IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 03.16

THE SMART TENNIS RACKET

Chips in clubs, bats, rackets, and balls are helping elite athletes sharpen their games P. 44

CLICK HERE TO ACCESS THE MARCH ISSUE OF **the institute**

IEEE

IEEE SPECTRUM

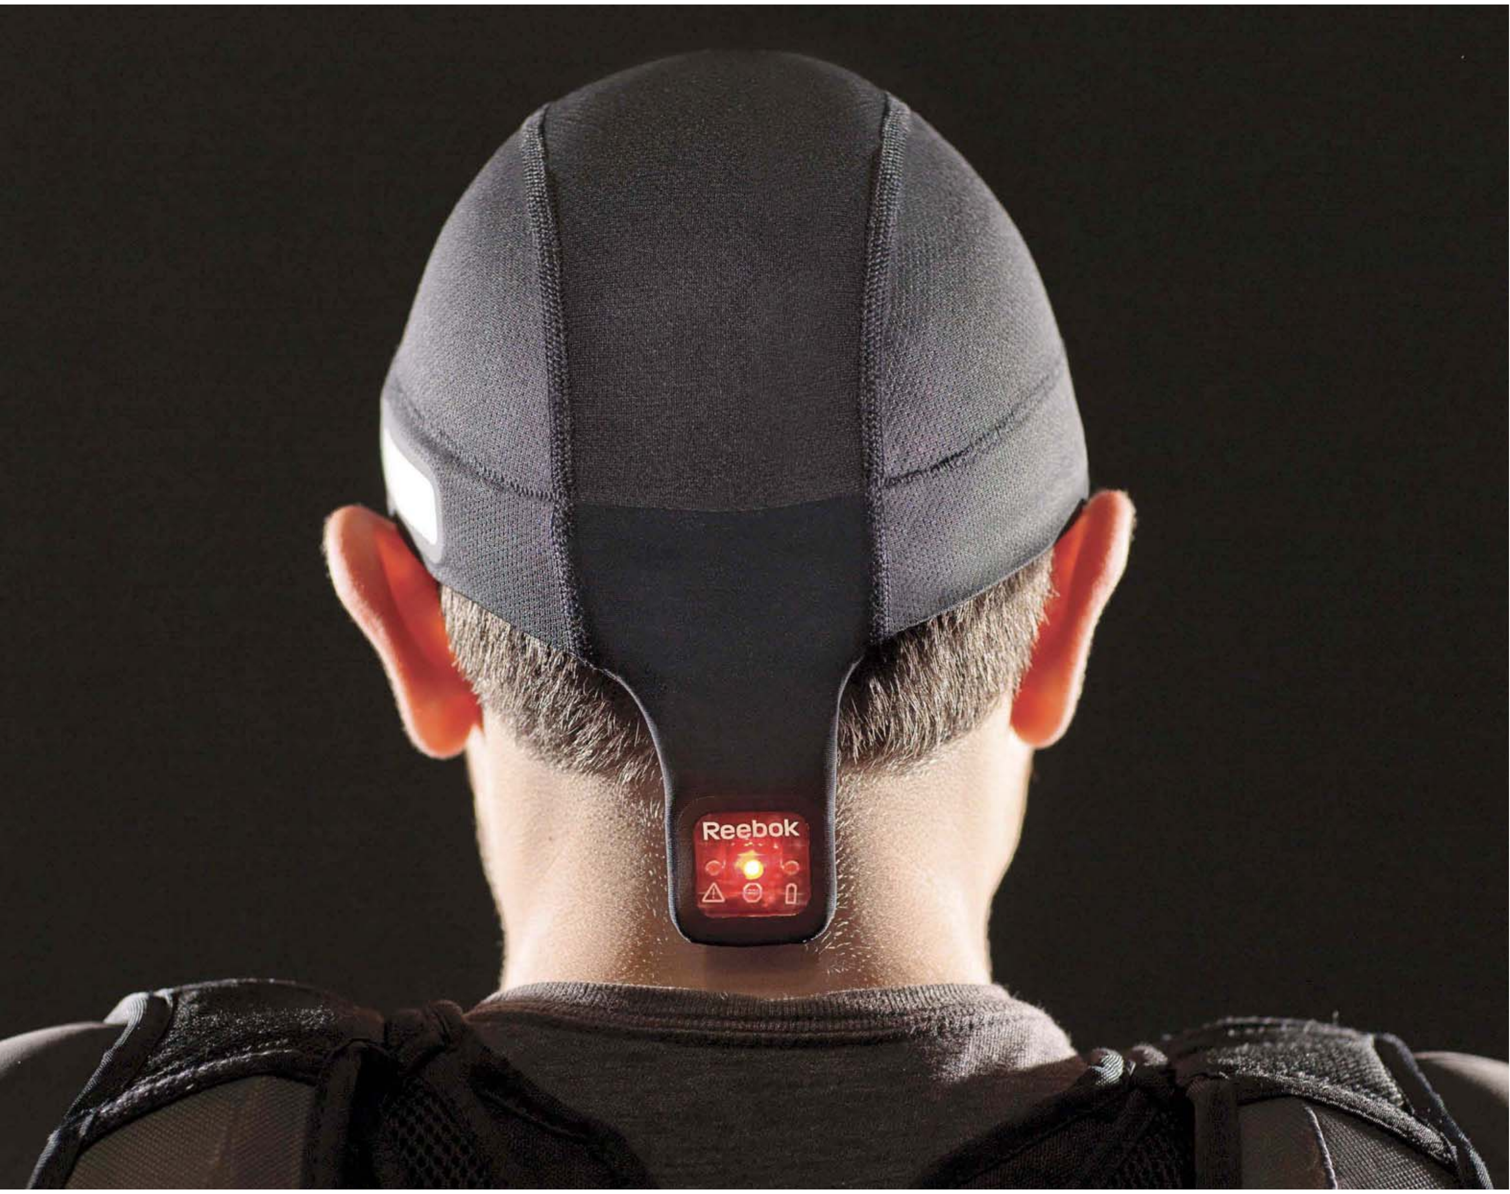

REEBOK

On the Cover **Photograph of Sloane Stephens by Dylan Coulter; graphics by John LePore**

There's a Better Way to Design 5G Wireless Systems

It starts with an integrated design approach for rapid prototyping.

In the race to 5G, researchers need to accelerate design cycles. With the LabVIEW Communications System Design Suite and NI software defined radio hardware, they can build 5G prototypes fast to decrease the time between an idea and a real-world application.

>> See how at ni.com/5g

800 891 8841

NATIONAL INSTRUMENTS™

DEPARTMENTS_03.16

IEEE SPECTRUM

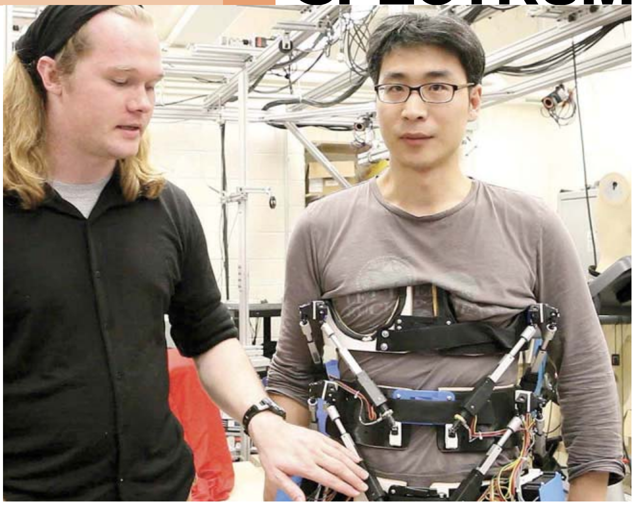

Online

Spectrum.ieee.org

Video: Robotic Spine Brace
A dynamic exoskeleton gathers data from the human spine that could provide clues for a more effective treatment of scoliosis. Watch our story on the technology here: http://spectrum.ieee.org/spinebrace0316

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Nonlinear Magnetic Materials With COMSOL Multiphysics—Available On Demand
- Grid Resiliency and Stability With Siemens Transformers—Available On Demand
- Signal Integrity: VNA Applications—Available On Demand

- **NEW PRODUCT RELEASE LIBRARY**
  http://spectrum.ieee.org/static/new-product-release-library

- **WHITE PAPERS**
  http://spectrum.ieee.org/whitepapers

The Institute

Available at theinstitute.ieee.org

- **GREEN ICT** The greenhouse-gas emissions produced globally by information and communication technology (ICT) account for more than 2 percent of the worldwide total—and that percentage is expected to double by 2020. This month we feature the efforts by IEEE and its members to make systems more energy efficient.

- **MEDAL OF HONOR** Life Fellow G. David Forney will receive IEEE's highest award this year, for pioneering contributions to error-correcting codes and reliable high-speed communications.

- **STARTUP STAR** Entrepreneur and IEEE Member Eileen Healy founded a pair of multimillion-dollar companies to address the growing needs of the mobile phone market. Read about her journey in the latest installment of our series on startups.

FROM LEFT: ELIZA STRICKLAND; RANDI KLETT; KRISTEN CLARK

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2016 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 53, issue No. 3, International edition. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

BACK STORY_

Hitting the Slopes—Literally

KAREN LIGHTMAN, executive director of the MEMS & Sensors Industry Group, knew of course that microelectromechanical systems were being used in all sorts of industrial and consumer products. But a ski accident in 2013 gave her a whole new perspective on them.

Lightman and her then 8-year-old daughter were taking a break at the side of a trail at a small ski resort in western Pennsylvania. They were waiting for the rest of their family to catch up when a skier, shooting down the mountain out of control, collided with another adult above them, causing both to tumble down the mountain. Lightman dove to try to shield her daughter from the impact, but she wasn't fast enough; they both hit their heads hard on the icy slope. Although they were wearing helmets, the impact was jarring. The MRIs revealed no damage, but Lightman quickly began showing symptoms of a concussion, which triggered a migraine headache and yet had no lasting effects. Her daughter seemed fine, she recalls, so she sent her to school the next day.

But shortly after she arrived at school, Lightman's daughter passed out on her desk. She was then diagnosed with a severe concussion. She ended up needing round-the-clock care for three months, and it took six months for her to fully recover.

"I began thinking about how sensors could have prevented this, or at least gotten her treatment sooner," Lightman said. "We should have sensors on all helmets; if we had data indicating that this had been a concussion-worthy hit, the emergency room doctors likely would have treated this differently. I also thought that we should have sensors on ski passes. If someone is skiing like a jerk, too fast or inebriated, they could be reprimanded or taken off the slopes before they hurt someone."

Now, she's encouraged that the technology is becoming more mainstream, Lightman says. It is moving into mouth guards, making it useful in sports like girls' lacrosse, in which helmets aren't typically worn. Indeed, Lightman spotted MEMS and sensor technology in all sorts of smart sports gear exhibited this past January at CES, in Las Vegas. In this issue's "Silicon Gets Sporty," she explains the trends behind the boom.

But the ski-jerk detector? We're still waiting for that one. ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. 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, Past Forward is in *IEEE Spectrum*, Vol. 53, no. 3 (INT), March 2016, p. 60, or in *IEEE Spectrum*, Vol. 53, no. 3 (NA), March 2016, p. 76.

03.16

IEEE SPECTRUM

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

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

EDITORIAL DIRECTOR, DIGITAL
Harry Goldstein, h.goldstein@ieee.org

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

SENIOR ART DIRECTOR
Mark Montgomery, m.montgomery@ieee.org

SENIOR EDITORS
Stephen Cass (Resources), cass.s@ieee.org
Erico Guizzo (Digital), e.guizzo@ieee.org
Jean Kumagai, j.kumagai@ieee.org
Samuel K. Moore (News), s.k.moore@ieee.org
Tekla S. Perry, t.perry@ieee.org
Philip E. Ross, p.ross@ieee.org
David Schneider, d.a.schneider@ieee.org

DEPUTY ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSOCIATE ART DIRECTOR Erik Vrielink, e.vrielink@ieee.org

SENIOR ASSOCIATE EDITORS
Rachel Courtland, r.courtland@ieee.org
Eliza Strickland, e.strickland@ieee.org
ASSOCIATE EDITOR
Celia Gorman (Multimedia), celia.gorman@ieee.org
VIDEO PRODUCER Kristen Clark, kristen.clark@ieee.org
ASSISTANT EDITOR Willie D. Jones, w.jones@ieee.org
SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

ADMINISTRATIVE ASSISTANT
Ramona L. Foster, r.foster@ieee.org

CONTRIBUTING EDITORS
Evan Ackerman, Mark Anderson, John Blau, Robert N. Charette, Peter Fairley, Tam Harbert, Mark Harris, David Kushner, Robert W. Lucky, Paul McFedries, Prachi Patel, Richard Stevenson, Lawrence Ulrich, Paul Wallich

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia
SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

PRODUCT MANAGER, DIGITAL Shannan Brown
WEB PRODUCTION COORDINATOR Jacqueline L. Parker
MULTIMEDIA PRODUCTION SPECIALIST Michael Spector

EDITORIAL / ADVERTISING CORRESPONDENCE
IEEE Spectrum
3 Park Ave., 17th Floor
New York, NY 10016-5997
EDITORIAL DEPARTMENT
TEL: +1 212 419 7555 FAX: +1 212 419 7570
BUREAU Palo Alto, Calif.; Tekla S. Perry **+1 650 752 6661**
ADVERTISING DEPARTMENT +1 212 705 8939

KAREN LIGHTMAN

CONTRIBUTORS_

David C. Brock

A historian of science and technology, Brock recently became director of the Center for Software History at the Computer History Museum. A few years back, while looking into the history of microcircuitry, he stumbled across the work of Dudley Buck, a pioneer of speedy cryogenic logic. He wrote about Buck in our April 2014 issue. In this issue, he explores what's happened after Buck, including a new effort to build a superconducting computer [p. 50]. This time, the draw is energy efficiency, not performance.

Al Geist

Geist is the chief technologist for the computer science and mathematics division at Oak Ridge National Laboratory, in Tennessee, where he's been studying ways that supercomputers can function in the face of inevitable faults. Today's leading strategy for recovering from errors won't work forever, he warns in "Supercomputing's Monster in the Closet" [p. 26]. When will it give out? "That's like asking when Moore's Law will end," he says. Nobody knows, but you still need to get ready.

Lucas Laursen

A regular contributor to *IEEE Spectrum*, Laursen is a science and technology journalist based in Madrid. For this issue, he interviewed Ronnie Nader, the prime force behind Ecuador's nascent space program [p. 23]. "Nader taught me how a small group of engineers are laying the groundwork for a self-sustaining space industry," says Laursen. "Probably the thing I least expected was how important he considers raising Ecuadorians' expectations" of what a developing nation can accomplish.

Emily Waltz

Waltz writes about biomedical technology from Nashville. In this issue, she explains the software behind a revolutionary new genetic engineering tool ["Software Helps CRISPR Live Up to Its Hype," p. 12]. "In the past anyone who wanted to do this kind of precise genetic engineering had to have serious expertise," says Waltz. But with this tool and the code that makes it work, just about any biology lab can do it. "It's a democratization of the technology. Everybody is pumped about it."

Winfried W. Wilcke

Wilcke has added many strings to his bow since graduating with a doctorate in nuclear physics. In 1983 he joined IBM, where he heads up nanoscale research at the company's lab in San Jose, Calif., dividing his time between cognitive computing and energy storage. With coauthor Ho-Cheol Kim, a materials scientist who leads the lab's advanced energy storage group, he writes about the latter in this issue: a promising new battery technology for electric vehicles [see "The 800-km Battery," p. 38].

03.16

IEEE

IEEE MEDIA
SENIOR DIRECTOR; PUBLISHER, IEEE SPECTRUM
James A. Vick, j.vick@ieee.org
ASSOCIATE PUBLISHER, SALES & ADVERTISING DIRECTOR
Marion Delaney, m.delaney@ieee.org
RECRUITMENT AND LIST SALES ADVERTISING DIRECTOR
Michael Buryk, m.buryk@ieee.org
BUSINESS MANAGER Robert T. Ross, r.ross@ieee.org
IEEE MEDIA/SPECTRUM GROUP MARKETING MANAGER
Blanche McGurr, b.mcgurr@ieee.org
IEEE MEDIA GROUP DIGITAL MARKETING SENIOR MANAGER
Ruchika Anand, r.t.anand@ieee.org
LIST SALES & RECRUITMENT SERVICES PRODUCT/ MARKETING MANAGER Ilia Rodriguez, i.rodriguez@ieee.org
REPRINT SALES +1 212 221 9595, EXT. 319
MARKETING & PROMOTION SPECIALIST Faith H. Jeanty, f.jeanty@ieee.org
SENIOR MARKETING ADMINISTRATOR Simone Darby, simone.darby@ieee.org
MARKETING ASSISTANT Quinona Brown, q.brown@ieee.org
RECRUITMENT SALES ADVISOR Liza Reich +1 212 419 7578
ADVERTISING SALES +1 212 705 8939

ADVERTISING PRODUCTION MANAGER Felicia Spagnoli
SENIOR ADVERTISING PRODUCTION COORDINATOR
Nicole Evans Gyimah
ADVERTISING PRODUCTION +1 732 562 6334

MANAGING DIRECTOR, PUBLICATIONS Michael B. Forster

IEEE BOARD OF DIRECTORS
PRESIDENT & CEO Barry L. Shoop, president@ieee.org
+1 732 562 3928 FAX: +1 732 465 6444

PRESIDENT-ELECT Karen Bartleson
TREASURER Jerry L. Hudgins **SECRETARY** Parviz Famouri
PAST PRESIDENT Howard E. Michel

VICE PRESIDENTS
S.K. Ramesh, Educational Activities; Sheila Hemami, Publication Services & Products; Wai-Choong "Lawrence" Wong, Member & Geographic Activities; Bruce P. Kraemer, President, Standards Association; Jose M.F. Moura, Technical Activities; Peter Alan Eckstein, President, IEEE-USA

DIVISION DIRECTORS
Maciej Ogorzalek (I); Hirofumi Akagi (II); Celia L. Desmond (III); William W. Moses (IV); Harold Javid (V); Rob Reilly (VI); Alan C. Rotz (VII); John W. Walz (VIII); Ray Liu (IX); Kazuhiro Kosuge (X)

REGION DIRECTORS
Ronald A. Tabroff (1); Timothy P. Kurzweg (2); James M.Conrad (3); Robert C. Parro (4); Francis B. Grosz Jr. (5); Thomas Coughlin (6); Witold M. Kinsner (7); Costas M. Stasopoulos (8); Antonio C. Ferreira (9); Ramakrishna Kappagantu (10)

DIRECTORS EMERITUS Eric Herz, Theodore W. Hissey

IEEE STAFF
EXECUTIVE DIRECTOR & COO James Prendergast
+1 732 502 5400, james.prendergast@ieee.org
CHIEF INFORMATION OFFICER Cherif Amirat
+1 732 562 6399, c.amirat@ieee.org
PUBLICATIONS Michael B. Forster
+1 732 562 3998, m.b.forster@ieee.org
CORPORATE ACTIVITIES Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
HUMAN RESOURCES Shannon Johnston, SPHR
+1 732 562 6343, s.johnston@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Eileen M. Lach, +1 212 705 8990, e.m.lach@ieee.org
CHIEF MARKETING OFFICER Patrick D. Mahoney
+1 732 562 5596, p.mahoney@ieee.org
EDUCATIONAL ACTIVITIES Jamie Moesch
+1 732 562 5514, j.moesch@ieee.org
CHIEF FINANCIAL OFFICER Thomas R. Siegert
+1 732 562 6843, t.siegert@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
Sheila Hemami, *Chair*; John Baillieul, Babak D. Beheshti, Herbert S. Bennett, Jennifer T. Bernhard, Eddie Custovic, Jean-Luc Gaudiot, Ron B. Goldfarb, Glenn Gulak, Lawrence Hall, Christofer Hierold, James M. Irvine, Hulya Kirkici, Carmen S. Menoni, Paolo Montuschi, Thrasos Pappas, Michael Pecht, Michael Polis, Sorel Reisman, Tariq Samad, Fred Schindler, Gianluca Setti, Gaurav Sharma, Curtis A. Siller, David Tian, Ravi Todi, H. Joel Trussell, Leung Tsang, Stephanie M. White, Steve 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

PHOTO-ILLUSTRATIONS BY **Gluekit**

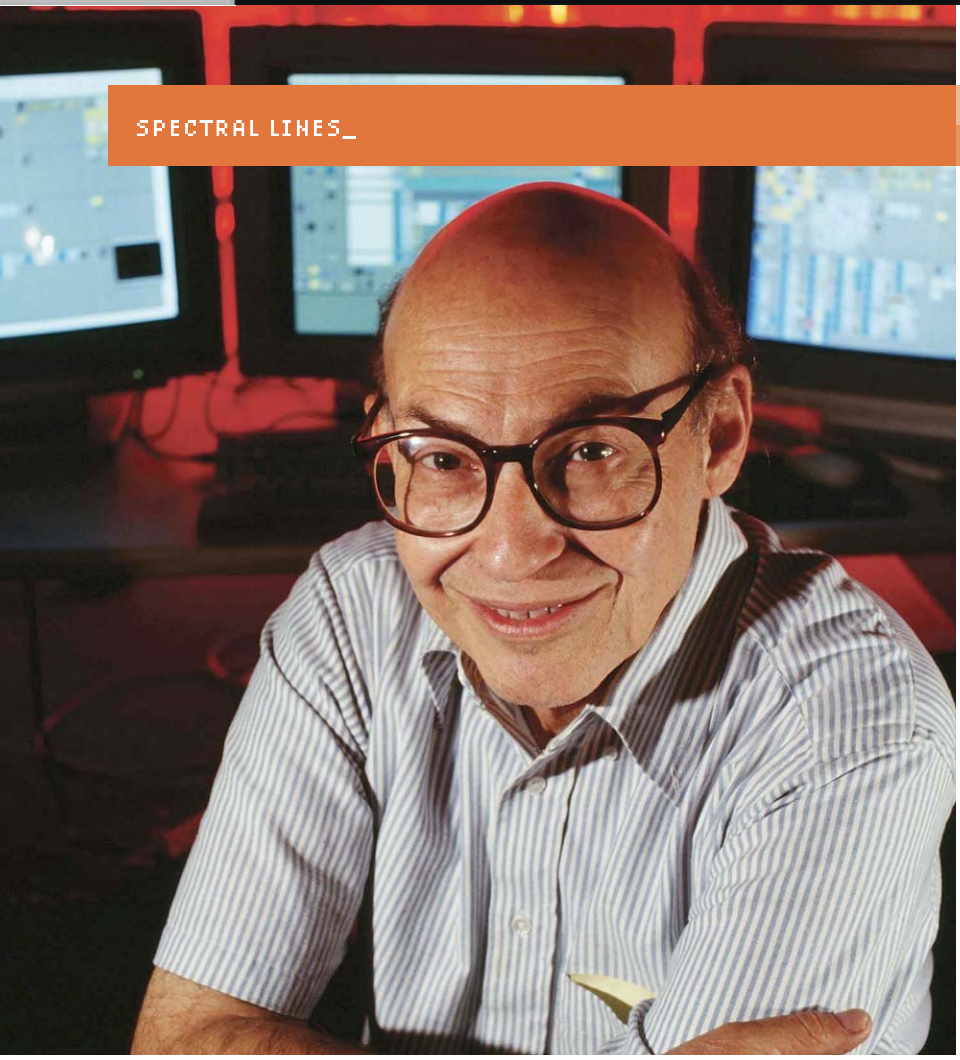

Marvin Minsky and the Pursuit of Machine Understanding

Making machines—and people—think

GENIUSES ARE MEAT MACHINES—that's how Marvin Minsky once characterized human beings—just like the rest of us, only much more efficient. And Minsky, who passed away in January, was certainly one of the most efficient meat machines this century or the last has ever seen.

There are the curriculum vitae facts of his genius: the degrees in mathematics from Harvard and Princeton; the cofounding, with John McCarthy, of the Computer Science and Artificial Intelligence (AI) Laboratory at MIT; the invention of the first head-mounted graphical display, the confocal microscope, and the first randomly wired neural network machine. He was a pioneering computer scientist, cognitive scientist, and roboticist, a fellow of IEEE and of the American Academy of Arts and Sciences, and the recipient of numerous honors and awards, among them the Turing Award, the IEEE Computer Society's Computer Pioneer Award, and the Franklin Institute's Benjamin Franklin Medal. He left his mark on every field that captured his interest, moving through several with seeming ease before finding his life's work: creating a theoretical framework in which to build machines that could understand as well as calculate.

But then there is the much-harder-to-characterize genius of his life as a public intellectual, provocateur, mentor, and friend.

In addition to thinking about how computers could be taught to learn as children do, bootstrapping their knowledge to learn more and more, he wrote about his efforts with verve and wit. Reading "Steps Toward Artificial Intelligence," written in 1961 and published in the *Proceedings of the Institute of Radio Engineers*, it's important to remember that the computers that sparked his inquiries then had mere kilobytes of memory. Minsky's best-known book, *The Society of Mind* (1986, Simon & Schuster), laid out how a human mind could emerge from elements, like neurons and neuronal pathways, that have no minds themselves. His articles, books, and interviews are a treasure trove of epigrams and ideas—funny, profound, confounding, and sometimes infuriating.

The most important part of Minsky's legacy just might be his students and his students' students. If you take a look at the Mathematical Genealogy Project, you'll find that he has 36 students and 1,123 descendants in this one knowledge domain alone. One of my colleagues quipped, "It's kind of like Mozart giving flute lessons."

Minsky wanted the AI that many find scary and threatening: He wanted to build computers that could actually think, not just machines that appear to think by virtue of their prodigious abilities to grind through masses of data. Steven Cherry, director of TTI/Vanguard, a think tank that sponsored meetings Minsky participated in, says, "He saw the developments of the last few years as steps in the wrong direction. Google and Facebook are exploiting their vast data sets, using deep learning. But Minsky saw this as achieving short-term gains at the expense of solving the real machine-intelligence problem."

Minsky wasn't naive about what could happen if machines achieved the capacity for humanlike learning, but he believed that humans would be able to deal with the challenges this would create. He worked on his magnificent obsession, not in tweets, or three-year research-grant sprints, or IPO-funding cycles, but over decades—by thinking one rich thought at a time. —SUSAN HASSLER

HANK MORGAN/SCIENCE FACTION/CORBIS


Discover
CST STUDIO SUITE 2016

NEWS

785 GW

ELECTRICITY CAPACITY THAT INDIA IS EXPECTED TO ADD BY 2040

ELIZA STRICKLAND

"BAREFOOT" MATRIARCHS TAKE ON INDIA'S ELECTRICITY GAP

Some 240 million people in India don't have electricity. Can illiterate village women solve the problem with off-grid solar power?

The women wielding soldering irons in this rustic solar engineering workshop may not know how to read or write, but they know their way around their circuit boards. "This is the shunt coil," Jansya Devi says, with a proud smile. "This is an eight-pin connector, and this is a drum coil."

These technical words were not in Devi's vocabulary a few months ago. She hails from Bihar, a state in eastern India known mainly for its rural poverty. Her village doesn't get electricity from the national power grid, she says, and after dark she typically does her housework by the light of a kerosene lantern and candles. She dreads the monsoon season, when high winds make it difficult to keep the wicks lit. She has a mobile phone, but to charge it she has to send it along with someone making a trip to the nearest town.

Her situation will soon change, thanks to the Barefoot College, a nonprofit school that trains "barefoot solar engineers" like Devi, using color-coded parts and »

SOLAR SCHOOL: Rural women learn to assemble solar-powered lanterns at the Barefoot College, in northwest India.

hands-on lessons. After six months at the peaceful campus in the northwestern state of Rajasthan, the new technicians will return to their home villages, bringing with them solar power equipment and know-how. They'll install solar panels, charging stations, and small LED lights in houses, and they'll stand ready to deal with breakdowns. While these systems offer only the most basic amenities of modern life, they also bring independence from India's dysfunctional national power grid.

The program initially enrolled men, but these students proved disinclined to return to their villages once they had marketable skills. When the school began recruiting grandmothers instead, the program took off. As of 2015, the college's graduates have brought light to some 20,000 houses in more than 300 villages across India, Barefoot administrators say. This very literal campaign of rural empowerment shows the untapped potential of women who are often passed over, says program manager Gloria Jonathan. "These women may be illiterate, but that doesn't mean they're uneducated," she says. "They have skills and they have intelligence."

LIGHTING UP LIVES: At a "night school" run by the Barefoot College [above], kids who work on farms during the day learn by the light of a solar-powered lantern. Women such as Jansya Devi [photo at right, seated on left] come to the Barefoot College from villages that don't get electricity from the national grid. After training, Devi will return home with solar panels, lights, and the skills to keep the systems running. Even in New Delhi, electricity distribution [top right] is a messy business. Barefoot College offers off-grid solar solutions [bottom right].

The Barefoot College believes that such women can solve a very big problem. In India, 240 million people don't have access to electricity, according to the 2015 report *India Energy Outlook*, from the International Energy Agency—and that's a low estimate. A 2014 World Bank report put the figure at 300 million people. Yet the founder and director of Barefoot College, Bunker Roy, doesn't sound daunted. His school's decentralized and off-grid approach to providing electricity is "the only answer" for India, he says. "Mahatma Gandhi said the ultimate solution for fighting poverty in India was not mass production but production by the masses," Roy says. "We have to apply the Gandhian model to solar-electrifying villages."

India's government, however, has a very different plan. Prime Minister Narendra Modi has pledged to give every citizen access to electricity by 2022. To meet this deadline, the government plans to rely heavily on solar power—but not off-grid projects run by local people. Last year, the Ministry of New and Renewable Energy announced an ambitious goal for building utility-scale solar facilities, including "ultra mega solar power projects" of 500 megawatts each, and establishing grid-connected rooftop solar projects in cities. With these two components, the government aims to reach 100 gigawatts of installed solar power capacity by 2022.

