IEEE Spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
8.12
SPECTRUM.IEEE.ORG

Building A Brain

An analog-inspired digital computer will succeed where others failed

Nuclear plant life extension: How long is too long?

How to keep technology from co-opting your life

IEEE

volume 49 number 8 international

ieee spectrum

Winner of the 2012 National Magazine Award for General Excellence

8.12

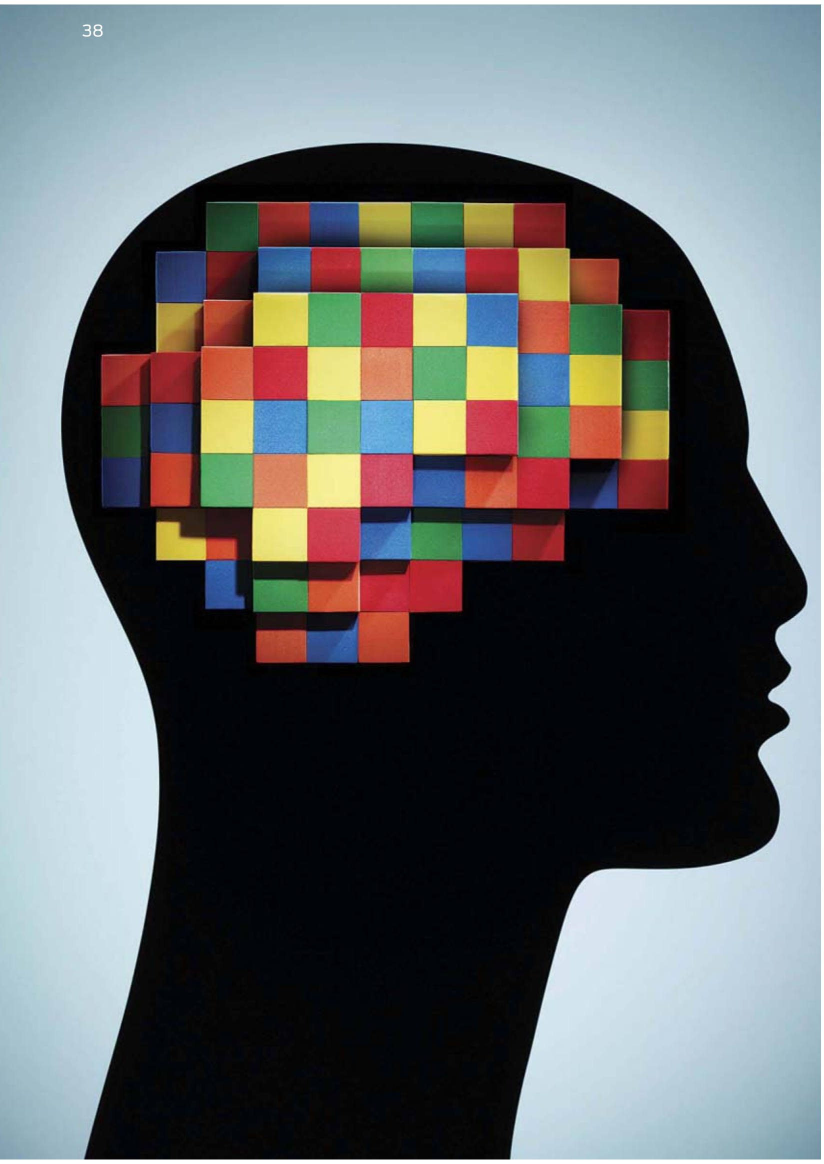

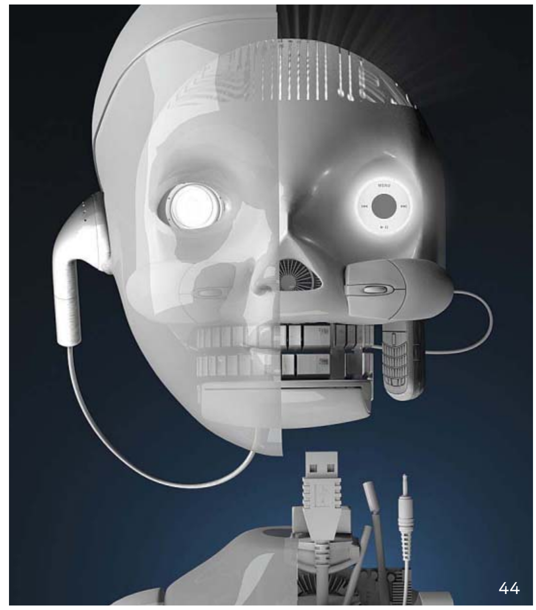

COVER AND ABOVE: PHOTOGRAPHER: DAN SAELINGER; PROP STYLIST: ARIANA SALVATO; TOP RIGHT: EDDIE GUY; BOTTOM RIGHT: EXELON NUCLEAR

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


Name
Dr. Dave Barrett
Job Title
Professor, Mechanical Engineering
Area of Expertise
Robotics
LabVIEW Helped Me
Bridge the gap between teaching theory and real-world design experience
Latest Project
Building a robotic tuna to swim across the Atlantic Ocean
NI LabVIEW
LabVIEW makes me better by making complex
CONTROL
simple and accessible

volume 49 number 8 international

ieee spectrum 8.12

12

UPDATE

OPINION

DEPARTMENTS

SPECTRUM.IEEE.ORG AVAILABLE 1 AUGUST

Did Bill Gates Steal the Heart of DOS?

The story of Gary Kildall, a computer pioneer who could have sold an operating system to IBM, is a matter of legend. Kildall didn't make the sale, Gates did, and the rest is history. But a cloud has hung over that history, with some people suggesting that the operating system sold by Gates was stolen from Kildall. Until recently, there's been no way to prove that claim true or false. But today, software forensic tools can answer the question once and for all.

THEINSTITUTE.IEEE.ORG AVAILABLE 8 AUGUST

BRINGING TECH HISTORY TO THE CLASSROOM
It's important for students to understand where their favorite gadgets came from and how today's technologies progressed from ideas to reality. That's why the IEEE History Center has developed a program to help preuniversity teachers incorporate the history of technology into their courses.

E-HEALTH COMES TO INDIA
The IEEE Humanitarian Technology Challenge recently launched a pilot program in Ahmedabad, India, called RFID Individual Tracking and Records Management. It uses RFID technology to track the health records of people in impoverished communities.

SPOTLIGHT ON PROCEEDINGS: QUALITY OF LIFE TECHNOLOGY
This month's *Proceedings of the IEEE* focuses on intelligent systems that augment body and mind functions for self-determination by older adults and people with disabilities.

ONLINE WEBINARS & RESOURCES

AVAILABLE AT http://spectrum.ieee.org/webinar

7 August: Rethink Your Systems-Engineering Approach to Combat Complexity in Product Development

23 August: Modeling Electromagnetic Structures in VSim EM From Vorpal

Modeling Magnetrons and TWTs With Vorpal

Applying Test-Driven Development for Embedded and Real-Time Development Using Model-Based Testing

Introducing COMSOL Multiphysics Version 4.3

Symbolic Computation Techniques for Advanced Mathematical Modeling

Modernist Circuit Topologies and Transistors for Smart-Grid Power Electronics

Embedded Meets Mobility—M2M Considerations and Concepts

Achieving Medical Device Software Compliance Under IEC 62304

ANSYS With Audio: A Chip Package System (CPS) Solution

Smart-Energy Technology and Application-Specific Solutions Training

Simulating Smart Materials

New product release library
http://spectrum.ieee.org/static/new-product-release-library

MasterBond white paper library
http://spectrum.ieee.org/static/masterbond-whitepaper-library

National Instruments resource library
http://spectrum.ieee.org/static/national-instruments-resource-library

CLOCKWISE FROM LEFT: JOHN MACNEILL; GLUEKIT; GETTY IMAGES

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2012 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of American Business Media, the Magazine Publishers of America, and 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

The Brain Maker

IN THE late 1970s, IEEE Fellow Steve Furber was a mild-mannered research fellow in aerodynamics at the University of Cambridge when he was snatched up by fledgling start-up Acorn Computers. There, in just a few years, he codesigned Britain's popular BBC Microcomputer System and the 32-bit ARM microprocessor, which debuted in 1985 and now forms the core logic of nearly every single mobile phone on the planet.

And he did all that before he realized what he really wanted to do with his life.

Furber left Acorn in 1990 for a professorship at the University of Manchester. He watched, delightedly, as ARM rose to dominate the low-power processor market. But he wasn't satisfied. "Processors were getting so much more efficient and powerful, but they still struggled to do things that 2-year-olds find very straightforward," he recalls.

For example, human brains excel at pattern recognition. We can recognize a face in shadow or in direct sun, in profile or from the front, and in blurred or cutoff photos. But in computer circuitry, matching an input pattern with a stored one is notoriously brittle. Unless there is an exact match, the memory will come up empty. In the late 1990s, Furber set out to see if he could do better. "But every time I tried to resolve it in circuitry, I ended up designing a match detector that looked like a neuron," Furber says. So he decided to devote himself full-time to brain modeling.

He used as his starting point the low-power ARM processor. By designing a novel 18-core chip and communication strategies to connect tens of thousands of those chips, he and his colleagues came up with a new kind of digital computer, called SpiNNaker, for Spiking Neural Network Architecture. As Furber writes in this issue, the massively parallel machine they are now building combines some of the best features of analog chips and digital supercomputers.

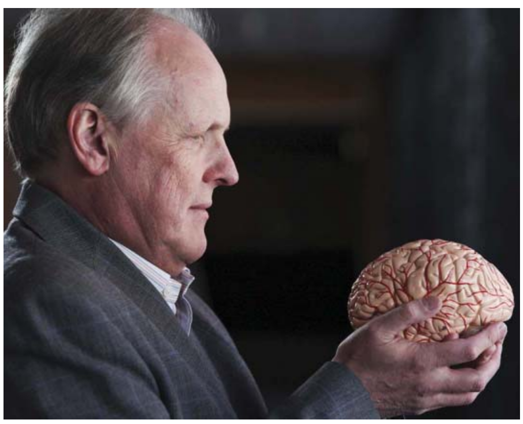

UNIVERSITY OF MANCHESTER

SpiNNaker's modest goal is to model neurons flexibly and in more or less real time. But Furber believes that models like SpiNNaker will one day help shed light on the big, profound issues: the origins and nature of self-awareness and consciousness. "I don't buy the argument that a brain can never understand itself," he says. ❏

CITING ARTICLES IN IEEE SPECTRUM

IEEE Spectrum publishes two editions. In the international edition, the abbreviation INT appears at the foot of each page. The North American edition is identified with the letters NA. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, The Data is in *IEEE Spectrum,* Vol. 49, no. 8 (INT), August 2012, p. 52, or in *IEEE Spectrum*, Vol. 49, no.8 (NA), August 2012, p. 64.

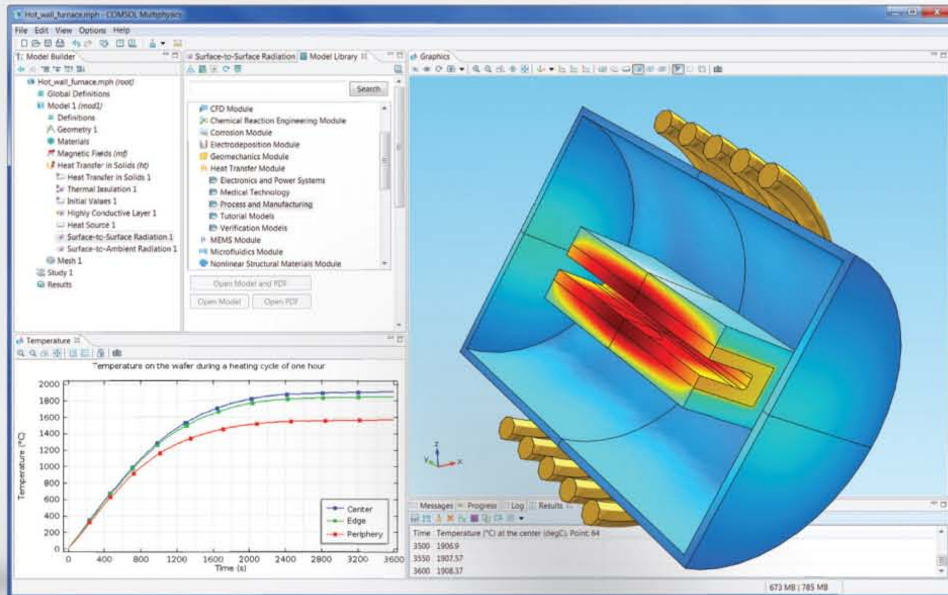

contributors

LEONARD J. BOND has worked on monitoring and prognostics for technologies like jet engines and nuclear reactors, including predicting the useful life span of these aging systems. He writes about the latter in "Old Reactors, New Tricks" [p. 24]. After 15 years as a laboratory fellow at the U.S. Department of Energy's Pacific Northwest National Laboratory, Bond recently moved to Iowa State University to direct the Center for Nondestructive Evaluation.

WILLIAM H. DAVIDOW is no ordinary technophobe, as he explains right at the start of his essay "Our Tools Are Using Us" [p. 44]. He trained as an electrical engineer at Dartmouth College, Caltech, and Stanford, then held positions at Intel, Hewlett-Packard, and General Electric before cofounding Mohr Davidow Ventures. He has written widely about how technology influences social institutions, most recently in *Overconnected: The Promise and Threat of the Internet*, which came out in 2011.

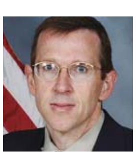

JAMES C. LYKE, a senior member of the IEEE, is technical advisor to the U.S. Air Force Research Laboratory's Space Electronics Branch, in New Mexico. In "Plug-and-Play Satellites" [p. 30], Lyke describes his group's efforts to design modular satellites that are less expensive and quicker to build than traditional orbiters. "Coming up with the satellite architecture was pretty hard," he says. "But convincing a risk-averse aerospace industry to even consider our approach has been even harder."

