONE AMONG major automakers, BMW had long lacked experience with electric drive. Then, last fall, the company unveiled its Mini E, an electric conversion of BMW's popular Mini Cooper hatchback. When the company announced a beta test of 500 units, thousands of drivers applied to lease one, at a whopping US $850 a month (close to twice the cost of leasing a BMW 5-series car).

Among other requirements, winners had to agree to keep the car in a garage, in which BMW would install a wall-mounted 240-volt charging unit. The battery pack goes where the rear seat would normally be, meaning that the Mini E is a two-seater. The 35-kilowatt-hour battery links together 5088 small off-the-shelf lithium-ion batteries (an approach similar to that used in the all-electric Tesla Roadster).

A 150-kilowatt (201 horsepower) electric motor drives the front wheels through a single-speed gearbox, replacing the standard Mini's engine and transmission. BMW quotes an acceleration time of 8.5 seconds from 0 to 100 kilometers per hour (62 miles per hour). That's respectable enough to support the Mini's sports-car image.

Longtime electric-drive pioneer AC Propulsion, of San Dimas, Calif., developed the battery pack, electric motor, and power electronics, as well as the wall charger. To make a production run of 500 cars possible, AC suspended all other work for several months and set up a global supply chain to ship electric motors from the United States and batteries and electronics from Asia to a BMW manufacturing facility in Munich.

Arriving simultaneously at the facility were Mini "gliders," completed cars minus their engines and transmissions. BMW workers assembled the components and tested the cars, the last of which was finished in November.

The car was instantly booked for test drives after its debut amid a crush of cameras at last fall's Los Angeles Auto Show—leaving dozens of disappointed journalists (including this one) out in the cold. Still, if BMW wants to get serious about electric cars, it will have to develop its own battery-management technology, negotiate with a myriad of new suppliers—and then ensure that it can produce and sell electric cars that maintain its traditional profit margins. That may prove to be harder than making a media splash with a limited-run car. ❑

RANGES: 200 KM (BATTERY), 600 KM (ELECTRIC + DIESEL), 400 KM (FUEL CELL)

CONCEPT: BLUEZERO

CLAIMED CO_2 EMISSIONS
32 g/km (E-Cell Plus in blended mode)

POWER PLANTS
70-kW (94 hp) continuous electric motor, powered by 35-kWh lithium-ion battery pack (E-Cell); 17.5-kWh lithium-ion battery pack, recharged by 50-kW (67 hp) 1-L three-cylinder turbodiesel engine (E-Cell Plus); or 80-kW (107 hp) continuous hydrogen fuel cell (F-Cell)

MORE
The Concept BlueZero is widely thought to preview the next B-Class, a compact five-seat hatchback Mercedes-Benz line not presently sold in the United States.

FUEL EFFICIENCY
4.5 L/100 km (E-Cell Plus in range-extending mode)

AQUADYNAMIC STYLING HIDES MULTIPLE POWER-TRAIN OPTIONS

MODEL: MERCEDES-BENZ CONCEPT BLUEZERO
AVAILABILITY: AUTO SHOWS GLOBALLY RELEASE: CONCEPT ONLY

MERCEDES-BENZ LIKES TO keep up with alternative propulsion technologies, and it's offering no less than three of them in its latest Concept BlueZero: battery electric (the E-Cell), extended-range electric (the E-Cell Plus), and fuel cell (the F-Cell).

The designers were able to fit this wide variety of power trains into a car just 422 centimeters long (barely 14 feet) thanks to a sandwich-like floor: Passengers sit above a flat chamber between the axles that can contain a tank for liquid fuel, a battery pack, or a fuel cell and its accompanying hydrogen cylinders. A 70-kilowatt (94 horsepower) electric motor drives the front wheels, regardless of the energy source powering the car.

In the E-Cell, the lithium-ion battery pack stores 35 kilowatt-hours, twice as much as the smaller packs in the other two variants. In the E-Cell Plus, the space saved by using a smaller battery helps to accommodate a range-extending engine, which sits under the rear seat.

Equally extreme is the styling, which draws heavily on the Bionic Car concept of 2005—known informally as the boxfish, or *kofferfisch* in German. With a rigid exoskeleton and very low aquadynamic resistance, the fish offered a model for designers to explore radical reductions in drag. The Bionic Car's coefficient of drag was a radically low 0.19. The full-size Concept F700 [see "Top 10 Tech Cars," *IEEE Spectrum*, November 2008] was the next Mercedes to explore biological models for vehicle shapes, and that trend is likely to continue, says Klaus Frenzel, manager of advanced design concepts at Daimler. The BlueZero's front and rear lights are, respectively, white and red C-shaped sweeps of light-emitting diodes. Its translucent roof panel is coated with thin-film solar cells, although they seem to have been put there largely for effect, because their power output is sufficient only to recharge a mobile phone.

BIO MOTIF: The sea-green cabin and smoothly contoured shapes emphasize the design's marine inspirations.

Daimler says the BlueZero's family of modular electric-drive components will be used in production vehicles this year and will be followed by low numbers of pure-battery vehicles next year. So if fish are the main inspiration, you might wonder: Will an electric eel be next? ❏

PRODUCTION: JETTA

US $21 990 ■ ONLY VW SELLS MASS-MARKET DIESELS IN U.S.

VW PROVES CLEAN DIESEL CAN BE AFFORDABLE

MODEL: 2009 VOLKSWAGEN JETTA TDI
AVAILABILITY: NORTH AMERICA, EUROPE
RELEASE: NOW

FOR ONCE, LAST year's "Green Car of the Year" wasn't a hybrid, a pure electric car, or a fuel-cell vehicle. It was simply a diesel, employing Rudolf Diesel's century-old idea of using high compression alone to ignite fuel mixed with air. And that diesel was fitted to a modestly priced compact sedan, no less. Could it be anything other than the latest incarnation of the beloved Volkswagen Jetta TDI?