Many experts are frankly skeptical of the government's ability to reach this goal, given that installed solar capacity in 2014 was less than 4 GW. Such a rapid scale-up will be "very difficult to achieve," states a report by Bloomberg New Energy Finance, adding that the 100-GW goal may be more "aspirational" than realistic. The International Energy Agency report agrees, listing challenges relating to land acquisition, solar panel manufacturing, and finances. The report projects that by 2020 solar capacity will reach only 40 GW.

How the country solves its electrification problem is a pressing concern

not only to rural Indians living in dark villages but to the whole world. India was estimated to account for 6 percent of global carbon dioxide output in 2014, and the country's rapid economic growth could double its emissions by 2030. At the Paris climate change conference last December, Indian officials alternated between making bold promises to invest in solar power and claiming their right to continue building coal-fired power plants. India currently generates 61 percent of its electricity from coal, according to government reports. If the Indian government can't meet its solar power pledges and provide clean electricity to all its citizens, the world's best hope may be that those citizens will provide it for themselves.

The Barefoot College began its pioneering work to bring off-grid solar power to India's villages in the 1990s. In the last decade it has been joined by similar philanthropic efforts as well as dozens of "social enterprise" companies, which operate microgrids or lease solar-powered lights to villagers in an effort to do good while making a profit.

The people who run these scrappy organizations have little faith in the government's pledge to provide reliable grid power to all Indians. Debajit Palit, who directs the "Lighting a Billion Lives" project at the nonprofit Energy and Resources Institute, in New Delhi, notes that previous administrations have made such pledges, and that target dates have been repeatedly pushed back. Much of the delay comes from India's chronically troubled electricity distribution companies, which have operated at a loss due mainly to low power prices and widespread electricity theft. These state-run utility companies together have accumulated US $65 billion (4.3 trillion rupees) in debts.

Last November, the central government approved a rescue plan for the distribution utilities that would clear their debts, freeing up resources to address the theft problem and improve infrastructure. It sounds good in theory, says Palit, "but the central government made the program, and the provincial governments have to implement it." He's seen enough mismanagement to question whether the state governments and companies will follow through with reforms.

Experts who study India's electricity gap say these companies' financial woes have not only stalled "last mile" efforts to electrify rural hamlets, they've also made grid power less attractive to everyone. The distribution companies often can't afford to buy power from generation stations, making long power cuts an everyday occurrence in some parts of the country. "Many villagers think, 'I won't get electricity during peak hours, so why should I connect?' " says Palit. He believes that off-grid solar kits, although they provide only a few watts of electricity, are a better solution for India's villages. "If I promise small, but keep my promise by providing a reliable supply, the villagers are happy," he says.

However, the financial models of off-grid solar efforts have also come under scrutiny. At Barefoot, for example, the bulk of the funding needed to keep its campus running comes from government and philanthropic sources. So scaling up its solar lighting program would require more and more largesse. But program manager Jonathan says that dependency doesn't extend to villages that already have their solar systems up and running, because each community pays its resident solar technician a small salary and funds her "rural electronics workshop." This money would otherwise be spent on kerosene and other fuels, Jonathan says. "We consider this a sustainable model," she says. "Even if we were to pull out, the community effort would still continue."

Still, providing electric light to those 240 million Indians who currently live without it will require scalable solutions.

GOING GLOBAL: The Barefoot College also runs an international version of its solar engineer program, bringing women like Erouby Kamwendo, of Malawi, to the campus for training.

NEWS

PHOTOGRAPHS BY Eliza Strickland

Anant Sudarshan, India director of the Energy Policy Institute at the University of Chicago, says the Barefoot College isn't the answer. "I don't think it's a great idea to throw money into something if we know it's not going to be viable, in the long run, without a continuous injection of funds," he says.

Sudarshan, who is based in New Delhi, has also studied social enterprise companies that offer off-grid solar power to rural villages. "Right now these companies are not doing very well, and they don't have much hope of doing very well," he says. His surveys of people living without electricity revealed that many weren't willing to pay for the tiny set of solar powered lights that such companies offer; instead they want enough power to run appliances such as televisions, electric sewing machines, and refrigerators. Even those villagers who are willing to pay for a few lights often consider this arrangement a stopgap solution while they wait for grid power. "There may be a sweet spot where these businesses can run profitably for a little while, but it's hard to find," Sudarshan says.

Selco India is one of the companies often praised for its attempts to use for-profit solar lighting projects to lift rural Indians out of poverty. But company cofounder H. Harish Hande explains that the solar power kits are only one piece of the services his company offers villagers—and not the most important selling point. His team also helps villagers secure bank loans for the initial solar setup, and then helps them start entrepreneurial efforts to make profitable use of their new power. "People are not interested in electricity; they're interested in what electricity can do," Hande says. "They're interested in running an electric sewing machine."

His full-service and customized approach may seem difficult to scale up, but Hande believes the plan is more logical than waiting for the government to provide grid electricity to villagers and expecting economic gains to follow. "The grid actually defeats development in many ways," he says. While the government's efforts position the poor as consumers, he says, they should instead be valued as producers. "That's the true way to solve poverty," he says, "and solar energy is a very powerful tool in making that happen."

Back at the Barefoot College, administrators have a similar vision for their solar engineering grandmas. When someone like Devi brings light to her dark village in Bihar, she's giving her community "a starting point," says project manager Jonathan. That village could theoretically run its equipment for 25 years before the solar panels give out, she says, but she guesses that young people growing up under those small lights will be dissatisfied with the meager power provided by the small solar kit. That doesn't bother her a bit, Jonathan says. "We see that as a good sign because then they've begun to think, 'We want more.' "

—ELIZA STRICKLAND

> "We consider this a sustainable model. Even if we were to pull out, the community effort would still continue"
>
> —Gloria Jonathan, Barefoot College

SOFTWARE HELPS CRISPR LIVE UP TO ITS HYPE

New algorithms make the gene-editing tool as easy as point-and-click

Biotechnologists are jumping at the chance to use the revolutionary gene-editing tool known as CRISPR. The molecular gadget can be programmed to accurately tweak the DNA of any organism, but scientists need software algorithms to hasten the programming process. Dozens of teams are developing such software, and each faces the task of keeping up with rapidly evolving science and an increasingly crowded field.

CRISPR—short for Clustered, Regularly Interspaced, Short Palindromic Repeats—is a genetic phenomenon found in microbes that scientists adapted to disable a gene or add DNA at precise locations in the genetic code. CRISPR isn't the first gene-editing tool on the block, but it is by far the simplest and cheapest, and since its adaptation four years ago, it has proliferated globally. Researchers can use it to knock out genes in animal models to study their function, give crops new agronomic traits, synthesize microbes that produce drugs, create gene therapies to treat disease, and potentially—after some serious ethical debate—to genetically correct heritable diseases in human embryos.

In less than four years, CRISPR "has transformed labs around the world," says Jing-Ruey Joanna Yeh, a chemical biologist at Massachusetts General Hospital's Cardiovascular Research Center, in Charlestown, who contributed to the development of the technology. "Because this system is so simple and efficient, any lab can do it." Traditional genome modification techniques involve shuttling DNA into cells without knowing where in the genome it will stick. Editing with CRISPR is like placing a cursor between two letters in a word processing document and hitting "delete" or clicking "paste." And the tool can cost less than US $50 to assemble. There are other genome-editing systems that are as precise as CRISPR, but they must be customized for every use and require far more expertise and resources to assemble.

As good as CRISPR is compared to its predecessors, the tool doesn't always work, says Jacob Corn, scientific director at the Innovative Genomics Initiative at the University of California, Berkeley. "We don't really understand why that is," he says. That's where software comes in. Algorithms can help researchers design their CRISPR tools in a way that is statistically more likely to succeed.

CRISPR systems are equipped with two main features: a short strand of programmable genetic code (called a guide RNA) and a protein (usually an enzyme called Cas9) that acts as a pair of molecular scissors. Once the complex is introduced into a cell, the guide RNA ushers Cas9 to a precise location in an organism's DNA sequence (or genome), sticks to it like Velcro, and lets the Cas9 snip the DNA. The cell's own machinery then repairs the cut, chewing up a bit of DNA or adding some in the process, thus disrupting the gene. Researchers can also intentionally introduce a piece of new genetic code to the site.

Guide RNAs find their targets in an organism's genome by looking for a DNA segment with a complementary code of molecules. The molecules are called bases and are represented by the letters A (adenine), T (thymine), G (guanine), and C (cytosine) [see illustration, "Smart Scissors"].

Scientists editing the genome are usually looking for a segment that controls a particular function–a gene. Those are generally hundreds to thousands of bases long. Guide RNAs, however, are only about 20 bases long, so scientists must choose a complementary 20-base segment within the gene to target. There are two main constraints to consider: The target has to be located near a landmark that the molecular scissors can recognize, and it must be unlike any other 20-base segment anywhere else in the genome.

SMART SCISSORS

The CRISPR system consists of an enzyme capable of chopping DNA and a piece of "guide RNA," which places the enzyme at the correct spot in the genome. The guide RNA is engineered to steer the enzyme to a particular spot in the DNA.

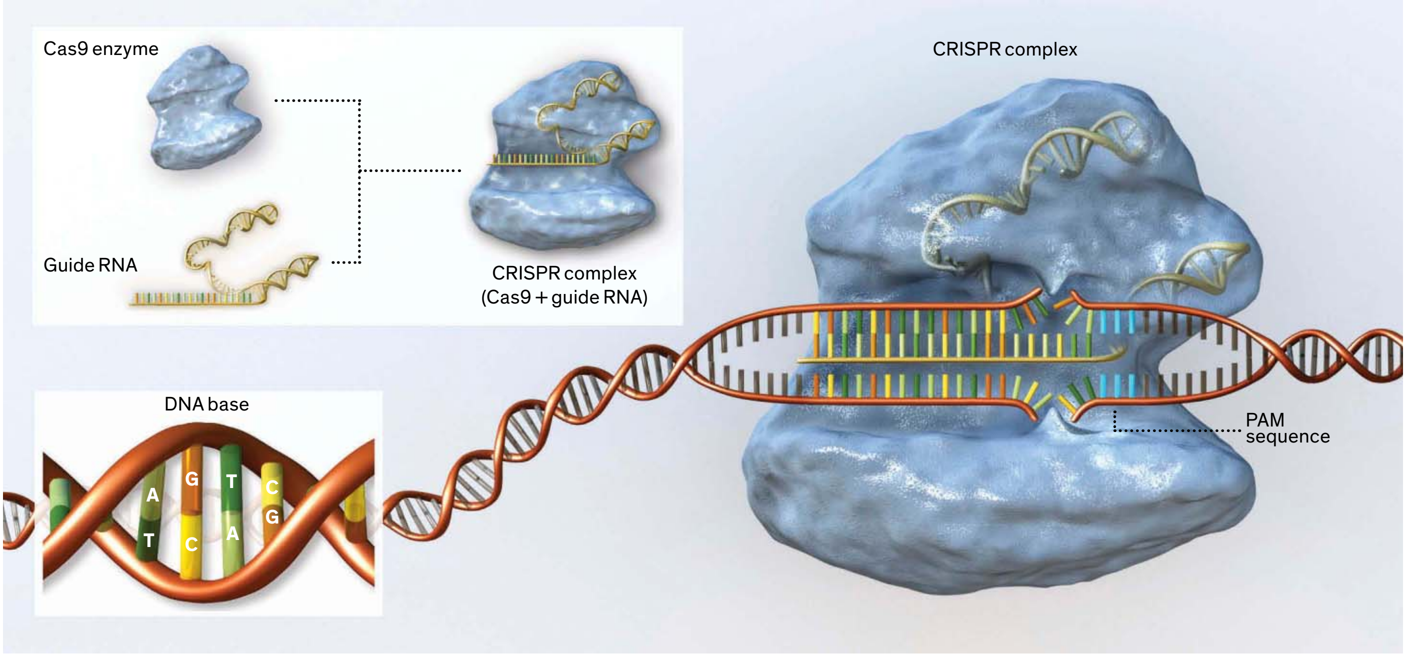

DNA encodes genes as a sequence of chemicals (symbolized as A, T, G, and C) that mate with their chemical complements (A with T, G with C) to form the rungs of the molecular helix.

Guide RNA ushers the CRISPR complex to a complementary site in the DNA, where the enzyme looks for landmarks called protospacer adjacent motifs, or PAMs. If the complex finds both a DNA match and a PAM, it snips the DNA strand so that either the gene's sequence will be disrupted or a new piece of DNA can be inserted at the site.

NEWS

TECHNICAL ILLUSTRATION BY Emily Cooper

The Cas9 enzyme's landmark is called a protospacer adjacent motif, or PAM. PAMs are easy to find in the genome–it's like looking for the word "the" in a book, and any complementary 20-base segment adjacent to a PAM can work as the target site.

Ensuring that the 20-base segment is unique, however, is tougher. With a genetic code having only four letters and most organisms' genomes ranging from millions to billions of base pairs, patterns are often repeated. Guide RNAs can get distracted by decoy segments, called off-target sites, and may end up mutating the wrong gene. Segments that differ from the target by just a couple of bases can trip up the tool. "You could scan across the whole genome by eye to try to find [off-targets], but it will take forever," says Cameron Ross McPherson, a data scientist at the Institut Pasteur in Paris, who codeveloped the CRISPR software Protospacer Workbench.

Algorithms can do the search rapidly with few inputs from the user. Harvard University's CHOPCHOP asks the user to enter the organism, the gene, and some optional advanced parameters. Within seconds, the algorithm finds within the target gene all the possible 20-base segments located near a PAM, ranks them based on their uniqueness in the genome and other parameters, and generates a list of guide RNAs to get you there. A search through the *spaw* gene in zebra fish, for example, yields 55 possible guide RNAs, most of which are unique sequences that differ from all other patterns in the genome by at least two bases.

Dozens of such software tools have popped up over the last two years, most of them available for free. A handful of companies, such as Benchling in San Francisco, offer more user-friendly interfaces than those of the free-to-the-public versions. But none of the software packages has emerged as a front-runner, says Michael Boutros at the German Cancer Research Center in Heidelberg, Germany, who codeveloped the CRISPR software E-CRISP.

Boutros says there is a lot of work to do. A list of 55 guide RNAs that theoretically might work is a helpful starting point, but it leaves researchers having to experiment, trial-and-error style, to see which one works the best. Algorithms that can instead predict with certainty that a particular guide RNA will work are needed.

To that end, biostatisticians are beginning to comb through experimental data to look for common patterns of successful guide RNAs that they can use to inform machine-learning-based prediction systems. But most of that data is scattered throughout small, individual studies. "Putting that all together will be a very powerful resource," and it presents an opportunity for computer engineers to get involved, says UC Berkeley's Corn. A few large data sets do exist already. A group from the Broad Institute in Cambridge, Mass., tested nearly 2,000 guide RNAs in human and mouse cells and recently published a set of rules for an improved algorithm.

Meanwhile, scientists are tinkering with Cas9 and other cutting proteins in an attempt to provide more options for CRISPR users. Some of those proteins improve the accuracy of the guide RNA. If they succeed, the need for software that predicts the accuracy of the tool may vanish, or at least evolve. "If we get to a place where we've absolutely, positively eliminated off-target effects, then great," says Corn. "Are we there yet? No."

–EMILY WALTZ

CRUNCH TIME FOR THE KILOGRAM

Superprecise standard masses are on their way to Paris for a critical weigh-in

K50 is going on a trip. Sometime soon, the golf-ball-size cylinder of platinum and iridium, which makes its home in Ottawa, will be carefully packed up and hand carried by a member of Canada's diplomatic corps on a plane to France, with special instructions that its case not be opened by customs or airport security.

This precious cargo will join a small clutch of similar objects that will trickle into Paris in the coming weeks. Their destination: the International Bureau of Weights and Measures (BIPM) just outside the city, where they'll be compared with one another to see just how well their estimated masses match up.

The exercise is a test, a dry run to see whether the world is ready for a shift in the way we define mass. Today, the International System of Units defines the kilogram by using a physical object: a supercentenarian cylinder called the International Prototype of the Kilogram (IPK). It is the progenitor of K50 and many other mass standard bearers. Although the IPK is stored in a vault and brought out only rarely, it is still a physical object, and so its mass can change. But because its mass is, by definition, 1 kilogram, any shifts that do occur essentially happen to the mass of everything else.

The world's metrologists aim to change this state of affairs in 2018 by fixing the kilogram to the Planck constant, a fundamental physical constant.

MASSES IN MOTION: Canada will send a cousin of this kilogram mass to Paris for a dry run of the unit's redefinition. It'll be joined by masses from around the world.

That shift would, at least in principle, allow any laboratory to "realize" the kilogram from scratch with a series of experiments and specialized equipment. But for that scheme to work, the kilogram derived by one laboratory must be the same as those derived by others.

"The question is, How would these independently realized kilograms agree?" says BIPM's Michael Stock, who is leading the pilot study. Stock expects roughly half a dozen groups to weigh in. Two will hail from the national metrology institutes of Germany and Japan, which maintain ultrapure silicon spheres whose characteristics were recently measured with great precision—one possible way of determining the kilogram. The remaining groups will represent institutes that are working on watt balances, which measure masses against the electrical force needed to resist their weight.

There are a number of watt balances and similar electromechanical balances under construction around the world. But to participate in this pilot study, a group must demonstrate that it can measure a kilogram with an uncertainty of 0.00002 percent or less, meaning it can't be off by more than 200 micrograms.

"We are ready to go," says Carlos Sanchez, discipline leader for electrical standards at Canada's National Research Council. As *IEEE Spectrum* went to press, he and his colleague Richard Green were working through a series of measurements that would calibrate K50 against masses measured in the group's watt balance. The apparatus recently achieved an uncertainty of just 18 micrograms, currently the lowest published and well below BIPM's requirement.

But a critical test will be how well other measurements agree. Previous experiments, which effectively flip the watt-balance measurement around so that masses are used to measure the Planck constant, have not been in perfect accord.

BIPM hopes to receive all the test masses by the end of March, although the list of participants and their expected arrival times was still being finalized at press time. The comparison measurements will take some months, and the results are not likely to be available before next year.

The outcome of the test is uncertain. So, too, is the fate of the kilogram in 2018. "If everything goes well, [the comparisons] would confirm that the future system would work. If it's not satisfactory, it could be an obstacle [to the redefinition]," Stock says. "So it's definitely going to be interesting."
—RACHEL COURTLAND

FEEL THE FORCE

You are soft. It's nothing to be ashamed of. So sensors designed to interact with you should be soft, too. One big problem is that a pressure sensor's accuracy falters when the sensor is bent. And bending is very important, because *you* bend. So engineers in Japan, led by Takao Someya, came up with nanofiber-based sensors that change their resistance under pressure. The fibers reorient themselves when bent, so sensors stay accurate even when curved around an 80-micrometer diameter.

For more, see http://spectrum.ieee.org/nanoclast0116

TOP: NATIONAL RESEARCH COUNCIL OF CANADA; BOTTOM: SOMEYA LABORATORY

PHOTOGRAPH BY Arca Space/Splash/Corbis

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

AIR PLANE

ELECTRIC SKATEBOARDS are all the rage right now. For some reason they're called hoverboards, although they don't, well...you know. But at last, there's a true hoverboard that rides on air like the one piloted by Marty McFly in *Back to the Future Part II* (1989). The ArcaBoard's batteries hold enough energy to keep its 36 electric fans running for between 3 and 6 minutes, depending on the rider's weight. A full recharge takes as little as 35 minutes. The US $20,000 deck, from Arca Space Corp. in Las Cruces, N.M., musters enough thrust to ferry a rider weighing 110 kilograms (about 242 pounds) at speeds approaching 20 kilometers per hour (about 12 miles per hour).

THE BIG PICTURE NEWS

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


Advance Your Career
With the Experts at Wiley-IEEE Press
Stuff You Don't Learn in Engineering School
Skills for Success in the Real World
CARL SELINGER
ISBN: 978-0-471-65576-3
Fast-Tracking Your Career
Soft Skills for Engineering and IT Professionals
Wushow "Bill" Chou
IEEE
IEEE PRESS
WILEY
ISBN: 978-1-118-52178-6
Ten Essential Skills for Electrical Engineers
BARRY L. DORR
IEEE
WILEY
ISBN: 978-1-118-52742-9
Negotiating Cultural Encounters
Narrating Intercultural Engineering and Technical Communication
EDITED BY Han Yu · Gerald Savage
WILEY
IEEE
ISBN: 978-1-118-06161-9
Information Overload
An International Challenge for Professional Engineers and Technical Communicators
EDITED BY Judith B. Strother · Jan M. Ulijn · Zohra Fazal
WILEY
IEEE
ISBN: 978-1-118-23013-8
Slide Rules
Design, Build, and Archive Presentations in the Engineering and Technical Fields
Traci Nathans-Kelly · Christine G. Nicometo
IEEE
WILEY
ISBN: 978-1-118-00296-4
ENGINEERING Your RETIREMENT
Retirement Planning for Technology Professionals
MIKE GOLIO
ISBN: 978-0-471-77616-1
ARE YOU AN EXPERT IN YOUR FIELD?
TO LEARN MORE ABOUT PUBLISHING WITH WILEY-IEEE VISIT:
HTTP://WWW.IEEE.ORG/PUBLICATIONS_STANDARDS/PUBLICATIONS/PRESS/AUTHORS/BOOKS.HTML
IEEE
VISIT: WWW.WILEY/IEEE.COM
WILEY

RESOURCES

735,405

NUMBER OF PEOPLE IN THE UNITED STATES WITH AMATEUR RADIO LICENSES IN 2015

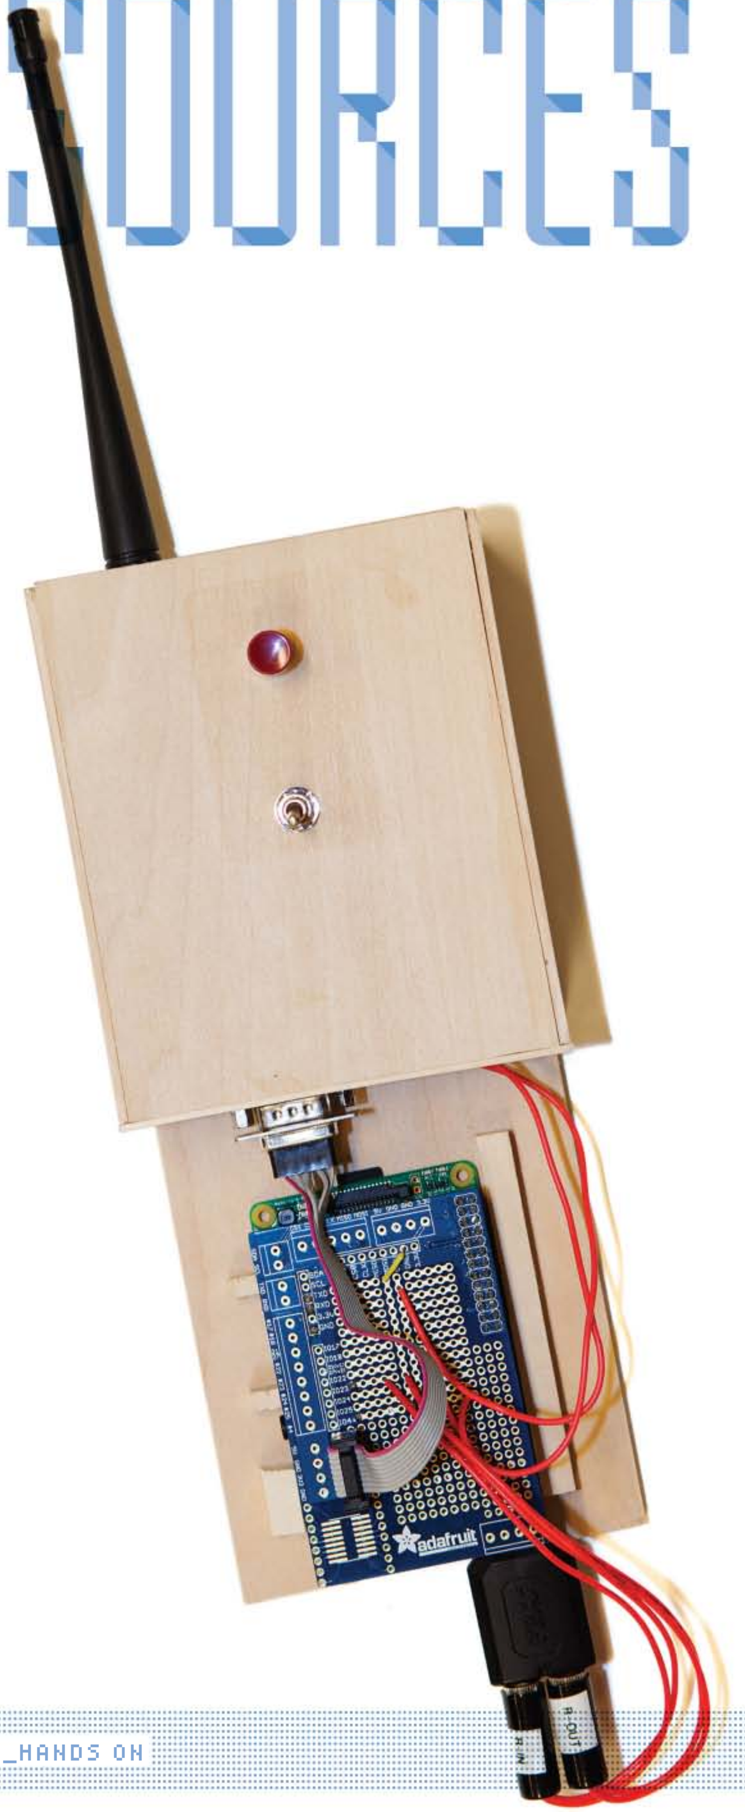

A HAM RADIO FOR MAKERS

THE RS-UV3 LETS YOU BUILD YOUR OWN ARDUINO- OR PI-BASED RADIO

RESOURCES_HANDS ON

PHOTOGRAPH BY Randi Klett

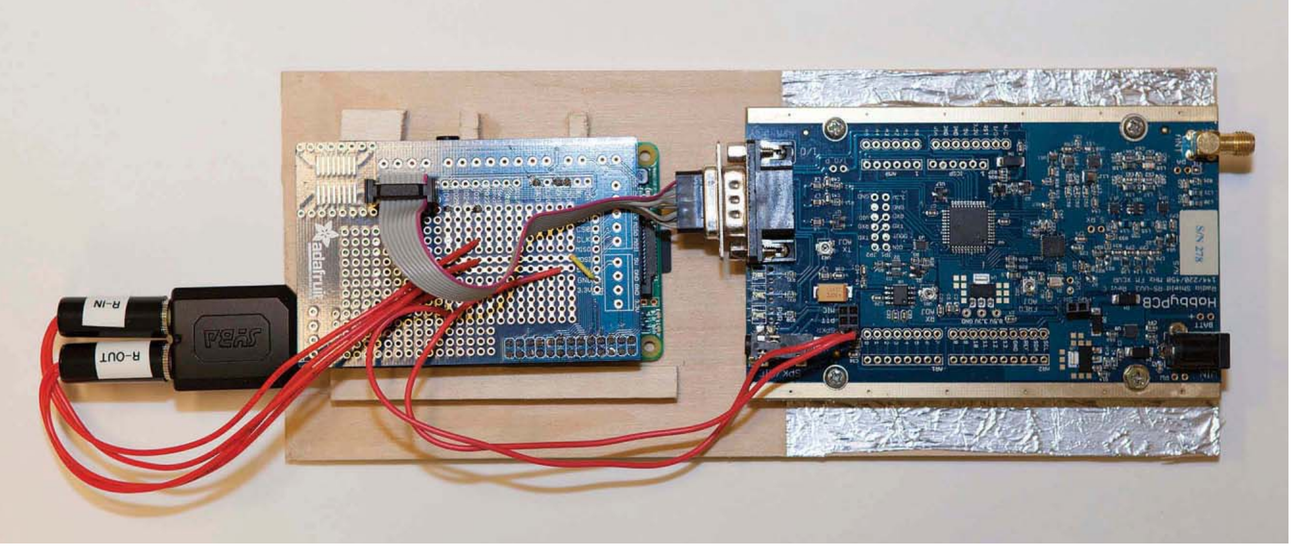

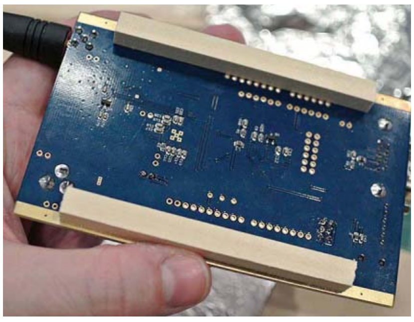

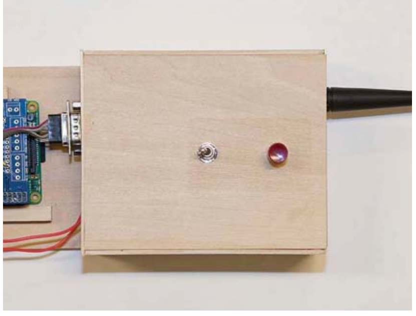

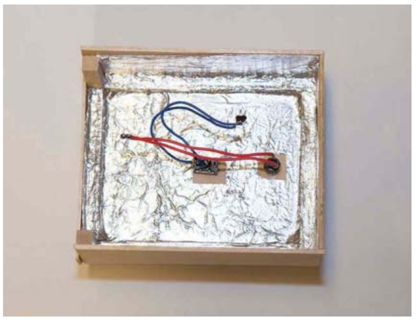

The RS-UV3 is a shot in the arm for amateur radio. Mobile phones and the Internet have made the basic act of talking to a faraway person an everyday experience. This means that much of the appeal of ham radio is now in things like emergency response; technically challenging exercises such as bouncing signals off satellites or ultralow-power long-distance contacts; and exploring a host of digital communications modes.