PAUL McFEDRIES has written *IEEE Spectrum*'s Technically Speaking column since 2002. In this issue he introduces "tweetation," a new kind of research citation, and other terms for gauging scholarly influence [p. 23]. This year HP Press published *Cloud Computing: Beyond the Hype*, the latest of his more than 70 books. His website, Word Spy, tracks emerging words and phrases. His "lexpionage," as he calls it, extends beyond technology: He recently wrote a post about "shtick lit," a genre in which an author takes on a stunt-like project in order to write about it.

PRACHI PATEL, a *Spectrum* contributing editor, writes in "The $10 000 College Degree" [p. 22] about a university in Texas that offers a four-year program in IT for about the cost of one year elsewhere. Her byline also appears this month on a story about a new technique that uses quantum dots to brighten the color of LCDs [p. 11]. Patel, who holds a master's in electrical engineering from Princeton, has written for *Discover* and the websites of *Scientific American* and *Technology Review*.

DAN SAELINGER, the photographic illustrator behind this month's cover, collaborated with prop stylist Ariana Salvato to turn the concept of a computer brain into a physical design of primary color blocks and wires. For another notable *Spectrum* cover, he shot—with both a pellet gun and a camera—a water-filled lightbulb, for our "Water vs. Energy" special issue in June 2010.

CLOCKWISE FROM TOP LEFT: PACIFIC NORTHWEST NATIONAL LABORATORY; KAREN HAMMOND; JOEL PREDD; GREG KRAUSS; KIRTLAND C4 TEAM; WILLIAM DAVIDOW


THE ULTIMATE IMU.
KVH
Series 1750 IMU
Fiber Optic Gyro IMU
Actual Size
Fiber Optic Gyro-based Inertial Measurement Unit with rugged COTS design, breakthrough performance in noise and bias stability, user-configurable outputs...and a great price.
Introducing the Series 1750 IMU –
THE DEVELOPER'S WEAPON OF CHOICE.
Need more reasons to make the switch?
Download the webcast and get all the details at:
kvh.com/1750IMU
KVH
KVH INDUSTRIES
World Headquarters: 50 Enterprise Center | Middletown, RI 02842-5279 U.S.A. | info@kvh.com +1 401.847.3327
©2011-2012 KVH Industries, Inc. KVH is a registered trademark of KVH Industries, Inc.

spectral lines

Herman Chinery-Hesse

ANDERS PETTERSSON/GETTY IMAGES

Africans Dial Up Innovation

Tired of hearing friends in Accra, Ghana, fret over the specter of home invasions, Herman Chinery-Hesse launched a text-messaging security service earlier this year that both addresses the problem of poor police response to property crimes and takes advantage of tight-knit neighborhoods.

If a house comes under attack, the owner and family members can sound an alarm from any of five mobile phones. The server automatically calls a private security service in the area—as well as up to 10 neighbors selected by the customer. Chinery-Hesse, who founded one of Africa's earliest software firms in 1991 [see "The African Hacker," *IEEE Spectrum*, August 2005], has even arranged with radio stations to interrupt normal programming with alerts on robberies in progress.

The unusual crime-fighting service could arise only in sub-Saharan Africa, where roughly half of the region's 850 million people now have cellphones. To a degree unmatched anywhere else in the world, the cellphone in Africa is an engine for innovation, spawning mobile technologies that a large swath of the population can take advantage of quickly and cheaply to tackle urgent needs.

The achievement is not Africa's alone. Global device makers, notably Nokia and Samsung, have found fertile ground for mobile phones as cheap as US $20 that boast both an FM radio and the capacity to support two SIM cards (neither feature is found in U.S. phones at any price). Indian communications company Bharti Airtel delivers the first borderless service in the world, allowing Africans from countries all over Africa to pay a single common rate.

Until recently, Africans were content to absorb, adapt, and assimilate new digital technologies. But boosted by declining costs of communications and rising prosperity, they are rebelling against the idea that the best they can do is import innovations from the rest of the world. "There's been a distinct shift in outlook," says Dorothy Gordon, who heads a technology incubator in Accra.

Direct experience of Silicon Valley helps, too. Chanda Chisala, who runs a software house in Lusaka, Zambia, credits the year he spent at Stanford University with expanding his horizons. With another Zambian code writer, he's designing an ambitious algorithm to organize news and information in "an unsung African version," Chisala says, of what Google offers.

By its number of patents or commercial innovations, Africa does not yet rate a spot on any global table. But there's a growing awareness that creating novel products and services is essential to Africa's health. This year, the Economic Commission for Africa launched the inaugural Innovation Prize for Africa, which is the first official effort on the continent to recognize the power of invention.

At a recent meeting on prospects for indigenous innovation, Ghana's minister for trade and industry, Hannah Tetteh, talked openly about how Africans can find pathways toward creating hits on the scale of Facebook or Twitter.

Such talk once would have seemed quixotic or even delusional. Not anymore, because of the smash success of M-Pesa, the original mobile-money system, introduced five years ago in Kenya and now used by tens of millions of people around Africa. M-Pesa, while designed in Britain and tested in India, satisfied an urgent African need: a safe way for people working in cities to send money quickly and cheaply to their families.

Given Africa's dismal history of sustaining complex systems, duplicating M-Pesa's success may prove impossible. But the spread of mobile money sends a powerful signal to African innovators that they can profit by drawing on their own distinctive life experiences.

Consider Esoko, an Accra software company, which makes a tool that lets farmers equipped with cheap, "handy" phones coordinate with customers and each other via text messaging. To explore the needs of their customers, Esoko programmers routinely visit farmers. "By meeting them, we better understand the real-life consequences of our design decisions," says Godwin Cudjoe, an Esoko code writer.

When Africans look inward at their own experience, says Chinery-Hesse, it "gives us confidence we can find our own innovation style."

African inventors see the potential to improvise in mobile technology in much the way Africans have long done in dance, music, and the visual arts. "Steve Jobs believed that the greatest tech innovations come from the intersection of arts and science," says Chisala. "If that's true, then Africa has an inherent advantage due to its highly developed, intuitive art culture. We just have to sort out the science part."

—G. Pascal Zachary

G. Pascal Zachary is a professor of practice at the Consortium for Science, Policy, and Outcomes at Arizona State University and a frequent contributor to IEEE Spectrum.

update

more online at spectrum.ieee.org

SEARCHING: As China's mobile phone users upgrade to smartphones, its Internet search firms are making their own operating systems and phones to keep up.
PHOTO: CRISTIAN BAITG/ ISTOCKPHOTO

Chinese Search Giant Goes Mobile

A new smartphone OS is Baidu's bid to own mobile search

China's 1 billion mobile phone users are ready for an upgrade. In the biggest mobile market in the world, consumers are switching in droves from feature phones to smartphones. As these consumers start relying on their phones for Web access, Baidu, the Internet company often called the Google of China, is making a serious play to capture the mobile search market. The company has introduced its own mobile operating system for smartphones and says it's now working on projects with about 20 phone makers. The OS, of course, puts Baidu search and other Baidu services front and center. China's other major Internet companies are also jockeying for a piece of the mobile OS market, but the stakes are higher for Baidu.

Baidu's imperative is obvious: The company dominates PC-based Internet searches in China but commands only 34 percent of mobile searches, according to Analysys International, a Chinese IT research firm. To stay relevant, the company must adapt.

In late 2011 Baidu first offered its OS, which it was then calling Baidu Yi, on a Dell smartphone. That phone hasn't become popular, but Wang Jing, Baidu's vice president of engineering, says it was a useful way to "test the water." Then, in May 2012, the company announced a collaboration with the Chinese electronics maker Changhong to sell a smartphone for 899 yuan (US $141) running the Baidu OS, which is now called the Baidu Cloud Mobile Platform.

Wang, who oversees Baidu's mobile Internet division, says the OS will be used on both high-end and low-end phones, but he acknowledges that the biggest market opportunity is in phones that cost less than 1000 yuan ($157). "That's the white-hot area, and it's caught a lot of attention," he says.

Baidu Cloud was designed to help manufacturers keep costs down so they can offer cheaper phones. "We provide 100 gigabytes of free cloud storage," says Wang. That free service allows manufacturers to save money by building in less memory,

update

he says. The company also shares with handset makers the revenue generated by mobile search ads.

Baidu Cloud is an Android "fork"—it's derived from Google's open-source Android OS but was adapted to meet local needs. Android-based phones dominate the Chinese smartphone market, with almost 70 percent market share, according to Analysys International. But some Google services, like YouTube, are blocked by the Chinese government, while others, like Google Maps, simply aren't optimized for China. Baidu Cloud's suite of services, including maps, a popular music service, and an e-reader platform, may give its OS an edge over foreign competitors. Wang says the company also provides training and support to help developers build apps, which results in more apps designed specifically for Chinese consumers.

Analysys telecom analyst Wang Ying says that some medium and small mobile phone manufacturers may be happy to use Baidu Cloud, but she doesn't expect the OS to dominate the smartphone market. "Lots of Internet players are joining the mobile phone market," she notes, including the social networking and gaming company Tencent and the e-commerce company Alibaba, which have both introduced their own phones and operating systems. "The whole smartphone OS market shows fragmentation," says Wang.

Other analysts say the proliferation of smartphones will benefit Baidu even if its OS doesn't catch on. "We're starting to see the stars align for Baidu when it comes to mobile search," says Michael Clendenin, the founder of the Shanghai-based research firm RedTech Advisors. He notes that 3G subscriptions are increasing and more Chinese websites are creating mobile versions. Meanwhile, Apple recently announced that it will integrate Baidu's search function into next-generation iPhones sold in China. "It's more important that Baidu be the default search on the mobile Web browser," says Clendenin. "They could totally fail with the Baidu Cloud OS and still be wildly successful at mobile search." —Eliza Strickland

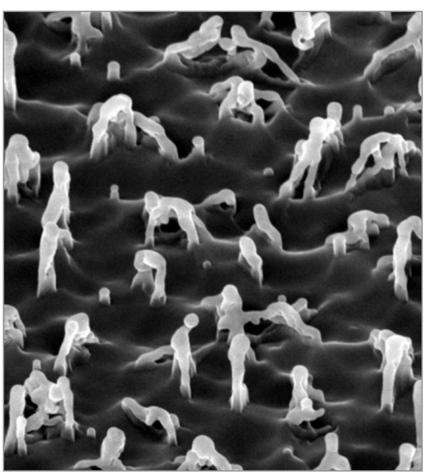

Microscopic Beach

A beach scene in black and white? That's what most people see in this scanning electron microscopy image by Lim Saw Sing of Infineon Technologies, in Germany [left]. The "beach" is actually a layer of polyimide that's been etched with reactive ions and coated in a fine film of metal, according to Lim. The image was one of the finalists in the fourth annual "Art of Failure Analysis" contest at the IEEE International Symposium on the Physical and Failure Analysis of Integrated Circuits, in July. Other images engineers found as they picked apart failed ICs include a forest scene made from the adhesion of carbon tape to a membrane and a pair of wings made from the fracture of a silicon sample.

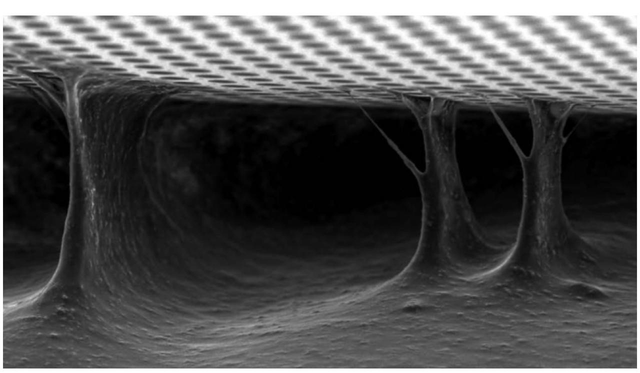

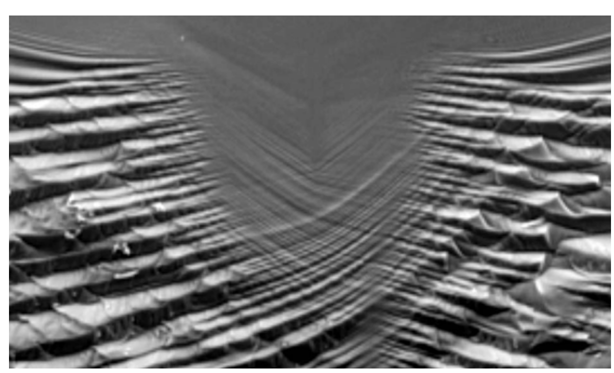

CLOCKWISE FROM TOP: LIM SAW SING; FOO FANG JIE; LIM CHAN WAY

"A lot can go wrong here today." —Google cofounder Sergei Brin, shortly before skydivers who were broadcasting video from Google's experimental augmented-reality glasses leaped from a blimp above San Francisco

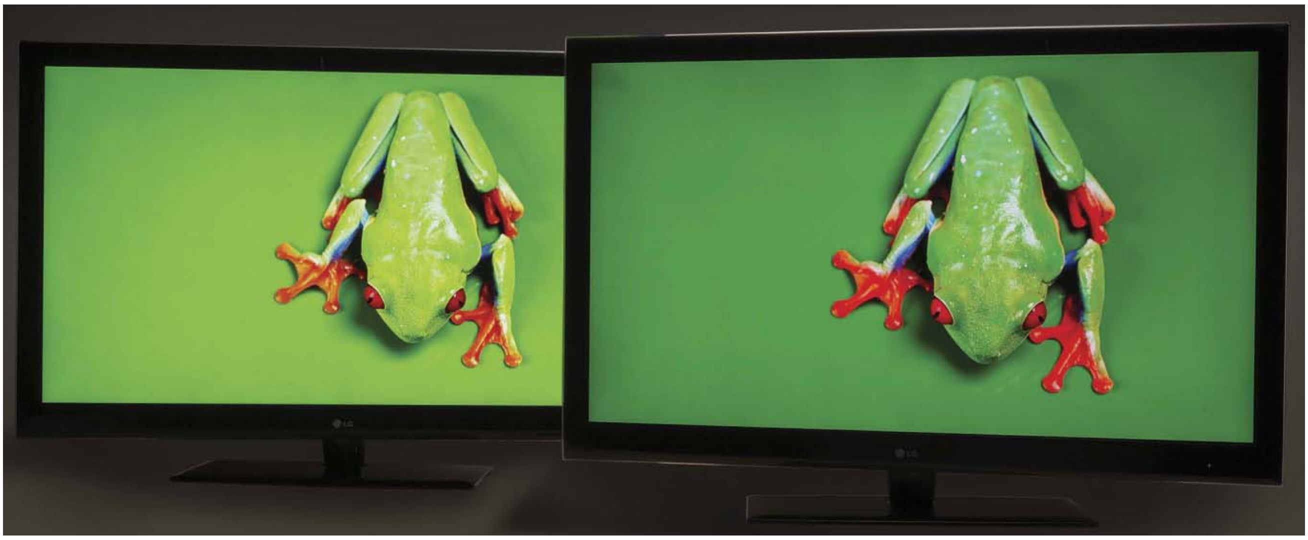

Quantum Dots Are Behind New Displays

They make LCDs brighter and could challenge OLEDs for future TV dominance

Liquid crystal displays dominate today's big, bright world of color TVs. LCDs are inefficient, though, and don't produce the vibrant, richly hued images of organic light-emitting diode (OLED) screens, which are expensive to make in large sizes.