Like most European carmakers, Volkswagen sells huge numbers of diesels in its home markets. Diesel fuel in much of Europe is less aggressively taxed than gasoline, in part because it delivers greater efficiency. But it has been a challenge to get those engines, with their relatively high emissions of particulate matter and nitrogen oxides, to comply with U.S. emissions laws. In California, the toughest state of them all, the current standard—known as Tier 2, Bin 5—temporarily halted the sale of diesel vehicles there after 2006.

Now a few diesels have returned, but only Volkswagen offers them in mass-market vehicles. The 2009 Jetta TDI, which went on sale in September 2008, carries the banner. At US $21 990, it costs about $4500 more than the gasoline model. VW's own fuel-efficiency tests indicate that fuel economy would, in practice, be 6.2 liters per 100 kilometers (38 miles per gallon) in the city and 5.3 L/100 km (44 mpg) on the highway, which are considerably better than the U.S. EPA's official estimates.

The Jetta has curbed emissions by replacing the previous mechanical fuel-injection system with a new device that uses piezoelectric crystals, which vibrate in response to electric signals, to better atomize the fuel and inject it at higher pressures. The exhaust passes through a storage catalyst, which temporarily holds the nitrogen oxides, while a filter holds soot and other particulates—until the engine control software can switch modes to burn the waste products off in a later combustion cycle.

Whether Diesel's invention becomes more than an oddity in the U.S. market depends largely on oil prices and the product mix coming out of refineries. More diesel fuel, at a cheaper price, will mean more diesel cars. Still, the Europeans are a lot less confident than they were just a year ago that diesels are the best answer to U.S. regulations. ❑

FIRST ELECTRIC CAR SOLD BY A MAJOR CARMAKER SINCE THE EV1

MODEL: MITSUBISHI i-MiEV
AVAILABILITY: JAPAN, UK
RELEASE: Q4 2009

Pure electric vehicles make sense in Japan more than anywhere else because so many commuters there go short distances in small cars and in stop-and-go traffic. So a lot of people will be following the progress of Mitsubishi Motors' i-MiEV electric minicar, which will go on sale to retail customers in Japan by the fourth quarter of 2009 at the latest.

Did we say small? The i-MiEV is adapted from a *kei* car, a class limited to 3.4 meters in length, 2 meters in height, 1.5 meters in width, 0.66 liter in engine size, and 47 kilowatts in power (63 horsepower).

The 330-volt, 16-kilowatt-hour lithium-ion battery pack was developed by Mitsubishi and GS Yuasa Corp., the only mass producer of large-format lithium-ion batteries in Japan. The 47-kW motor generates 180 newton meters of torque, enough to go from 0 to 100 kilometers per hour (62 miles per hour) in 13 seconds; the claimed top speed is 130 km/h (80 mph).

There are three driving modes: Standard, Eco, and "B." The Eco mode, in comparison with Standard, limits the motor's output to just 18 kW, stretching the range of a single charge. Range varies with usage and driving style, of course. The B mode adds more regenerative braking on downhill stretches and when the car is coasting, to recharge the pack more aggressively. ❑

FAR ENOUGH: Mitsubishi says the i-MiEV can go 160 kilometers on a standard Japanese road and 120 km on a standard U.S. road.

HIGHEST ELECTRIC SPEED: 75 KM/H (47 MPH) ▪ 1100-KM RANGE (CITY)

PRODUCTION: **FUSION**

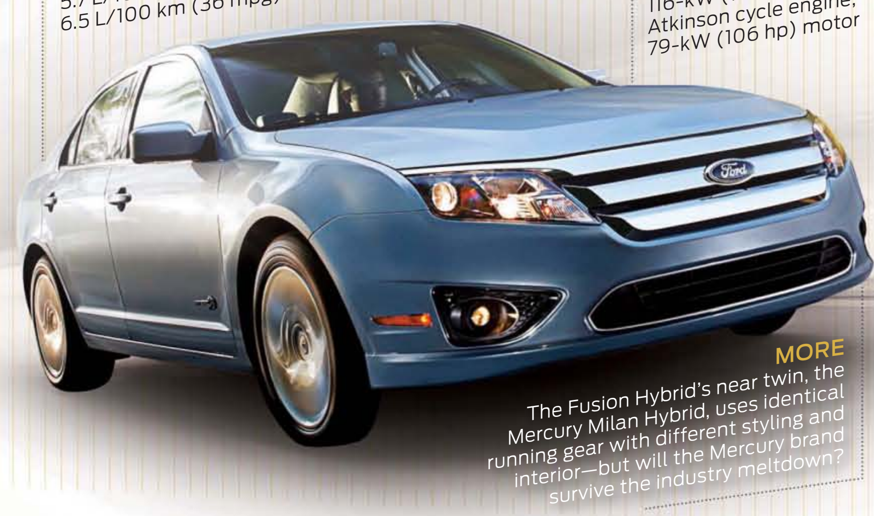

FORD'S LATEST TROUNCES TOYOTA ON MILEAGE

MODEL: 2010 FORD FUSION HYBRID
AVAILABILITY: NORTH AMERICA
RELEASE: NOW

AFTER THE WELL-RECEIVED Escape Hybrid launched in late 2004, Ford's hybrid program seemed to stall. The company regularly updated the small sport utility and added hybrid versions of its Mercury and Mazda platform mates. But it couldn't raise combined sales beyond a modest 25 000 or so per year—and no new hybrids appeared over the next four years.

Now Ford has unveiled a hybrid version of its Fusion midsize sedan. It is pitted directly against Toyota's popular Camry Hybrid, also a hybrid adaptation of an existing gasoline-engine car. And in testing by the U.S. Environmental Protection Agency, the Fusion Hybrid handily beat the 2009 Camry Hybrid by several miles per gallon.

The Fusion Hybrid is the first vehicle to use Ford's second-generation hybrid system. It uses a 1.3-kilowatt-hour nickel-metal-hydride battery with 20 percent more power than the Escape's in a package that's 30 percent smaller, meaning that the Fusion Hybrid sacrifices little of the trunk space available in the standard version.