In some ways, trying out such digital modes has never been easier. Free desktop programs like Fldigi can work with the audio tones used in a smorgasbord of communications schemes, from the 1930s-era radio-fax Hellschreiber protocol to today's complete bulletin-board systems. But linking the computers running such software to radios is often surprisingly fiddly in the age of painless USB and Bluetooth. Except for high-end rigs, connecting a computer to a ham radio typically involves navigating legacy interfaces and connectors and can call for specialized additional equipment, a turnoff for makers who might otherwise be interested in the possibilities of radio.

But the US $90 RS-UV3 radio shield from HobbyPCB is an FM transceiver that's welcoming by design and built for makers from the ground up. It's not the only radio shield available for the Arduino that works in the UHF/VHF frequency bands, but the RS-UV3 is the most flexible one I've spotted when it comes to interfaces. As well as a legacy ham radio interface, the shield provides multiple ways to connect it to an Arduino, PC, or Raspberry Pi—and to two of those at once if required. By itself, the RS-UV3 can transmit with only 0.25 watt, but HobbyPCB plans to sell an add-on amplifier for more powerful transmissions.

(A quick legal aside: Using the RS-UV3 to transmit requires an amateur radio license. Depending on your country, you may also require a license to receive, so check your local regulatory regime.)

Unlike standard ham radios, the RS-UV3 can't operate as a stand-alone device. First you have to choose how to turn it on, either by soldering a jumper across the power line, or as I did, by wiring in a connector to which I attached a switch. The onboard processor then boots up in what's known as simplex mode (where the same frequency is used for both reception and transmission) at a default frequency of 146.52 megahertz. Plugging in a handheld speaker/microphone with a push-to-talk (PTT) button via a jack on the RS-UV3 will let you talk back and forth on that frequency. Alternatively, you could wire an electret microphone, small speaker, and PTT button directly to the shield.

BOTTOM LEFT AND OPPOSITE PAGE: STEPHEN CASS; THIS PAGE: RANDI KLETT (3)

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

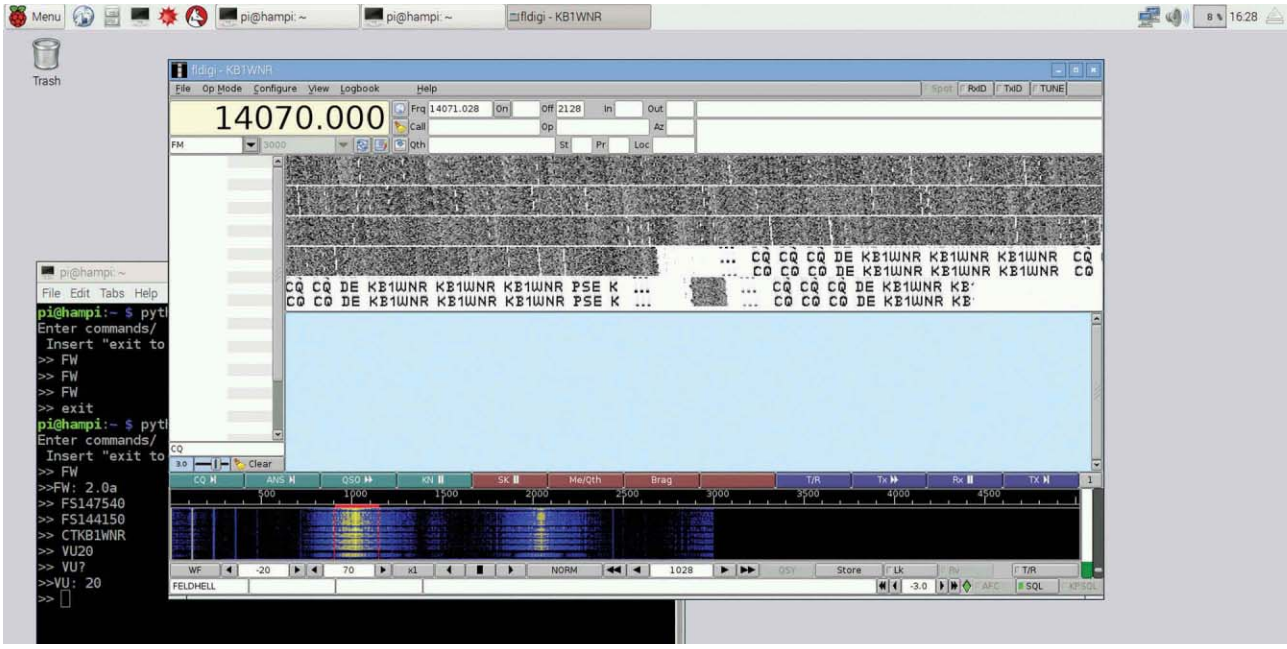

RASPBERRY PI RADIO: The RS-UV3 was screwed to wooden supports [bottom left] and mounted. A USB adapter allows audio to pass between the RS-UV3 and a Pi, while commands are passed via a DB-9 connector [top left]. To reduce interference, I housed the RS-UV3 in a foil-lined enclosure [bottom right and middle]. Signals are decoded with Fldigi [above].

But most people who have a radio want to tune it to more than one frequency. You accomplish this, and other functions such as adjusting the volume, by setting up a serial connection with the RS-UV3's processor and sending it commands. (The setup is reminiscent of the old dial-up modem Hayes command codes.)

You can establish the serial connection in a number of ways, such as mounting an Arduino directly to the RS-UV3 using the provided through-holes and then wiring two of the Arduino's input/output pins to a header on the shield. Another option is to use a header intended for FTDI serial cables and connect it to a PC's USB port.

Or you can use the shield's versatile DB-9 connector. Along with serial transmit and receive lines, the connector also provides audio input and output to the transceiver and a PTT control line. These serial lines use 3.3 volts, rather than the typical 5 V. This is handy because 3.3 V is the operating voltage of the Raspberry Pi. Consequently, you can connect the RS-UV3 directly to the general-purpose input/output header (GPIO) of the Pi without much in the way of interface electronics.

And thanks to the upgrade in processing power that came with the release of the Raspberry Pi 2, the Pi now has enough horsepower to run Fldigi in addition to controlling the RS-UV3. So I did just that.

First, I built a simple hardware interface between the Pi's GPIO port and the RS-UV3's DB-9 connector using an old prototyping "hat" from Adafruit I had lying around. I also created a wooden enclosure (as is my wont) for the shield—lined with tinfoil to reduce radio frequency interference—with a power switch and a PTT button on top and a slot to hold the Pi.

I used the Pi's configuration tool so that the Raspbian operating system wouldn't reserve the serial pins on the GPIO for its own use. I tweaked a Python script I found online, written by Fabio Varesano, to create a basic command-line terminal to send commands to the RS-UV3.

Then I turned my attention to using Fldigi to listen to and generate signals for the transceiver. I wired up two mono audio jacks to the RS-UV3's audio input and output feeds. Unfortunately, while the Pi has a built-in socket for audio output, it doesn't have one for audio input. So, following the advice of Lior Elazary's website, I purchased a Syba USB audio adapter for the Pi and plugged the jacks into that.

Installing Fldigi from its source code was straightforward (I followed Jeffrey Kopcak's online instructions, although without the remote access steps).

Firing everything up (with the Pi connected to a keyboard, mouse, and monitor), I was able to set the RS-UV3 to a suitable frequency and communicate via a number of digital modes to a nearby test rig (which consisted simply of me holding up a handheld ham radio to a laptop running Fldigi through its built-in loudspeakers and microphone in the quiet of *IEEE Spectrum*'s offices late at night). Success! I now have a dedicated digital radio rig, for less than $150 all told.

Future plans will require obtaining an amplifier, but I hope to build on this basic setup to create a complete packet radio bulletin-board system that will fit in a relatively small box. Just in case that whole mobile phone/Internet thing has a hiccup. **—STEPHEN CASS**

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

CES AND THE FUTURE OF VR

DIVERSITY IN INNOVATION WAS THE ORDER OF THE DAY

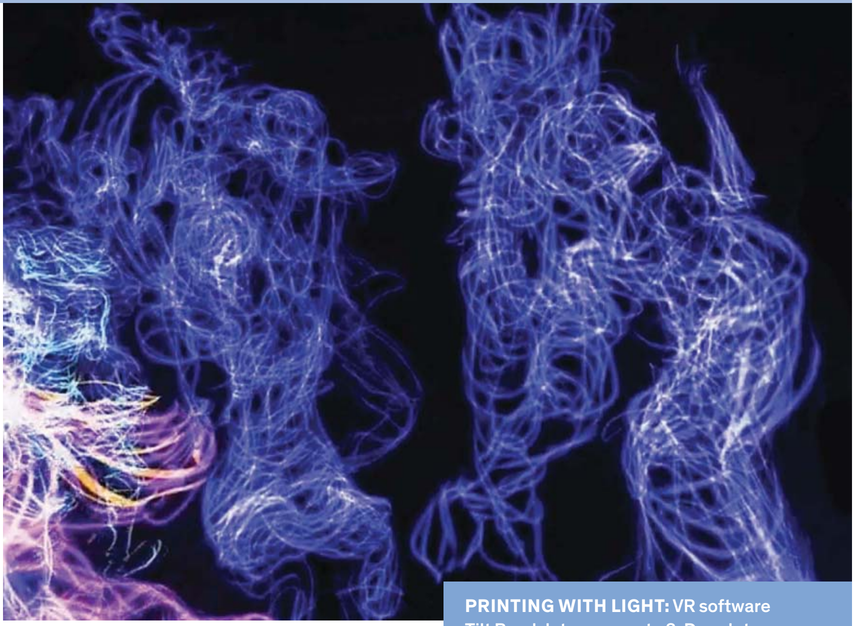

PRINTING WITH LIGHT: VR software Tilt Brush lets you create 3-D sculptures with hand controllers.

There was a certain amount of media hand-wringing surrounding January's CES, the big kahuna of consumer electronics trade shows, held annually in Las Vegas. Where was the big game-changing product? CES was so over.

It's true that it's been a couple of years since, say, smart watches arrived on the scene, and a full five years since the last really big new product category—tablets—was introduced. But new categories tend to leave an innovation hangover: In a rush to grab market share, funding and engineering talent tend to get focused on me-too products. In contrast, this year's exhibitors were presenting a lot of different ideas pushing in a lot of different directions.

The result can be bemusement—for example, there was a US $300 "game console" for dogs, a machine for writing names in coffee foam, and a wristband that lets you give yourself electric shocks to break bad habits. But there were also a lot of ideas with the potential for a big impact, such as a wearable two-way voice translator that doesn't rely on an Internet connection, a design for a two-dimensional tactile display for the visually impaired (the makers hope to have a working prototype later this year), and smart-fabric ideas that could help push fashion tech forward.

But the most interesting emerging trend was in virtual reality. Not so much in the hardware, the basic outlines of which have pretty much revolved around smartphone-display-based headsets for several years now. What was interesting was that people are coming up with engaging content and applications for VR.

For example, the Lowe's chain of home improvement stores was showing off a system it is testing in half a dozen stores. Customers can design a kitchen or bathroom and see how it would look to stand in the room, with an employee making adjustments in real time. When I tested it out, the result was good enough that I kept catching myself trying to lean against a purely virtual countertop. Another company, with the aid of a chair that can tilt a few centimeters forward and back and side to side, provided a simulated carnival ride that was somatically close enough to the real thing to produce literal white knuckles.

But the killer app was something I almost missed, coming across it outside the main halls in the last minutes of the show. This was a demo of HTC's Vive VR system—which uses two pole-mounted locating beacons and hand controllers in addition to a headset—running Google's Tilt Brush software.

Those of us over 40 will recall the era of purely character-based interfaces. Remember then the psychological impact the first time you actually got to use MacPaint or Microsoft Paint and you spent an hour just trying out each type of brush?

Using Tilt Brush was like that, only more so. The system creates a 3-D space—the boundaries are demarcated in your view with a light-grid pattern in best holodeck tradition—in which you are free to walk around. One hand controller becomes a universal palette, which you rotate with your wrist to call up things like a color picker, brush selection, and so on. The other controller becomes your brush: Pull its trigger and "paint" begins to appear at your hand's location.

The tracking is exquisite: As your hand traces a path through space, the brush follows perfectly. Sketch out a shape, and then walk around it to look at it from another side. It's incredibly intuitive (and a lot easier than when we all tried to write our names with a mouse in MacPaint). Despite my limited artistic abilities, within a few moments I had painted a rough 3-D volcano, complete with glowing lava spilling down the sides. Were it not for the line of folks waiting their turn, I could have spent hours in there.

The system is expected to ship later this year, with pricing to be announced. Apart from its value to artists of all stripes, unless the system is a complete flop, it should only be a matter of time before it interoperates with existing computer-aided design software, bringing 3-D design out of the realm of high-end outfits. **—STEPHEN CASS**

This article is based in part on material published online as part of IEEE Spectrum*'s live CES coverage.*

GOOGLE

RONNIE NADER
HE'S LEADING ECUADOR'S CHARGE INTO ORBIT

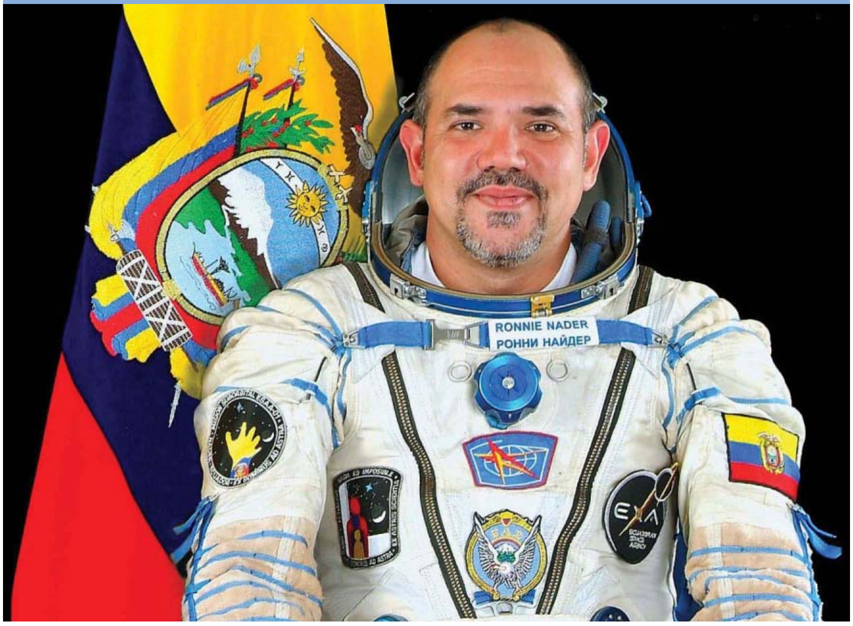

ECUADORIAN CIVILIAN SPACE AGENCY

Ronnie Nader is practically a one-man space program. Nader, a systems engineer and Ecuador's only astronaut candidate, completed four years of cosmonaut training in Moscow in 2007, subsequently helped establish Ecuador's own "vomit comet" zero-gravity training program, and managed the design, construction, launch, and operations of the country's first two orbiting satellites, in 2013.

Of course, Nader is not really alone. The Ecuadorian Civilian Space Agency (EXA), which he founded in Guayaquil in 2007, counts on about a dozen other engineers. Nor is EXA alone in the developing world: Thirty-five developing countries have some kind of space program, the World Academy of Sciences reports.

EXA's first satellite, a 2.1-kilogram CubeSat called Pegaso, suffered a collision with a decommissioned Soviet satellite. Happily, the second, Krysaor, was able to pick up Pegaso's erratic radio signal and help it continue streaming video back home. Here, Nader explains to *IEEE Spectrum* contributor Lucas Laursen how EXA approaches its missions.

Lucas Laursen: Why should Ecuador have a space program now?

Ronnie Nader: Many people think that developing countries should occupy themselves first with the most basic problems and, after those problems are resolved, see to things in space. But that is to assume that these scientific and technological questions are a luxury. We simply don't think that way. You have to be able to do multiple things at once. In developed countries, poverty isn't beaten, nor is crime. Yet those places don't stop developing science or technology.

L.L.: What breakthroughs made it possible for Ecuador to launch its first satellites?

R.N.: In 2010, the satellite ground station was the first enabler. Previously we innocently assumed that there was a way to take the signal from a satellite to the Internet. There wasn't. So we built our own.

Another big difference is the Internet. We have to connect the Internet to space. That was the idea with Pegaso, which emitted video through the Internet. You have to bring space to the people. Our satellite even tweeted.

L.L.: What commercial prospects does spacefaring technology have in Ecuador?

R.N.: The problem with the small CubeSat is that if you buy one part from one place, you have to buy the rest from there too. So we designed almost everything. We bought solar panels, bought raw aluminum to build the frame, had to modify the camera, design and order battery cells from China, make the PCB boards, and made the tooling to make the machines.

We finance ourselves with what we can sell, including data from climate observations. We have seven years of data, so that's valuable now, and sales provide quite a bit.

We haven't yet gotten to the spin-off phase. But in 2011 we exported titanium parts for a British satellite. [We are also going to provide to American CubeSat clients] solar arrays, a radiation shield, titanium infrastructure, and batteries.

L.L.: What value does Ecuador get from your space program that it can't get another way?

R.N.: The ultraviolet-exposure data, from whom could we have bought that? UV-exposure data from the Pegaso and Krysaor are now included in local weather reports. Local firefighters are also benefiting—they came to us to ask for help identifying wildfires.

I would also say that one cannot buy a people's sense of capability. One cannot fire the imagination of a whole generation of future engineers by buying a satellite.

L.L.: When will Ecuador's first astronaut fly, and on what hardware?

R.N.: I continue training, and we're looking to train a backup. For financial reasons, obviously, we can't pay [US] $70 million to the Russians and send a cosmonaut to the [International] Space Station. The Space Station is a private club, although they say quite the contrary. Our approach to the problem is to partner with private companies, such as Xcor, Blue Origin, or PoSSUM, that have advanced suborbital technology. It's cheaper than Virgin Galactic, and you can't do science with eight tourists shouting and floating around.

An extended version of this interview is available online.

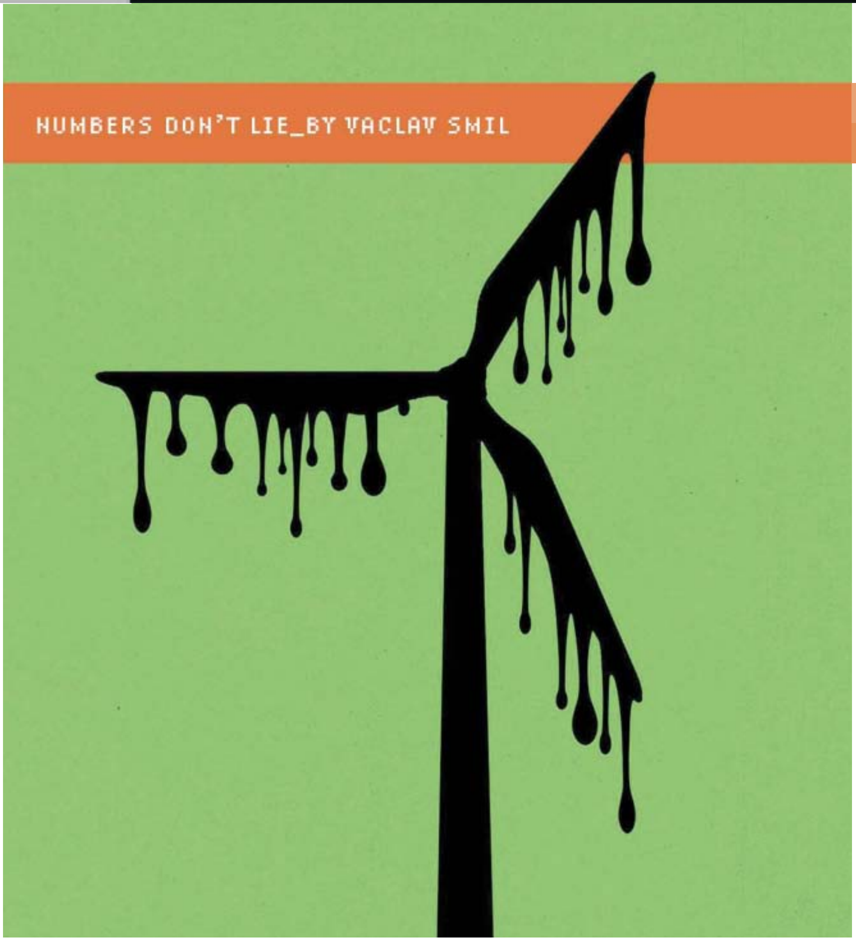

NUMBERS DON'T LIE_BY VACLAV SMIL

OPINION

WHAT I SEE WHEN I SEE A WIND TURBINE

WIND TURBINES ARE THE MOST VISIBLE SYMBOLS OF the quest for renewable electricity generation. And yet, although they exploit the wind, which is as free and as green as energy can be, the machines themselves are pure embodiments of fossil fuels. ● Large trucks bring steel and other raw materials to the site, earth-moving equipment beats a path to otherwise inaccessible high ground, large cranes erect the structures, and all these machines burn diesel fuel. So do the freight trains and cargo ships that convey the materials needed for the production of cement, steel, and plastics. For a 5-megawatt turbine, the steel alone averages 150 metric tons for the reinforced concrete foundations, 250 metric tons for the rotor hubs and nacelles (which house the gearbox and generator), and 500 metric tons for the towers. ● If wind-generated electricity were to supply 25 percent of global demand by 2030 (forecast to reach about 30 petawatt-hours), then even with a high average capacity factor of 35 percent, the aggregate installed wind power of about 2.5 terawatts would require roughly 450 million metric tons of steel. And that's without counting the metal for towers, wires, and transformers for the new high-voltage transmission links that would be needed to connect it all to the grid. ● A lot of energy goes into making steel. Sintered or pelletized iron ore is smelted in blast furnaces, charged with coke made from coal, and receives infusions of powdered coal and natural gas. Pig iron is decarbonized in basic oxygen furnaces. Then steel goes through continuous casting processes (which turn molten steel directly into the rough shape of the final product). Steel used in turbine construction embodies typically about 35 gigajoules per metric ton.

To make the steel required for wind turbines that might operate by 2030, you'd need fossil fuels equivalent to more than 600 million metric tons of coal.

A 5-MW turbine has three roughly 60-meter-long airfoils, each weighing about 15 metric tons. They have light balsa or foam cores and outer laminations made mostly from glass-fiber-reinforced epoxy or polyester resins. The glass is made by melting silicon dioxide and other mineral oxides in furnaces fired by natural gas. The resins begin with ethylene derived from light hydrocarbons, most commonly the products of naphtha cracking, liquefied petroleum gas, or the ethane in natural gas.

The final fiber-reinforced composite embodies on the order of 170 GJ/t. Therefore, to get 2.5 TW of installed wind power by 2030, we would need an aggregate rotor mass of about 23 million metric tons, incorporating the equivalent of about 90 million metric tons of crude oil. And when all is in place, the entire structure must be waterproofed with resins whose synthesis starts with ethylene. Another required oil product is lubricant, for the turbine gearboxes, which has to be changed periodically during the machine's two-decade lifetime.

Undoubtedly, a well-sited and well-built wind turbine would generate as much energy as it embodies in less than a year. However, all of it will be in the form of intermittent electricity—while its production, installation, and maintenance remain critically dependent on specific fossil energies. Moreover, for most of these energies—coke for iron-ore smelting, coal and petroleum coke to fuel cement kilns, naphtha and natural gas as feedstock and fuel for the synthesis of plastics and the making of fiberglass, diesel fuel for ships, trucks, and construction machinery, lubricants for gearboxes—we have no nonfossil substitutes that would be readily available on the requisite large commercial scales.

For a long time to come—until all energies used to produce wind turbines and photovoltaic cells come from renewable energy sources—modern civilization will remain fundamentally dependent on fossil fuels. ■

ILLUSTRATION BY Dan Page

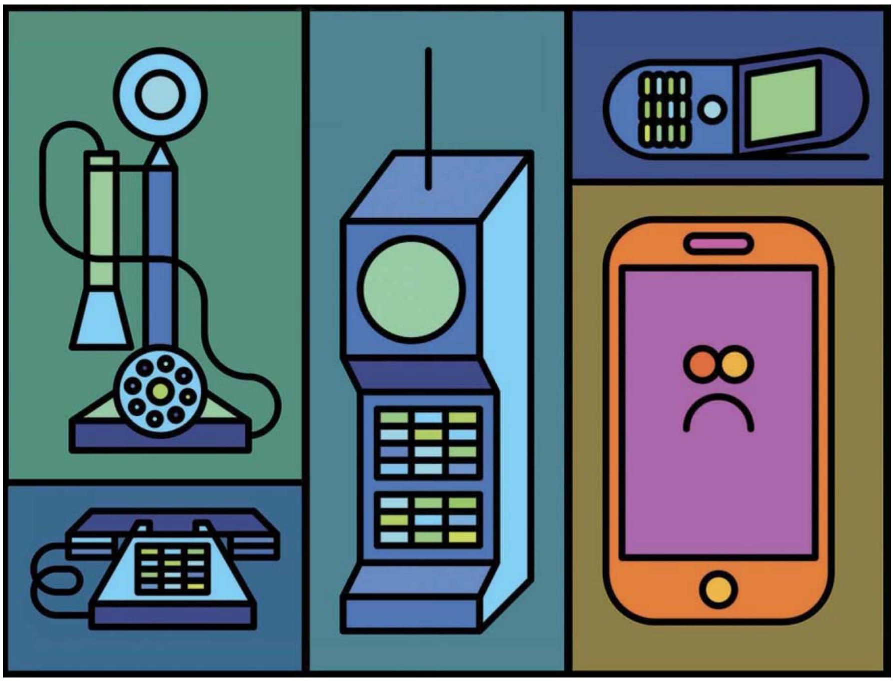

THE END OF THE SMARTPHONE?

It's time to start thinking about what's next

RECENTLY I WAS LISTENING TO A TALK—OSTENSIBLY about the future–in which the smartphone was pictured as an evolutionary end point that would last indefinitely. I thought, *Nothing is forever, and forever comes very quickly in technology*. I started to worry about what would come after the smartphone, and what the end of the smartphone's dominance would mean to the electronics industry. • The smartphone has been an incredible success, both in technological and business terms. I often think of it as the pinnacle of engineering brilliance. There is an entire ecosystem of parts, systems, software, even design philosophy, that has built up around it. Anything that can be made from smartphone technology can be made quickly, easily, and inexpensively. It has been the most successful electronics product ever. • But there are dark clouds on the horizon. The good news is that practically everyone has one. The bad news is that practically everyone has one. Moreover, the new models seem to offer diminishing advances in performance and features that are easily apparent to users. The ads for Apple's iPhone often say that the new model is its best iPhone ever. Well, *duh*, of course. Imagine if the ads said that the new model wasn't quite as good as the previous one! • I remembered my father's old rotary-dial telephone from my youth. He still had it in the 1990s, and several times I suggested that he should get a new telephone. "What for?" he would say. "The new ones do the same thing." And I would think, *He has a point there*.

The standard landline telephone lasted almost a century. In electronics, this probably does qualify as being forever. Meanwhile, other things–turntables, tape decks, cassette players, beepers, and so forth–came and then were suddenly gone as if they had never existed (although turntables are making something of a comeback to cater to the neonostalgic crowd). Perhaps, like my father's phone, the smartphone will endure for decades. But also, like that old telephone, it may be something everyone will have, and there will be no compelling reason to buy a new one.

Certainly, the smartphone can be further improved. We need better battery life, faster and wireless charging, and I'd like to see the camera have optical zoom. But on the other hand, users may not see the need for better displays and faster processors. Meanwhile, the supporting infrastructure will see evident improvements–gigabit wireless speeds, greater coverage, better integration with Wi-Fi, and cheaper service plans. However, my concern here is with the smartphone itself. I think of it as the little engine that could–pulling the entire electronics industry in its frantic rush to keep up with the acceleration required by Moore's Law. What if that little engine that could runs out of steam?