But a handful of start-up companies have been plugging away at another display technology that could enhance LCDs and maybe unseat OLEDs: quantum dots. These light-emitting semiconductor nanocrystals shine pure colors when excited by electric current or light and promise rich, beautiful displays that would be inexpensive and easy to manufacture.

In June, quantum-dot developer Nanosys announced that it was working with 3M to commercialize a quantum-dot film that could be integrated into the back of ordinary LCD panels. The film could cut the display's power consumption by half and enable LCDs to generate 50 percent more colors within the range set by the National Television System Committee. Nanosys's CEO, Jason Hartlove, says that major LCD manufacturers are now testing the film, and a 17-inch notebook computer incorporating the technology should be on shelves within six months.

"We've designed this technology to give LCD manufacturers a competitive counterweight to OLEDs," he says. "[It] drops in and is compatible with existing LCD manufacturing methods."

New LCD TVs typically have a strip of white LED backlights along the panel's edge. Light guides spread the illumination evenly across the panel; the light then passes through a series of optical films that colorize, polarize, and diffuse it. The resulting red, green, and blue hues are mixed at different intensities to produce the display's color spectrum. But with today's technology, "the broad-spectrum white has a bunch of pinks, yellows, and oranges," Hartlove says. "So when you create a red from a color filter, you let in a bunch of different reddish colors."

Nanosys's quantum-dot film converts some of the light from highly efficient blue LEDs to spectrally pure red and green, resulting in a broad color gamut and more-lifelike images.

QD Vision, based in Lexington, Mass., is developing its own LCD backlight technology, which literally and figuratively has an edge over Nanosys's, says CTO Seth Coe-Sullivan. Instead of using a large film of quantum dots, its product would integrate the quantum dots into the LED array on the panel's edge. "So we use 1/100 of the material, which is not inexpensive," he says. That advance required engineering dots that could operate efficiently at the higher temperatures at the backlights.

QD Vision is also working on a quantum-dot-based display that emits light, the way OLED

BRIGHT IDEA: The 47-inch HDTV on the right has been modified with Nanosys's quantum-dot film to enhance the performance of its backlight. *PHOTO: NANOSYS*

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

update

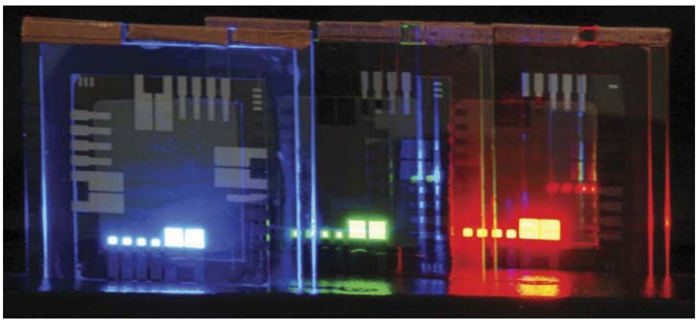

BRIGHT DOTS: QD Vision is creating a display from quantum-dot LEDs.
PHOTO: QD VISION

displays work, rather than transmitting it as LCDs do. OLED pixels contain a thin layer of light-emitting organic semiconductor sandwiched between two electrodes. QD Vision swaps quantum dots for the organic material.

Quantum dots can easily be suspended in solvents to make inks that can be stamped or printed without losing their glowing prowess, so the displays should be cheap to make. What's more, quantum dots are also intrinsically more efficient than organic semiconductors.

QD Vision demonstrated a 4-inch-diagonal prototype of the display last year, and the company has since teamed up with LG Display to commercialize the technology. Samsung, meanwhile, is also pursuing quantum-dot LED displays. Last year, researchers at the Samsung Advanced Institute of Technology reported 4-inch-diagonal color displays made on glass and plastic.

Quantum dots, both in LCD backlight films and as light emitters, might still find OLEDs to be tough competition. OLED displays have been relegated to small screens, but this year LG and Samsung unveiled the first 55-inch OLED TV prototypes. Samsung claims that it will launch a commercial version later in 2012 in South Korea, though at a price of US $9000—or 10 times as expensive as LCD sets—but the cost should drop as the technology improves. "There [have] been decades of work on organic emissive material, and there have been a lot of improvements," says Paul Semenza, a senior analyst at research firm DisplaySearch. "It's not like quantum dots are competing against a new thing."

Semenza says that the success of quantum-dot display technology will depend in part on developments in OLEDs and LEDs over the next few years. The LCD backlight approach shows promise because it's easy and shows visible improvements, but "as white LED backlights improve, the benefit from doing this might diminish," he says. OLEDs, meanwhile, are already beautiful but costly to manufacture. "If quantum-dot displays are easy to manufacture [and] have a lower cost process and better efficiency, they could be big." —Prachi Patel

A version of this article appeared online in June 2012.

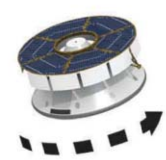

1. About 20 minutes from landing and spinning at 2 revolutions per minute, the spacecraft takes one last fix on the sun and stars for orientation.

2. The cruise stage separates about 1500 kilometers above the surface, 17 minutes from landing.

3. Thrusters are fired to stop the robot's rotation. Two 75-kilogram tungsten masses are ejected from one side of the spacecraft. Once the craft is in the atmosphere, the resulting offset in its center of mass will create a 20-degree angle of attack, providing lift.

NASA Mars Science Laboratory Spacecraft

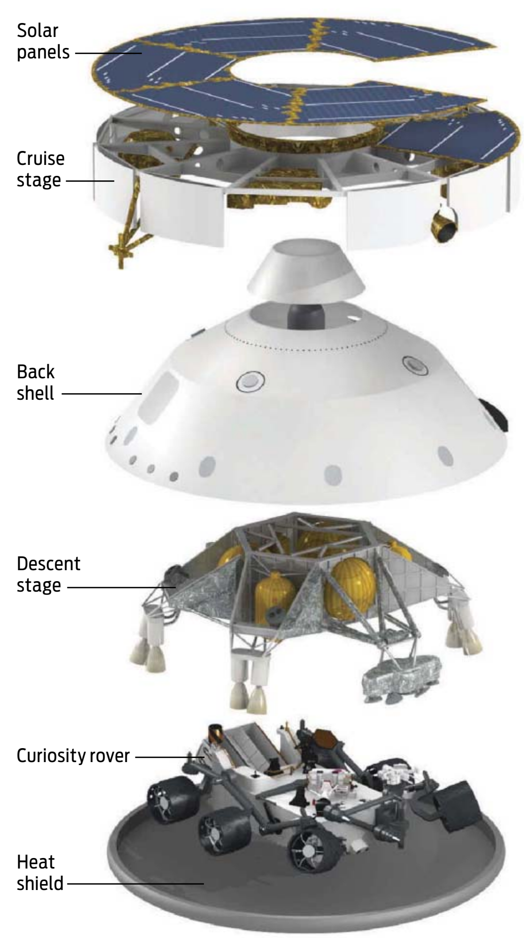

Curiosity's 1-Ton Touchdown

The Mars rover will rely on dead reckoning and radar to land on the Red Planet

On 6 August, NASA's robotic rover Curiosity is scheduled to reach the end of its eight-month journey to Mars and begin its high-stakes descent down to the surface of the Red Planet. In 7 nail-biting minutes, the spacecraft must transform a 5.9-kilometer-per-second approach into a soft landing near the foot of the Martian equatorial mountain Aeolis Mons. And with a 14-minute communications lag between Mars and Earth, Curiosity's computer will have to perform the feat all on its own.

Two sensor systems, both stored in the spacecraft's descent stage, will prove crucial to safely landing the hefty, 900-kilogram robotic rover. One of those is the spacecraft's Descent Inertial Measurement Unit (DIMU), a Honeywell-made device containing three gyros and three accelerometers. For more than half the landing, until about 2 minutes from touchdown, Curiosity will be flying blind, relying entirely on DIMU data to estimate the spacecraft's position, orientation, and speed, and by extension, its likely altitude.

NASA Jet Propulsion Laboratory's Steven Lee, who is the guidance, navigation, and control systems manager for the mission, likens this stage of the descent to riding in the passenger seat of a car with your eyes closed. You can use your inner ear and sense of touch to approximate the location of the car, assuming you know its starting point. The spacecraft essentially uses its "vestibular system to get in the vicinity," Lee says.

But to get the precision needed for the last moments of landing, the spacecraft's descent stage will need to open its eyes. About 8 km above the planet's surface, the robot's heat shield will drop away, revealing six independent pulsed-Doppler radar antennas. Distance readings from this JPL-designed, long-range radar system will be used to signal when the parachute should be deployed, when retro-rockets should fire, and when the rover should be lowered down from the descent stage. If all goes well, Curiosity will make the most precise landing yet attempted on the Red Planet, and 14 or so minutes later, a host of engineers will toast the achievement.

—Rachel Courtland

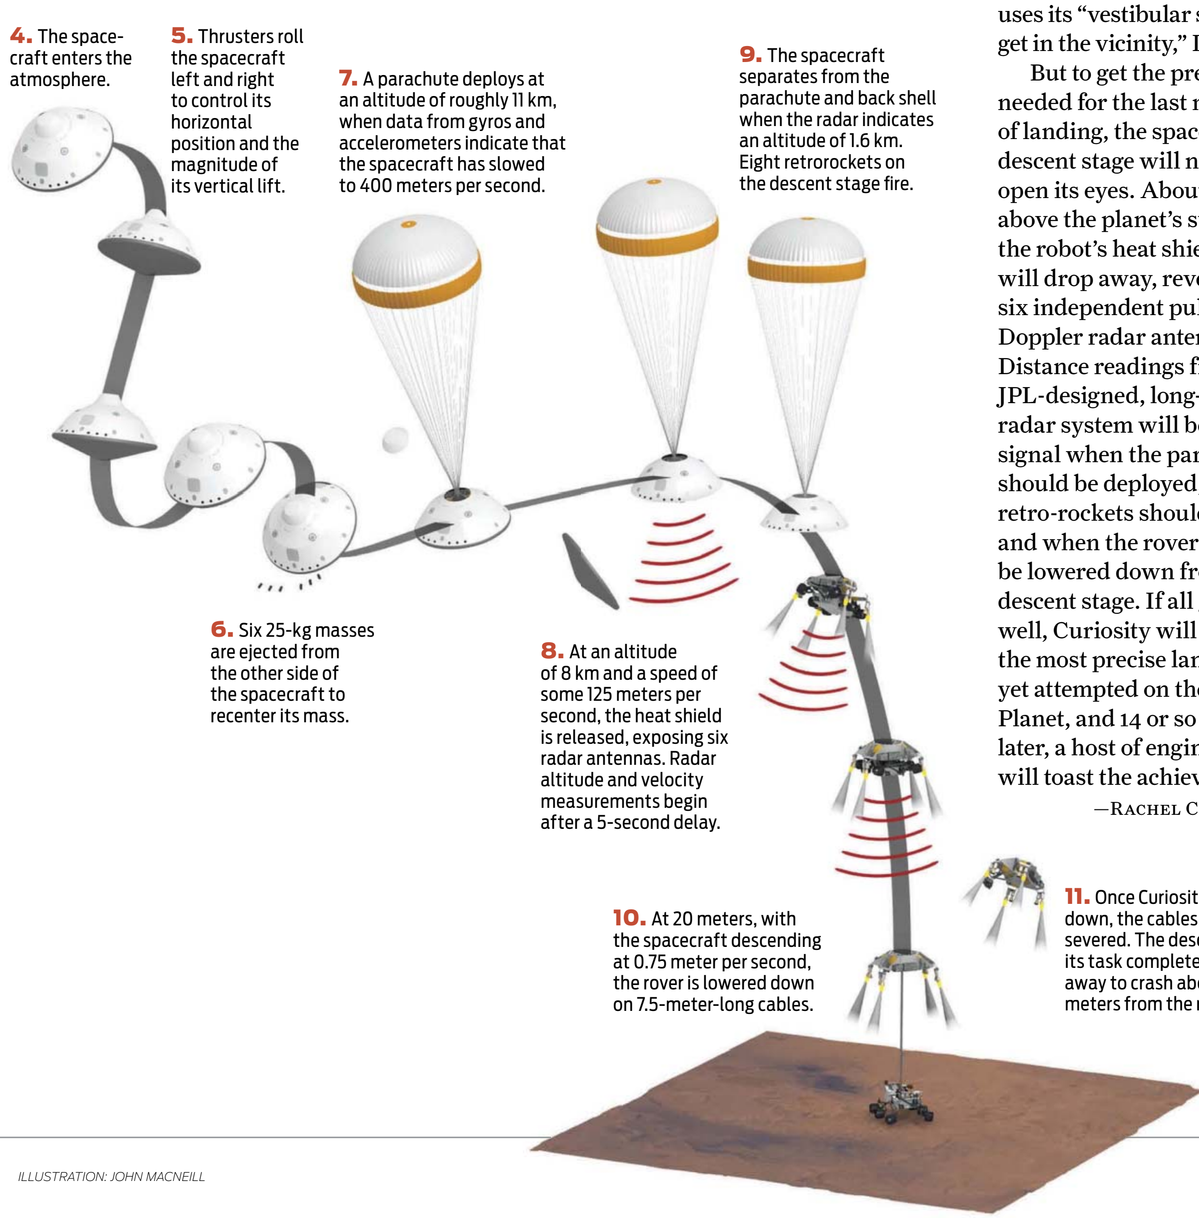

Illustration: John MacNeill

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

update

Declarations of Cyberwar

What the revelations about the U.S.-Israeli origin of Stuxnet mean for warfare

Mouths went agape when *New York Times* reporter David Sanger wrote in June that anonymous sources within the United States government admitted that the United States and Israel were indeed the authors of the Stuxnet worm and related malware. Those two countries had long been suspected of creating the code that wrecked centrifuges at Iran's Natanz uranium enrichment facility. But never before had a government come so close to claiming responsibility for a cyberattack.