Ford engineer Gil Portalatin said that improvements in control logic allowed the design team to be very aggressive in determining when the load is light enough to justify shutting off the gas and letting the electric motor propel the car all by itself. Such shutoffs happen about twice as often as in the early-model Escape.

The new system adds a variable-voltage controller that can boost the voltage to the battery, letting the traction motor and generator increase the battery recharge rate while operating at their most efficient speeds. In addition, new control logic for the regenerative brakes scavenges up to 94 percent of the braking energy and feeds it to the battery.

The result is an impressively high speed for operation under electric power alone: 75 kilometers per hour (47 miles per hour). It also allows up to 3 km of continuous electric-only driving and a range of 1100 km of around-town driving on a tank of gasoline.

The Fusion also offers Ford's popular Sync, a voice-activated digital entertainment and mobile phone system; real-time traffic and weather data delivered through the satellite radio; and a system that uses radar sensors in the rear quarter panels to detect other vehicles in the car's blind spots and alert drivers to crossing traffic when backing out of a parking space. ❑

QUINTESSENTIAL HYBRID GETS MAKEOVER

MODEL: 2010 TOYOTA PRIUS **RELEASE:** JUNE 2009 (NORTH AMERICA)

ECO SYMBOL: The new Prius is the third generation of the world's most popular hybrid and the shining green symbol of Toyota itself.

The Prius has become the definitive image of the hybrid-electric vehicle, with more than 1 million sold in 44 countries. Toyota now has a new version of the car, and it hopes to sell 180 000 units in the first full year in the United States.

The new Prius improves incrementally on the 2004–2009 version. It is more spacious, somewhat more luxurious, and has a larger engine (1.8 liters versus 1.5), yet consumes slightly less fuel: 4.7 liters per 100 kilometers (50 miles per gallon) on the combined U.S. urban and highway cycle, against the previous 5.1 L/100 km (46 mpg). This is the first engine Toyota has ever designed that has no mechanically powered accessory drives: water pumps, air-conditioning compressors, and the like. All of them run off the car's electrical system, to cut down on parasitic losses.

The new Prius offers three driving modes. The EV-Drive mode propels the Prius at low speeds on battery power alone. It can go up only to 1.6 km this way, but sometimes that's enough (indeed, this is jocularly known as the cheating-husband button, for those who may use it to slink home in the wee hours without waking their families). The Eco mode cuts performance for better fuel economy, while the Power mode gives quicker throttle response and more aggressive acceleration than standard.

Still, Prius buyers await the plug-in version, fitted with a larger battery pack containing lithium-ion cells. The company has promised to offer the car within 18 months. ❑

PRODUCTION: F3DM

US $21 700 ■ HOPES TO ENTER EUROPE BY 2010, U.S. BY 2011

WORLD'S FIRST PRODUCTION PLUG-IN HYBRID

MODEL: 2009 BYD F3DM
AVAILABILITY: CHINA
RELEASE: Q3 2009

WHILE GENERAL MOTORS constantly hypes its plans to sell the Chevrolet Volt extended-range electric car late next year, a Chinese battery company has quietly beaten GM to the punch. In December, BYD Automotive, of Shenzhen, started selling the F3DM, the world's first mass-produced plug-in hybrid-electric vehicle.

The compact sedan will be limited to China, however, until it can meet Western countries' standards for safety and comfort. The company is selling the F3DM (for "dual mode") to government fleets and will offer it to consumers this summer. It goes for roughly 150 000 yuan (US $21 700)—a little more than half the 280 000-yuan cost of a Toyota Prius but more than twice the price of the gasoline-powered F3, one of China's best-selling cars. BYD says it hopes ultimately to sell 2000 F3DMs a month.

The battery pack contains 100 of BYD's own 3.3-volt lithium-ion cells, which are based on iron phosphate. That chemistry has less energy than the cobalt alternative, but it's far less prone to internal short circuits of the sort that have famously caused laptops to burst into flames.

Unlike any other hybrid around, the F3DM has three distinct modes of operation: all-electric; series hybrid, which uses the engine to recharge the battery; and parallel hybrid, in which the gasoline engine drives the vehicle directly, with assistance from the electric motor.

Local journalists who tested the F3DM say that the transition between the three modes is rough, that the power steering vibrates, and that battery recharging is accompanied by a roar. On the other hand, journalists say the car accelerates very well under electric power.

A year ago, BYD Automotive showed the F6DM—a larger midsize plug-in hybrid sedan—at the Detroit auto show, quoting a 97-kilometer (60-mile) electric range. The company hopes to sell cars in Europe by 2010, says BYD board member Tony Mampilly, and in the United States by 2011. Well...within two to four years, says Micheal Austin, of BYD America.

But thus far, no Chinese carmaker has a firm date for U.S. sales. Industry experts dismiss or laugh at the idea of U.S. sales in the near term. Chinese makers will need years of experience elsewhere, says Aaron Bragman of the analyst firm IHS Global Insight, before they can meet stringent U.S. crash safety and emissions standards. ❑

WORLD'S FIRST HYBRID WITH LITHIUM-ION BATTERY

MODEL: 2009 MERCEDES-BENZ S400 BLUEHYBRID **AVAILABILITY:** EUROPE, NORTH AMERICA

The European love affair with diesel engines can be summed up succinctly: fuel that costs only about €1 per liter. But now, worried about pursuing the wrong track, European carmakers are introducing their first hybrids.

So here comes the new Mercedes-Benz S400 BlueHybrid (in Mercedes-speak, "blue" means "green"), on sale in Europe and planned for the United States as a 2010 model. A 15-kilowatt (20 horsepower) electric motor contributes oomph when needed, restarts the engine, and assists with initial torque to move the car away from a stop.

The battery is the first in any production hybrid to be based on lithium-ion technology. The cells, built in France by Johnson Controls–Saft, together can store 0.7 kilowatt-hours. Because lithium-ion cells pack about twice as much energy into a given volume as nickel-metal hydride, Mercedes was able to fit them into the same space that had previously housed the car's standard lead-acid 12-volt starter battery.