I know that I'm asking unanswerable questions, and my own response is only to propose a dream goal–that we invent another electronic device that everyone wants and needs. Maybe the smartphone will drive the market, but maybe not. Either way, I feel the need for some new device that I just can't do without. I look forward to wanting that new gadget. ■

ILLUSTRATION BY **Greg Mably**

SUPERCOMPUTING'S MONSTER IN THE CLOSET

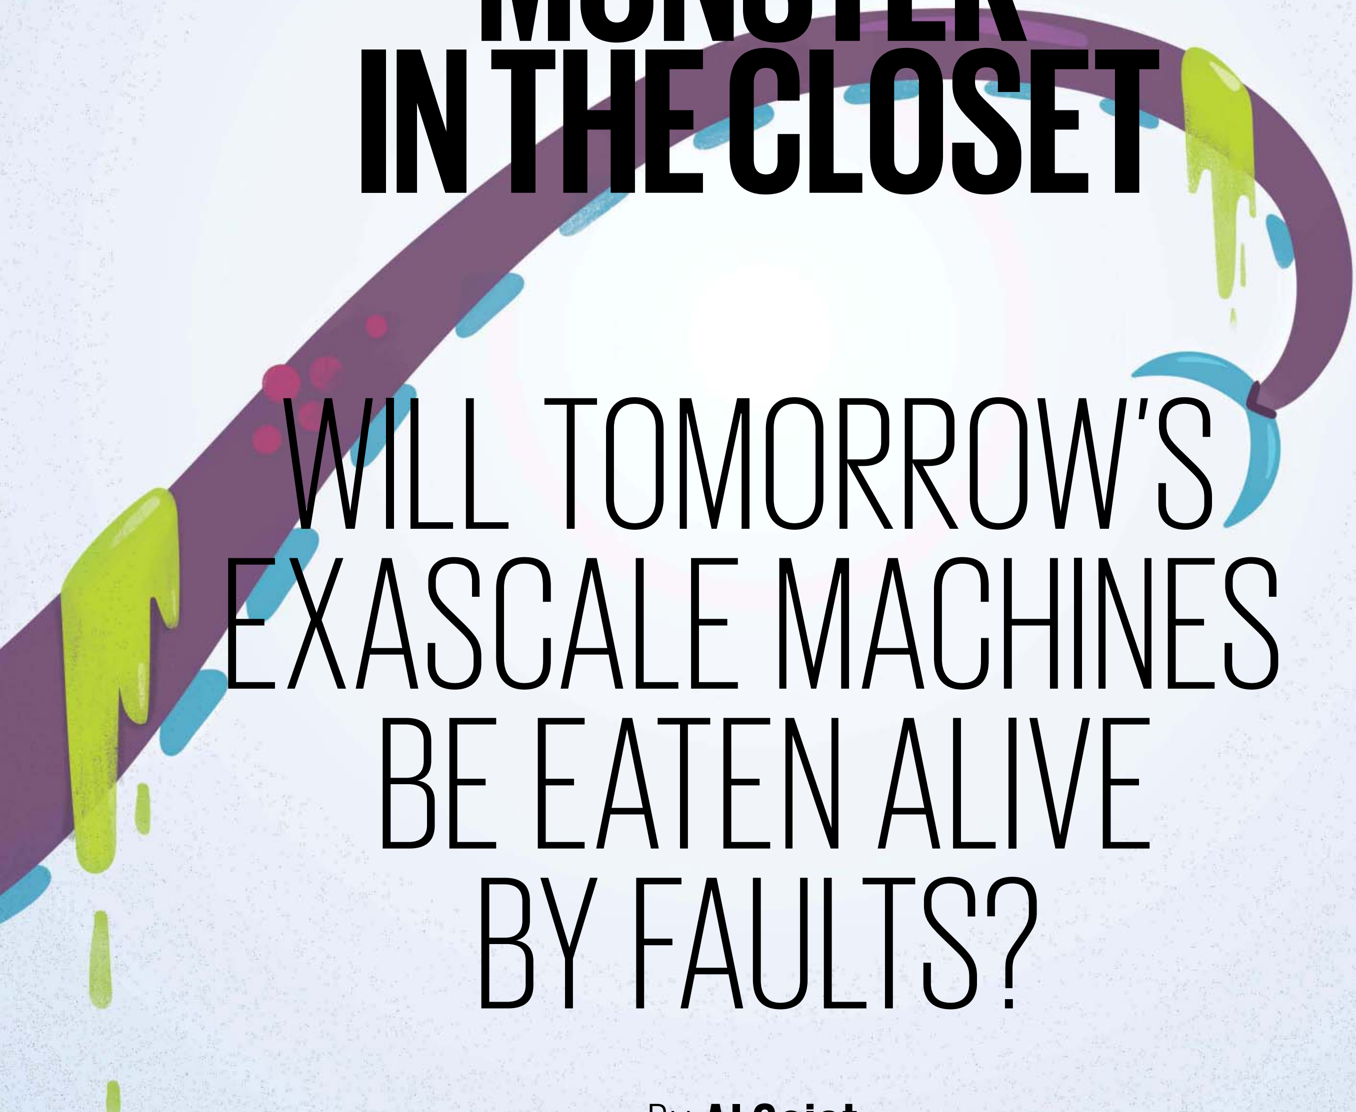

WILL TOMORROW'S EXASCALE MACHINES BE EATEN ALIVE BY FAULTS?

By **Al Geist**

Illustrations by **Shaw Nielsen**

As a child, were you ever afraid that a monster lurking in your bedroom would leap out of the dark and get you? My job at Oak Ridge National Laboratory is to worry about a similar monster, hiding in the steel cabinets of the supercomputers and threatening to crash the largest computing machines on the planet. ¶ The monster is something supercomputer specialists call resilience—or rather the lack of resilience. It has bitten several supercomputers in the past. A high-profile example affected what was the second fastest supercomputer in the world in 2002, a machine called ASCI Q at Los Alamos National Laboratory. When it was first installed at the New Mexico lab, this computer couldn't run more than an hour or so without crashing.

The ASCI Q was built out of AlphaServers, machines originally designed by Digital Equipment Corp. and later sold by Hewlett-Packard Co. The problem was that an address bus on the microprocessors found in those servers was unprotected, meaning that there was no check to make sure the information carried on these within-chip signal lines did not become corrupted. And that's exactly what was happening when these chips were struck by cosmic radiation, the constant shower of particles that bombard Earth's atmosphere from outer space.

To prove to the manufacturer that cosmic rays were the problem, the staff at Los Alamos placed one of the servers in a beam of neutrons, causing errors to spike. By putting metal side panels on the ASCI Q servers, the scientists reduced radiation levels enough to keep the supercomputer running for 6 hours before crashing. That was an improvement, but still far short of what was desired for running supercomputer simulations.

An even more dramatic example of cosmic-radiation interference happened at Virginia Tech's Advanced Computing facility in Blacksburg. In the summer of 2003, Virginia Tech researchers built a large supercomputer out of 1,100 Apple Power Mac G5 computers. They called it Big Mac. To their dismay, they found that the failure rate was so high it was nearly impossible even to boot the whole system before it would crash.

The problem was that the Power Mac G5 did not have error-correcting code (ECC) memory, and cosmic ray–induced particles were changing so many values in memory that out of the 1,100 Mac G5 computers, one was always crashing. Unusable, Big Mac was broken apart into individual G5s, which were sold one by one online. Virginia Tech replaced it with a supercomputer called System X, which had ECC memory and ran fine.

Cosmic rays are a fact of life, and as transistors get smaller, the amount of energy it takes to spontaneously flip a bit gets smaller, too. By 2023, when exascale computers—ones capable of performing 10^{18} operations per second—are predicted to arrive in the United States, transistors will likely be a third the size they are today, making them that much more prone to cosmic ray–induced errors. For this and other reasons, future exascale computers will be prone to crashing much more frequently than today's supercomputers do. For me and others in the field, that prospect is one of the greatest impediments to making exascale computing a reality.

JUST HOW MANY SPURIOUS BIT flips are happening inside supercomputers already? To try to find out, researchers performed a study in 2009 and 2010 on the then most powerful supercomputer—a Cray XT5 system at Oak Ridge, in Tennessee, called Jaguar.

Jaguar had 360 terabytes of main memory, all protected by ECC. I and others at the lab set it up to log every time a bit was flipped incorrectly in main memory. When I asked my computing colleagues elsewhere to guess how often Jaguar saw such a bit spontaneously change state, the typical estimate was about a hundred times a day. In fact, Jaguar was logging ECC errors at a rate of 350 per minute.

In addition to the common case of a single cosmic ray flipping a single bit, in some cases a single high-energy particle cascaded through the memory chip flipping multiple bits. And in a few cases the particle had enough energy to permanently damage a memory location.

ECC can detect and correct a single-bit error in one word of memory (typically 64 bits). If two bits are flipped in a word, ECC can detect that the word is corrupted, but cannot fix it. The study found that double-bit errors occurred about once every 24 hours in Jaguar's 360 TB of memory.

The surface area of all the silicon in a supercomputer functions somewhat like a large cosmic-ray detector. And as that surface area grows, the number of cosmic-ray strikes also grows. Exascale systems are projected to have up to 100 petabytes of memory—50 times as much as today's supercomputers—resulting in that much more real estate for a cosmic-ray particle to hit.

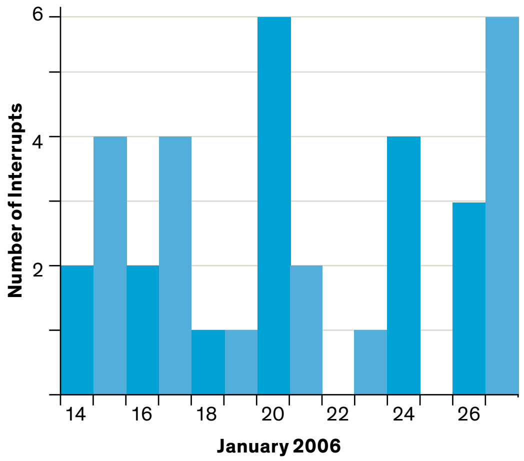

FAILURE NOT OPTIONAL: Modern supercomputers are so large that failures are expected to occur regularly. In 2006, the Red Storm supercomputer at Sandia National Laboratories typically suffered a handful of system interruptions each day, for example.

DATA SOURCE: LOS ALAMOS NATIONAL LABORATORY

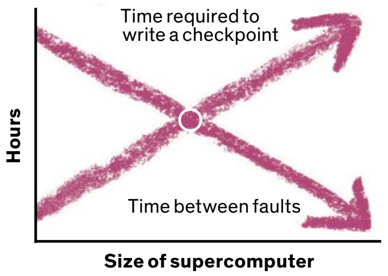

A LOOMING CRISIS: As systems get larger, the time it takes to save the state of memory will exceed the time between failures, making it impossible to use the previous "checkpoint" to recover from errors.

But resilience is not all about bit flips and cosmic rays. Even the simplest components can cause problems. The main resilience challenge for Jaguar was a voltage-regulator module. There were 18,688 of them, and whenever one failed, a board carrying two of the machine's 37,376 hex-core processors powered off.

Two lost processors wasn't the issue—Jaguar would automatically detect the malfunction and reconfigure the system to work without the problematic board. But that board also contained a network-communication chip, which all other such boards in the system depended on to route messages. When this board powered down, the system would continue to run a while, but it would eventually hang, requiring a reboot of the entire supercomputer to reset all the board-to-board routing tables. While today's supercomputers do dynamic routing to avoid such failures, the growing complexity of these computing behemoths is increasing the chances that a single fault will cascade across the machine and bring down the entire system.

Supercomputer operators have had to struggle with many other quirky faults as well. To take one example: The IBM Blue Gene/L system at Lawrence Livermore National Laboratory, in California, the largest computer in the world from 2004 to 2008, would frequently crash while running a simulation or produce erroneous results. After weeks of searching, the culprit was uncovered: the solder used to make the boards carrying the processors. Radioactive lead in the solder was found to be causing bad data in the L1 cache, a chunk of very fast memory meant to hold frequently accessed data. The workaround to this resilience problem on the Blue Gene/L computers was to reprogram the system to, in essence, bypass the L1 cache. That worked, but it made the computations slower.

SO THE WORRY IS NOT THAT THE monster I've been discussing will come out of the closet. It's already out. The people who run the largest supercomputers battle it every day. The concern, really, is that the rate of faults it represents will grow exponentially, which could prevent future supercomputers from running long enough for scientists to get their work done.

Several things are likely to drive the fault rate up. I've already mentioned two: the growing number of components and smaller transistor sizes. Another is the mandate to make tomorrow's exascale supercomputers at least 15 times as energy efficient as today's systems.

To see why that's needed, consider the most powerful supercomputer in the United States today, a Cray XK7 machine at Oak Ridge called Titan. When running at peak speed, Titan uses 8.2 megawatts of electricity. In 2012, when it was the world's most powerful supercomputer, it was also the third most efficient in terms of floating-point operations per second (flops) per watt. Even so, scaled up to exaflop size, such hardware would consume more than 300 MW—the output of a good-size power plant. The electric bill to run such a supercomputer would be about a third of a billion dollars per year.

No wonder then that the U.S. Department of Energy has announced the goal of building an exaflop computer by 2023 that consumes only 20 MW of electricity. But reducing power consumption this severely could well compromise system resilience. One reason is that


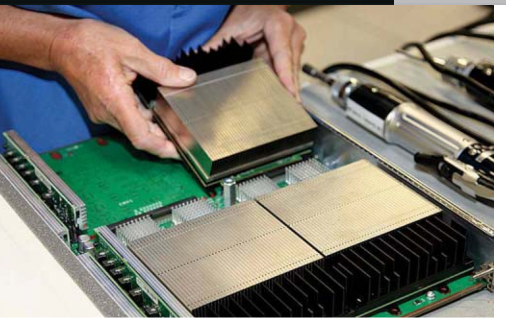

REDUCE, REUSE, RECYCLE: When your supercomputer starts showing its age, you have to do something or else the cost of the electricity to run it won't be worth the results you obtain. But that doesn't mean you need to throw it out. In 2011 and 2012, Oak Ridge National Laboratory upgraded its Jaguar supercomputer, first installed in 2005, transforming it into a far more capable machine called Titan [see table below]. The effort, as shown in these photos, was extensive, but it made Titan No. 1 in the world for a time.

OAK RIDGE NATIONAL LABORATORY (6)

the power savings will likely have to come from smaller transistors running at lower voltages to draw less power. But running right at the edge of what it takes to make a transistor switch on and off increases the probability of circuits flipping state spontaneously.

Further concern arises from another way many designers hope to reduce power consumption: by powering off every unused chip, or every circuit that's not being used inside a chip, and then turning them on quickly when they're needed. Studies done at the University of Michigan in 2009 found that constant power cycling reduced a chip's typical lifetime up to 25 percent.

Power cycling has a secondary effect on resilience because it causes voltage fluctuations throughout the system—much as a home air conditioner can cause the lights to dim when it kicks on. Too large of a voltage fluctuation can cause circuits to switch on or off spontaneously inside a computer.

Using a heterogeneous architecture, such as that of Titan, which is composed of AMD multicore CPUs and Nvidia GPUs (graphics processing units), makes error detection and recovery even harder. A GPU is very efficient because it can run hundreds of calculations simultaneously, pumping huge amounts of data through it in pipelines that are hundreds of clock cycles long. But if an error is detected in just one of the calculations, it may require waiting hundreds of cycles to drain the pipelines on the GPU before beginning recovery, and all of the calculations being performed at that time may need to be rerun.

SO FAR I'VE DISCUSSED HOW HARD it will be to design supercomputer hardware that is sufficiently reliable. But the software challenges are also daunting. To understand why, you need to know how today's supercomputer simulations deal with faults. They periodically record the global state of the supercomputer, creating what's called a checkpoint. If the computer crashes, the simulation

Machine Makeover: A supercomputer is reborn

	Jaguar, XT5 partition (2011)	Titan (2012)
Compute nodes	18,688	18,688
Log-in and I/O nodes	256	512
Memory per node	16 gigabytes	32 GB + 6 GB
Number of Opteron CPU cores	224,256	299,008
Number of Nvidia GPU accelerators	0	18,688
Total system memory	300 terabytes	710 TB
Total system peak performance	2.3 petaflops	20+ petaflops

Source: AnandTech, 31 October 2012

can then be restarted from the last valid checkpoint instead of beginning some immense calculation anew.

This approach won't work indefinitely, though, because as computers get bigger, the time needed to create a checkpoint increases. Eventually, this interval will become longer than the typical period before the next fault. A challenge for exascale computing is what to do about this grim reality.

Several groups are trying to improve the speed of writing checkpoints. To the extent they are successful, these efforts will forestall the need to do something totally different. But ultimately, applications will have to be rewritten to withstand a constant barrage of faults and keep on running.

Unfortunately, today's programming models and languages don't offer any mechanism for such dynamic recovery from faults. In June 2012, members of an international forum composed of vendors, academics, and researchers from the United States, Europe, and Asia met and discussed adding resilience to message-passing interface, or MPI, the programming model used in nearly all supercomputing code. Those present at that meeting voted that the next version of MPI would have no resilience capabilities added to it. So for the foreseeable future, programming models will continue to offer no methods for notification or recovery from faults.

One reason is that there is no standard that describes the types of faults that the software will be notified about and the mechanism for that notification. A standard fault model would also define the actions and services available to the software to assist in recovery. Without even a de facto fault model to go by, it was not possible for these forum members to decide how to augment MPI for greater resilience.

So the first order of business is for the supercomputer community to agree on a standard fault model. That's more difficult than it sounds because some faults might be easy for one manufacturer to deal with and hard for another. So there are bound to be fierce squabbles. More important, nobody really knows what problems the fault model should address. What are all the possible errors that affect today's supercomputers? Which are most common? Which errors are most concerning? No one yet has the answers.

And while I've talked a lot about faults causing machines to crash, these are not, in fact, the most dangerous. More menacing are the errors that allow the application to run to the end and give an answer that looks correct but is actually wrong. You wouldn't want to fly in an airliner designed using such a calculation. Nor would you want to certify a new nuclear reactor based on one. These undetected errors–their types, rates, and impact–are the scariest aspect of supercomputing's monster in the closet.

GIVEN ALL THE GLOOM AND DOOM I've shared, you might wonder: How can an exascale supercomputer ever be expected to work? The answer may lie in a handful of recent studies for which researchers purposely injected different types of errors inside a computer at random times and locations while it was running an application. Remarkably enough, 90 percent of those errors proved to be harmless.

One reason for that happy outcome is that a significant fraction of the computer's main memory is usually unused. And even if the memory is being used, the next action on a memory cell after the bit it holds is erroneously flipped may be to write a value to that cell. If so, the earlier bit flip will be harmless. If instead the next action is to read that memory cell, an incorrect value flows into the computation. But the researchers found that even when a bad value got into a computation, the final result of a large simulation was often the same.

Errors don't, however, limit themselves to data values: They can affect the machine instructions held in memory, too. The area of memory occupied by machine instructions is much smaller than the area taken up by the data, so the probability of a cosmic ray corrupting an instruction is smaller. But it can be much more catastrophic. If a bit is flipped in a machine instruction that is then executed, the program will most likely crash. On the other hand, if the error hits in a part of the code that has already executed, or in a path of the code that doesn't get executed, the error is harmless.

There are also errors that can occur in silicon logic. As a simple example, imagine that two numbers are being multiplied, but because of a transient error in the multiplication circuitry, the result is incorrect. How far off it will be can vary greatly depending on the location and timing of the error.

As with memory, flips that occur in silicon logic that is not being used are harmless. And even if this silicon is being used, any flips that occur outside the narrow time window when the calculation is taking place are also harmless. What's more, a bad multiplication is much like a bad memory value going into the computation: Many times these have little or no affect on the final result.

So many of the faults that arise in future supercomputers will no doubt be innocuous. But the ones that do matter are nevertheless increasing at an alarming rate. So the supercomputing community must somehow address the serious hardware and software challenges they pose. What to do is not yet clear, but it's clear we must do something to prevent this monster from eating us alive. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/resilience0316