The origins of the most sophisticated cyberattacks ever undertaken may now be clear, but exactly where such attacks fit in the universe of war and foreign policy—and what the international community would consider a proper response to them—is still the subject of debate.

A particularly important question is what sort of cyberattack is the equivalent of a traditional armed attack. Efforts to answer that question have culminated in the *Manual on International Law Applicable to Cyber Warfare* (also known as the *Tallinn Manual*), which will be published later this year.

The *Tallinn Manual* is a nonbinding yet authoritative restatement of the law of armed conflict as it relates to cyberwar. It offers attackers, defenders, and legal experts guidance on how cyberattacks can be classified as actions covered under the law, such as armed attacks. "The term 'armed attack' has a precise meaning in international law: Not all 'cyberattacks' rise to the level of an armed attack," says Bret Michael, a professor of computer science and electrical engineering at the U.S. Naval Postgraduate School, who has been serving as a technical expert to the group drafting the *Tallinn Manual*.

Despite this progress, the inter-

national community is just at the beginning of what could be a long process, says Charles Barry, a senior research fellow at the National Defense University's Institute for National Strategic Studies, in Washington, D.C. He predicts that it will take "another 20 to 50 years to get traction on cyberrules."

What is certain, say observers, is that going forward, conventional warfare will almost always be complemented by cyberwarfare aimed at knocking out an opponent's communications and intelligence-gathering capabilities. "Actually, that's already being done," says Michael.

Cyberattacks can aid in military campaigns, but can the threat of a military response serve as a cyberdeterrent? "That's downright silly, because it's difficult, bordering on impossible, to identify a cyberattacker beyond a shadow of a doubt," says Larry Constantine, a professor in the mathematics and engineering department at the University of Madeira, in Portugal.

However, identification beyond a shadow of a doubt might not really be needed to escalate a cyberattack into an armed conflict. In June at CyCon 2012, a NATO-sponsored cyberconflict conference in Tallinn, Estonia, U.S. Air Force Lt. Col. Forrest Hare told attendees that attribution is a political, not a legal, concept. The three standards of proof used in criminal law—"beyond a reasonable doubt," "clear and compelling," and "preponderance of the evidence"—don't apply to military and intelligence operations. Michael adds that the difficulty of reliably tracing an attack to its source does not preclude the use of other sources to weave together what he calls "a clear mosaic of responsibility." Showing who funded the activity or provided the actors with guidance may be enough.

And there is already a deterrent in the form of the law of armed conflict, says Michael. It holds military commanders or their civilian superiors who order attacks that amount to a war crime as criminally responsible.

In the meantime, governments can try to take heart in the belief that there are few nations capable of fielding a cyberweapon with the sophistication of Stuxnet. But Jeffrey Voas, a computer scientist in the computer security division at the U.S. National Institute of Standards and Technology, in Gaithersburg, Md., notes that if an attack doesn't require stealth, the code doesn't have to be nearly as artful. And there are tens of thousands of people who could pull off a less sophisticated strike, says Constantine, who built his own Stuxnet-like malware in 2003 to prove a point. In other words, powerful cyberattacks are within the range of many states, so long as they don't care if they get caught. —Willie D. Jones