The 3.5-liter V6 engine has been converted to operate on the Atkinson (or five-stroke) cycle, the standard for hybrid systems. In an Atkinson engine, the piston moves through strokes of different lengths so that the combusting fuel-air mixture expands to a greater volume on the power stroke than it occupies on the intake stroke. The result is greater efficiency. The drawback—low torque at low engine speeds—doesn't matter in a hybrid, which can use the electric motor to augment low-end torque. ❑

AT LAST, LITHIUM: The S400 BlueHybrid, the first production hybrid fitted with a lithium-ion battery, cools the battery pack with the car's air-conditioning system.

THE FIRST CAR THAT AVOIDS REAR-ENDERS

MODEL: 2010 VOLVO XC60
AVAILABILITY: GLOBAL
RELEASE: NOW

Volvo, which has long emphasized safety, is going beyond anything it's done before with a system that in some cases substitutes its own judgment for that of the driver. The new XC60 can slow down and stop automatically, to avert or mitigate a low-speed collision with the car in front.

An infrared laser sensor behind the windshield feeds data to a processor that figures out how far ahead the next car is and how fast it's going. That way, the system can calculate the braking force necessary to stop before hitting that vehicle. If the system notices that the car is closing in on the car in front, it warns the driver by flashing a red light at the base of the windshield and sounding an alarm.

If the driver fails to react in time, the system applies the brakes automatically. The system brakes hard, and only at the last minute, to keep drivers from relying on it in traffic.

An adaptive cruise control, incorporating radar, lets the driver program the car to maintain a certain distance from the vehicle up ahead. Another system warns against an impending collision while putting a little tension on the disc-brake calipers so they'll take effect faster when the driver makes a panic stop.

Such features can be viewed as first steps toward vehicles that drive themselves. Perhaps one day, when a driver wants to fool with the navigation system, a patient spouse will smile sweetly and say, "Oh, honey, just let the car drive." ❑

WATCH THEIR BACK: Volvo's CitySafety braking system uses laser sensors to keep a safe distance from the car ahead.

US $20 500 ■ SPEEDOMETER DISPLAYS EFFICIENCY

PRODUCTION: INSIGHT

POWER PLANT
73-kW (98 hp) 1.3-L four-cylinder engine; 10-kW (13 hp) electric motor

MORE
The original Honda Insight, a teardrop-shaped two-seater, delivered real-world fuel economy near 3.4 L/100 km (70 mpg), but the company sold only 18 000 from 1999 to 2006.

CLAIMED FUEL EFFICIENCY
5.9 L/100 km (40 mpg) city, 5.5 L/100 km (43 mpg) highway

MINIMAL HYBRID, MAXIMAL PAYBACK

MODEL: 2010 HONDA INSIGHT
AVAILABILITY: GLOBAL **RELEASE:** NOW

FOR 45 YEARS, Honda has offered small, cleverly engineered cars, which for the past 15 years have achieved the highest average fuel economy in the United States. But while it debuted its first hybrid in 1999—the short-lived Insight—Honda never had a breakout like Toyota's Prius.

Now the company has launched its first four-seater designed, like the Prius, as a hybrid from the ground up. The vehicle, for which Honda resurrected the Insight name, incrementally improves Honda's "mild hybrid" approach, which incorporates a relatively small electric motor and uses it less frequently than do typical (or "full") hybrids. The Prius, for example, can go 1 to 2 kilometers on electricity alone, under some circumstances; the Insight can't.

Honda's formula combines a small internal combustion engine—which shuts off when the car comes to a stop—with the smallest battery-and-motor combination capable of moving the car from a resting state while the engine takes a second or two to restart itself. Of course, the motor assists the engine under heavy loads, too—half the point of a hybrid system is to exploit a motor's ability to provide full torque from the moment it switches on.

This "mild" strategy costs much less to implement than the full-hybrid setup, letting Honda call the Insight "the lowest-priced hybrid offered in the United States." It is priced at US $20 500, compared to a likely $25 000 or more for the 2010 Toyota Prius, for which prices had not been announced at press time.

Under the Insight's hood is a 1.3-liter four-cylinder engine producing 73 kilowatts (98 horsepower). The heart of the hybrid system is a lightweight, ultrathin electric motor between the engine and transmission. It puts out just 10 kW (13 hp) and is powered by a flat nickel-metal-hydride battery pack under the rear deck that holds 0.58 kilowatt-hour of energy (not quite half the 1.3 kWh of the Toyota Prius pack). Honda claims the Insight's system is 19 percent smaller and 28 percent lighter than the previous generation used in its Civic Hybrid. The Insight's motor can, by itself, occasionally power the car at urban speeds, though only if the conditions are just right.

A system Honda calls EcoAssist accumulates data on your driving style, so you can delve into your history to improve your mileage. Drivers can further enhance their fuel economy by selecting Econ mode, which sets the control logic so that the car accelerates more slowly and backs off the gasoline engine more quickly. ❑

Shrinking Possibilities

Continued from page 25

light into the pattern. If you can't, your image will be blurry.

Because of tricks like water-immersion lithography—which increases resolution by replacing the standard air gap between the lens and the wafer surface with a liquid—chipmakers have been able to print sharp images within these parameters. But that ability has come at a great cost. ASML, Nikon, and Canon have all pushed water-immersion lenses as far as they can. Efforts to switch from a wavelength of 193 nm to 157 nm have failed because of difficulties with optical materials.

Most industry experts agree that features of complementary metal-oxide-semiconductor (CMOS) silicon transistors will continue scaling to below 20 nm. But that won't happen without extreme ultraviolet lithography.