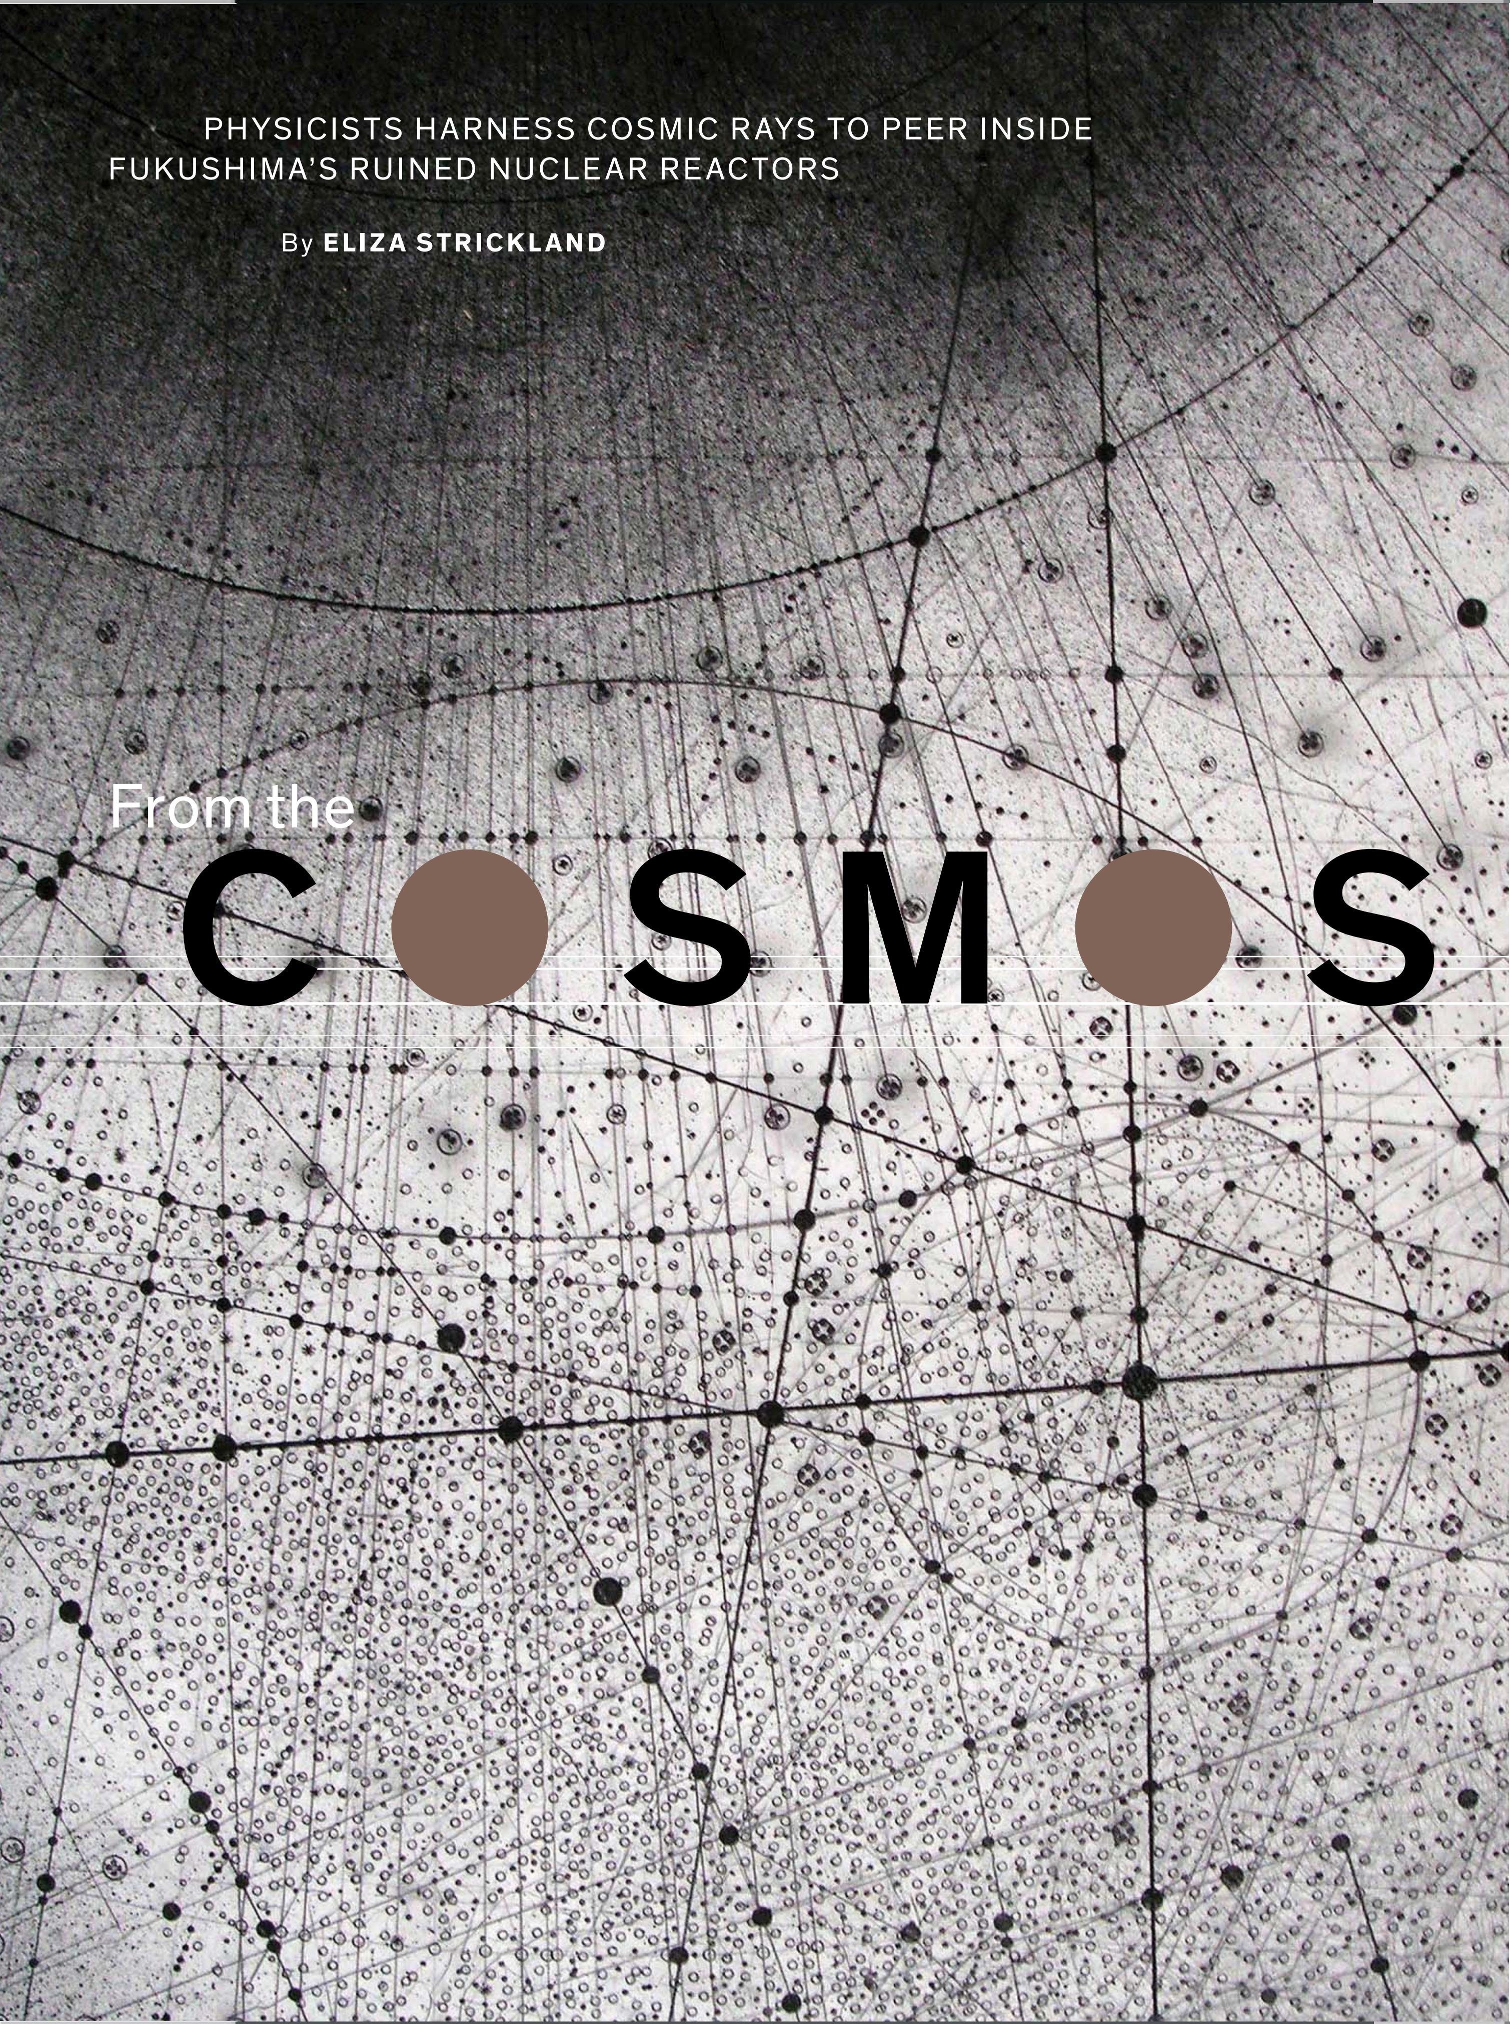

PHYSICISTS HARNESS COSMIC RAYS TO PEER INSIDE FUKUSHIMA'S RUINED NUCLEAR REACTORS

By **ELIZA STRICKLAND**

From the COSMOS

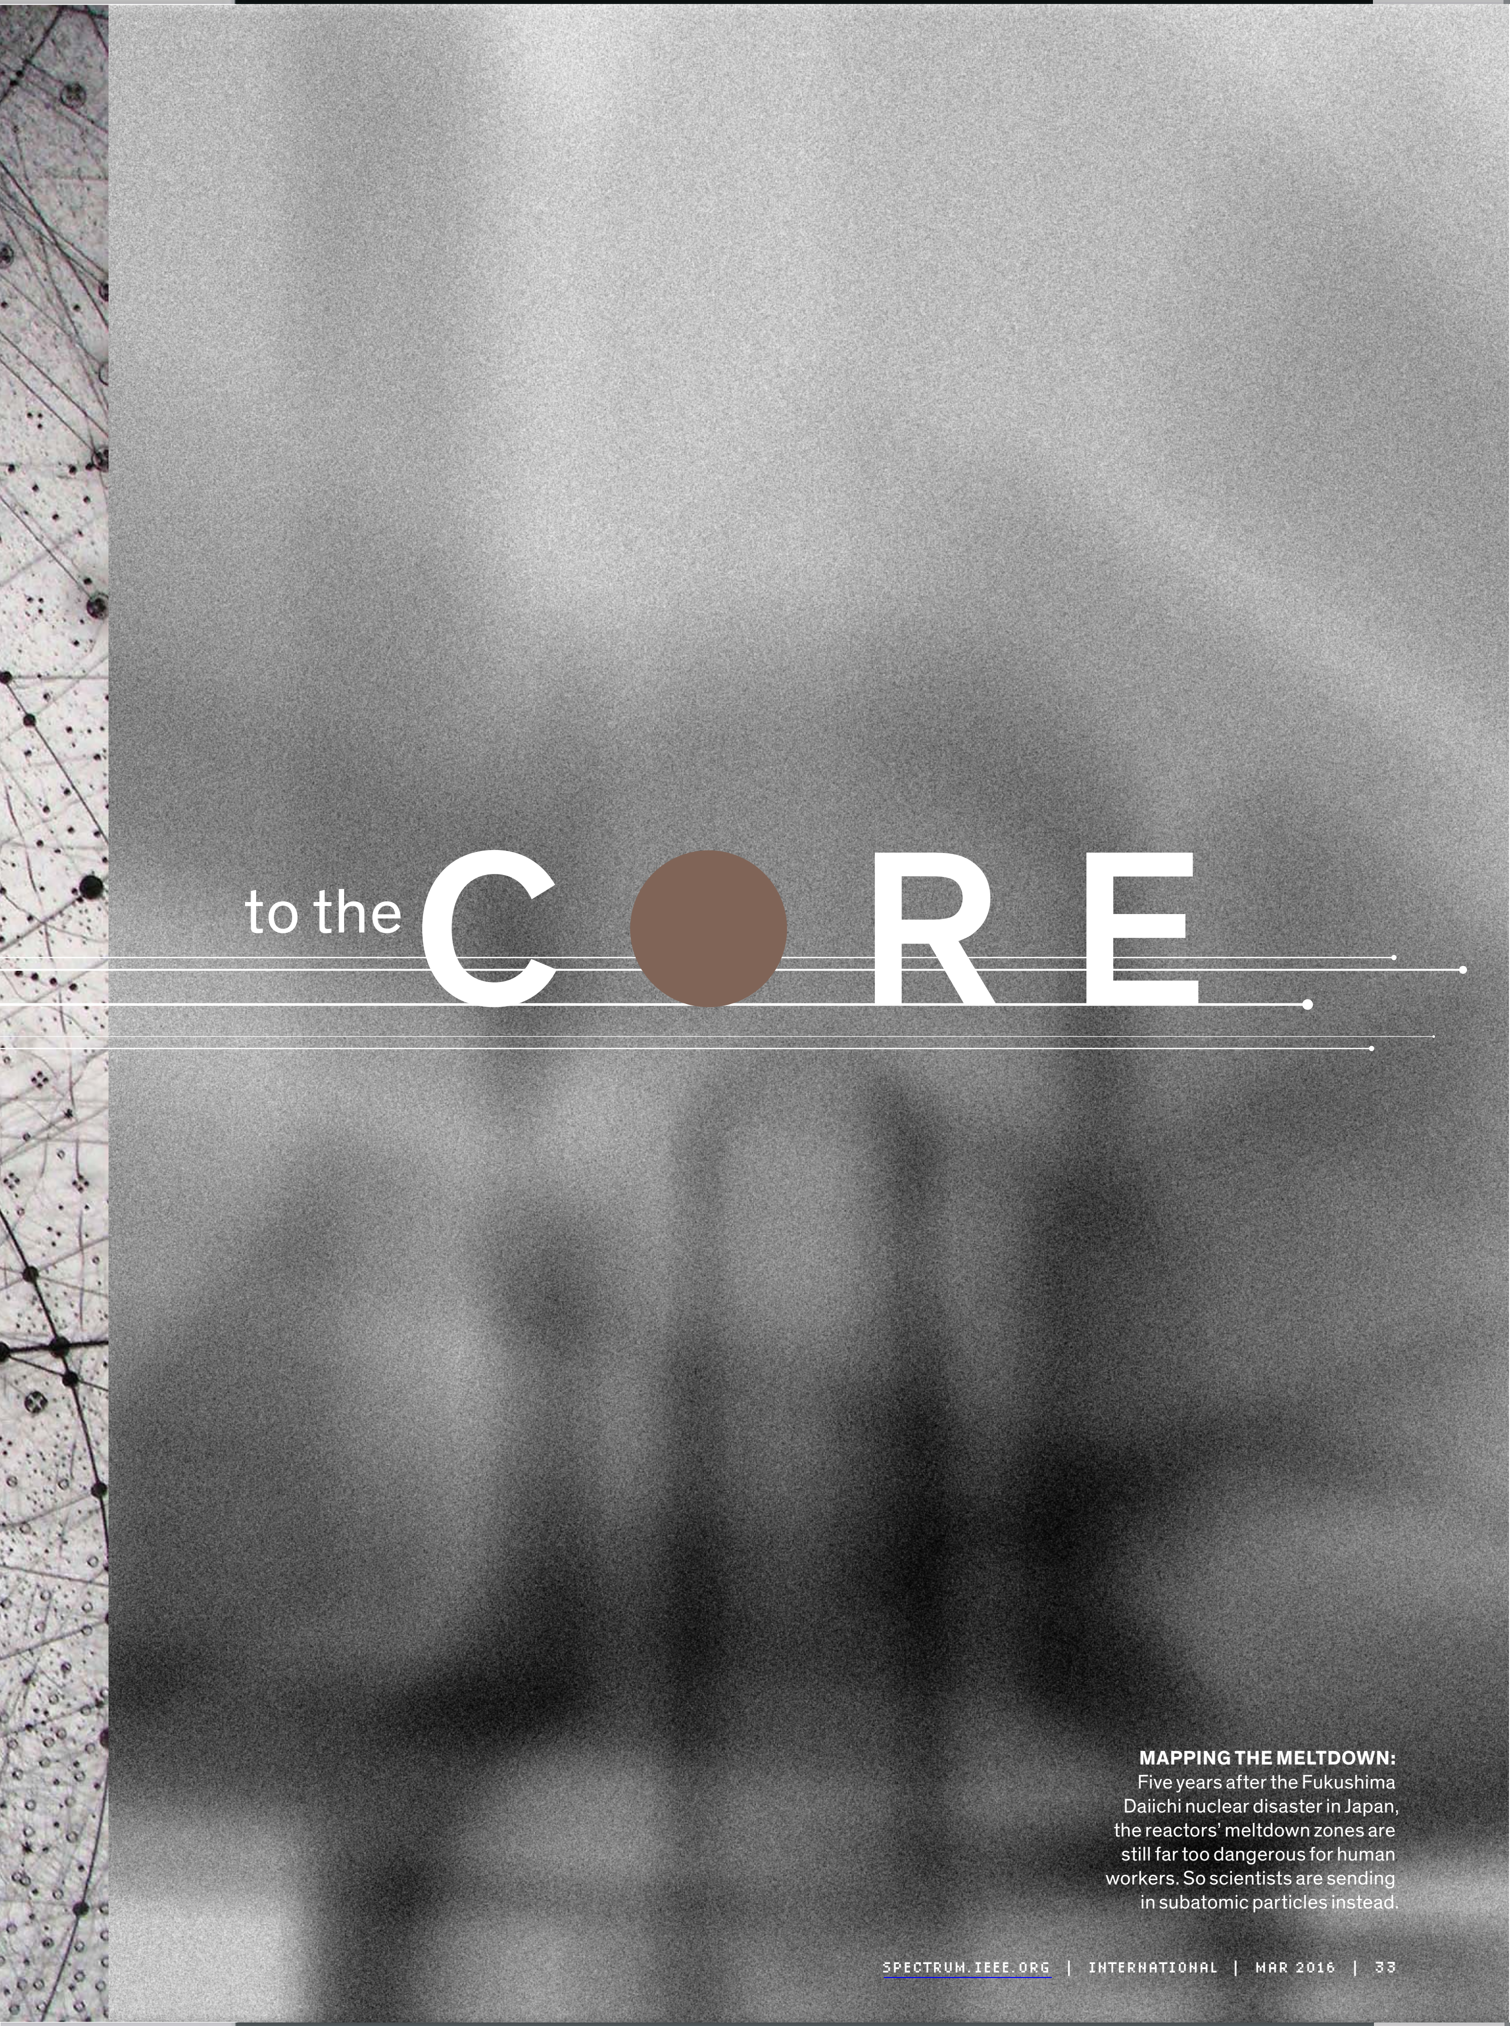

to the CORE

MAPPING THE MELTDOWN: Five years after the Fukushima Daiichi nuclear disaster in Japan, the reactors' meltdown zones are still far too dangerous for human workers. So scientists are sending in subatomic particles instead.

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

THEY COME FROM OUTER SPACE.

Some are born within the cozy confines of our solar system, surging forth when our sun flares up, as it routinely does, with geysers of plasma. Others have traveled from unfathomably distant reaches beyond our Milky Way galaxy, where stars at the end of their lives went supernova with mighty and sustained blasts.

They are cosmic rays: streams of electrically charged subatomic particles that perpetually bombard Earth. When they hit the thick blanket of gas that surrounds and protects our planet, they collide with atoms and split into even tinier fragments that rain down to the ground. Physicists call these assortments of particles "air showers."

One plentiful particle in these showers is the energetic muon. And here on Earth these infinitesimal specks of matter might be the answer to a very big problem called Fukushima.

FIVE YEARS AGO THIS MONTH, the Japanese nuclear power plant Fukushima Daiichi was torn apart by meltdowns or explosions in four reactors. *IEEE Spectrum* has previously told the story of that horrific accident, the second worst in nuclear history: Shifting tectonic plates in the Pacific Ocean triggered a massive tsunami that battered Japan's coast and inundated the nuclear plant, knocking out the power systems required to pump cooling water over the reactor cores. Deprived of coolant, the fuel rods in the cores of three of the plant's six reactors melted down.

Workers struggled courageously for months to stabilize the wrecked facility, trying to establish a continuous flow of water into the reactor vessels to keep the melted cores cool—no easy task when the perforated vessels were leaking like sieves. It took four months, until July 2011, for Tokyo Electric Power Co. (TEPCO), the utility company that runs the plant, to install the pumps and water recycling systems necessary for the government to declare the plant "stable." In December 2011, the authorities agreed that the plant had achieved "cold shutdown," meaning that reactor conditions made it impossible for nuclear fission reactions to start up again in the cores.

Those milestones may make it sound as if the Fukushima crisis has been resolved. It hasn't been. Should the circulation systems at Fukushima Daiichi break down tomorrow, water levels in the leaky reactor vessels would drop, exposing the nuclear fuel. In a flash, Japan would again be in the midst of a nuclear disaster.

To eliminate such a possibility, TEPCO has pledged to remove all nuclear fuel, including every drop of the melted stuff, and to completely dismantle the power plant. The company estimates this mammoth decommissioning project will likely take 40 years. Why such a long timeline? Because some of Fukushima's challenges require the invention, design, testing, and implementation of entirely new technologies.

The main problem lies in and under the three reactors that experienced meltdowns. (A fourth reactor building was shattered by an explosion, but its nuclear fuel had been removed earlier for routine maintenance and wasn't damaged.) During the first terrible days of the accident, the uranium fuel rods inside Fukushima Daiichi's reactors 1, 2, and 3 overheated, slumped, oozed, and finally dripped down to the bottom of the reactor vessels. Some of this radioactive sludge probably melted through those 14-centimeter-thick steel vessels and spilled onto concrete pads below, making the cleanup task much more formidable.

Or maybe it didn't—and that's the issue. To this day, TEPCO doesn't know exactly where this deadly material ended up. The monitoring equipment inside the reactors was destroyed in the accident, and workers can't go near the crucial areas to conduct inspections. TEPCO hasn't yet been able to take measurements in the most dangerous areas, and has already found hot spots around those concrete pads where radiation reaches the lethal level of 5 sieverts per hour. TEPCO's technology partners are busily developing robots that can go in to survey the scene, but rubble-filled buildings with narrow hatches and stairs make it tough for bots to get around.

Lake Barrett knows a lot about cleaning up a nuclear disaster: He managed the decommissioning of the Three Mile Island reactor after its 1979 accident, the worst nuclear accident the United States has ever seen. He's now a consultant for TEPCO and regularly flies to Japan to advise the company. What will it take to defuel Fukushima's reactors? "It starts with investigative diagnostics," Barrett says. "They'll make fundamental engineering decisions based on the diagnostics." He notes that TEPCO's timetable calls for the company to decide on a defueling procedure in 2017—an optimistic goal, he says, given the current state of knowledge. "We're coming up on five years and we don't have a lot of information yet," he explains.

Until TEPCO figures out the exact location of every radioactive glob of melted metal, the overseers can't make a plan to remove the fuel from the reactors and proceed with the decommissioning. They're stymied. And that's where the muons come in. An imaging technology that makes use of these humble particles may be the investigators' best hope for getting a look at the lethal hearts of Fukushima's ruined reactors.

THE CRUCIAL QUESTION: To clean up and dismantle the Fukushima Daiichi nuclear reactors, experts must first figure out the exact locations of all the melted globs of nuclear fuel.

PREVIOUS PAGES: LEFT: EMMA MCNALLY; RIGHT: IRID/TEPCO. THIS PAGE: PALLAVA BAGLA/CORBIS

TWO WAYS TO LOOK INSIDE

Subatomic particles called muons can penetrate all matter, including Fukushima's concrete and steel reactor buildings and its dense uranium fuel. By analyzing the muons' flight paths, scientists can map the substances they passed through.

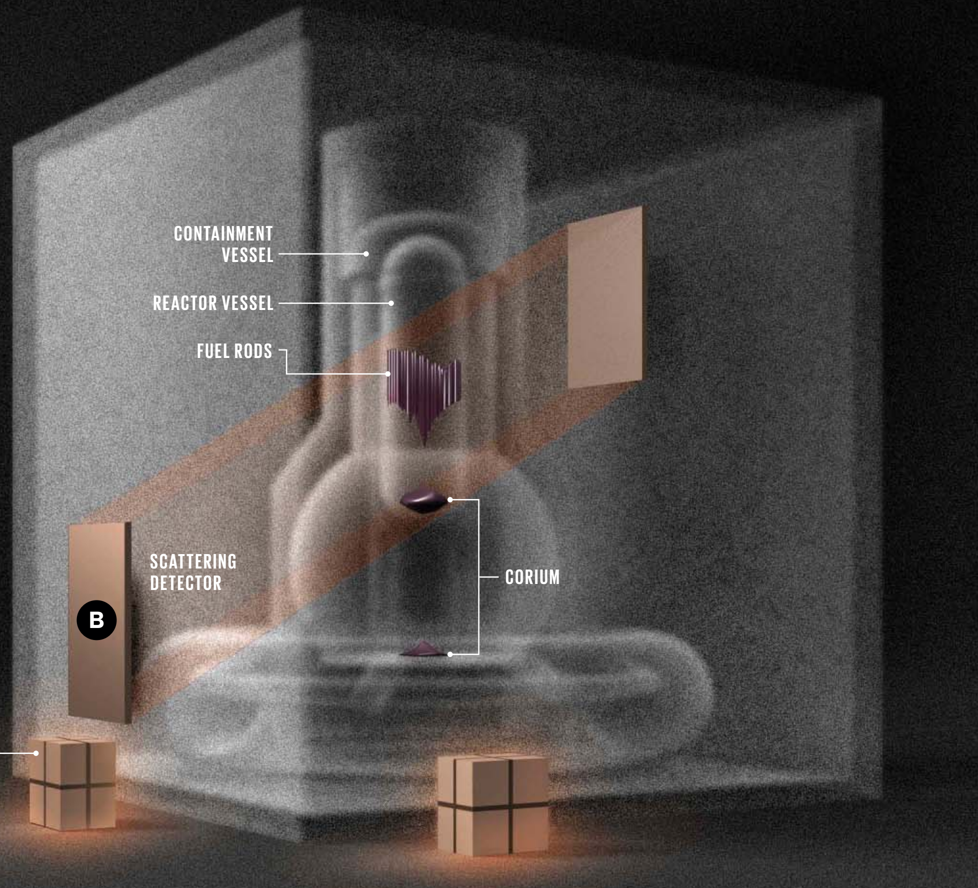

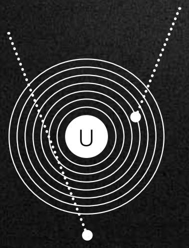

A MUON TRANSMISSION As it passes through an atom, a muon can change an electron's orbit and lose energy in the process. If it's a low-energy muon, that loss may stop the traveling particle in its tracks. Atoms of bulky elements like uranium have more electrons and therefore stop more muons. The detectors count incoming muons and construct an image of the materials they encountered during their flight.

B MUON SCATTERING The whizzing muon can also be attracted or repelled by an atom's nucleus, which alters its trajectory. A uranium atom's hefty nucleus exerts a stronger force on the muon than do the nuclei of other substances. Two detectors will record the paths of muons before and after they pass through a Fukushima reactor vessel, and will use the scattering pattern to produce an image of the materials inside.

A BLAND OFFICE PARK in the suburbs of San Diego seems an unlikely place to encounter ardent fans of subatomic particles. But inside the offices of Decision Sciences, an 11-year-old company headquartered there, you can't throw a piece of chalk without hitting a physics Ph.D. CEO Gene Ray, for example, picked up a doctorate in theoretical physics before embarking on a 40-year business career. "But I haven't used it much," he demurs.

So when it comes time to talk cosmic rays, the CEO happily yields to his vice president of analytics, Konstantin Borozdin, who trained in experimental nuclear physics in Moscow and worked for a decade as an astrophysicist at Russia's Space Research Institute. The muons produced from cosmic rays, Borozdin explains, aren't usually objects of interest for theoretical physicists. The muons that shower the Earth from the upper atmosphere don't interact strongly with other particles, so they pass right through most matter. Look at your hand for one second; statistics say a muon has whizzed through it.

For most physicists, they're like static in the signal. "In the underground detectors that look for very rare particles, muons are a nuisance," Borozdin says. "They try to shield from them."

Borozdin got interested in muons when 9/11 blew away America's peaceful haze. By that time he was on staff at Los Alamos National Laboratory, in New Mexico, investigating cosmic rays. Los Alamos, the storied birthplace of the atomic bomb, received a special call from the U.S. government after the 9/11 attack. "They asked us to look for technologies that could address nuclear terrorism," Borozdin remembers. "They wanted ways to detect nuclear weapons that someone might try to smuggle into the country." Borozdin and his colleagues realized that those nuisance particles, the ubiquitous muons, could be used in imaging systems that would spot nuclear materials.

ILLUSTRATION BY Emily Cooper

To explain the principles at work, let's start with high school physics. Recall that an atom is mostly empty space, with a hefty nucleus surrounded by a cloud of diminutive electrons. Most of the time, a passing muon streaks through the atom's space unimpeded, but occasionally it interacts with the atom's constituent particles. In one kind of interaction, the muon, which carries either a positive or negative electric charge, either pulls or pushes a negatively charged electron out of its normal orbit. That process causes the muon to lose a few electron volts of energy and slow down.

Some decades back, physicists came up with a basic imaging method to take advantage of these electron interactions, focusing on low-energy muons that lose enough energy to halt them in their tracks. This technique got an exotic tryout in 1968, when a team of physicists installed a detector inside the base of an Egyptian pyramid. The detector counted incoming muons and traced their paths back through the pyramid to determine whether they'd traveled through solid limestone, where densely packed atoms would cause more low-energy muons to stop, or the air of a secret chamber, which would allow more to whiz through. After several months of analyzing the data, the physicists got their answer: no hidden chamber (and hence no hidden treasure).

But there's another and more subtle signal that can be teased out of the interplay of particles. If a speeding muon passes close enough to the nucleus, which carries a positive charge, the muon is either repelled or attracted. That interaction alters the direction of the muon's flight. You may remember that each element in the periodic table is defined by the number of protons in its nucleus; helium has 2 protons, for example, while uranium has 92. Now imagine two muons, one flying through a helium-filled balloon and another zipping through a chunk of uranium ore. The muon passing through uranium experiences a stronger repulsion or attraction from the hefty nucleus, and is pushed or pulled farther off course.

Borozdin's Los Alamos group saw potential here. By analyzing how muons scattered after passing through a material, they could determine what that material was. In a breakthrough 2003 report, published in *Nature*, they suggested that detectors based on muon scattering could, for example, distinguish "a block of uranium concealed inside a truck full of sheep."

Decision Sciences was soon formed to turn the concept into a commercial technology. Borozdin helped the company develop its scanning system, and he officially joined the staff in 2014. The system uses two detectors: one above the object of interest to register the incoming muons' trajectory, and one below to record the scattering. Decision Sciences began marketing its scanners to cargo ports and borders that need to screen for nuclear materials, extolling the system's superiority to existing technologies that use X-rays or gamma rays: Muon-imaging systems can look through thick layers of lead and steel, and they don't make use of a dangerous radiation source.

Then, in March 2011, Fukushima Daiichi melted down.

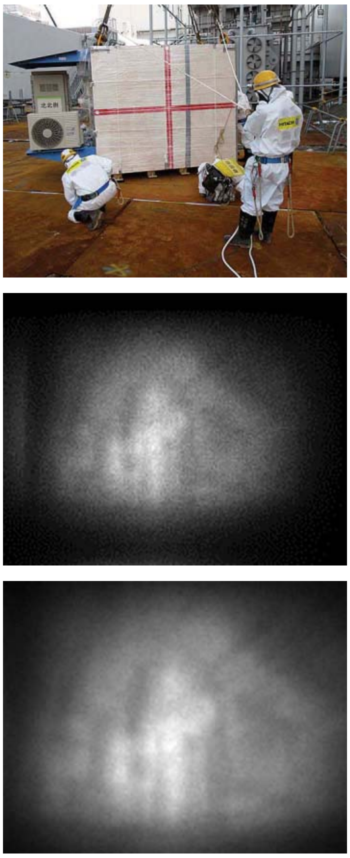

FIRST TRYOUT: The muon transmission detectors installed next to reactor 1 [top photo] produced shadowy images of the building's interior after collecting data for one month and three months [bottom two images].

IRID/TEPCO (3)

AFTER BRINGING A MEASURE of stability to the crippled reactors, TEPCO has been striving to characterize the mess they have become. Roboticists at Japan's major tech companies are developing unique shape-shifting bots that can survive the site's lethal radiation and navigate the reactor buildings' complex layouts. Last year TEPCO sent the first robots into the containment vessels, the concrete and steel structures that surround the reactor vessels.

TEPCO has deemed these robot excursions a success (including one in which a bot got stuck inside), and plans to continue them. But no robot has yet approached the centers of the containment vessels, where thick concrete pads sit directly under the reactor vessels. If nuclear fuel melted through the reactor vessels (a near certainty in reactors 1 and 3), that's where it landed.

Barrett, TEPCO's decommissioning consultant, says critical questions must be answered about the breach of the reactor vessels: "Was it a slower flow, or a catastrophic blowout? How much water was on the floor of the containment vessel when it breached through? Did it flow horizontally, like the lava flow from volcanoes?" What's more, in the superheated conditions during the meltdown, the nuclear fuel must have mixed with other disintegrating reactor components to form an unpredictable substance that nuclear specialists wryly refer to as corium, he says. "Uranium, steel, lead, debris, concrete, handrails, electric cables, aluminum–it's going to be a real agglomeration."

Muon imaging may be the best way to look inside Fukushima's deadliest spaces and find answers. In 2011, two teams of physicists came to this conclusion and started building muon detectors specifically for Fukushima Daiichi. Japanese researchers at the High Energy Accelerator Research Organization (known as KEK) constructed a system that uses muon transmission imaging, the same type of system used in the Egyptian pyramid. In a separate effort, Toshiba,

PROTECTING PORTS: At its R&D center, Decision Sciences demonstrates how its muon-scattering imaging system can be used to screen for smuggled nuclear materials at borders.

DECISION SCIENCES

a major contractor for TEPCO, asked Decision Sciences to build an instrument based on its muon-scattering system.

In February 2015, TEPCO debuted the KEK system at reactor 1, setting up the detector outside the reactor building and letting it collect muon data for about 90 days. A TEPCO spokesman says the images verified the company's working assumptions about the behavior of reactor 1's fuel load–namely, that all of it melted and flowed downward. The problem is that the images reveal the interior of the reactor building only in broad strokes. Barrett says it's hard to imagine that these fuzzy images will help engineers make serious decisions about defueling operations. "The data coming out of it is pretty crude," he says. Nevertheless, TEPCO plans to use the KEK system to study reactor 2 soon.

The Decision Sciences system currently sits in storage at a Toshiba research center in Japan, where it has been calibrated and could be used whenever TEPCO wishes. The system has already been tested on a small-scale mock-up of a Fukushima reactor at Los Alamos, and on a working research reactor at the Toshiba facility. The results of those tests, say Borozdin and Ray, bode well for the real trial of the technology at reactor 2, which TEPCO initially scheduled for mid-2015.

Yet the device wasn't tested at Fukushima then, and instead still waits in storage. Surprisingly, TEPCO won't specify a timeline for its use. As this article goes to press, a TEPCO spokesman would say only that the KEK device will be given priority at reactor 2, and that the Decision Sciences system will be installed "depending on the status of on-site work."

The Decision Sciences system's two detectors can't be placed directly above and below the reactor vessel, so they'll be situated at a slant: one detector in the reactor building above and to the side of the vessel, the other outdoors, below and on the opposite side. Enough muons fly down from the atmosphere at an angle, the Decision Sciences physicists say, to provide meaningful results. They estimate that approximately 500 to 1,000 muons will speed through both detectors each second, which could provide, within a week, a useful image of the reactor core or the puddle of corium in the containment vessel. The longer the system is left in place, the more fine-grained the results will be.

Borozdin recalls his tour of the wreckage of Fukushima Daiichi in 2012. During a discussion of the reactor buildings' conditions, the site director mentioned the 40-year decommissioning plan that TEPCO had mapped out. "He said that if they knew the location of the fuel, this road map could be reduced by at least 10 years," Borozdin says. That's a bold statement, which TEPCO won't now confirm. It's also a rare instance of good news emerging from the Fukushima site.

IN THE ENORMOUS TASK of decommissioning Fukushima, one question gives way to more: Once the nuclear fuel has been located, TEPCO must determine how it can be safely broken into bits, hauled out, and packed into protective canisters. The company's provisional plans call for a "wet" defueling procedure that mimics the method used at Three Mile Island: Each containment vessel would be filled with water to the brim, providing more protection for workers as they drill down to break up the corium.

However, Barrett notes that the condition of the containment vessels is another big unknown. Damaged by the accident's meltdowns and explosions, these vessels are obviously leaking: TEPCO must continuously add water through input pipes and pump water out of the reactor buildings' basements to keep the radioactive mess submerged. Before TEPCO can plan a wet defueling, it must determine whether all the leaks in the containment vessels can be identified and patched. Only then could the containment vessels be flooded. "Until we get the diagnostics done, it's hard to know which defueling procedure will work," Barrett says. He's counting on new robots–transformer bots, pipe-crawling bots, underwater bots, and more–to conduct the necessary inspections.

Barrett and many other experts now suspect that these inspections will show the wet defueling option to be infeasible, which means TEPCO will have to develop a plan for a "dry" fuel removal. Such a procedure might require remotely operated tools that could attack the corium from the side and cut it into pieces, along with water sprays to keep the radioactive dust down.

That method would be extraordinarily complicated, and it's never been done before. But where others might see impossibility, Barrett sees another challenge. "We put a man on the moon; we can certainly defuel these things," he says. "We just have to develop the technology to do it." ■

POST YOUR COMMENTS at http://spectrum.ieee.org/fukushima0316

The 800-km Battery

By WINFRIED W. WILCKE & HO-CHEOL KIM

Lithium-ion batteries are played out.