BRIAN STAUFFER

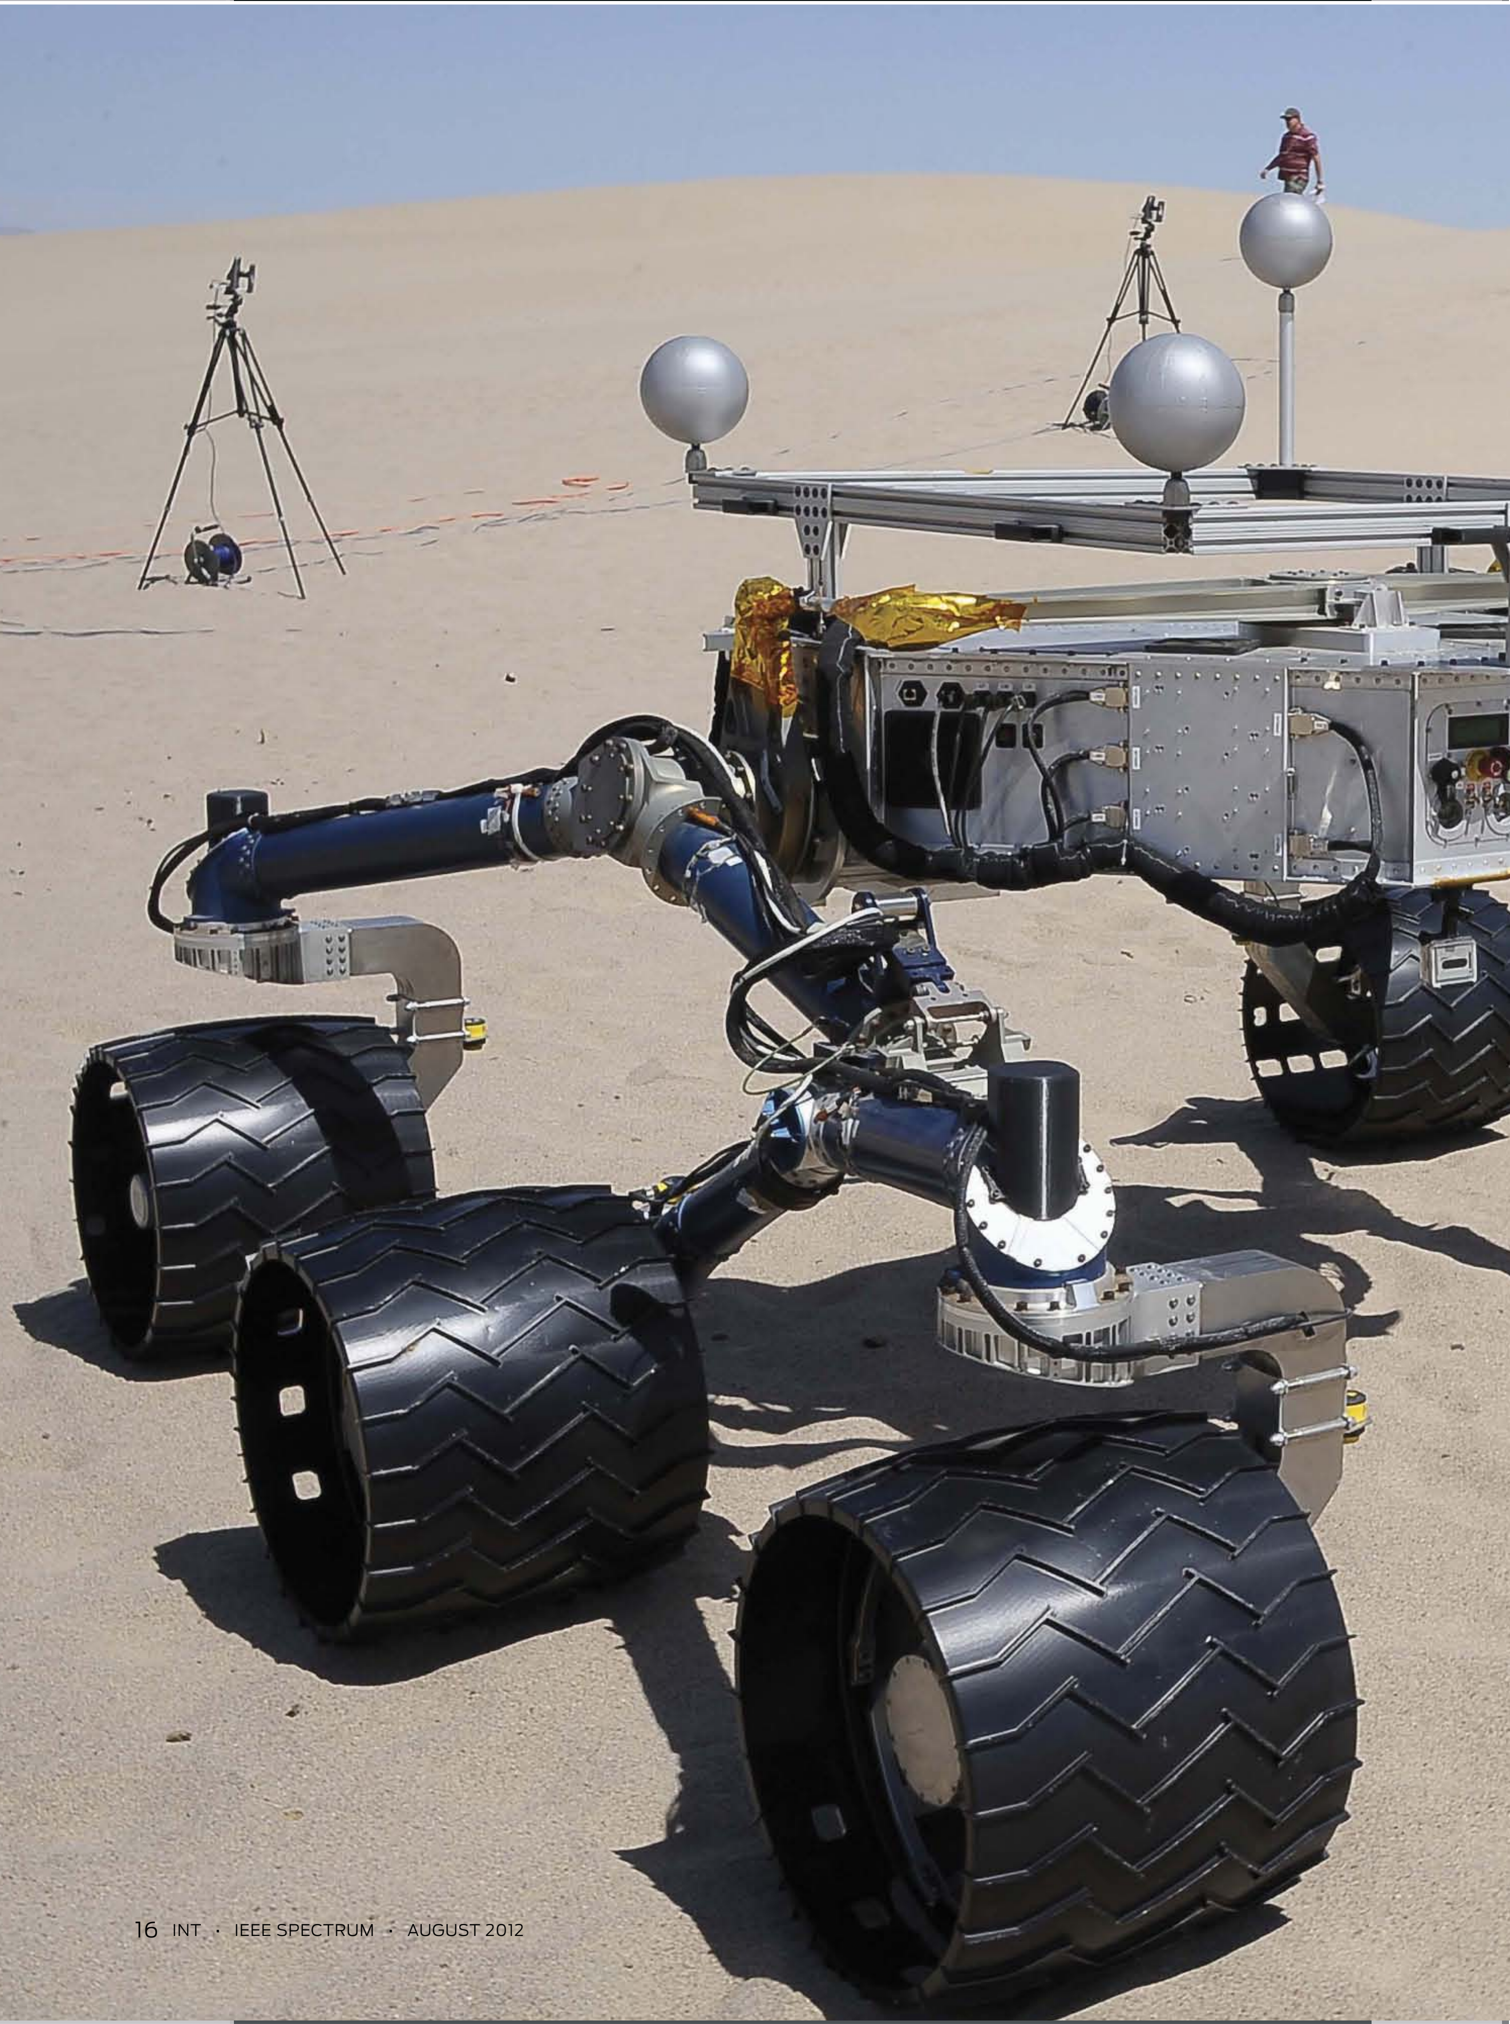

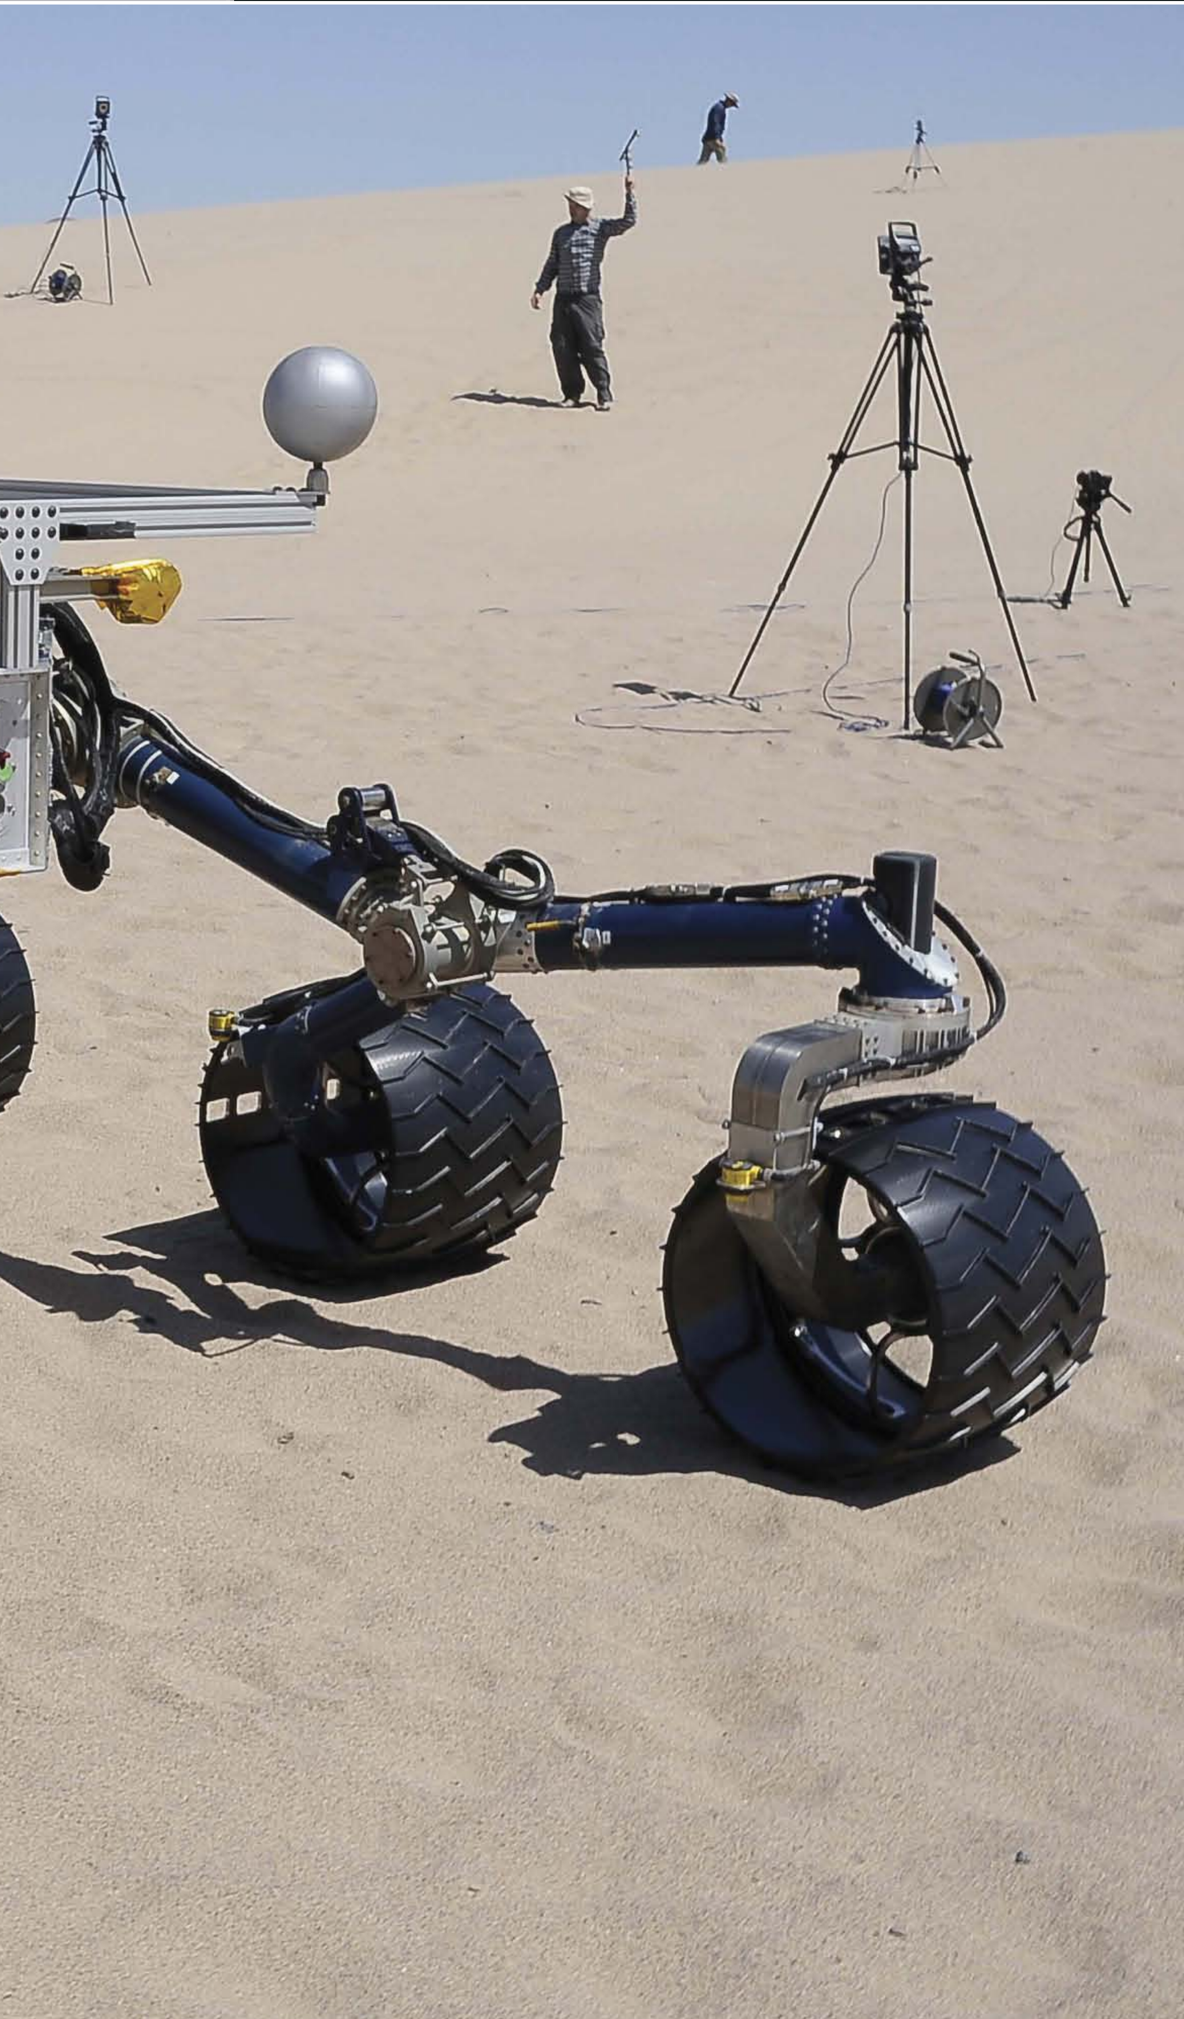

the big picture

DRIVING ON MARS

Even as the Mars rover Curiosity made its way toward the Red Planet, researchers from the Mars Science Laboratory were busily refining their ability to manipulate the vehicle remotely. In this photo, a replica of the rover is being put through its paces in the Mojave Desert near Baker, Calif. The real thing will use its array of instruments, such as a high-resolution imager that can resolve features as tiny as 12.5 micrometers, to study whether Mars has what it takes to support microbial life. The vehicle is powered by a radioisotope thermoelectric generator that will deliver enough energy to keep the mission going for nearly two Earth years. Curiosity is expected to descend into a target area about twice the size of Manhattan on or about 6 August.

PHOTO: GENE BLEVINS/REUTERS

hands on

Khan Academy Unplugged

How to run Salman Khan's educational website off-line

When a friend recently asked for volunteers to help tutor inmates at a local prison, I asked, "Why not use the Khan Academy?"

This website contains a wealth of video tutorials that span a wide range of topics and complexity, ideal for inmates seeking high school equivalency diplomas. The math offerings are probably the most complete—ranging in difficulty from basic addition to linear algebra and calculus—but the site has lessons in astronomy, art history, economics, chemistry, biology, physics, and more.

The academy has received considerable media coverage. So it's easy to find out lots about it—how it began with a few simple videos Salman Khan made to help tutor a cousin in math, how his lessons unexpectedly went viral on YouTube, how Bill Gates took an interest and funded the expansion of his work, and how the academy has become a pillar in some technologically progressive schools. Khan and a small team of programmers are now on a mission to provide "a free world-class education to anyone anywhere." And they seem well on their way.

What's harder to discover is how the Khan Academy works under the hood. In particular, the website includes not just videos but exercises generated on the fly to test whether you've understood a lesson. It also provides sophisticated tools for teachers, allowing these "coaches," as Khan labels them, to monitor how students are doing with their self-paced learning. So it's a rather complex Web application. And that made for a problem: Prisoners are not allowed Internet access.

Fortunately, the open-source Khan Academy provides an off-line-capable version of its site. Go to http://khanacademy.org and scroll to the bottom of the home page. Clicking on Downloads takes you to a short description of the academy's off-line server and a link to the download page, where you'll find three off-line-capable packages—for Windows, Linux, and Mac OS. Apart from some difficulties I had getting this to run on my 64-bit Windows 7 machine, there's a fundamental problem with the version offered there (revision 8051): When you attempt the exercises off-line, an error message reports that access to Google.com is needed.

While the ability to watch Khan's videos off-line would be a valuable thing in itself, it seemed a shame to miss out on Khan's computer-generated exercises. Happily, with a little digging I discovered that earlier versions of the Khan Academy Web application will run the exercises off-line. Return to the downloads page and search "All downloads" (leave the search box blank). You'll then see links to several previous versions.

I grabbed the file KhanAcademy-1482.zip, unzipped it, and soon had a version running that could serve exercises and coaching tools off-line. To get the videos, you'll need to follow the instructions in the Read.me file. This automates the download of some 2100 videos—fewer than the current crop (around 3200) but still an ample supply. It took me more than a day to download the 23.5-gigabyte collection.

So how does the off-line version work? The Khan Academy uses Google App Engine, a cloud service, to host its website. That is, Google serves the Khan Academy website, created dynamically by the Khan Academy's software running on Google's machines. But Google also provides a version of its app engine that can run locally—used normally for development purposes to test Web-application code before deploying it on Google's infrastructure. The off-line version of the Khan Academy runs just on this local development app server. Once everything is installed, you view it by

JESSE LEFKOWITZ

pointing your browser to http://localhost:8080.

There are some complications. For one, the Web app sometimes tries to fetch things from the Internet. In particular, the home page contains an embedded YouTube video player to show media coverage of the Khan Academy. Without an Internet connection, there's just a big blank space on the home page instead.

That looks bad. So I replaced the HTML for that YouTube video player with a couple of lines that link to an image from a "PBS NewsHour" story on the Khan Academy. The image I used is a screen shot from the opening moments of that news story. Now you see newscaster Judy Woodruff, rather than a bunch of white space, on the home page. And because not all of the news-report videos advertised on the home page come in that 23.5-GB download, I added a disclaimer line to the home page noting that only some of the media coverage can be viewed off-line. Another infelicity was a set of links to YouTube, Twitter, and Facebook at the top of the pages. As these social networking sites aren't available, I removed them from the relevant HTML template.

My biggest disappointment with the off-line version was that the exercise dashboard didn't show Khan's lovely "knowledge map," a graphical representation of how his math lessons are conceptually organized, which helps you keep track of what you've mastered and what you have next to tackle. The dashboard even indicated that the knowledge map was unavailable in off-line mode. But I found that if I navigated away from the exercise dashboard and then clicked the back button on my browser, the knowledge map would mysteriously materialize.

I don't understand the mechanism behind this odd behavior, but I figured I could nevertheless take advantage of it. In place of the notification that the knowledge map isn't available off-line, I put a link to show the knowledge map. That link sends you to a dummy page I created, which includes nothing but a bit of JavaScript that simulates hitting the back button. Users still don't see the knowledge map when they first load the exercise dashboard, but when they click on the new show-the-knowledge-map link, the map appears.

These modifications really weren't necessary, but they make the off-line version look a bit sharper. And they demonstrate the power of having access to the source code. You don't need to understand how everything works—I certainly don't—to be able to customize even a complex Web application to meet your needs.

I was trying to come up with an educational resource for inmates, but others are using the Khan Academy off-line in very different settings, such as for teaching children in remote African villages that don't have Internet access. It's easy to imagine more prosaic uses, too—you might just want to use it as a backseat "game" during family car trips. Whatever your purpose, an off-line Khan server is easy to set up, and with a little hacking you can tailor it to your liking.