Pitch Counts

Year	Node	Half-pitch	Gate length*
2009 [a]	32	52	29
2007 [a]	45	68	38
2005 [b]	65	90	32
2004 [b]	90	90	37
2003 [b]	100	100	45
2001 [c]	130	150	65
1999 [c]	180	230	140
1997 [d]	250	250	200
1995 [d]	350	350	350
1992 [d]	500	500	500

* Here, gate width is defined as the physical gate length, which in recent years became smaller than the printed gate length.

[a] ITRS data 2008 update [b] ITRS data 2006 [c] ITRS data 2001
[d] ITRS data 1997

Note that each year skipped is identified on the ITRS as between nodes.

ROAD MAP: Before 1985, chips were classified fairly empirically. Node, half-pitch, and physical gate length are given for advanced logic such as microprocessor (MPU). Memory (DRAM) half-pitch is smaller by approximately one generation. In 2005, the International Technology Roadmap for Semiconductors (ITRS) retired the term "node," as it was no longer a meaningful descriptor.

Extreme ultraviolet lithography uses a wavelength of 13.5 nm, right near the point where the deep ultraviolet becomes X-rays. If we could harness light at that wavelength, we could continue shrinking features without many of the resolution enhancement tricks we have developed to push 193-nm lithography to the limit. ASML Holding, where I am chief scientist, recently introduced an EUV lithography system that can produce chips with features smaller than 30 nm. Nikon reportedly has a similar tool in development.

ASML expects to ship its first commercial EUV production lithography systems next year. We have already installed two EUV development tools—one at the Albany NanoTech Complex, in New York state, and the other at IMEC, in Leuven, Belgium. These machines produce patterns with a 28-nm half-pitch, better than the so-called 22-nm node.

These EUV developments are welcome, but they shouldn't be interpreted as proof that EUV has arrived. The systems are experimental, capable of turning out chips at a rate of a few wafers per hour, much slower than would be needed for a commercial system. EUV still faces significant technical challenges. Consider the contortions we have to go through inside the box to get 13.5-nm radiation. It can't be done with any traditional light source. Instead, we use a big carbon dioxide laser to vaporize liquid tin droplets [see illustration, "Seeing the Light"].

First we boil the tin, and then we drip the liquid tin in carefully timed droplets, synchronized to the firing of the CO_2 laser so that it hits each tin drop as it falls. When the laser hits the droplet, the tin is vaporized, and 13.5-nm photons are released. A spherical reflector mirror takes this radiation and channels it into the optical system. Quite a bit of the original radiation is lost in the process because you can capture only what is collected in the reflector, and the reflector does not surround the tin droplets completely. The upshot is that only a small percentage of the radiation created at the 13.5-nm wavelength reaches the photoresist to make the pattern.

And that's just the first challenge. At 13.5 nm, your optics can't be made of glass, because glass—and air, and just about everything else—absorbs 13.5-nm radiation. You need a good vacuum to prevent EUV light from being absorbed by stray molecules. You need to use mirrors rather than lenses, and that brings up the next issue: About 30 percent of the light that hits that mirror is absorbed. Finally, not all the radiation emitted by your CO_2 laser and tin setup is 13.5-nm radiation—some of it is infrared and is lost as heat. EUV's critics say that too much power is necessary to yield the number of photons required to expose wafers at the chip fab.

The mirrors in our EUV lithographic system are based on Bragg reflection, a concept used for optical fibers and other waveguides. To create a strong Bragg reflector, you start with a rigid substrate, then coat it with several dozen alternating layers of molybdenum and silicon. The layers must be spaced with a uniform thickness of about half the wavelength to reinforce the reflected wave as it bounces off the many dozen layers. With this staggering technique, it's possible to make a surface that's about 70 percent reflective (the other 30 percent is absorbed by the mirror surface).

The EUV development tools use a different light source. They still use the tin droplets, but instead of a CO_2 laser, they vaporize the tin drops with an electrical discharge—a small lightning bolt, basically. These have proved rather inefficient.

One of the many technical challenges of making EUV work is figuring out how good the vacuum has to be to prevent EUV light from being absorbed by stray molecules.

Who are the potential customers for EUV? Anyone who wants to stay on the road map implied by Moore's Law: big memory makers that have to shrink relentlessly, the largest microprocessor companies, and chip foundries.

But the majority of semiconductor companies don't and won't need to be at the cutting edge. Consider the chip-foundry business, which accounts for a large and growing share of the overall chip business worldwide.

Taiwan is home to several chip foundries, the two biggest being Taiwan Semiconductor Manufacturing Co. and United Microelectronics Corp., both headquartered in Hsinchu. A chip foundry fabricates anyone's chips, on a contract basis. Unlike integrated device makers like Intel, Samsung, and Toshiba, which design and build all their own chips, some semiconductor companies do none of their own fabrication. Others have evolved from a past where they made their own chips to becoming "fab-lite," where they retain some facilities to develop the initial technology but then send their volume business to foundries. This allows these companies to do rapid development under their own control and avoid the cost of massive fabs for volume. More and more companies in the United States and Europe are taking these approaches because they believe that building a new fab (which can cost about $5 billion) is an impossible expense in the current economic climate. In the United States alone, AMD, Freescale, and Texas Instruments all recently went fab-lite.

Soon only a few foundries, many of them in the Far East, will remain to run all these companies' production lines. Others include Chartered Semiconductor in Singapore and Semiconductor Manufacturing International Corp. in Shanghai, as well as IBM Microelectronics in East Fishkill, N.Y., which makes processors for its own equipment but has also become a dominant foundry for game processors. AMD recently spun off its manufacturing operations to start Globalfoundries, a joint venture with the government of Abu Dhabi.

About half the foundry customers need chips that can be fabricated with the previous generation of process technology—currently 65 nm. Another 40 percent of foundry customers use even larger, earlier nodes. At most, only 10 percent have orders that require the most up-to-date chipmaking technology, to keep up with the likes of Intel and Samsung. These customers include companies that create graphics processors, FPGAs, and phone chips. Many foundry customers are focusing on design innovation, which may or may not require being at the forefront of chip-fab technology.