Next up: lithium-air

ILLUSTRATION BY Elias Stein

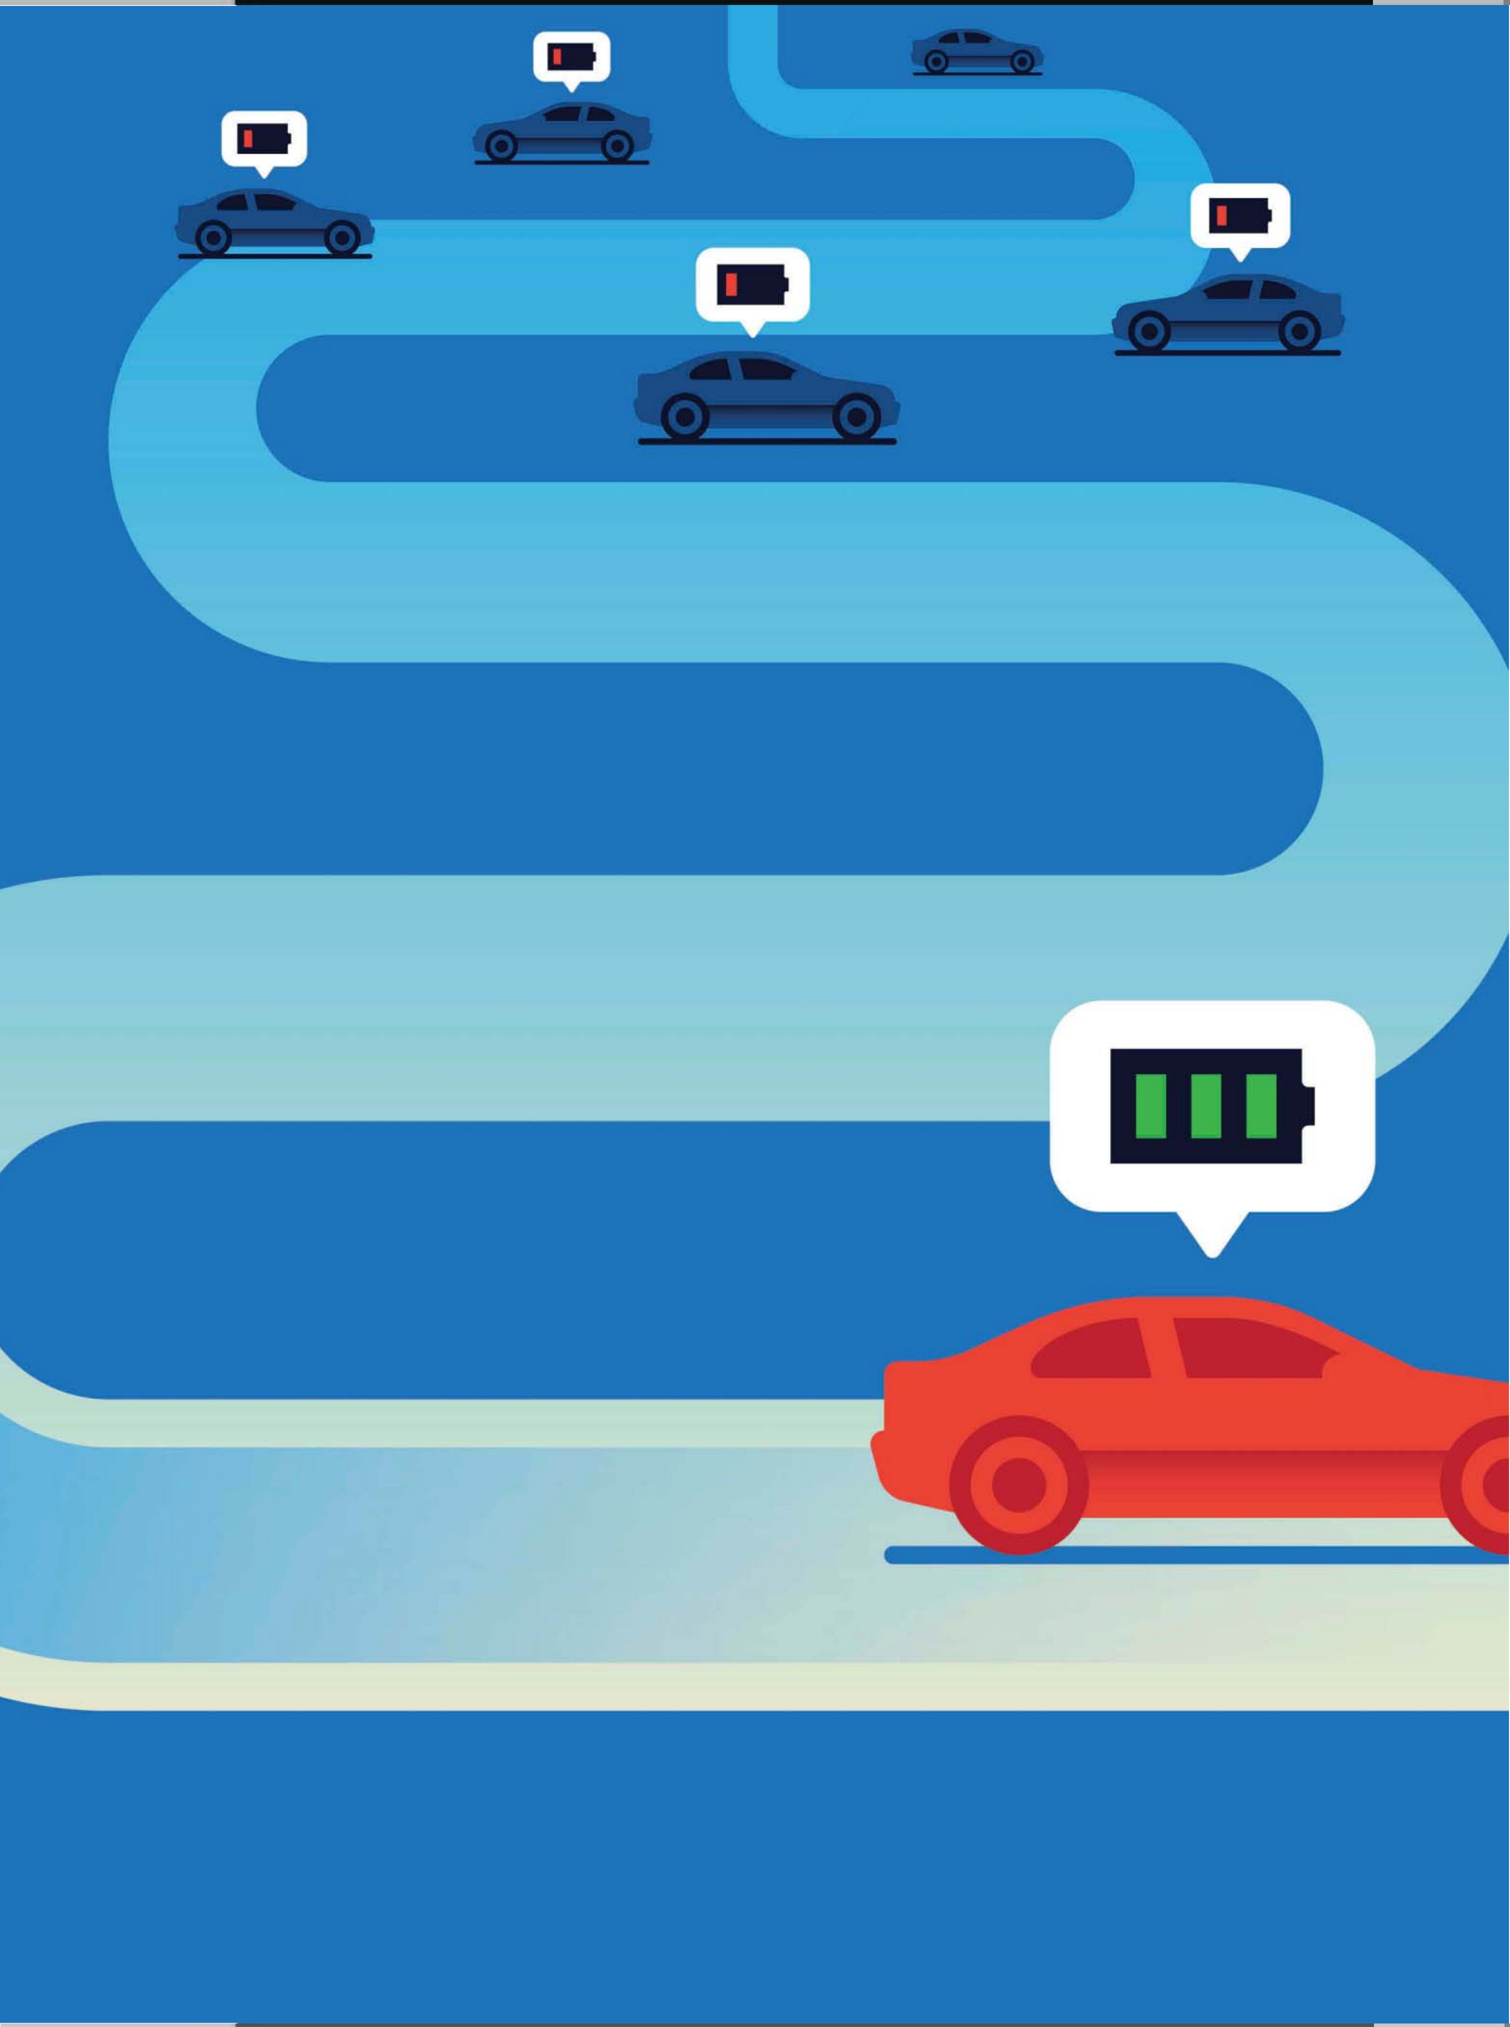

PROPOSITION: Electric cars will remain mostly niche products until they have a range of 800 kilometers, or roughly 500 miles, with an affordable battery.

That's as far as most people would want to drive in a day, and then they have all night to recharge.

That's how we came up with a figure of 800 km—or a nice round 500 miles—as the goal for our R&D project, Battery 500. It began in 2009 at the IBM Almaden Research Center, in San Jose, Calif., and has grown since then into a multinational partnership with commercial and academic participants in Europe, Asia, and the United States. It is based on metal-air technology, which packs far more energy into a battery of a given mass than today's state-of-the-art technology, the lithium-ion battery. We are still years away from commercialization, but we have made enough progress to predict that these batteries could be used in cars in the foreseeable future. Why are we so confident? Read on.

ELECTRIC MOTORS ARE IDEALLY suited for powering cars. They're lightweight and extremely powerful, they achieve efficiencies in excess of 90 percent, they don't need complex transmissions, and they churn out torque in just the right way, providing full rotational force starting with zero rpms. Internal-combustion engines, by contrast, don't produce high torque until they're spinning at thousands of rpms.

But even though they're propelled by a near-ideal mechanism, electric cars have a huge drawback, which is the low energy content of the batteries. Gasoline packs about 13,000 watt-hours per kilogram; the best production lithium-ion cells store only about 250 Wh/kg. Add the mass of the ancillary battery equipment—including the bus bars, cooling system, and battery management system—and the energy density of the entire system drops by half, giving the batteries a pitiful *1 percent* of the raw energy density of gasoline.

This huge gap between the energy densities of gasoline and batteries seemed to make it impossible to build competitive electric cars, but the success of the Tesla Model S has shown that it can be done. One major factor in favor of the electric car is the high efficiency with which it converts battery power to motive power at the wheels—about six times as efficiently as the average for gasoline-fueled cars in the United States. Also, electric car makers put the biggest, heaviest battery they can reasonably fit into their designs. Even so, the ranges fall far short of the 500-mile target. The upshot is that electric-car batteries need to attain at least twice the energy density of Li-ion cells to achieve a range of 800 km.

Cost is at least as important as energy density. Today's battery cells run from US $200 to $300 per kilowatt-hour, which means that—given an average range of 5 or 6 km/kWh—a car with an 800-km range will require a 150-kWh battery costing $30,000 to $45,000. By comparison, the base price of a BMW 2 Series car is $33,000. The price per kilowatt-hour must fall to at most $100 if the technology is to gain a serious foothold. At that price point, the car's much lower operating costs in energy use and maintenance, together with the sheer pleasure of driving such a responsive machine, will assure success in the marketplace.

But how do we get that 800-km-range battery? Well, start with the current state of the art, the lithium-ion battery.

A conventional, or "intercalation," Li-ion battery is a sealed system that has one electrode made of graphite (the anode) and an opposing electrode (the cathode) typically made from various oxides of transition metals, such as cobalt, nickel, or manganese. Both electrodes are immersed in a liquid organic electrolyte that contains dissolved lithium salts. In this electrolyte, lithium ions travel from one electrode to the other, the direction of travel depending on whether the battery is charging or discharging. In between the electrodes, immersed in the electrolyte, there's a porous polymeric separator, which prevents the electrodes from short-circuiting. The ions insert themselves between the atomic layers of the electrode material. This process is known as intercalation, and it is reversible—that is, it allows recharging.

If the electrodes are connected through an external circuit, lithium ions will migrate from the negative electrode to the positive electrode while electrons flow through the attached external circuit, thus discharging the battery. An externally applied voltage will reverse the ion flow, recharging the battery. The battery capacity depends on how much material is available for intercalation. In other words, the battery capacity is related to the volume and hence, the mass of the anode and cathode.

Metal-air batteries, however, employ a true electrochemical reaction rather than intercalation. For clarity, we'll assume that the metal is lithium. During discharge, the metallic lithium anode releases lithium ions; these travel through the electrolyte and combine with oxygen at the cathode, forming lithium peroxide (Li_2O_2). As in a conventional Li-ion battery, electrons flow through the external load circuit to compensate for the Li-ion flow inside the battery. The lithium peroxide accumulates on the surface of the porous carbon cathode, where the three participants in the reaction (lithium ions, electrons, and oxygen) meet and react. Since the reaction occurs on a surface, the volume or mass of the cathode material does *not* matter so long as it has a large surface. This is the main reason why this battery type has such high energy density.

Recharging reverses this order of events: An externally applied voltage breaks up the lithium peroxide, the oxygen diffuses back out into the environment, and the metal ions migrate back to the anode, where they acquire electrons and thus convert back to bulk metal.

This general principle can be applied to several different metals. Lithium-air, sodium-air–an interesting new contender [see sidebar, "Sodium: Less Energy, More Stability"]–and potassium-air are all possible systems because they lend themselves to recharging. Heavier metals, such as zinc, magnesium, iron, or aluminum have proved very hard to recharge, so we will not consider them here.

OUR WORK HAS FOCUSED on lithium and sodium. Let's start with lithium, which has the greater energy storage capacity. A lot of extraneous chemical reactions can mess things up in these cells. To understand these side reactions, we precisely measure the gases consumed and produced during the cycling of these cells. For this, we used a sophisticated differential electrochemical mass spectrometer that we built at the IBM Almaden Research Center. It

IN THE ELECTRODE —OR ON IT

A lithium-ion battery puts lithium ions *inside* the electrode; a lithium-air battery compounds lithium and oxygen *on top* of the electrode, allowing lighter electrodes

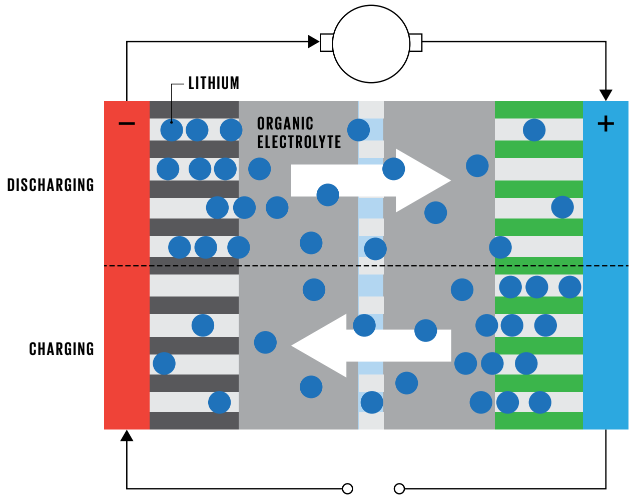

A LITHIUM-ION BATTERY shunts lithium ions from one electrode to the other, the direction of travel depending on whether the battery is charging or discharging. The ions fit between the atomic layers of the electrode, in a process called intercalation. Battery capacity depends on the volume—and hence, the mass—of the electrode.

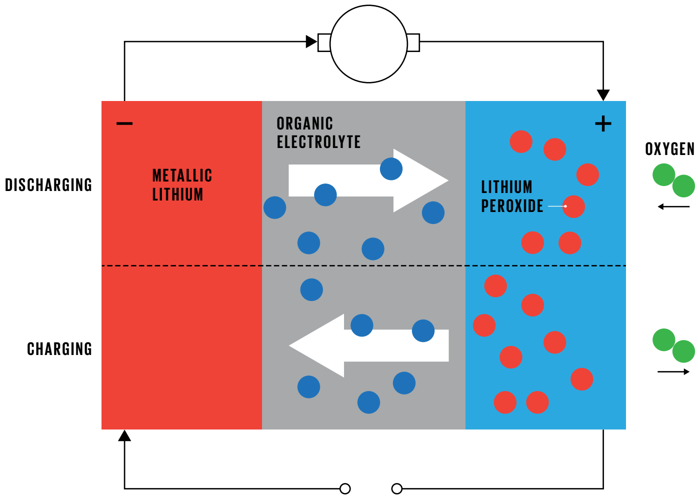

A LITHIUM-AIR BATTERY instead reacts lithium ions with oxygen on the surface of the electrode, forming lithium peroxide (Li_2O_2). Recharging reverses the reaction. In this chemistry, the capacity of the battery depends not on the volume of the electrode but on its surface area. That is why even a very light electrode can store a lot of energy, producing high energy density.

features eight stations to measure gas in parallel experiments.

It was this instrument that really gave us key insights. For example, it showed that the early lithium-air cells released far less oxygen on recharge than they consumed during discharge. (For most experiments we use dry oxygen rather than air.)

In an ideal cell, the amount of oxygen consumed during discharge should exactly equal the amount released on recharge. Our finding, therefore, was bad news because it implied that much of the oxygen released during the (desired) breakup of the Li_2O_2 on recharge was attacking components in the cell itself, notably the electrolyte. The cells weren't recharging–they were self-destructing!

With the help of our IBM sister laboratory, in Zurich, we traced the source of this parasitic reaction through experiments and computer simulations. We determined that the main problem was our organic electrolyte–it was breaking down. Since then we've gone a long way toward solving the problem. When our latest cells recharge, our new electrolyte (we'll describe it shortly) allows them to release most of the oxygen they take in during discharge. We also carefully monitored the hydrogen and water produced during the cycling of the cells, because their presence indicates that we're still seeing at least one other parasitic reaction.

We have been able to achieve 200 discharge-charge cycles, although so far we can do so only by limiting the discharge to less than the theoretical maximum.

Here are some of our key findings:

ANODE: Unlike the graphite anodes in a standard Li-ion cell, our metallic lithium anodes change their surfaces dramatically during recharging by growing mossy or treelike structures, called dendrites. These dendrites are dangerous because they provide conducting pathways between anode and cathode that can short-circuit the cell.

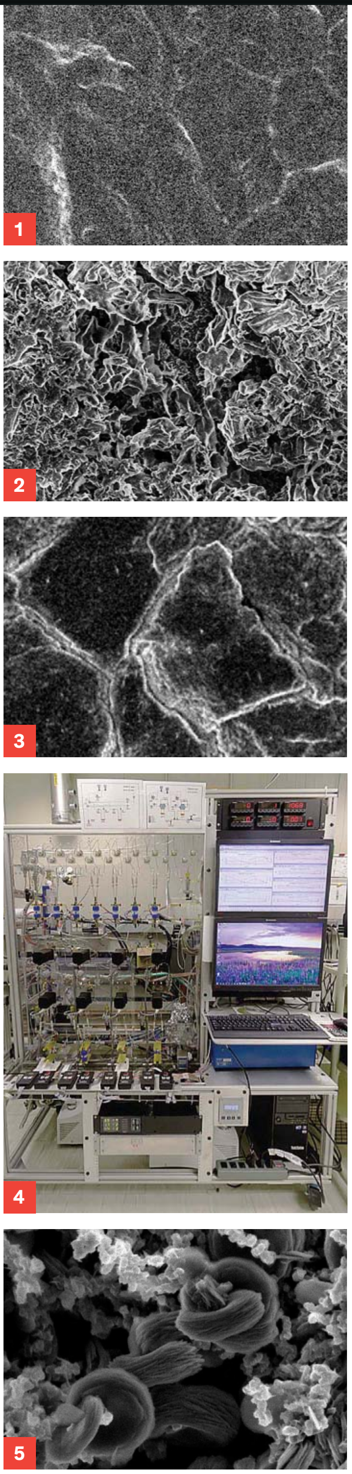

NOT SO ROUGH: With a conventional separator between the electrodes, a lithium anode in its pristine state is smooth [1], but after 10 discharge-charge cycles, its surface is rough, showing the growth of dangerous dendrites [2]. When the authors used a nanoporous separator, the surface stayed smooth through the same number of cycles [3]. The authors studied reactions in the lithium-air battery by channeling gases through a mass spectrometer [4]. The main reaction, between lithium and oxygen, produced lithium peroxide, which appears as doughnut-like structures in this electron micrograph [5].

IBM (5)

We've had great success in limiting the formation of dendrites by putting a special separator between the anode and the source of the lithium ions. This separator consists of a layer–we've tried both organic and inorganic materials–that contains nanometer-scale pores. That's small enough to uniformly distribute the flowing ions' current and thus suppress dendrite formation. With this nanoporous separator, the metal remains smooth for hundreds of cycles, whereas with a standard separator it forms dendrites after only a few cycles. Another membrane that combines ion-conducting glasses with a polymer matrix works even better.

Fortunately, an electric car's big battery requires only hundreds, not thousands, of full cycles. For example, a car with the 500-mile range needs to be fully recharged only 400 times for a lifetime range of 200,000 miles (roughly 320,000 km).

ELECTROLYTE: The improved electrolyte solvent molecules we use can still be broken apart by oxygen and other compounds generated during the operation of a cell. We haven't yet discovered any single solvent that's stable enough for commercially useful lithium-air cells, but we have found a cocktail of solvents that works pretty well.

CATHODE: We add traces of $LiNO_2$ (lithium nitrate) to our carbon cathode to minimize the undesirable catalytic effect that accelerates the breakdown of the electrolyte during charging and releases carbon dioxide. Even so, this reaction still requires that the applied charging voltage must be higher–by up to 700 millivolts–than the battery's operating voltage. Such a high overvoltage reduces electric efficiency–that is, the fraction of the energy pumped into the battery during charging that's returned during discharge. Although it's much better than what you get with plain carbon cathodes (over 1,200 mV), it's still too high for practical use. We've had similar results when we replaced carbon with metal oxides.

CATALYSTS: The pros and cons of deliberately using catalysts in metal-air batteries are subject to much scientific debate. Catalysts often lead to an apparent reduction of overvoltage, but one has to be extraordinarily cautious about claiming a net benefit, because catalysts generally accelerate the destruction of electrolytes. Also, our theoretical studies indicate that the activation energy of the lithium-oxygen reaction–in both directions–is very low. Thus catalysts should not be needed.

AIR PREPARATION: We've called these devices lithium-air cells, but in fact we've mostly been using dry oxygen gas. The emphasis here is on "dry" because we need only remove the water vapor and the carbon dioxide from the air, not the nitrogen, to make it usable. To do that at scale in a commercial battery, we'll need to put in a substantial engineering effort to create an air-cleaning system that's sufficiently light, efficient, and reliable to retain the energy advantage of this technology.

Another outstanding engineering task is how to scale up to much larger cells and integrate them into a multicell module and pack, including a tailor-made battery management system. Our original cells measure roughly 13 millimeters in length and 76 mm in diameter; we are testing versions that measure 100 by 100 mm.

The whole project was motivated by the desire to achieve high specific energy density–that is, the energy per unit mass. Where are we now?

The lithium-oxygen reaction has a *theoretical* (specific) energy density of 3,460 Wh/kg, which is much higher than the theoretical limit of any lithium-ion intercalation chemistry. The *practical* energy is much lower than the theoretical values for both intercalation and metal-air chemistries due to the inert mass contributed by those cell components that take no part in the reaction. These include the electrolyte, the cell housing, current collectors, and the separator. Furthermore, a lithium-air battery would also include the inert mass of the machinery needed to prepare ambient air for use in the cell. It is these engineering problems that make the practical development of lithium-air batteries for cars such a challenge.

It is too early to quote a practical energy density for the lithium-oxygen technology, let alone lithium-air technology. These numbers depend on engineering details, and the project is still focused on the basic science of the materials and chemistry. However, early results are encouraging. For example, we have measured a specific energy density of 15 kWh/kg of the raw carbon cathode material (5,700 milliampere hours multiplied by 2.7 volts per gram of carbon black). But as we pointed out earlier, the practical energy density will be *greatly* reduced by the mass of all the other components in the cell. Our current best estimate of what is practically attainable is around 800 Wh/kg at the cell level.

The first practical metal-air batteries might be used in buses, trucks, and other large vehicles that can more easily accommodate the mass of the air-cleaning machinery. But the most profound change will come when the technology reaches the family car, freeing it from the "range anxiety" of today's electric cars–and freeing us all from our dependence on oil and the many problems it causes. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/800kmbattery0316

SODIUM: LESS ENERGY, MORE STABILITY

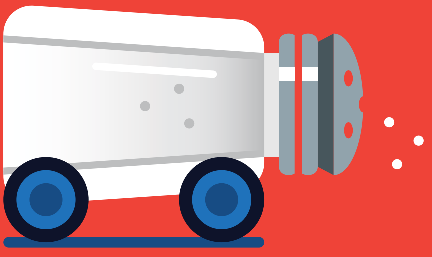

SODIUM-AIR BATTERIES ARE another interesting possibility, despite having an energy density lower than that of the lithium-air chemistry.

The lower energy reflects the nature of the reaction, which uses only one electron and thus generates a superoxide (NaO_2) instead of sodium peroxide (Na_2O_2). This reduces energy density immediately by a factor of two. The reaction's theoretical specific energy is approximately 1,100 watt-hours per kilogram.

On the other hand, sodium-air batteries charge up more efficiently than lithium-air batteries because they have a very low overpotential—less than 20 millivolts as opposed to 700 mV. As a result, it's possible to keep the operating voltage under 3 volts, which protects components from destructive oxidation, notably the electrolyte destruction observed in the lithium-air system. We have proved this by measuring efficiencies above 98 percent. These results make for good stability during cycling: After 50 cycles, the cell's capacity is essentially unchanged.

There are a few technical challenges to overcome. For instance, because of the nature of the oxidation, the sodium-air battery sucks in twice as much air as its lithium-air equivalent, requiring an airflow comparable to that of a piston engine of the same power. Then there's the high chemical reactivity of sodium metal, which you may remember seeing demonstrated in high school, in which a small piece of sodium reacts violently with water.

Lithium is comparatively rare, and it isn't cheap. But sodium is as common as table salt, and it's not expensive. The materials of a sodium-air cell would likely cost less than a tenth as much as those in a lithium-ion battery. In the long run, lithium-metal batteries promise the best performance, but given the combination of stability, low cost, and still-impressive specific energy, the sodium-air technology might serve to bridge the gap between today's batteries and those of the more distant future. *—W.W.W. & H.-C. Kim*

ILLUSTRATION BY Elias Stein

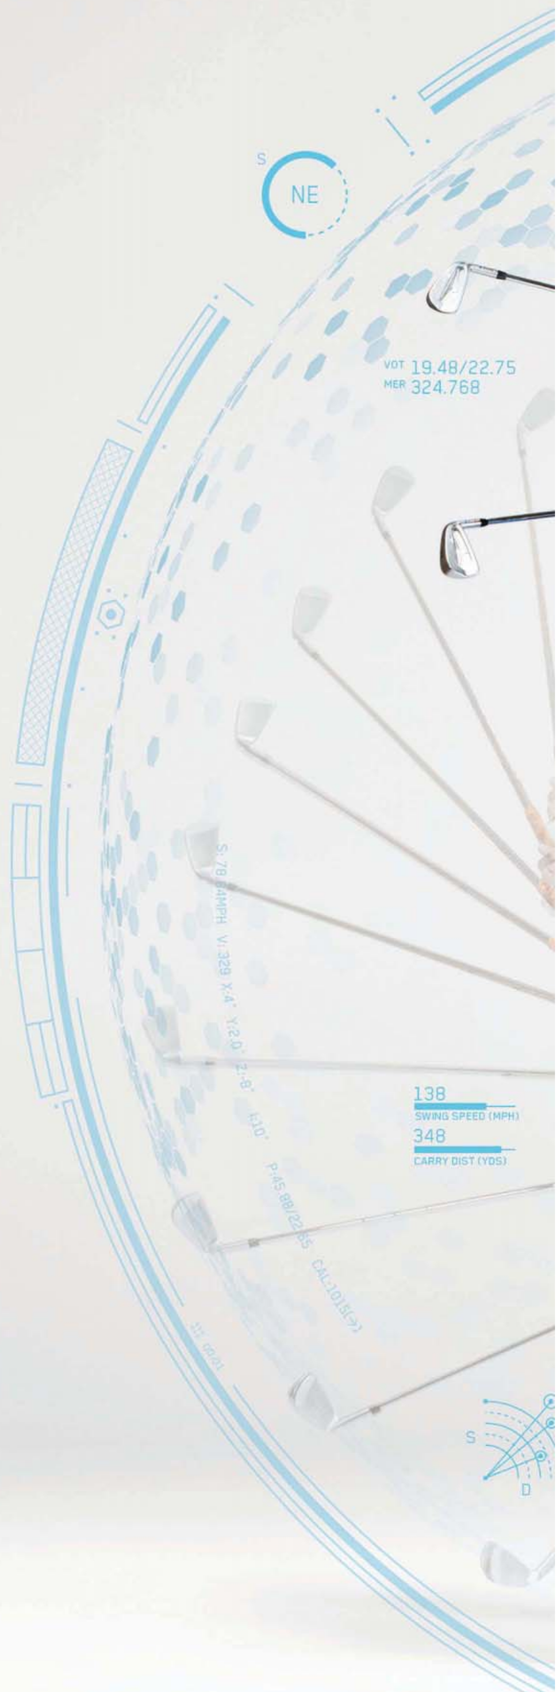

SILICON GETS SPORTY

Next-generation sensors make golf clubs, tennis rackets, and baseball bats smarter than ever

By Karen Lightman

HUNTER MAHON PHOTOGRAPH: DYLAN COULTER; DATA GRAPHICS: JOHN LE PORE


S: 78.64
V: 329
X: 4" Y: 2.0" Z:-8"
I: 10"
P: 45.88 / 22.65
CAL: 1015
POS: Y44.782
VMAX
+2.1
20 40 60 80
FACE POS
22
35
72.3 MPH
Z: -003.3°
Y: -022.1°
X: -001.3°

A golfer stands in the dreaded sand trap, carefully considering how to balance his weight as he eyes the ball. He takes a few practice swings. If he swings too deeply, he'll hit the ground and lose another stroke. It's a tough shot, but he swings without hesitation. Embedded in his club are microelectromechanical system (MEMS) devices–tiny machines with elements about the thickness of a human hair. These devices aren't going to swing the club for him, but he's been using them to analyze his swing and practice this shot. Maybe this time he'll make it.

The wild popularity of smartphones and wearables has been driving down the cost of MEMS devices, including accelerometers, gyroscopes, magnetometers, and pressure sensors. These minuscule chips help to count steps, track calories burned, and monitor heart rate. Such data are useful, sure, but while these devices may nudge users to be more active, they don't actually improve a swing, a punch, or a kick. To do so means moving sensors off the wrist and into sports gear–and that's quickly happening. Indeed, you can now buy sensor-based equipment that can boost your performance, not only for golf but also for tennis, baseball, boxing, and soccer.

MEASURING THE OFTEN-COMPLEX motion of an athlete takes many more metrics than just tracking steps. Consider the sport of sculling: To gauge the efficiency of a rower as his oar moves through the phases of catch, drive, release, and recovery, you need to track the movement of his legs, back, and arms. If you want to analyze a baseball player as he whips the bat around, you need to consider rotational angles and swing speed.

The current boom in sensor-laden sports equipment is an outgrowth of the dramatic drop in sensor prices and the development of technology that makes it easy to integrate data from multiple sensors. In the past, synthesizing the data outputs from multiple sensors in real time was a nightmare. Figuring out how to make sense of streams of data flying into a general-purpose proces-

THE GEAR

A sampling of sensor-laden sports equipment shows the range of athletic activity being measured, analyzed, and affected by this technology. The products listed are representative examples, but they're rarely the only brands available in each sports category. More products, from both established companies and startups, are being launched regularly.