—David Schneider

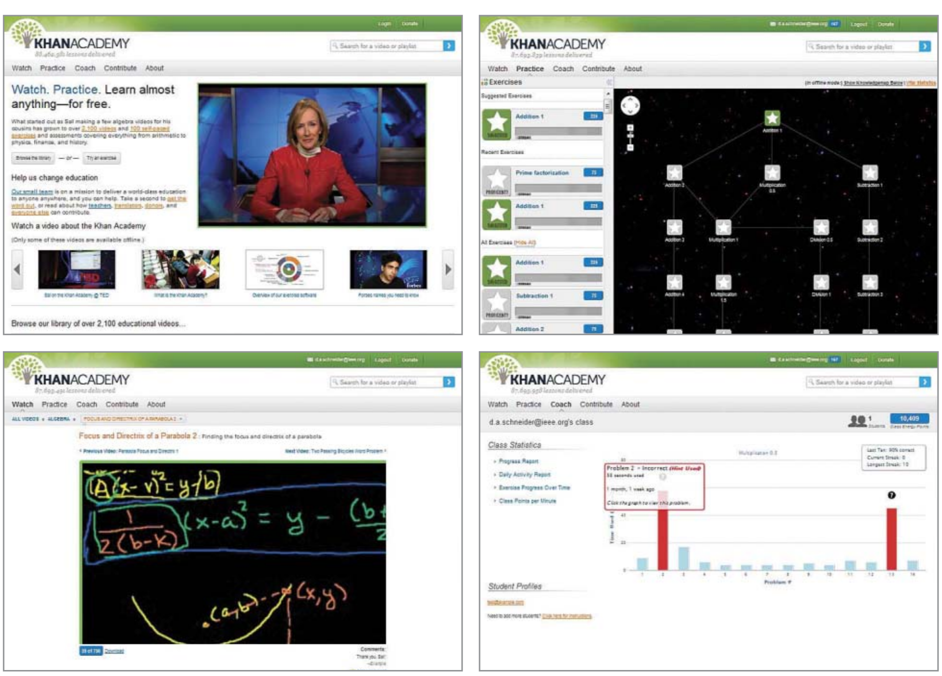

ACT LOCALLY: A little hacking makes the Khan Academy home page [top left] look presentable when hosted on your local machine, and a simple fix allows the site's helpful knowledge map to appear [top right]. Students can use the off-line version to watch Khan's educational videos and to do practice exercises [bottom left], and teachers can access the various monitoring tools [bottom right].
IMAGES: DAVID SCHNEIDER

geek life

POCKET STUDIOS

Hollywood pros are starting to use smartphones to record music and film

PHONE UPGRADE: *Olive*, a full-length movie, was shot by Patrick Gilles [above] and Hooman Khalili using a Nokia smartphone fitted with professional film lenses. Chris Price [left] recorded his latest album using an iPhone, performing in locations outside the studio to create unusual background ambiances.

PHOTOS: TOP, DAVID PAUL MORRIS/BLOOMBERG/GETTY IMAGES; LEFT: KYLE SAFIEH

WHEN LOS ANGELES songwriter Chris Price got his first iPhone last year, he filmed himself playing a couple of songs. "I was so struck by the quality of the sound, I ended up doing the whole record that way," he says.

The result, *Homesick*, which was released in June in CD, vinyl, and online digital formats, is part of a growing trend of albums and movies being recorded and filmed on iPhones.

"I think my goal was to legitimize smartphones as part of the new frontier of DIY recording," says Price, who began experimenting with an updated '60s sound after leaving his old label, Geffen Records, two years ago. "It's getting harder for young new artists to gather the thousands of dollars needed to rent studio space and hire an engineer and producer."

Price also found that as expenses fell, creativity rose. "When you're given less to work with," he says, "your brain works overtime." The phone's portability freed him up in other ways as well.

"A phone is a mobile studio, so you're not limited to a single room with the same acoustics for every song," Price says. "I could pick out different spots around the city for their sonic qualities," like reverb in a downtown tunnel. He used 4Tracks, an app that records in high-quality WAV files, and mixed them in GarageBand on his computer. The result was a throwback to analog's warmer, imperfect aesthetic.

He's not the only artist to discover that small is beautiful. Earlier this year, Alabama rockers One Like Son released an iPhone-recorded album, *Start the Show*, produced with FourTrack, MultiTrack DAW, AmpKit, ThumbJam, and GuitarJack apps. In 2010, Gorillaz recorded an album, *The Fall*, on the iPad, using 20 different apps. The electronic duo Nuclear O'Reilly is believed to have created the first iPhone album, *Phoning It In*, in 2009.

A full-length movie, *Olive*, starring Gena Rowlands with music by Ben Lear (son of TV mogul Norman Lear), premiered last December and is making its way through the festival circuit. With a 12-member crew, San Francisco radio personality Hooman Khalili and guitarist Pat Gilles made the US $430 000 film in just over six weeks (including editing). "There's something to be said about being first, ahead of the curve," says Khalili.

The pair used a Nokia N8 smartphone, which shoots in high definition, affixed to a 35-mm Carl Zeiss lens to enable depth of field. They hired a cellphone hacker to turn off the pixel-degrading auto-focus and auto-zoom but ultimately had to color-correct in postproduction. "We couldn't control the digital iris. It more than washed out; it actually changed the color," says Khalili. "Indoor shots were especially tricky. We had to trick the camera quite a bit and put as much light on the actors as possible."

There's been enough fervor to spawn the iPhone Film Festival, which started last year and holds its next festival in October.

"People are using iPhone films as entrée to the traditional filmmaking industry and as a way to keep their skills sharp. It's equivalent to the musical jam session," says festival cofounder Ruben Kazantsev. "Another advantage is having angles and shots not available with a traditional camera. For one film, *Goldilocks*, the director put an iPhone in a Ziploc bag into a wine glass and filmed the wine pouring in. You'd never be able to get that shot with any other camera."

—SUSAN KARLIN

A version of this article appeared online in June 2012.

SORTING AND EDITING: Using Lightroom's side-by-side sorting [top left], you can quickly compare images and assign them ratings, color flags, and keywords; AfterShot Pro can compare up to six images side by side [top right], though its sorting process isn't as quick. Lightroom's before-and-after views [bottom left] provide great feedback as you use its top-notch global and selective editing tools to fine-tune your picture. Though AfterShot doesn't have Lightroom's before-and-after view, it does use layers [bottom right] to help track and preserve your edits.

SALLY WIENER GROTTA & DANIEL GROTTA (4)

ADOBE LIGHTROOM AND COREL AFTERSHOT PRO

Two programs can help you get your photo house in order

ONCE, IT was an effort to take pictures and get them developed, yet easy to keep and show them. Today, it's the other way around—pictures are easy to take, and as a result, we're swimming in them. Drowning, even. Every time we shoot another thousand or so at a simple birthday party, we're left with a major task—sorting out the best ones, then editing and sharing them.

Fortunately, help is on the way: The workflow software the pros use is now filtering out to the masses in the form of a US $100 program, Corel AfterShot Pro. The most popular professional software, Adobe Photoshop Lightroom, a mature product ($150 for version 4), is smoother and more versatile. But the scrappy newcomer has a nice selection of photo tools that will satisfy many users.

Here's how it works. You know that somewhere among those hundreds of digital pictures you took at the birthday party lurk a few gems—if only you could find them. Both Lightroom and AfterShot enable fast comparing and sorting. Lightroom's two-up side-by-side display compares a "select" image to a candidate. You can quickly zoom into details, swap out a new select or candidate, and rate each using one to five stars and color flags (plus add keywords, if you didn't do that automatically on import). Then, you can view, say, all the photos with three stars or higher. AfterShot displays up to six images side by side, but swapping out images requires more steps.

Editing in Lightroom or AfterShot Pro can also be quite fast, with top-notch exposure, color, tone, and detail tools. If, for instance, you take a series of photos under the same lighting, using similar exposure settings, you can modify one picture and apply those edits to all the others.

Both programs edit "nondestructively," leaving your original image untouched. Here's where AfterShot has an advantage: It uses layers to track edits and extends its nondestructive editing to special effects filters. Lightroom doesn't have layers, and though it supports important third-party filters, the filters aren't nondestructive.

Finally, what use are photos without sharing and printing? This time, Lightroom has the edge, with more output options: Web galleries, photo books, slideshows, prints with multiple-size images on a single sheet (if you wish), and easy uploads to Facebook and YouTube are all possible. AfterShot does an excellent job of saving a series of photos (even while you continue to edit other pictures), custom or template-based printing (including multiple-size images per sheet), and creating nice Web galleries.

—DANIEL GROTTA & SALLY WIENER GROTTA

A version of this article appeared online in June 2012.

File Formats: Both Lightroom and AfterShot Pro are known as RAW workflow programs. If you, like most serious photographers, shoot RAW, you'll want to be sure the software you choose supports your camera. We found that AfterShot Pro doesn't support as many RAW file formats as Lightroom does, but Corel is adding more formats on an ongoing basis.

Both programs work well with TIFF and JPEG images, though AfterShot Pro supports such files only in RGB, and only up to 40 megapixels. In addition, AfterShot doesn't support video files, while Lightroom does.

education

DANETTE HOLLEY PHOTOGRAPHY/GETTY IMAGES

THE $10 000 COLLEGE DEGREE

A Texas university offers a low-cost bachelor's in IT. But what's it worth?

COME THE summer of 2016, Texas A&M University–San Antonio will graduate its first group of students with bachelor's diplomas in information technology, at the bargain price (for the United States) of US $9700.

When the university's president, Maria Hernandez Ferrier, announced the bachelor's of applied arts and sciences IT degree in March, listeners were astounded. Just last year, critics scoffed and education officials scratched their heads at Governor Rick Perry's call to his state's public higher-ed institutions to develop a four-year bachelor's degree for no more than $10 000. That's within range of a single year's tuition and fees at sister school Texas A&M–College Station or the University of Texas–Austin.

The San Antonio university's intent is clear: an affordable degree designed to make disadvantaged kids highly employable. Still, critics worry that it will result in a subpar education.

The new degree relies on collaboration with community colleges and a stripped-down version of the college IT curriculum. Texas high school students will take dual credit courses starting their junior year, followed by networking and security courses at a community college, which gets them an associate's degree in applied sciences. Next will come nearly 70 hours of college junior- and senior-level business communication and IT courses at the university. At the end, students will emerge prepared for entry-level jobs in network administration, security, and support, says Carolyn Green, director of Texas A&M–San Antonio's center for information technology and cybersecurity. "They could be hired from the associate's program, but with the bachelor's they're in a position to advance in their jobs or go to graduate school."

There is no pretense that the degree is equivalent to a traditional four-year computer science or IT degree. For the price, it's next to impossible to include the math, programming, and computer engineering courses, electives, and lab work that those degrees require, Green says.

A major criticism of the $10 000 degree proposal is that it will undermine the notion of a bachelor's degree. Even with all the cost-cutting measures, "I do not see how they can meet a $10 000 target without severely reducing the quality of the degree, probably by using a lot of online instruction, packaged learning, and low-cost and possibly outdated instructional material," says Peter Hugill, a professor at Texas A&M–College Station. "I suspect many students will rarely see a real live faculty member."

Other schools offer low-cost IT or computer science degrees. Public universities in many states, including Florida and New Mexico, offer in-state tuition at less than $20 000 for four years. And undergraduate tuition is free or waived at a handful of private institutions such as Berea College in Kentucky and the more selective Cooper Union in New York City.

Even schools that cost over $40 000 are worth their price for engineering and technology degrees because of the more-than-twofold jump in salary the degree brings over a high school diploma, says Nicole Smith, a senior economist at the Georgetown University Center on Education and the Workforce. But, she says, college isn't just about getting a job; it's a multifaceted experience that builds social and problem-solving skills.

Nevertheless, Smith adds, cheap college degrees are a must in the current economy, even at the expense of that quintessential college experience. Many states can't subsidize education, and many students want an education that can guarantee a job. "The proof will be in the pudding," she says, about the $10K degree in Texas. "Once you get this degree, can you get a job that will pay wages?"

Green says that big San Antonio–area employers—the United Services Automobile Association, the U.S. Air Force, and energy and banking companies—are already aware of and excited about the new degree program.

The crucial message here is that higher education is not one size fits all, according to Steve Moore, a Texas A&M University System spokesman. "There's a realization now that not everybody can be educated in the same way, nor do they need to be educated in the same way to be effective in the workforce."

—PRACHI PATEL

A version of this article appeared online in June 2012.

 BY PAUL McFEDRIES

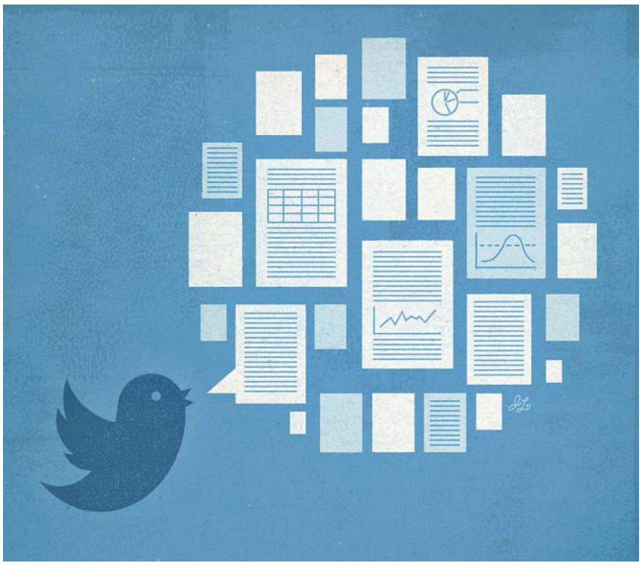

Measuring the Impact of Altmetrics

There is a growing movement within the scientific establishment to better measure and reward all the different ways that people contribute to the messy and complex process of scientific progress.
—Samuel Arbesman, *Wired*

How do you measure the influence of a journal or scientist? Until recently that question was largely settled. For a journal, you could turn to the **impact factor** (or IF), which determines the relative importance of a journal within its field by looking at how many times its articles get cited in other journals relative to the total number of articles it publishes. **PageRank** (predating and loosely related to the famous PageRank algorithm used by the Google search engine) is a kind of IF measure that gives greater weight to journals with high impact; a similar measure is the **Eigenfactor score** created by the evolutionary biologist Carl T. Bergstrom. For an individual scientist, you could calculate his or her ***h*-index** (in which *h* of the scientist's total number of papers have received at least *h* citations).