The standard sedative to EUV anxiety is the assurance that the gap will be bridged by double-patterning lithography, a technique of last resort that improves the resolution possible with 193-nm light by making two or more exposures, slightly shifted with respect to each other, and with two or more different masks [see "Seeing Double," *IEEE Spectrum*, November 2008].

But double patterning is more cumbersome than many people realize. The trouble comes down to design restrictions, cost, and yield. To understand why, let's go back to the foundry customer. With double patterning, these companies face a horde of new design restrictions. For example, with double patterning at 193-nm wavelengths, the only things that are easy to print are parallel or perpendicular lines and spaces. It's quite difficult to print holes or elbows, and if you've ever looked at an actual (but unexceptional) circuit design, it's absolutely lousy with zigs and zags.

So, first, companies painstakingly rework their designs to meet these onerous restrictions (no elbows! no zigs! no zags!). Then they've got to split the design into two or more parts. There is just no easy way to do that. Different electronic design automation companies are working on that problem, but none of them has been able to create the vaunted "black box" method, an ideal tool that will send the design on its way, no brain required.

Double patterning also demands more process steps, which again adds to the cost by cutting effective throughput, increasing fab cycle time, and adding additional defects. If your original process could yield 100 wafers an hour, using double

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

patterning will leave you with much less. If you'd been turning a profit at 100 per hour, there's no guarantee you will at 50.

Then there's depreciation. Most fab owners keep their tools running 24 hours a day because they depreciate at the staggering rate of several thousand dollars per hour. That's one of the major problems with double patterning. Cutting your throughput means you're losing thousands of dollars an hour. Pretty soon you're talking about real money.

Memory-chip manufacturing doesn't require as many masks as logic does. Where memories have only a few wiring layers, state-of-the-art logic has a complicated stack of as many as 10 metal layers. Consequently, a set of masks needed to make a logic chip is going to cost more than a set for a memory chip. Both processes require the use of expensive masks. And masks need an inspection tool to insure against infinitesimal flaws. Yet another tool is needed to repair these flaws, another expense that neither logic nor memory manufacturers can avoid.

Memory makers don't care as much about the cost of the masks because they don't use as many. For them, a single mask set will make up to 100 million chips. A $1 million mask set, even doubled, works out to about a penny per chip. But for logic chipmakers, the economics are completely different, because the vast majority of mask designs may be used to make only a few hundred wafers. That's because of the punishingly rapid evolution required to keep developing the latest features for cellphones or digital cameras.

If we meet our targets for EUV throughput, EUV is preferable to double patterning because it lets layers be exposed with single masks. Double patterning begins to fail at about 20 nm. Lens designers at Carl Zeiss (ASML's supplier) believe we can build EUV optics capable of reaching at least an 11-nm half-pitch.

The principal logic makers are spending hundreds of millions of dollars on R&D every year to keep scaling their transistors, in speed as well as in physical size. Memory chipmakers, on the other hand, are now in glorious pursuit of the "Grand Unified Memory"—one technology that will do it all, easily taking over for the NAND flash memories you find in your MP3 player and USB drives, and the DRAM that dominates applications for high-speed (but power-hungry) computation. The contenders for this crown include phase-change memories, resistive RAM, and spin torque transfer magnetic RAM (STT-MRAM), but these find themselves approximately in the same place as EUV—each concept has been proved and even demonstrated, but it's not quite ready for prime time.

Is it worth it? That's the billion-dollar question. Scaling won't continue forever, if only because we will eventually be down in the atomic realm.

Right now we're at 34-nm features. Let's assume no features can be smaller than the spacing between one atom and the next, which in silicon is 0.546 nm. Within a few years, progress won't depend so much on making transistor parts smaller. Instead, it will increasingly depend on new transistor designs and on materials that will make the transistors and the chips they reside on unrecognizable. In the next few generations of shrink, the industry needs EUV lithography to continue Moore's Law economics.

The enormous capital investments required by optical lithography and silicon manufacturing mean that no new technologies will easily displace these workhorses. To take over, any contender technology must build on, and incorporate, the incumbents. ❑

ieee spectrum online | Tech Insiders Webinar Series

The Robotics Continuum Webinar Series

ROBOTICS 101 - 16 APRIL 2009

Explore the ubiquitous world of robotics, focusing on unmanned vehicle applications in academia and industry. Learn about how the next generation of embedded technologies is helping roboticists design, prototype and deploy sophisticated robotic architectures.

Moderator: Bob Malone, IEEE

Presenter: Anu Saha, Academic Product Manager, National Instruments

UPCOMING WEBINARS:

- **Sense, Think, Act** – 18 June 2009
- **Robots in Academia** –17 Sept. 2009
- **The Future of Robotics** –12 Nov. 2009

Be a part of the Robotics Continuum!

Share your best practice or idea as it relates to Robotics as a video post on YouTube. Just email the link to your video to RoboticsWebcast@ni.com by 23 April, 2009.

Your video could be selected to be featured alongside the next Robotics Continuum webinar on IEEE Spectrum Online! To learn learn more, please visit ni.com/robotics.

Sponsored by

NATIONAL INSTRUMENTS™

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

IEEE
Celebrating 125 Years
of Engineering the Future

Internships

with the

The Microsoft® Medical Media Laboratory *M3L*

Washington, District of Columbia

Microsoft's Health Solutions Group, in conjunction with Microsoft® Research, is pleased to offer three to six month internships for individuals holding or pursuing graduate level degrees.

The M3L explores the use of information technology in healthcare, and in particular, the opportunities afforded by surface computing, multi-media displays, mobile computing, unified communications, human-computer interaction and decision support.

Interns will have the opportunity to design and build novel solutions in one of these areas and examine the impact of this technology in a real-world healthcare environment.

Stipend and expenses paid.