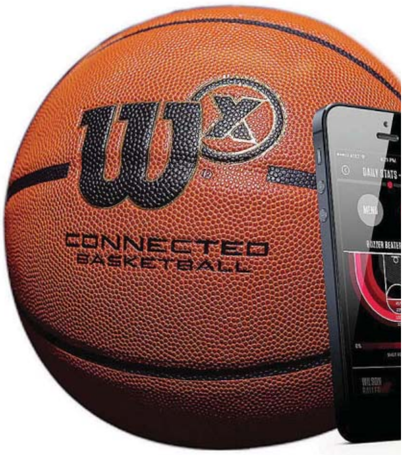

Wilson X Connected Basketball

TENNIS

Manufacturer: Babolat
Equipment: Babolat Pure Drive Play racket
Sensors on board: InvenSense motion-sensing system embedded in its handle
What it tracks: Type of shot made (forehand, backhand, serve, and smash), location where the tennis ball connects with the racket's strings, overall playing time versus active playing time, and ball speed during play and when serving
How it works: Players can connect the racket via Bluetooth to a mobile device to share data with a coach for postgame analysis or over social media.
Commercial status: Announced at French Open 2013; available now

US $200

BOXING

Manufacturer: Athletec
Equipment: Corner
Sensor on board: Tri-axis accelerometer
What it tracks: Punch types (jab, cross, hook, uppercut), speed, power, hit/miss ratio, blocks, combinations, punch rate, pacing
How it works: A small device clips onto the boxer's hand wraps and fits under the boxing gloves. An onboard microprocessor analyzes the accelerometer data to detect punches and extract key features from each punch's waveform. The device then sends this data via Bluetooth to a smart device for display and storage to build a picture of the boxer's performance in each round, session, week, and month of training.
Commercial status: Announced in January 2016; will begin shipping in August 2016

$70 (for a pair)

BASEBALL, SOFTBALL

Manufacturer: Diamond Kinetics
Equipment: SwingTracker
Sensors on board: Tri-axis accelerometer and tri-axis gyroscope
What it tracks: Fifteen different swing metrics, including power, speed, efficiency, and distance the bat travels in the hitting zone
How it works: Mounts to the bottom of a bat, capturing 11,000 data points per second to analyze swing data in real time; shares that information with coaches via Bluetooth to a mobile device.
Commercial status: Announced 2014; available now

$150

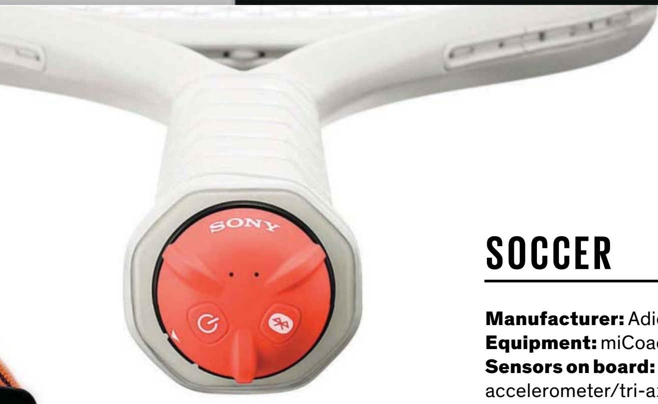

Sony Smart Tennis Sensor

Athletec Corner

SwingTracker

CLOCKWISE FROM TOP LEFT: WILSON SPORTING GOODS; SONY; ATHLETEC; DIAMOND KINETICS

TENNIS

Manufacturer: Sony Corp.
Equipment: Sony Smart Tennis Sensor
Sensors on board: Three-axis motion-tracking sensor and vibration sensor
What it tracks: Shot type (serve, smash, volley forehand, volley backhand, topspin forehand, topspin backhand, slice forehand, and slice backhand), ball speed, swing speed, ball spin, and ball impact spot; will match shot data to video recorded on smart device
How it works: Attaches to the grip end of a compatible tennis racket.
Commercial status: Released in Japan May 2014; in United States 2015

$200

SOCCER

Manufacturer: Adidas
Equipment: miCoach Smart Ball
Sensors on board: Tri-axis accelerometer/tri-axis gyroscope combo
What it tracks: How hard the ball has been struck, speed, spin, and flight path
How it works: Soccer ball with integrated sensors tracks the dead-ball kicking used for penalties, shooting, corners, long passes, and goal kicks; transmits data via Bluetooth to an iPhone or Android device. Coaches can use this data to train players.
Commercial status: Announced May 2014; available now

$200

RUNNING

Manufacturer: Bodytech
Equipment: Lumo Run running shorts and capris
Sensors on board: Nine-axis inertial measurement unit (accelerometer, gyroscope, magnetometer), pressure sensor
What it tracks: Measures 14 different running metrics, including pelvic rotation, tilt and drop, braking, bounce, ground contact time, cadence, and stride length
How it works: A removable sensor in the waistband provides real-time audio cues via Bluetooth to headphones and data to a smart device for detailed postrun analysis.
Commercial status: Announced October 2015; shipping in March 2016

$150 (shorts), $170 (capris)

BASKETBALL

Manufacturer: Wilson Sporting Goods Co.
Equipment: Wilson X Connected Basketball
Sensor on board: Tri-axis accelerometer
What it tracks: Shots made and shots missed
How it works: An algorithm developed by a machine learning system detects shots made and missed based on accelerometer data. Requires standard-height hoop with a net, communicates to a smart device via Bluetooth technology; nonreplaceable battery tracks at least 100,000 shots.
Commercial status: Available in United States now, Canada by end of March 2016

Coming: Wilson X Connected Football
What it tracks: Throw distance, speed, spiral efficiency, and whether a football has been caught or dropped

$200

sor—so-called sensor fusion—presented a real hurdle to developing sports-training products that used such data, even though the need was clear.

In the past couple of years, though, a number of companies have come out with off-the-shelf technology for sensor fusion. Now, rather than develop an algorithm from scratch that measures linear acceleration, for example—which is useful in measuring the speed of a kicked or hit ball—a design engineer can use a third-party linear-acceleration algorithm to measure motion across the x, y, and z axes.

Sensor-fusion hardware and software have also taken on another tricky task: managing power consumption, which is key if you want to avoid having to frequently charge your smart balls, bats, boxing gloves, tennis rackets, or golf clubs. These systems may, for example, turn off a gyroscope, which is notoriously power hungry, when not in use. They may allocate work to the most efficient sensors, like an accelerometer or an accelerometer-magnetometer combination instead of a gyro, if it can do the job.

While both sensor manufacturers and the companies that make tools to process sensor data agree that sensor fusion is a good idea, they have tackled the implementation problem in different ways, creating systems that range from very proprietary to essentially open.

On the proprietary end of that spectrum, the makers of sensors and MEMS—such as Analog Devices, Bosch Sensortec, Freescale, InvenSense, Kionix, and STMicroelectronics—are running their own sensor-fusion algorithms on the chips they build that include multiple sensors and a microcontroller, for example Bosch Sensortec's BNO055. Some of these manufacturers also provide application programming interfaces that can run on other hardware but work only with one company's brand of sensors. The combo sensor-microcontroller chips allow for data batching, which means that they'll store sensor readings in a buffer until the application processor is awake. This power-saving mechanism is a big benefit of sensor fusion.

In a more flexible approach, Hillcrest Labs, PNI Sensor, and QuickLogic have developed stand-alone sensor-fusion coprocessors that support chips from a variety of sensor makers. These hardware-software systems poll the sensors for data, push it to their own low-power coprocessors, and manage the data using proprietary algorithms.

Finally, my organization, MEMS & Sensors Industry Group, has rolled out an open-source option that can help designers jump-start development of systems that require the use of sensor fusion. Our Accelerated Innovation Community, which began in late 2014, makes some of the most popular sensor-fusion algorithms available for free to design engineers.

WITH CHEAP SENSORS AND simple sensor fusion on the market, the current wave of smart sports gear is not likely to crest anytime soon.

Sensors have already been inserted into just about every type of sports equipment that can be swung, punched, or worn, and more are in development. Noel C. Perkins of the University of Michigan and William Clark of the University of Pittsburgh developed one recent example, a smart baseball bat add-on. The prototype is a small, sensor-laden circuit board affixed to the knob end of a bat. It contains MEMS chipsets that measure the three components of the bat's acceleration as well as rotation rates around three orthogonal axes. A smartphone collects the data and analyzes the swing time, speed, control, and other factors that relate to a successful hit.

The University of Michigan has licensed the technology to Clark's new company, Diamond Kinetics, which in 2014 rolled out its first commercial product, SwingTracker, a lightweight sensing accessory that attaches temporarily to the knob of the bat. The technology, which Perkins originally developed to improve his fly fishing, can and likely will be applied to other sports.

Sensors in athletic clothing, already in limited use by elite athletes, will soon become accessible to ordinary consumers. These can be used to analyze body motion for multiple sports. Such wearable gear is still in the nascent stages of evolution. But more than one startup has already begun to sell sensor-equipped workout clothes. California-based MAD Apparel, for example, claims that its US $400 Athos Upper Body Package analyzes in real time how hard an athlete's muscles are working by using a shirt with 14 embedded biosensors and a hardware "core" positioned over the heart.

Another maker of smart clothing is the Dutch company Xsens Technologies, which Fairchild Semiconductor purchased in April 2014. Xsens is adapting its full-body motion-tracking suits (used to make animated movies and video games) to sports applications. Independent researchers are using Xsens's bodysuits for such sports as ski jumping, shot put, archery, baseball pitching, rowing, and even wingsuit flight.

And, in another sign that smart clothing is moving toward the mass market, Under Armour chief digital officer Robin Thurston has said his company is developing smart clothing and expects to have commercial products within the next two to three years.

For many sports, the remaining challenge lies not in collecting data but in making that information accessible and meaningful. Both professional athletes and casual sports enthusiasts want to ensure that they are maximizing the effects of their time at the gym. That's coming.

Within a few years, our MEMS-enabled golfer should have plenty of help for his game: His smart eyewear will give him a hole-by-hole analysis of the course and will record and instantly play back each shot while MEMS motion sensors in his shoes will help him position his feet. His clothes will have sensors sewn into the fabric so he can adjust his hips and legs appropriately. He'll swing his MEMS-based golf club to strike a MEMS-enabled golf ball, connecting just right. The ball will land on the green, transmitting its short distance to the hole immediately to his smartphone. At long last, he'll *really* enjoy the sport of golf. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/sportsensors0316

THE HITS OF CONCUSSION DETECTION

I can't describe the use of sensors in athletics without mentioning where they will have the biggest impact—literally. I'm talking, of course, about concussion detection. The numbers are staggering—and growing, particularly for children who play sports. According to the Centers for Disease Control and Prevention, diagnoses of concussion or traumatic brain injury rose a whopping 57 percent in children 19 years old and younger between 2001 and 2009. There were more than 248,000 sports-related concussions in children in 2009 alone. ¶ Today's push for more extensive concussion analysis stems from a Simbex Study funded by the National Institutes of Health. More than eight years of detailed data from football practices and games at all levels, from Pop Warner Little Scholars to the National Football League, showed that concussions occur far more often and have longer-lasting effects than previously believed. ¶ The Head Impact Telemetry System, or HITS, developed for that study is now the property of the football helmet manufacturer Riddell, which acquired HITS in 2004. HITS uses multiple single-axis MEMS accelerometers to record the location, magnitude, duration, and direction of impacts. Players and coaches can upload and evaluate the data either right there on the playing field or after the game.

HIT HARD: David Bruton Jr. of the Denver Broncos lies on the ground after a hit in 2014; he left the game with a reported concussion.

"The HITS system showed that there are actually two events in each impact–a high *g*-force event on initial contact and then a low *g*-force event due to the percussive reaction of the head," says Rob O'Reilly, a senior member of the technical staff at Analog Devices. "At first, we were adjusting the same single-axis technology used in automotive airbag deployment sensing, and a lot of people thought we just couldn't do it."

While that technology worked reasonably well, Analog Devices introduced the ADXL377, a three-axis MEMS accelerometer able to measure up to 200 *g*'s of force instantly and accurately while withstanding peak impacts of up to 10,000 *g*'s. Intended specifically for contact sports of all kinds, that accelerometer is small enough to fit easily inside an earpiece. The military is already using the company's ADXL375, a 200-*g* digital-output accelerometer, in a body-worn blast gauge for sensing the likelihood of traumatic brain injury from explosions. Data from the use of this sensor will improve the technology and likely drive a broader adoption in youth sports.

A handful of companies are building accelerometers and gyroscopes to gauge the possibility of brain injury into products worn by nonhelmeted athletes playing such sports as soccer and rugby. X2 Biosystems in Seattle makes technology-laden sports mouth guards, and Triax Technologies of Norwalk, Conn., is selling sensor-studded headbands and skullcaps. Reebok markets the US $100 Checklight skullcap, which uses accelerometers incorporated into a mesh of stretchable electronics to measure the severity of an impact. Even U.S. Women's National Team soccer star Abby Wambach is in the game. Having sustained some serious concussions because of her affinity for heading the ball, Wambach has partnered with Triax on the launch of its head-impact monitoring products.

What's still debatable among athletes, coaches, and parents with regard to concussion-detection equipment is its accuracy. According to recent Washington State University research, some of the concussion-detection products (such as smart mouth guards) used for nonhelmeted sports aren't fast enough to measure hard hits and don't accurately measure angular hits, which may be more dangerous than straight-on impacts.

This doesn't invalidate the merit of concussion-detection equipment for nonhelmeted (or helmeted) athletes. It just means that engineers must work more closely with sports-equipment manufacturers to produce products with greater accuracy. *—K.L.*

FOOTBALL, ICE HOCKEY*

Manufacturer: Reebok
Equipment: Reebok Checklight skullcap for head-impact detection
Sensors on board: MEMS accelerometers embedded in MC10's conformal electronics
What it tracks: Potential severity of head impacts
How it works: MEMS accelerometers track direct accelerations to the head; red and yellow indicator lights display the severity of impact.
Commercial status: Announced July 2013; available now

** and other helmeted sports*

US $100

TOP: DOUG PENSINGER/GETTY IMAGES; BOTTOM: REEBOK

TODAY, SILICON MICROCHIPS underlie every aspect of digital computing. But their dominance was never a foregone conclusion. Throughout the 1950s, electrical engineers and other researchers explored many alternatives to making digital computers. ¶ One of them seized the imagination of the U.S. National Security Agency (NSA): a superconducting supercomputer. Such a machine would take advantage of superconducting materials that, when chilled to nearly the temperature of deep space–just a few degrees above absolute zero–exhibit no electrical resistance whatsoever. This extraordinary property held the promise of computers that could crunch numbers and crack codes faster than transistor-based systems while consuming far less power. ¶ For six decades, from the mid-1950s to the present, the NSA has repeatedly pursued this dream, in partnership with industrial and academic researchers. Time and again, the agency sponsored significant projects to build a superconducting computer. Each time, the effort was abandoned in the face of the unrelenting pace of Moore's Law and the astonishing increase in performance and decrease in cost of silicon microchips. »

The NSA's FROZEN DREAM

AFTER SIX DECADES, WILL THE SPY AGENCY'S VISION OF SUPERCONDUCTING SUPERCOMPUTERS FINALLY BECOME A REALITY? *By* David C. Brock

PHOTO-ILLUSTRATION BY **Bryan Christie Design**

Now Moore's Law is stuttering, and the world's supercomputer builders are confronting an energy crisis. Nuclear weapon simulators, cryptographers, and others want exascale supercomputers, capable of 1,000 petaflops–1 million trillion floating-point operations per second–or greater. The world's fastest known supercomputer today, China's 34-petaflop Tianhe-2, consumes some 18 megawatts of power. That's roughly the amount of electricity drawn instantaneously by 14,000 average U.S. households. Projections vary depending on the type of computer architecture used, but an exascale machine built with today's best silicon microchips could require hundreds of megawatts.

The exascale push may be superconducting computing's opening. And the Intelligence Advanced Research Projects Activity, the U.S. intelligence community's arm for high-risk R&D, is making the most of it. With new forms of superconducting logic and memory in development, IARPA has launched an ambitious program to create the fundamental building blocks of a superconducting supercomputer. In the next few years, the effort could finally show whether the technology really can beat silicon when given the chance.

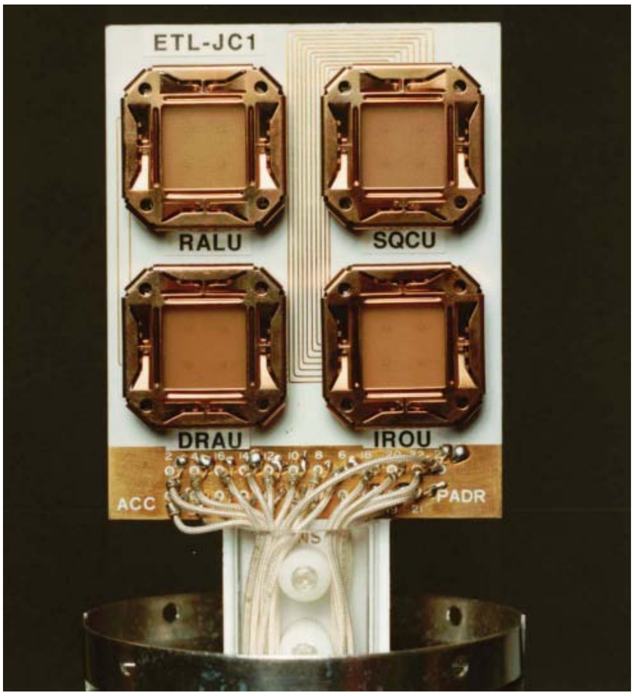

AIST

THE NSA'S DREAM of superconducting supercomputers was first inspired by the electrical engineer Dudley Buck. Buck worked for the agency's immediate predecessor on an early digital computer. When he moved to MIT in 1950, he remained a military consultant, keeping the Armed Forces Security Agency, which quickly became the NSA, abreast of new computing developments in Cambridge.

Buck soon reported on his own work–a novel superconducting switch he named the cryotron. The device works by switching a material between its superconducting state–where electrons couple up and flow as a "supercurrent," with no resistance at all–and its normal state, where electrons flow with some resistance. A number of superconducting metallic elements and alloys reach that state when they are cooled below a critical temperature near absolute zero. Once the material becomes superconducting, a sufficiently strong magnetic field can drive the material back to its normal state.

In this, Buck saw a digital switch. He coiled a tiny "control" wire around a "gate" wire, and plunged the pair into liquid helium. When current ran through the control, the magnetic field it created pushed the superconducting gate into its normal resistive state. When the control current was turned off, the gate became superconducting again.

Buck thought miniature cryotrons could be used to fashion powerful, fast, and energy-efficient digital computers. The NSA funded work by him and engineer Albert Slade on cryotron memory circuits at the firm A.D. Little, as well as a broader project on digital cryotron circuitry at IBM. Quickly, GE, RCA, and others launched their own cryotron efforts.

Engineers continued developing cryotron circuits into the early 1960s, despite Buck's sudden and premature death in 1959. But liquid-helium temperatures made cryotrons challenging to work with, and the time required for materials to transition from a superconducting to a resistive state limited switching speeds. The NSA eventually pulled back on funding, and many researchers abandoned superconducting electronics for silicon.

Even as these efforts faded, a big change was under way. In 1962 British physicist Brian Josephson made a provocative prediction about quantum tunneling in superconductors. In typical quantum-mechanical tunneling, electrons sneak across an insulating barrier, assisted by a voltage push; the electrons' progress occurs with some resistance. But Josephson predicted that if the insulating barrier between two superconductors is thin enough, a supercurrent of paired electrons could flow across with zero resistance, as if the barrier were not there at all. This became known as the Josephson effect, and a switch based on the effect, the Josephson junction, soon followed.

IBM researchers developed a version of this switch in the mid-1960s. The active part of the device was a line of superconducting metal, interrupted by a thin oxide barrier cutting across it. A supercurrent would freely tunnel across the barrier, but only up to a point; if the current rose above a certain threshold, the device would saturate and unpaired electrons would trickle across the junction with some resistance. The threshold could be tuned by a magnetic field, created by running current through a nearby superconducting control line. If the device operated close to the threshold current, a small current in the control could shift the threshold and switch the gate out of its supercurrent-tunneling state. Unlike in Buck's cryotron, the materials in this device always remained superconducting, making it a much faster electronic switch.

As explored by historian Cyrus Mody, by 1973 IBM was working on building a superconducting supercomputer based on Josephson junctions. The basic building block of its circuits

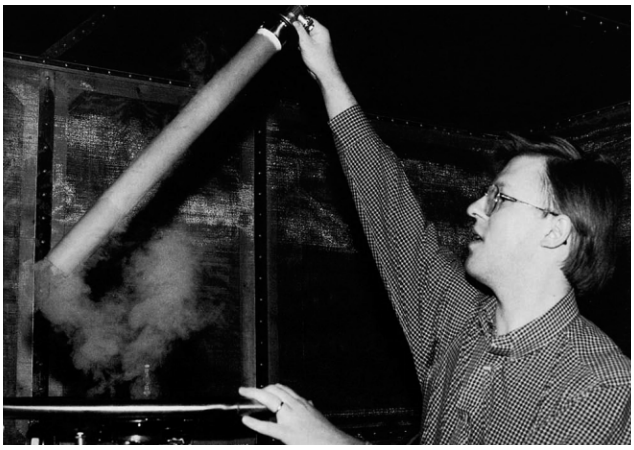

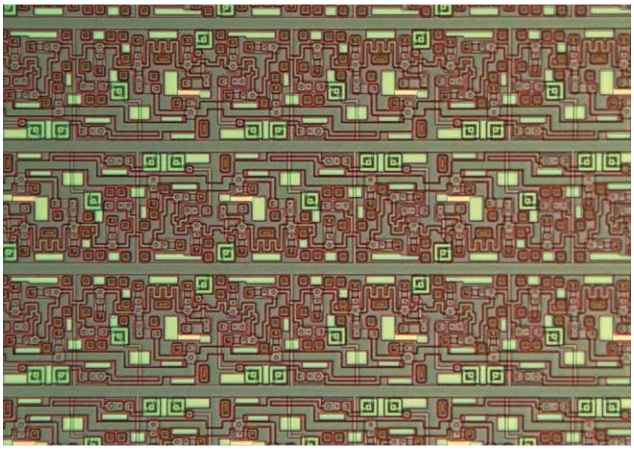

COOL RUNNING: The ETL-JC1, a superconducting computer developed in Japan, included four Josephson-junction-based integrated circuit chips, for logic and memory [opposite]. A later effort, the Hybrid Technology Multi-Threaded project, used a new form of superconducting logic; team member Dmitry Zinoviev is pictured here [above left] holding a flask plug for a dewar of liquid helium in front of shielding for sensitive superconducting circuits. Long-time superconducting-electronics manufacturer Hypres is making low-power circuits [above right] for Raytheon BBN Technologies, as part of a memory project for the present-day IARPA program.

LEFT: JUDY CONLON/NASA; RIGHT: HYPRES

was a superconducting loop with Josephson junctions in it, known as a superconducting quantum interference device, or SQUID. The NSA covered a substantial fraction of the costs, and IBM expected the agency to be its first superconducting-supercomputer customer, with other government and industry buyers to follow.

IBM's superconducting supercomputer program ran for more than 10 years, at a cost of about US $250 million in today's dollars. It mainly pursued Josephson junctions made from lead alloy and lead oxide. Late in the project, engineers switched to a niobium oxide barrier, sandwiched between a lead alloy and a niobium film, an arrangement that produced more-reliable devices. But while the project made great strides, company executives were not convinced that an eventual supercomputer based on the technology could compete with the ones expected to emerge with advanced silicon microchips. In 1983, IBM shut down the program without ever finishing a Josephson-junction-based computer, super or otherwise.

Japan persisted where IBM had not. Inspired by IBM's project, Japan's industrial ministry, MITI, launched a superconducting computer effort in 1981. The research partnership, which included Fujitsu, Hitachi, and NEC, lasted for eight years and produced an actual working Josephson-junction computer—the ETL-JC1. It was a tiny, 4-bit machine, with just 1,000 bits of RAM, but it could actually run a program. In the end, however, MITI came to share IBM's opinion about the prospect of scaling up the technology, and the project was abandoned.

CRITICAL NEW DEVELOPMENTS emerged outside these larger superconducting-computer programs. In 1983, Bell Telephone Laboratories researchers formed Josephson junctions out of niobium separated by thin aluminum oxide layers. The new superconducting switches were extraordinarily reliable and could be fabricated using a simplified patterning process much in the same way silicon microchips were.

Then in 1985, researchers at Moscow State University proposed a new kind of digital superconducting logic. Originally dubbed resistive, then renamed "rapid" single-flux-quantum logic, or RSFQ, it took advantage of the fact that a Josephson junction in a loop of superconducting material can emit minuscule voltage pulses. Integrated over time, they take on only a quantized, integer multiple of a tiny value called the flux quantum, measured in microvolt-picoseconds.

By using such ephemeral voltage pulses, each lasting a picosecond or so, RSFQ promised to boost clock speeds to greater than 100 gigahertz. What's more, a Josephson junction in such a configuration would expend energy in the range of just a millionth of a picojoule, considerably less than consumed by today's silicon transistors.

Together, Bell Labs' Josephson junctions and Moscow State University's RSFQ rekindled interest in superconducting electronics. By 1997, the U.S. had launched the Hybrid Technology Multi-Threaded (HTMT) project, which was supported by the National Science Foundation, the NSA, and other agencies. HTMT's goal was to beat conventional silicon to petaflop-level supercomputing, using RSFQ integrated circuits among other technologies.

It was an ambitious program that faced a number of challenges. The RSFQ circuits themselves limited potential computing efficiency. To achieve tremendous speed, RSFQ used resistors to provide electrical biases to the Josephson junctions in order to keep them close to the switching threshold. In experimental RSFQ circuitry with several thousand biased Josephson junctions, the static power dissipation was negligible. But in a petaflop-scale supercomputer, with possibly many billions of such devices, it would have added up to significant power consumption.

The HTMT project ended in 2000. Eight years later, a conventional silicon supercomputer—IBM's Roadrunner—was touted as the first to reach petaflop operation. It contained nearly 20,000 silicon microprocessors and consumed 2.3 megawatts.

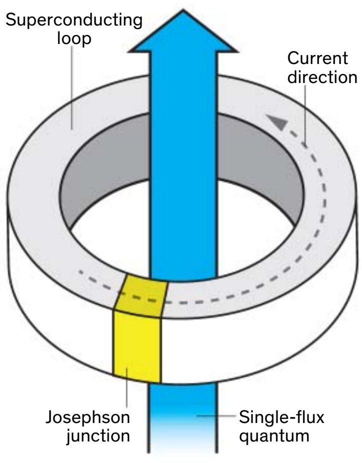

SINGLE-FLUX QUANTUM: Current in a superconducting loop [left] containing a Josephson junction—a nonsuperconducting barrier—generates a magnetic field with a tiny, quantized value.

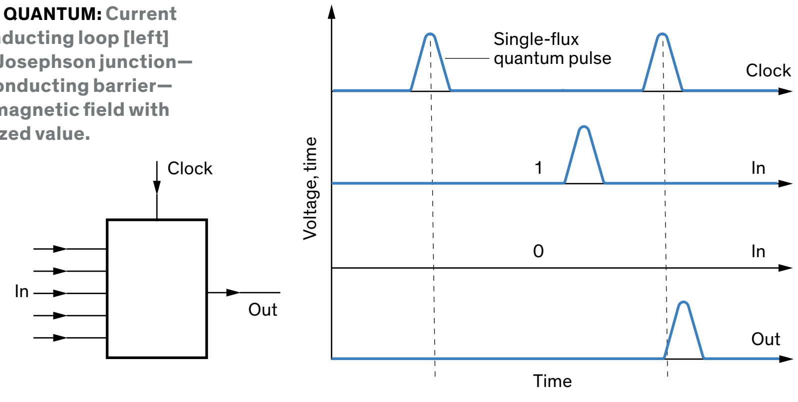

ON THE MOVE: Flux ejected from the loop through a Josephson junction can take the form of tiny voltage pulses [above]. The presence or absence of a pulse in a given period of time can be used to perform computations.

FOR MANY RESEARCHERS working on superconducting electronics, the period around 2000 marked a shift to an entirely different direction: quantum computing. This new direction was inspired by the 1994 work of mathematician Peter Shor, then at Bell Labs, which suggested that a quantum computer could be a powerful cryptanalytical tool, able to rapidly decipher encrypted communications. Soon, projects in superconducting quantum computing and superconducting digital circuitry were being sponsored by the NSA and the U.S. Defense Advanced Research Projects Agency. They were later joined by IARPA, which was created in 2006 by the Office of the Director of National Intelligence to sponsor intelligence-related R&D programs, collaborating across a community that includes the NSA, the Central Intelligence Agency, and the National Geospatial-Intelligence Agency.

Nobody knew how to build a quantum computer, of course, but lots of people had ideas. At IBM and elsewhere, engineers and scientists turned to the mainstays of superconducting electronics–SQUIDs and Josephson junctions–to craft the building blocks. A SQUID exhibits quantum effects under normal operation, and it was fairly straightforward to configure it to operate as a quantum bit, or qubit.

One of the centers of this work was the NSA's Laboratory for Physical Sciences. Built near the University of Maryland, College Park–"outside the fence" of NSA headquarters in Fort Meade–the laboratory is a space where the NSA and outside researchers can collaborate on work relevant to the agency's insatiable thirst for computing power.

In the early 2010s, Marc Manheimer was head of quantum computing at the laboratory. As he recently recalled in an interview, he saw an acute need for conventional digital circuits that could physically surround quantum bits in order to control them and correct errors on very short timescales. The easiest way to do this, he thought, would be with superconducting computer elements, which could operate with voltage and current levels that were similar to those of the qubit circuitry they would be controlling. Optical links could be used to connect this cooled-down, hybrid system to the outside world–and to conventional silicon computers.

At the same time, Manheimer says, he became aware of the growing power problem in high-performance silicon computing, for supercomputers as well as the large banks of servers in commercial data centers. "The closer I looked at superconducting logic," he says, "the more that it became clear that it had value for supercomputing in its own right."