Lately, however, scholars have become increasingly disenchanted with these and similar **bibliometric** indicators that use such values as total number of articles published or total number of citations. They complain that traditional measures of scientific impact are too slow and too narrow to accurately reflect science in the Internet age.

Enter, then, the new field of **article level metrics** or, as it is increasingly known, **altmetrics**. This blend of *alternative* and *metrics* refers to tools based on bookmarks, links, blog posts, tweets, and other online measures that presumably indicate ways that readers have been influenced by an article—in short, how much "buzz" the paper is generating online.

Extremely astute readers may recall an earlier column of mine [see "The Coming Data Deluge," *IEEE Spectrum*, February 2011] that took note of researchers using "syndromic surveillance" to predict flu outbreaks based on an analysis of Google searches for flu-related terms. This is part of the emerging field of **infodemiology** (that is, information-based epidemiology), which is part of a broader field called **infoveillance**, the monitoring of online health information. If Google searches can show us the influence (no pun intended) of a flu virus on a population, why can't we use similar online data to judge the influence of a researcher or a scientific article?

Much of the altmetric scholarship has focused on Twitter and what Gunther Eysenbach, a researcher at the University of Toronto, has called **tweet metrics**. Although the prudent neologism collector must be on guard against Twitter-based coinages [see "All A-Twitter," *Spectrum*, October 2007] that are just silly (an adjective that can be rightfully applied to the vast majority of them), exceptions sometimes cry out to be made. To wit, I offer you the **tweetation**, a mash-up of *tweet* and *citation* that refers to a Twitter post that links to a scholarly article.

Another of Eysenbach's creations is the **TW*n*** score, which measures the number of tweets within *n* days of publication. This is the basis of the **twimpact factor**.

Then there's the **tweeted half life** (THL), which is the number of days after publication that it takes for an article to generate 50 percent of the tweetations that occur within a defined TW*n* period, say 30 days. If the article's TW30 is 100—that is, it generated 100 tweets in its first 30 days—and it generated 35 tweets on day 0 (the publication date), 10 tweets on day 1, and 8 tweets on day 2, then its THL is 2, because it was on day 2 that it surpassed 50 tweets.

This is all part of what researchers are calling **scientometrics 2.0**, where data mining techniques are brought to bear on massive social media databases and other online storehouses to search for fresh indicators of scholarly impact. Will they replace traditional measures such as the impact factor? Almost certainly not. The goal is merely to drag the concept of scientific influence into a century characterized by the rapid dissemination of information and near-universal social media. Tweet on. ❑

JESSE LEFKOWITZ


OYSTER CREEK NUCLEAR GENERATING STATION, NEW JERSEY
42 YEARS OLD

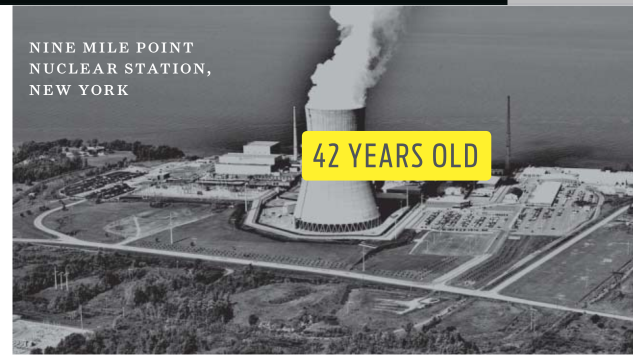
NINE MILE POINT NUCLEAR STATION, NEW YORK
42 YEARS OLD

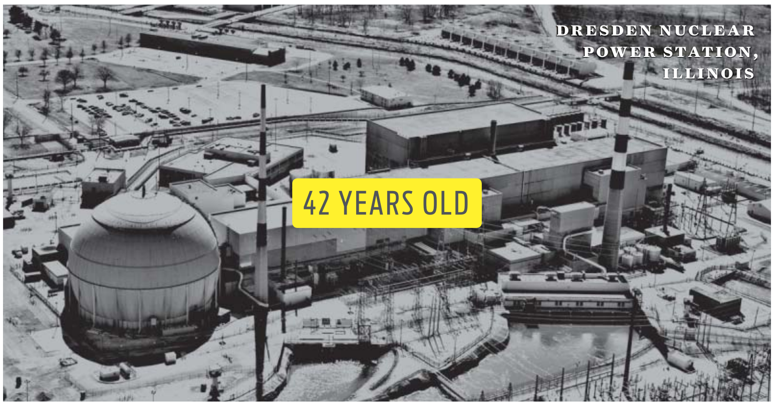
DRESDEN NUCLEAR POWER STATION, ILLINOIS
42 YEARS OLD


H.B. ROBINSON STEAM ELECTRIC PLANT, SOUTH CAROLINA
41 YEARS OLD


MONTICELLO NUCLEAR GENERATING PLANT, MINNESOTA
41 YEARS OLD

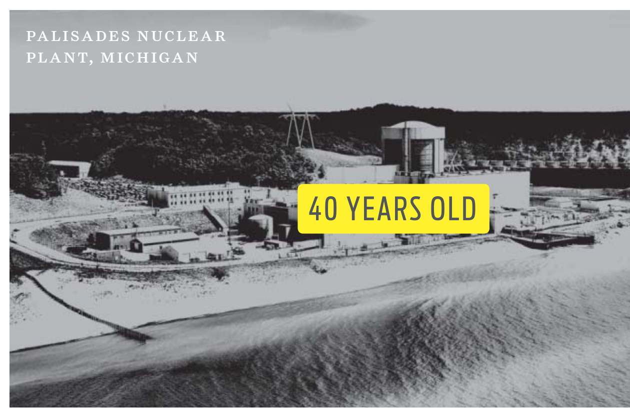
PALISADES NUCLEAR PLANT, MICHIGAN
40 YEARS OLD

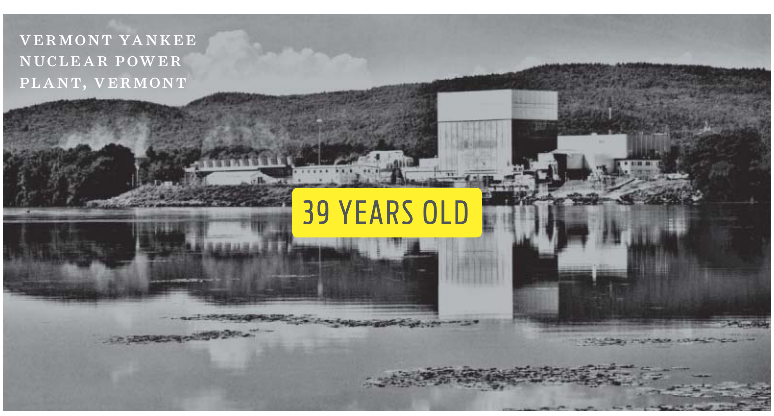
VERMONT YANKEE NUCLEAR POWER PLANT, VERMONT
39 YEARS OLD

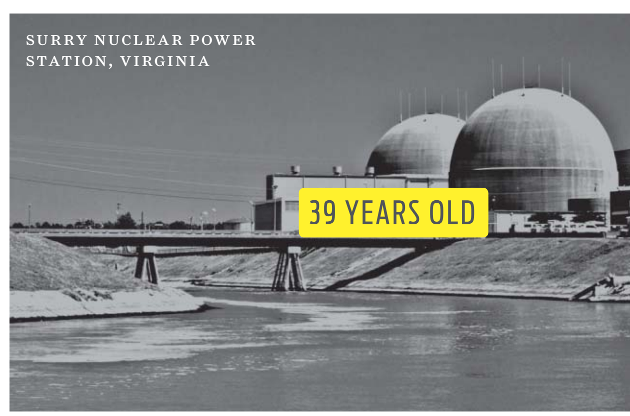
SURRY NUCLEAR POWER STATION, VIRGINIA
39 YEARS OLD


TURKEY POINT NUCLEAR GENERATING STATION, FLORIDA
39 YEARS OLD


39 YEARS OLD
OCONEE NUCLEAR STATION, SOUTH CAROLINA

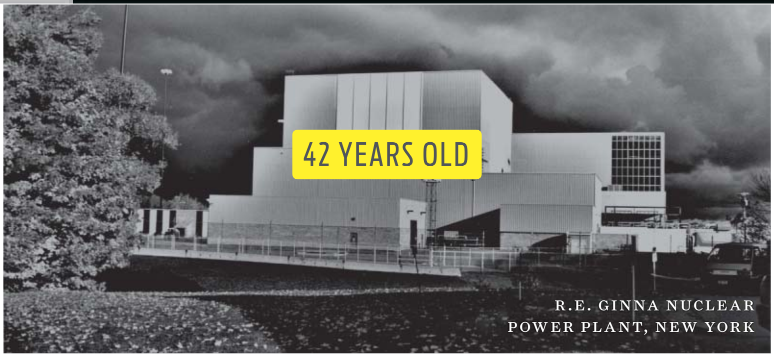

R.E. GINNA NUCLEAR POWER PLANT, NEW YORK

OLD REACTORS, NEW TRICKS

Can new monitoring techniques keep nuclear power stations operating safely for 80 years—twice their expected life spans?

BY LEONARD J. BOND

POINT BEACH NUCLEAR PLANT, WISCONSIN

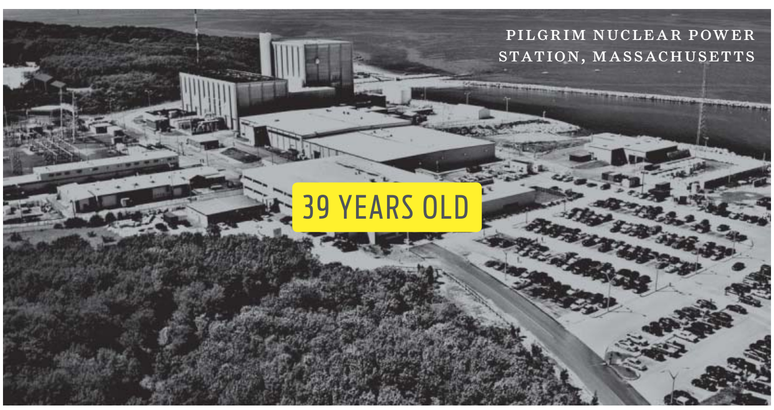

PILGRIM NUCLEAR POWER STATION, MASSACHUSETTS

QUAD CITIES NUCLEAR POWER STATION, ILLINOIS

PRAIRIE ISLAND NUCLEAR GENERATING PLANT, MINNESOTA

STAYING POWER: These nuclear power stations, the 15 oldest in the United States, were originally licensed to operate for 40 years. The Nuclear Regulatory Commission has since extended their licenses to 60 years, and researchers are contemplating another round of extensions.

PHOTOS: EXELON NUCLEAR; CONSTELLATION ENERGY/AP PHOTO; ENTERGY NUCLEAR; DANIEL ACKER/LANDOV; CAROLINA POWER & LIGHT CO.; FPL; KAREN BLEIER/AFP/GETTY IMAGES; ENTERGY NUCLEAR (3); DOMINION ENERGY; EXELON NUCLEAR; FPL; DUKE ENERGY; NUCLEAR MANAGEMENT CO.

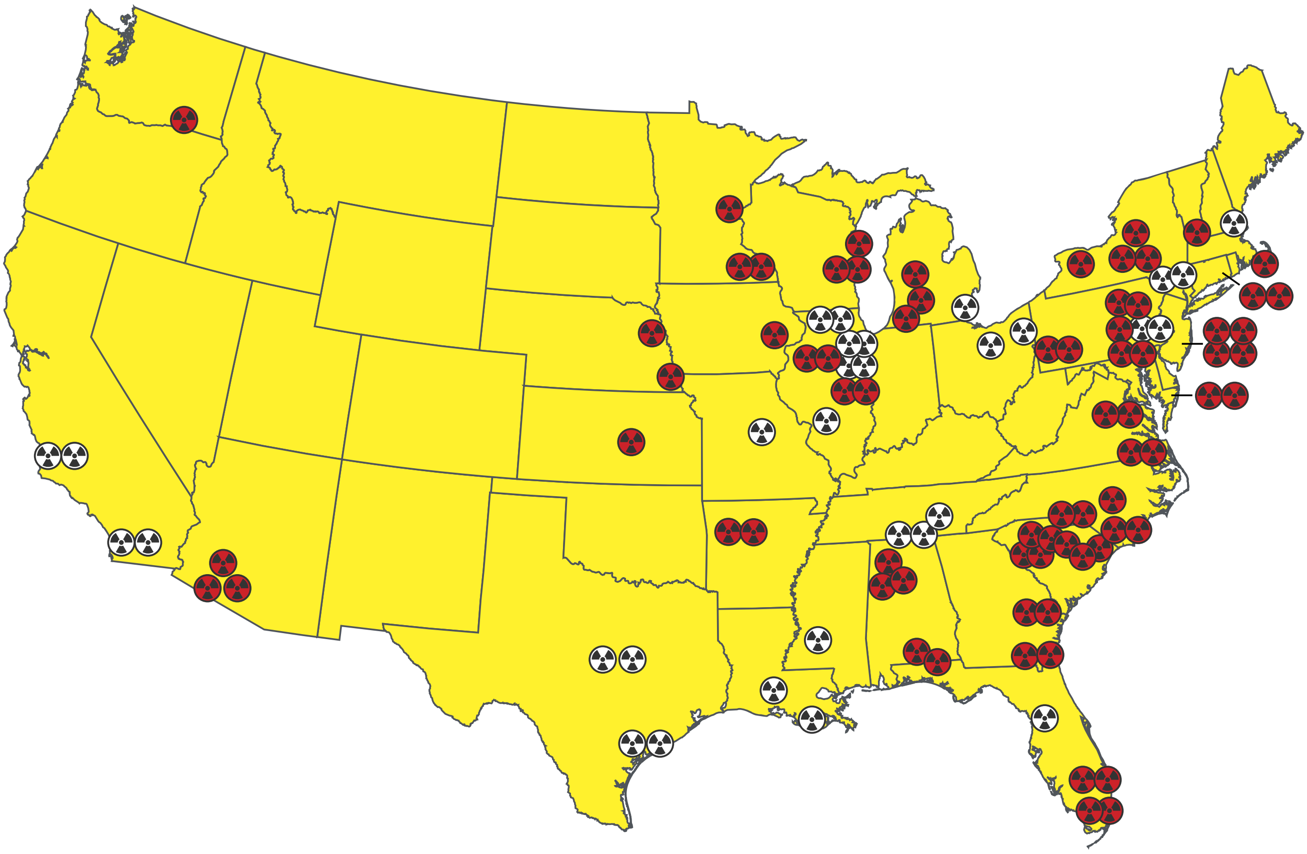

IN FEBRUARY 2002, during a routine inspection at Ohio's Davis-Besse nuclear power station, inspectors found three cracks in the lid of the reactor's pressure vessel, the mighty steel cylinder that encloses the radioactive core. One crack was in the housing of a mechanism that drives control rods into the reactor core to manage the nuclear reaction. The flaws needed to be repaired, but there was no sense of urgency—that is, until workers began fixing the crack in the control rod mechanism and they felt a wiggle. A wiggle that was all wrong.