Contact: Hank Rappaport, MD
Email: M3L_int@microsoft.com

More information, including an on-line application, can be found at **http://research.microsoft.com/en-us/jobs/intern/medicalmedia.aspx**

Microsoft®

UNIVERSITY of DAYTON

Endowed Chair
in Sensor Exploitation and Fusion
Department of Electrical and Computer Engineering

The Department of Electrical and Computer Engineering at the University of Dayton invites nominations and applications for the Ohio Research Scholar's Endowed Chair in Sensor Exploitation and Fusion. This Ohio Research Scholar position is one of several newly created endowed chairs that are part of a multi-university, statewide initiative formed under the Ohio Research Scholars Program (ORSP).

See ***http://www.ohiochannel.org/your_state/third_frontier_project/rfp.cfm?rfp_id=94444*** for more information about ORSP. The goal of the ORSP is to strengthen and increase the number of University-led research clusters within the state of Ohio. The ORSP will achieve such a missive through aggressive recruitment and attraction of talent to Ohio.

The successful candidate will have an established record of innovative research in wide area surveillance techniques, which support the communication and processing of data from heterogeneous sensor suites to provide real-time surveillance capabilities. The Chair will possess a doctorate in electrical or computer engineering and an established record of disciplinary research and extramural funding commensurate with appointment at the rank of Full Professor. The Chair will lead collaborative research, mentor a diverse population of students, and expand our academic program in areas **relating to the development of algorithms and structures that allow intelligent interpretation of sensed data from multiple sensors. This will include remote and in-situ sensors, varying resolutions, compression schemes, representation methods and timing.** The Chair will forge partnerships with industry, research centers and government labs and advance the regional, national and international stature of the University in this field. The successful candidate will have demonstrated ability to work with researchers and colleagues from diverse cultural backgrounds.

Because of the need for collaborations with AFRL, U.S. citizenship or permanent resident status is highly desirable. Send a cover letter, curriculum vitae, and the names and contact information of four references to: Dr. Don Moon, Endowed Chair Search Committee, Electrical and Computer Engineering Department, University of Dayton, Dayton, OH 45469-0232. Applications may also be sent electronically to ***Don.Moon@notes.udayton.edu***. Review of applications will begin on April 1, 2009, and continue until the position is filled. UD offers a competitive salary, attractive fringe benefits and a generous start-up package.

The University of Dayton, a comprehensive Catholic University founded by the Society of Mary in 1850, is an Equal Opportunity/Affirmative Action employer. Women, minorities, individuals with disabilities and veterans are strongly encouraged to apply. The University of Dayton is firmly committed to the principle of diversity.

Sign me up!

What's New @ IEEE

10 FREE NEWSLETTERS TO CHOOSE FROM

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

IEEE

THE HONG KONG POLYTECHNIC UNIVERSITY
香港理工大學

DEPARTMENT OF COMPUTING

The Department invites applications for Professors/Associate Professors/Assistant Professors in Database and Information Systems/Biometrics, Computer Graphics and Multimedia/Software Engineering and Systems/Networking, Parallel and Distributed Systems. Applicants should have a PhD degree in Computing or closely related fields, a strong commitment to excellence in teaching and research as well as a good research publication record. Applicants with extensive experience and a high level of achievement may be considered for the post of Professor/Associate Professor.

Please visit the website at ***http://www.comp.polyu.ed.hk*** for more information about the Department. Salary offered will be commensurate with qualifications and experience. Initial appointments will be made on a fixed-term gratuity-bearing contract. Re-engagement thereafter is subject to mutual agreement. Remuneration package will be highly competitive. Applicants should state their current and expected salary in the application. Please submit your application via email to ***hrstaff@polyu.edu.hk***. Application forms can be downloaded from ***http://www.polyu.edu.hk/hro/job.htm***. **Recruitment will continue until the positions are filled.**

Details of the University's Personal Information Collection Statement for recruitment can be found at ***http://www.polyu.edu.hk/hro/jobpics.htm***.

UNIVERSITY of DAYTON

Endowed Chair in Wide Area Surveillance

Department of Electrical and Computer Engineering

The Department of Electrical and Computer Engineering at the University of Dayton invites nominations and applications for the Ohio Research Scholar's Endowed Chair in Wide Area Surveillance. This Ohio Research Scholar position is one of several newly created endowed chairs that are part of a multi-university, statewide initiative formed under the Ohio Research Scholars Program (ORSP).

See ***http://www.ohiochannel.org/your_state/third_frontier_project/rfp.cfm?rfp_id=94444*** for more information about ORSP. The goal of the ORSP is to strengthen and increase the number of University-led research clusters within the state of Ohio. The ORSP will achieve such a missive through aggressive recruitment and attraction of talent to Ohio.

The successful candidate will have an established record of innovative research in wide area surveillance techniques, which support the communication and processing of data from heterogeneous sensor suites to provide real-time surveillance capabilities. The Chair will possess a doctorate in electrical or computer engineering and an established record of disciplinary research and extramural funding commensurate with appointment at the rank of Full Professor. The Chair will lead collaborative research, mentor a diverse population of students, and expand our academic program in areas relating to **high performance computer architecture and technology including high speed input/output and processing hardware, firmware and software**. The Chair will forge partnerships with industry, research centers and government labs and advance the regional, national and international stature of the University in this field. The successful candidate will have demonstrated ability to work with researchers and colleagues from diverse cultural backgrounds.

Because of the need for collaboration with AFRL, U.S. citizenship or permanent resident status is highly desirable. Send a cover letter, curriculum vitae, and the names and contact information of four references to: Dr. Don Moon, Endowed Chair Search Committee, Electrical and Computer Engineering Department, University of Dayton, Dayton, OH 45469-0232. Applications may also be sent electronically to ***Don.Moon@notes.udayton.edu***. Review of applications will begin on April 1, 2009, and continue until the position is filled. UD offers a competitive salary, attractive fringe benefits and a generous start-up package.

The University of Dayton, a comprehensive Catholic University founded by the Society of Mary in 1850, is an Equal Opportunity/Affirmative Action employer. Women, minorities, individuals with disabilities and veterans are strongly encouraged to apply. The University of Dayton is firmly committed to the principle of diversity.