MANHEIMER PROPOSED a new direct attack on the superconducting supercomputer. Initially, he encountered skepticism. "There's this history of failure," he says. Past pursuers of superconducting supercomputers "had gotten burned… so people were very cautious." But by early 2013, he says, he had convinced IARPA to fund a multisite industrial and academic R&D program, dubbed the Cryogenic Computing Complexity (C3) program. He moved to IARPA to lead it.

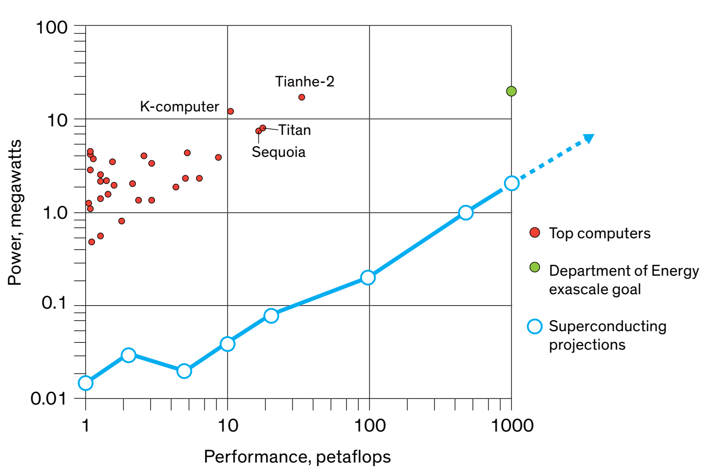

GO FOR POWER: Performance demands power. Today's most powerful supercomputers consume multiple megawatts [red], not including cooling. Superconducting computers, cryocoolers included, are projected to dramatically drop those power requirements [blue].

TOP RIGHT: HYPRES; BOTTOM: IEEE TRANSACTIONS ON APPLIED SUPERCONDUCTIVITY, VOL. 23, #1701610; MARC MANHEIMER

The first phase of C3—its budget is not public—calls for the creation and evaluation of superconducting logic circuits and memory systems. These will be fabricated at MIT Lincoln Laboratory—the same lab where Dudley Buck once worked.

Manheimer says one thing that helped sell his C3 idea was recent progress in the field, which is reflected in IARPA's selection of "performers," publicly disclosed in December 2014.

One of those teams is led by the defense giant Northrop Grumman Corp. The company participated in the late 1990s HTMT project, which employed fairly-power-hungry RSFQ logic. In 2011, Northrop Grumman's Quentin Herr reported an exciting alternative, a different form of single-flux quantum logic called reciprocal quantum logic. RQL replaces RSFQ's DC resistors with AC inductors, which bias the circuit without constantly drawing power. An RQL circuit, says Northrop Grumman team leader Marc Sherwin, consumes 1/100,000 the power of the best equivalent CMOS circuit and far less power than the equivalent RSFQ circuit.

A similarly energy-efficient logic called ERSFQ has been developed by superconducting electronics manufacturer Hypres, whose CTO, Oleg Mukhanov, is the coinventor of RSFQ. Hypres is working with IBM, which continued its fundamental superconducting device work even after canceling its Josephson-junction supercomputer project and was also chosen to work on logic for the program.

Hypres is also collaborating with a C3 team led by a Raytheon BBN Technologies laboratory that has been active in quantum computing research for several years. There, physicist Thomas Ohki and colleagues have been working on a cryogenic memory system that uses low-power superconducting logic to control, read, and write to high-density, low-power magnetoresistive RAM. This sort of memory is another change for superconducting computing. RSFQ memory cells were fairly large. Today's more compact nanomagnetic memories, originally developed to help extend Moore's Law, can also work well at low temperatures.

The world's most advanced superconducting circuitry uses devices based on niobium. Although such devices operate at temperatures of about 4 kelvins, or 4 degrees above absolute zero, Manheimer says supplying the refrigeration is now a trivial matter. That's thanks in large part to the multibillion-dollar industry based on magnetic resonance imaging machines, which rely on superconducting electromagnets and high-quality cryogenic refrigerators.

One big question has been how much the energy needed for cooling will increase a superconducting computer's energy budget. But advocates suggest it might not be much. The power drawn by commercial cryocoolers leaves "considerable room for improvement," Elie Track and Alan Kadin of the IEEE's Rebooting Computing initiative recently wrote. Even so, they say, "the power dissipated in a superconducting computer is so small that it remains 100 times more efficient than a comparable silicon computer, even after taking into account the present inefficient cryocooler."

For now, C3's focus is on the fundamental components. This first phase, which will run through 2017, aims to demonstrate core components of a computer system: a set of key 64-bit logic circuits capable of running at a 10-GHz clock rate and cryogenic memory arrays with capacities up to about 250 megabytes. If this effort is successful, a second, two-year phase will integrate these components into a working cryogenic computer of as-yet-unspecified size. If that prototype is deemed promising, Manheimer estimates it should be possible to create a true superconducting supercomputer in another 5 to 10 years.

Such a system would be much smaller than CMOS-based supercomputers and require far less power. Manheimer projects that a superconducting supercomputer produced in a follow-up to C3 could run at 100 petaflops and consume 200 kilowatts, including the cryocooling. It would be 1/20 the size of Titan, currently the fastest supercomputer in the United States, but deliver more than five times the performance for 1/40 of the power.

A supercomputer with those capabilities would obviously represent a big jump. But as before, the fate of superconducting supercomputing strongly depends on what happens with silicon. While an exascale computer made from today's silicon chips may not be practical, great effort and billions of dollars are now being expended on continuing to shrink silicon transistors as well as on developing on-chip optical links and 3-D stacking. Such technologies could make a big difference, says Thomas Theis, who directs nanoelectronics research at the nonprofit Semiconductor Research Corp. In July 2015, President Barack Obama announced the National Strategic Computing Initiative and called for the creation of an exascale supercomputer. IARPA's work on alternatives to silicon is part of this initiative, but so is conventional silicon. The mid-2020s has been targeted for the first silicon-based exascale machine. If that goal is met, the arrival of a superconducting supercomputer would likely be pushed out still further.

But it's too early to count out superconducting computing just yet. Compared with the massive, continuous investment in silicon over the decades, superconducting computing has had meager and intermittent support. Yet even with this subsistence diet, physicists and engineers have produced an impressive string of advances. The support of the C3 program, along with the wider attention of the computing community, could push the technology forward significantly. If all goes well, superconducting computers might finally come in from the cold. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/superconducting0316

Faculty Positions in Robotics and Mechatronics

NAZARBAYEV UNIVERSITY

Nazarbayev University is seeking highly qualified full-time faculty members at the Assistant, Associate and Full Professor levels to join its rapidly growing program in Robotics and Mechatronics in the School of Science and Technology. Successful candidates must have a Ph.D. degree from a reputed research university, a demonstrated ability for research, excellent English-language communication skills and a commitment to graduate and undergraduate teaching and program development.

Launched in 2010 as the premier national university of Kazakhstan, NU's mandate is to promote the emergence of Astana as the research and educational center of Eurasia. The strategic development of this English-language university is based on the Western model via partnerships with top ranking world universities. The university is committed to be a world-class research university and has made significant investments in research infrastructure, equipment and personnel.

Applications are particularly encouraged from candidates with expertise and research interests in the areas of biorobotics, unmanned aerial vehicles, motion planning, and mechanical systems design. Exceptional candidates with research interests in the broader field of Robotics and Mechatronics are also encouraged to apply.

Benefits include an internationally competitive salary, international health care coverage, free housing (based on family size and rank), child educational allowance, and home-leave travel twice per year.

Applicants should send a detailed CV, teaching and research statements, and list of publications to **sst@nu.edu.kz**. Review of applications will begin immediately but full consideration will be given to applications submitted no later than May 1st, 2016. Successful appointments are expected to begin on August 1st, 2016. For more information please visit **http://sst.nu.edu.kz**.

csem

Alpnach
Landquart
Muttenz
Neuchâtel
Zurich

The high-tech activities of CSEM at both Swiss and International level range from applied research to the development of industrial products. The high degree of expertise of our staff is our strength. CSEM is a private, non-profit Swiss company engaged in applied research, including technology, strategy and innovation consulting, and focused on generating value for a sustainable world.

Our Micro & Nanosystems Division in Neuchâtel and its high-tech clean room are at the peak of current technology. Building on recent changes within the management of the division, we are seeking a person with natural personal and leadership skills to take up the position of

Division Head Micro & Nanosystems

Your mission

In close collaboration with the Director of the MEMS programme, you are responsible for defining and communicating the strategic direction of the division and its areas of activity, together with CSEM's management board and scientific board. On the one hand, your job will be to guide and manage a team of around 60 people, directing and motivating them to achieve their objectives. On another, you are in charge of the division's operations, including setting and monitoring its budget. You are responsible for the smooth running of the division's business activities and portfolio of projects, as well as for project acquisition and negotiation of the most important contracts. You work closely with the other divisions within CSEM, as well as with our academic and industrial partners in Switzerland and Europe. Furthermore, you will develop and maintain strong contact with local political authorities.

Your profile

You are an EPF Engineer (or equivalent) with a PhD in microtechnics, physics or material science, and a minimum of 5 years' experience in an industrial management position. As recognized expert in micro & nanosystems, you master both the high level overview as well as in-depth specifics. Appreciated as an inspiring and motivating leader, you are able to listen and instil confidence in your team. You enjoy taking on new challenges, keep cool under pressure, and have a clear, results-oriented approach to your work. You are fluent in French and English; knowledge of German is a plus.

Working at CSEM

We offer a stimulating, multidisciplinary environment with the opportunity to work with leading Swiss and international companies. You will benefit from excellent social security conditions, and have opportunities to evolve within a multicultural company with a strong employee-driven culture.

We look forward to receiving your application at jobs@csem.ch, or by mail at our headquarters. Preference will be given to professionals applying directly.

CSEM SA
Human Resources ¦ Jaquet-Droz 1 ¦ P.O. Box ¦ CH-2002 Neuchâtel
Phone : +41 32 720 5111 ¦ jobs@csem.ch ¦ www.csem.ch

MISSOURI S&T University of Science & Technology

Faculty Position in Smart Cyber-Physical Systems
Computer Science
Electrical and Computer Engineering
Engineering Management and Systems Engineering
Missouri University of Science and Technology

The Departments of Computer Science, Electrical and Computer Engineering and Engineering Management and Systems Engineering at the Missouri University of Science and Technology are seeking a **tenure-/tenure-track faculty member** (open rank) in the broad area of **SMART CYBER-PHYSICAL SYSTEMS**. This is intended to build on existing research strengths in cyber-physical systems, communications (e.g. extreme connectivity, reliability, dynamic reconfiguration), smart environments (e.g., smart grid, smart transportation, smart healthcare, disaster response and extreme event management), decision-making, data mining and analytics, and security to address key challenges facing smart and connected communities.

Missouri S&T is located in Rolla and offers abundant outdoor recreational opportunities in the beautiful Ozark Mountain region. Information about the university and the departments can be found at: **https://cs.mst.edu/** for CS; **http://ece.mst.edu/** for ECE; and **http://emse.mst.edu/** for EMSE.

Candidates must submit an electronic application at **jobs.mst.edu**. They must include Position Reference Number 00065395. For further details please see the following site using job ID 18533: **https://myhr.umsystem.edu/psp/tamext/ROLLA/HRMS/c/HRS_HRAM.HRS_CE.GBL?SiteId=10**

Missouri S&T is an AA/EEO employer and does not discriminate on the basis of race, color, religion, national origin, sex, sexual orientation, gender identity, gender expression, age, disability or status as a protected veteran. Females, minorities, and persons with disabilities are encouraged to apply. The university participates in E-Verify. For more information on E-Verify, please contact DHS at: 1-888-464-4218.

PAST FORWARD_BY EVAN ACKERMAN

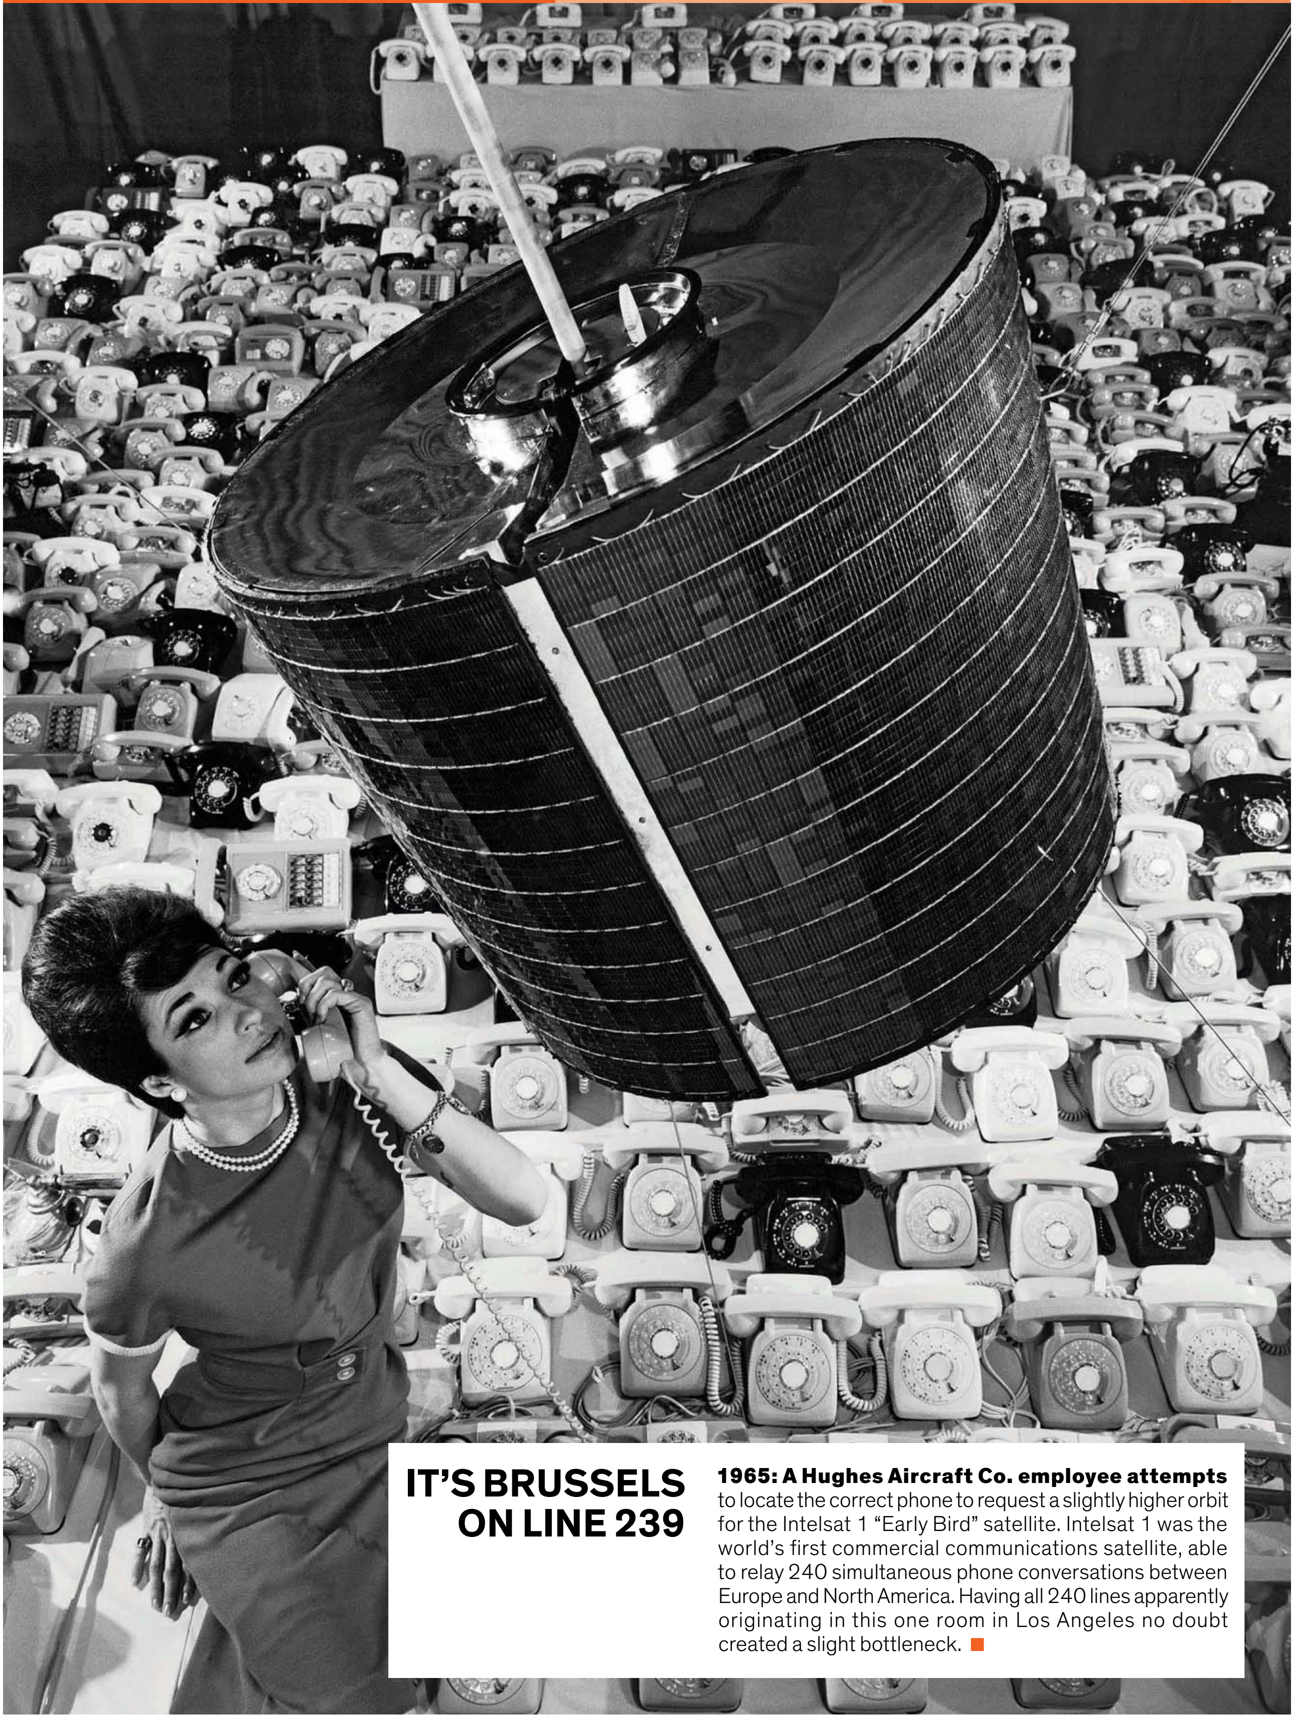

IT'S BRUSSELS ON LINE 239

1965: A Hughes Aircraft Co. employee attempts to locate the correct phone to request a slightly higher orbit for the Intelsat 1 "Early Bird" satellite. Intelsat 1 was the world's first commercial communications satellite, able to relay 240 simultaneous phone conversations between Europe and North America. Having all 240 lines apparently originating in this one room in Los Angeles no doubt created a slight bottleneck. ■

UNDERWOOD ARCHIVES/GETTY IMAGES

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

Recruit the perfect candidate
for your faculty position!

IEEE MEDIA
Recruitment Services

Place your faculty position in IEEE Spectrum and you're sure to reach the most qualified candidates!

Over 28% of IEEE members are deans, professors and instructors who work in universities and colleges or independent & university research, test & design labs – and they all rely on IEEE Spectrum to help them advance their careers.

This is why Engineering and Computer Science departments of many worldwide academic institutions choose IEEE members to enhance their faculty and research lab staff – because they are the most qualified and sought after candidates.

Contact us at employerjobsite@ieee.org

to learn more about our Special Recruitment Advertising Promotion.

Don't delay – Recruit the perfect professor!

IEEE

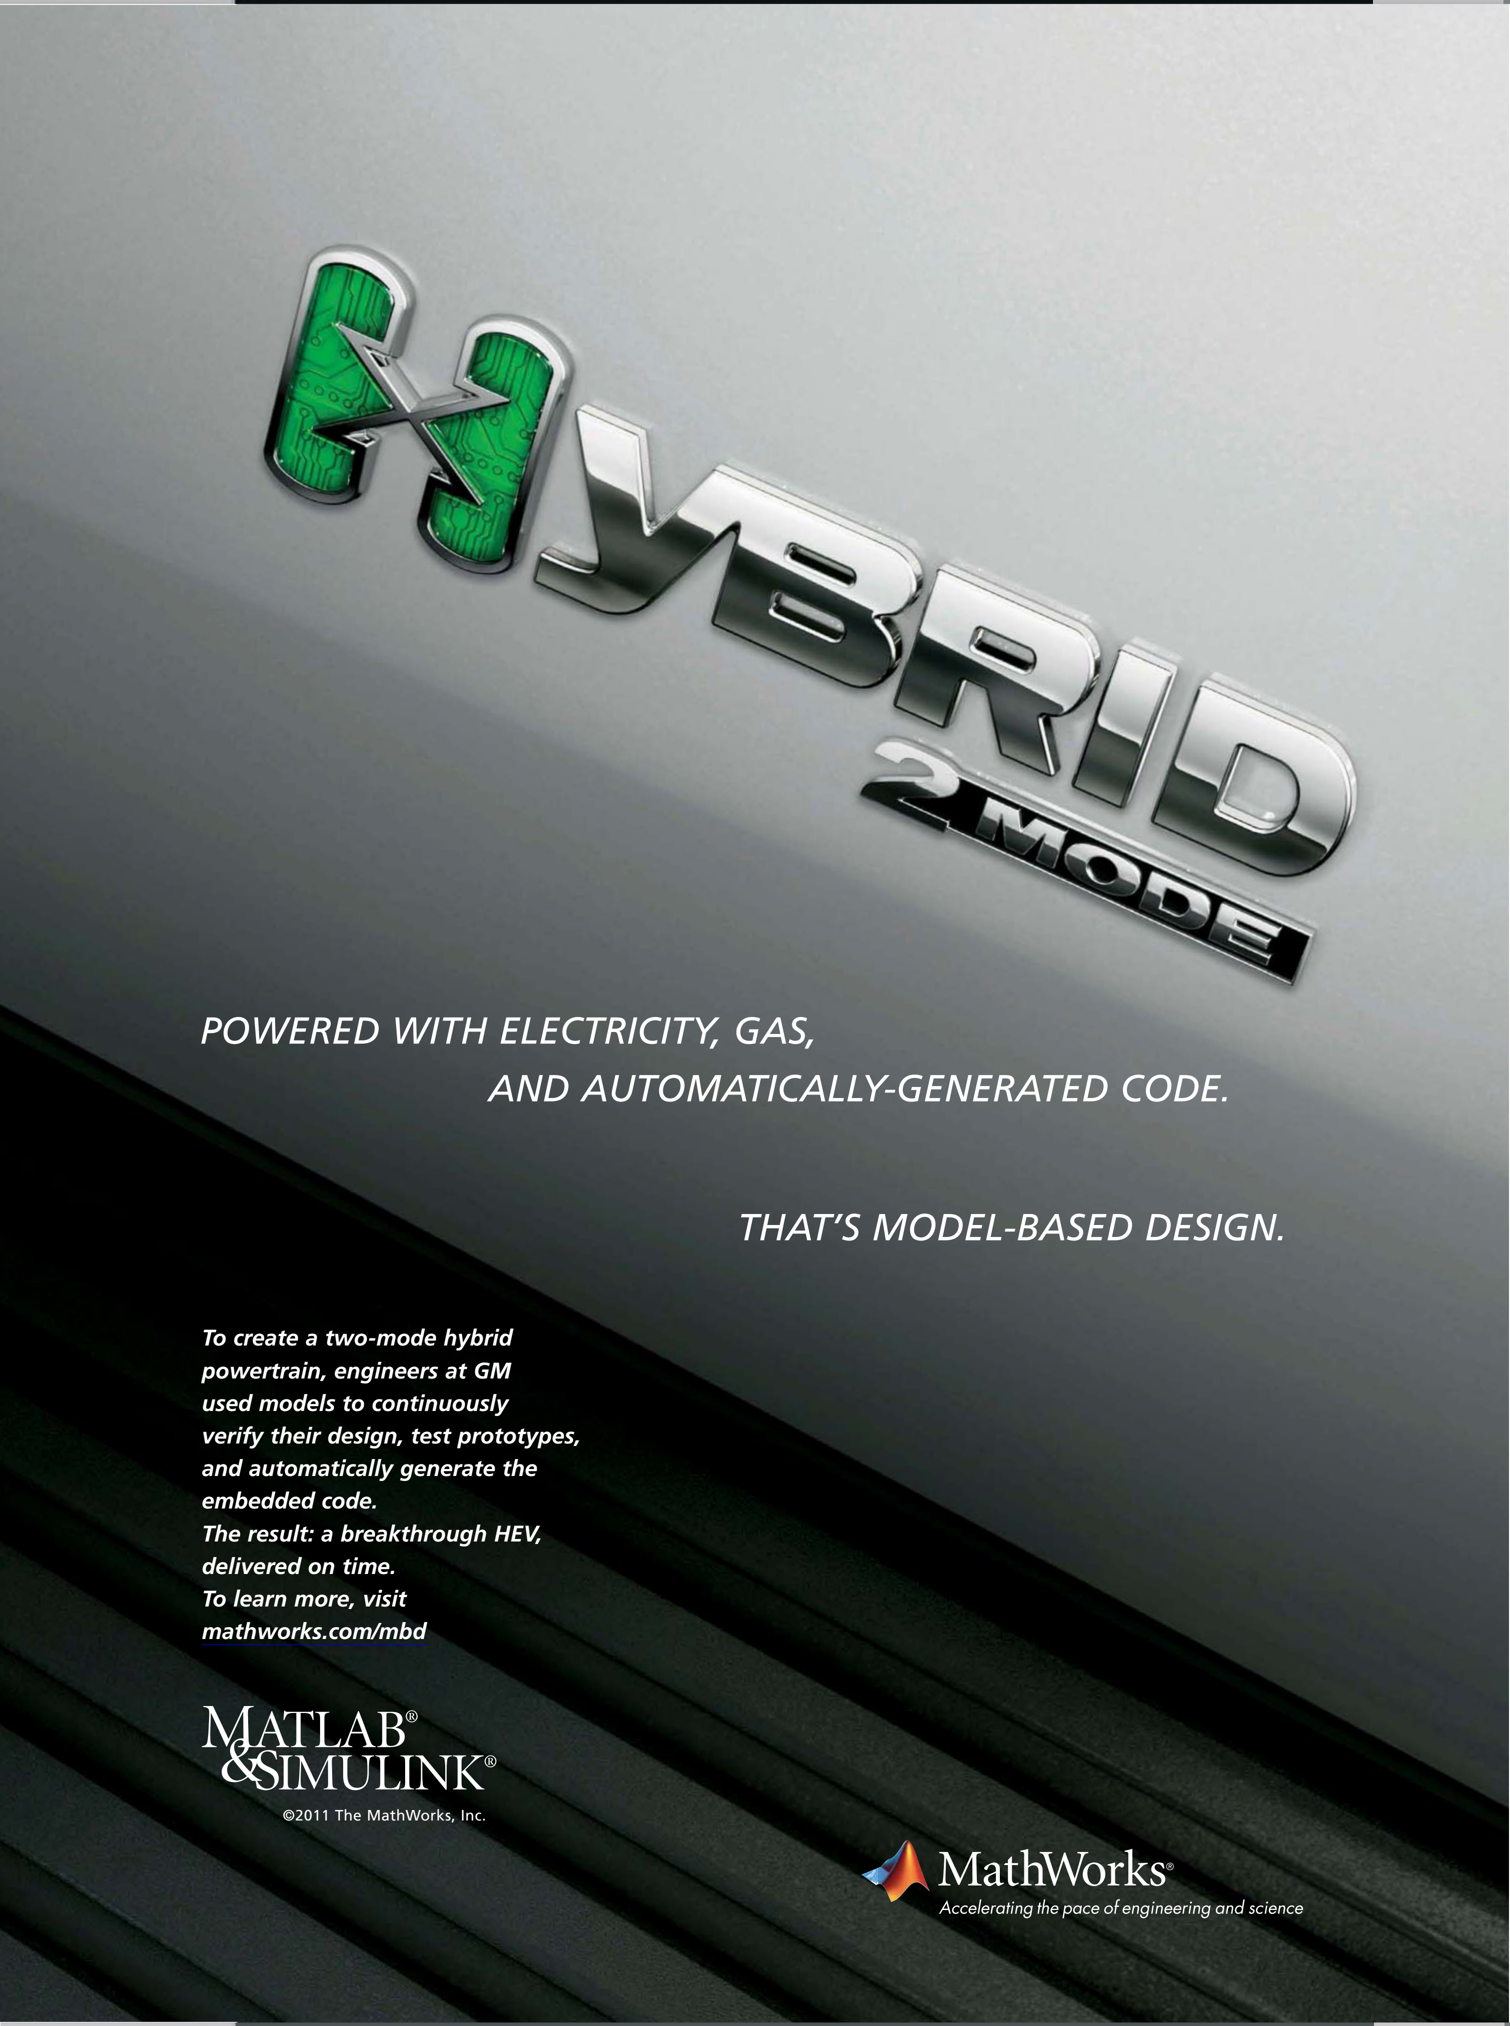
HYBRID
2 MODE
POWERED WITH ELECTRICITY, GAS,
AND AUTOMATICALLY-GENERATED CODE.
THAT'S MODEL-BASED DESIGN.
To create a two-mode hybrid powertrain, engineers at GM used models to continuously verify their design, test prototypes, and automatically generate the embedded code.
The result: a breakthrough HEV, delivered on time.
To learn more, visit mathworks.com/mbd
MATLAB® & SIMULINK®
©2011 The MathWorks, Inc.
MathWorks®
Accelerating the pace of engineering and science