DP DAKOTA POWER

A Unique Opportunity to Innovate the Next Generation Electric Drive Systems

Join the Dakota Power R&D team to design and build the lightweight electric drive systems to power the next generation of electric vehicles. Dakota Power is developing unique high efficiency technologies to prepare globally competitive electric drive systems for both military and civilian applications.

Dakota Power seeks innovative electrical engineers and computer engineers with experience in optimizing magnetic fields in motors and electric power systems. Additionally, computer science graduates are needed with experience in finite element modeling and system simulation.

Dakota Power, LLC is located in Rapid City in the beautiful Black Hills of South Dakota. A generous compensation and benefit package is provided to the members of the Dakota Power R&D team.

For more information visit
www.dplwed.com

DAKOTA POWER, LLC IS AN EQUAL OPPORTUNITY EMPLOYER

the data

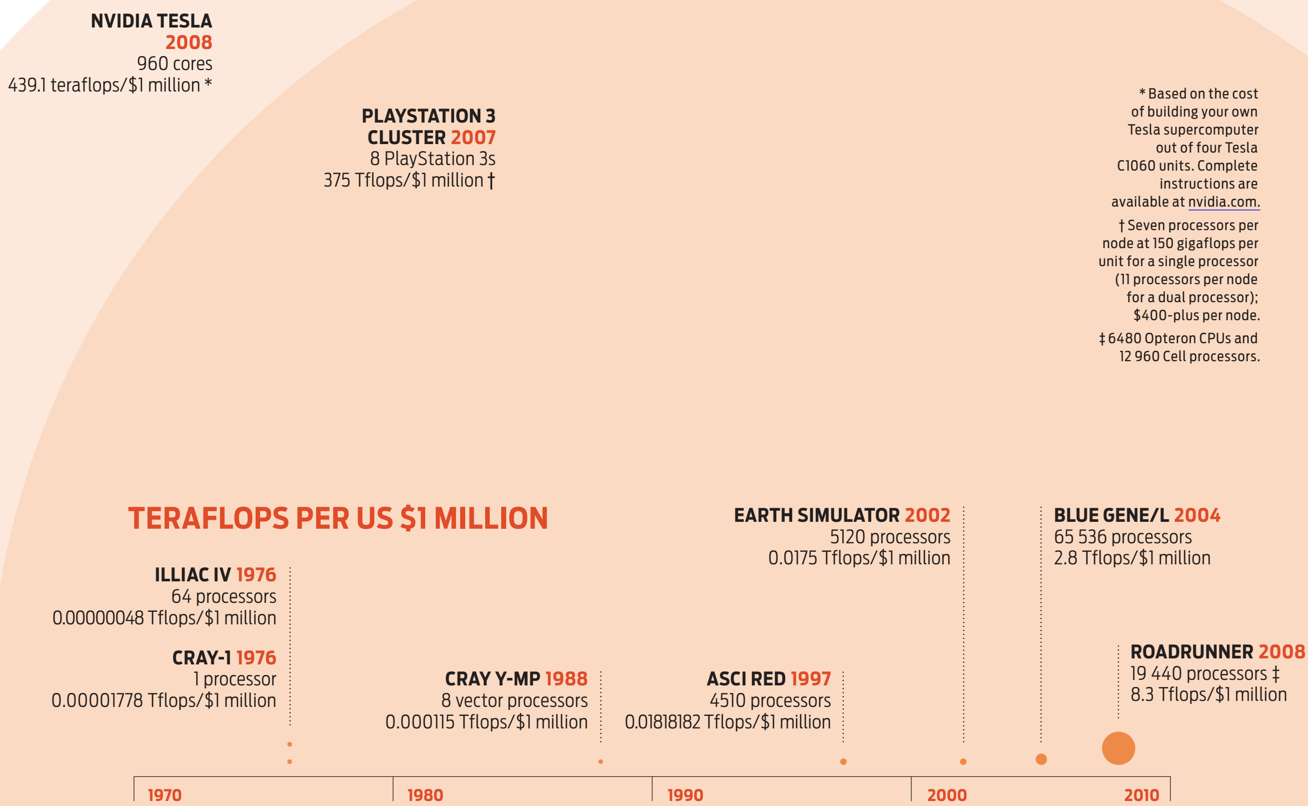

* Based on the cost of building your own Tesla supercomputer out of four Tesla C1060 units. Complete instructions are available at nvidia.com.

† Seven processors per node at 150 gigaflops per unit for a single processor (11 processors per node for a dual processor); $400-plus per node.

‡ 6480 Opteron CPUs and 12 960 Cell processors.

The Supercomputer Goes Personal

VIDEO GAMERS and graphic artists have funded a revolution in high-performance computing. The same massively parallel chips that figure out exactly the right color for each pixel in *Street Fighter IV* are now being used to calculate protein folding, predict climate change, and crack the encryption of hitherto-secure Web sites.

In October 2007, astrophysicist Gaurav Khanna, of the University of Massachusetts Dartmouth, lashed together eight Sony PlayStation 3s to continue the simulations of black holes he used to perform using rented time on traditional supercomputers funded by the National Science Foundation. He got about 1.2 teraflops; the hardware cost in PS3s at, say, Best Buy would be about US $3200. Twenty years earlier you would have needed more than 500 Cray X-MPs, at a cost of at least $78 billion, to break the teraflops boundary.

Today, graphics titan Nvidia advertises its new workstation, the Tesla, as a "personal supercomputer." It clusters four Nvidia C1060 processing boards, each of which unites 240 graphics cores to process instructions at nearly teraflops speeds. We calculate it as about 17 percent more cost-effective than Khanna's PS3 solution, and a lot more elegant. Of course, neither is perfect. For one thing, the Tesla and the PS3 do single-precision floating-point calculations using four 8-bit bytes. The Roadrunner, by contrast, uses 64-bit floating integers. But is the greater accuracy worth an additional $117 722 400? —*Paul Wallich*