The control-rod housing moved slightly, which should have been impossible, as it was supposed to be surrounded on all sides by the 15-centimeter-thick steel of the reactor vessel. When workers investigated, they found a cavity roughly the size of an American football in the steel next to the housing. This void left less than one centimeter of metal protecting the pressurized interior of the reactor vessel, with its radioactive core. If the vessel had ruptured while the reactor was in operation and at pressure, the water that cooled the core would have gushed out through the hole. Such a serious "loss of coolant" accident might have led to serious core damage. To fix the vessel, the plant's owners installed a new lid at an estimated cost of US $600 million.

Investigations by the U.S. Nuclear Regulatory Commission (NRC) and the plant's owner determined that a tiny fissure probably appeared in the control rod mechanism as early as 1990. By around 1995, acidic water from inside the reactor was leaking through the crack and eroding the surrounding steel of the pressure vessel; it ate away at the steel for seven years before workers discovered the metal loss. Nuclear researchers are acutely aware that this kind of slow, steady degradation becomes more likely as nuclear power stations age. Every day of operation, the rugged steel and concrete that make up a reactor's containment structures are bombarded with radiation and stressed by both high temperature and high pressure. Given enough time, these forces can potentially weaken even the toughest materials.

In the aftermath of Japan's Fukushima Dai-ichi nuclear accident, governments all over the world are reevaluating the safety of their nuclear power plants. In the United States, where nuclear power supplies 20 percent of the country's electricity, attention is focused on the aging of the country's 104 active nuclear reactors, which are 32 years old on average. (Four Westinghouse AP-1000 reactors now being built in the United States are the country's first new nuclear construction projects in decades.) When

AN AGING INDUSTRY: Of the United States' 104 commercial nuclear power reactors, 73 [in red] have received license extensions allowing them to operate for 60 years. The other 31 [in white] are still operating under their original 40-year licenses.
Source: U.S. Nuclear Regulatory Commission

these reactors went online, regulators granted them licenses to operate for 40 years, a conservative estimate of their life span. Now these plants are being awarded license extensions; 73 reactors have already received approval to operate until they're 60 years old, and 10 of those reactors have already entered this new era of extended operation.

But that's not the end of the story. Operators are performing major "midlife" refurbishments that can cost $1 billion per plant. Meanwhile, regulators and nuclear researchers are studying these aging plants to find an answer to one of the most important questions now facing the industry: Would it be safe and economically sound to keep these plants running until they reach 80 or more years of age?

That question, on which billions of dollars will depend in coming years, is also being asked in Europe, Asia, and former Soviet states. Although license periods and practices vary across countries, it's in the United States, with its large concentration of aging plants, that regulatory and industrial decisions will establish guidelines for reactors around the planet. If the U.S. nuclear industry demonstrates that refurbished nuclear power stations can operate until they're octogenarians or even longer, other countries will likely follow its example.

Current plans for managing aging reactors include periodic inspections of the components that are most difficult to replace: the pressure vessel, the concrete containment structure that surrounds it, and the main pipes and cables that connect to it. Over the past 15 years at Pacific Northwest National Laboratory (PNNL), in Richland, Wash., my colleagues and I have sought new types of online monitoring and nondestructive testing technologies that can provide early warnings of degrading materials. Our goal has been to transition from the current "find and fix" approach to one we call "model and predict."

INSIDE A NUCLEAR power station, fierce forces are at work. In the pressurized water reactors (PWRs) and boiling water reactors (BWRs) that generate power in the United States, the nuclear cores consist of rods of uranium dioxide. Inside this radioactive material, a nuclear fission reaction produces energy and many forms of radiation, including gamma rays and neutrons. The extremely high radiation levels are reduced by about a factor of 20 by the steel walls of a reactor pressure vessel, and then to safe levels by the massive reinforced concrete containment structure that jackets the vessel.

Both reactor types use water as the coolant. In a PWR, water enters the reactor core at about 275 °C and is heated as it flows upward through the core to a temperature of about 315 °C. The water remains liquid due to high pressure, usually around 15.5 megapascals (about 150 times the atmospheric pressure at sea level). In a BWR, the cooling water is maintained at about 7.6 MPa so that it boils in the core at about 285 °C. In both cases steam is produced to drive turbines that generate electricity.

CORRODED STEEL: At the Davis-Besse nuclear power station in Ohio, an acidic leak gradually created a cavity in the lid of the reactor's pressure vessel. The fissure developed over about 12 years.
PHOTOS: TOP, NUCLEAR REGULATORY COMMISSION/AP; BOTTOM: NUCLEAR REGULATORY COMMISSION

High temperature, high pressure, and radiation all stress a reactor's components. Inside a reactor, neutrons bombard the pressure vessel's steel walls; over a period of years, that bombardment can cause reactions that displace atoms in the material and produce impurities and tiny voids. These microscopic phenomena can reduce the metal's toughness and its ability to resist cracking.

The NRC and the nuclear industry, working with the Electric Power Research Institute, are now determining how to measure and monitor the aging of a reactor's key components. The major concerns are embrittlement and cracking in the reactor pressure vessel and its piping; degradation of the concrete containment; aging cables; and corrosion in buried water pipes. At the moment we just don't know which of these problems will be the most critical in any given plant. After all, no one has ever before operated a commercial-scale nuclear reactor for six or seven decades. We have entered a new era in the atomic age.

During the past 30 years, many parts of plants have been replaced or refurbished, including turbines, some major piping, and pressure vessel lids. But the central components of a nuclear plant—the pressure vessel itself and its reinforced concrete-and-steel containment—were never designed for replacement. The pressure vessel of a typical 1-gigawatt power plant weighs about 300 metric tons and is more than 12 meters tall. Most analysts believe that it would be easier to build a new plant than to cut into the containment to extract and replace a pressure vessel.

So how do you determine whether a vessel or another major component is robust enough to last another 20 years?

IF YOU WANT TO KNOW what's happening in an aging reactor, to really understand how its thick steel and tough concrete are faring after years of relentless bombardment, the best thing to do may be to listen to it. Nuclear researchers are now testing acoustic and ultrasonic monitoring techniques drawn from the civil and aerospace engineering communities. The same methods used to monitor the structural integrity of a bridge or an airplane may work for a nuclear pressure vessel as well.

One promising technique was demonstrated decades ago in an operational nuclear plant. In 1989, inspectors at the Limerick Generating Station, in

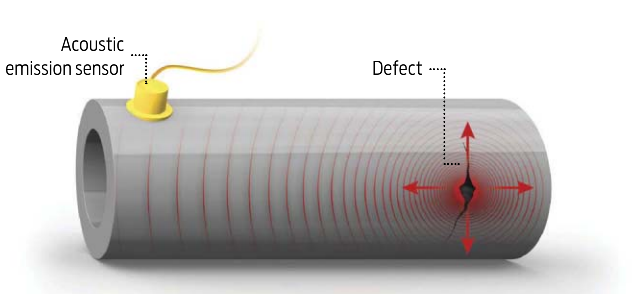

ACOUSTIC EMISSION MONITORING: When a crack grows in metal, the rupture releases tiny pulses of acoustic energy. Sensors detect these waves and can monitor a developing flaw.

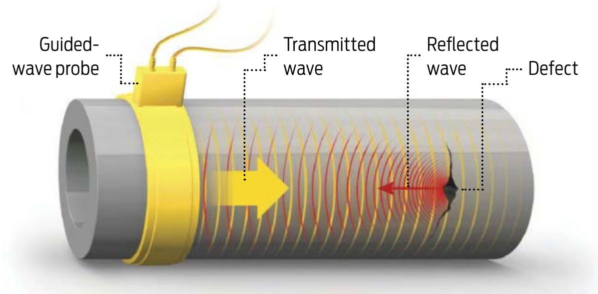

GUIDED WAVES: To check for flaws in a pipe or other metal structure, ultrasonic waves are pulsed through the material. Any flaw reflects part of the wave back toward the sensor.

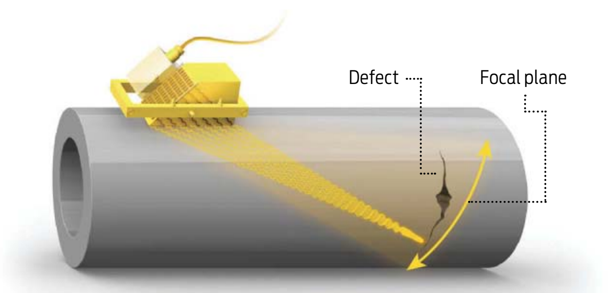

PHASED ARRAY: In this technique a group of transmitters release separate ultrasonic waves, which interact to form one larger wave front. By controlling the timing and amplitude of the individual pulses, researchers can steer the wave front to scan a structure for flaws.

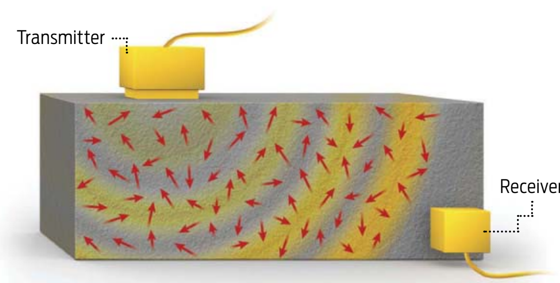

DIFFUSE FIELD: To monitor a coarse-grained material like concrete, a single ultrasonic pulse is introduced into the material. Receivers listen for the tiny echoes produced by the wave's interactions with all the grains. The composite signal creates a distinct signature for that material, which will change if the material degrades.

EMILY COOPER (4)

LISTENING FOR TROUBLE Researchers are experimenting with acoustic tools to monitor nuclear power stations' critical parts, like steel pressure vessels, pipes, and concrete containment structures. These tools can probe interiors without damaging the material.

Pennsylvania, found a tiny crack in the welding around a pressure vessel pipe that brought cooling water into the bottom of the reactor. The operators concluded that the flaw didn't pose a threat, but they wanted to see if it was possible to monitor crack growth in an operating plant. They turned to a technique called acoustic emission monitoring, which is used to check on metallic structures like pipelines and wind-turbine blades. This method relies on the fact that when a crack grows, acoustic energy is released in tiny pulses—much the same way an earthquake sends out seismic waves. Once the acoustic system was installed, operators could listen for the ultrasonic waves that would indicate a growing fracture.

The acoustic system was kept in place for three years, during which time researchers listened in as one part of the crack grew to a depth of 12 millimeters. The system also detected the growth of minuscule cracks that wouldn't have been noted by traditional monitoring methods, and researchers deemed the technology demonstration a success. In the decades since, fossil-fueled power plants and petrochemical facilities have installed acoustic emission systems to monitor vessels and pipes. However, nuclear power stations in the United States have been slow to adopt this proven technology.

With advances in both computer hardware and processing software, acoustic emission systems are now little larger than a laptop and are capable of displaying data nearly in real time. At PNNL, my colleagues and I recently tested acoustic emission monitoring along with another technique for metal monitoring that makes use of "guided waves." In this technique, transducers generate ultrasonic waves with specific frequencies, which propagate through a structure such as a metal pipe or the walls of a pressure vessel. Because the ultrasonic waves are scattered and reflected by discontinuities in a material, they can provide clear indications of cracks or corrosion. This technique could be particularly useful because it wouldn't require inspectors to strip off insulation to inspect pipes, like the all-important cooling pipes that circulate water through the reactor core.

In our recent laboratory study, we tested these two monitoring methods on a fatigued stainless steel pipe. The acoustic emission monitoring detected signals caused by the formation of a crack before we visually confirmed that tiny fissure. After we knew the crack was there, we monitored it with the guided-wave technique. When the waves encountered a crack, they bounced back to the sensor; by monitoring those received signals we were able to follow the growth of the defect from a starter notch that was 2.45 mm deep and 47.7 mm long to a fissure that measured 68 mm long. This may not seem like dramatic growth, but such a crack would be a serious cause for concern at an operating nuclear power plant.

Guided-wave technology, which is rapidly maturing, is now regularly used for pipe testing in the oil and gas industry. In the nuclear industry, regulators are working to standardize the monitoring procedures. To use the technology inside an active plant, however, operators must overcome challenges like high temperatures—it can hit 200 °C inside a light-water reactor's primary piping. That's far too hot for the most common type of transducers, which use piezoelectric materials to convert electricity into ultrasonic waves in the transmitters (and vice versa in the receivers). To get around this problem, some researchers are testing more rugged piezoelectric materials. Others are experimenting with different ways to generate the waves—for example, a laser pulse that heats and expands a pipe's surface to create waves that ripple outward.

Two other ultrasonic techniques show potential for long-term deployment. A kind of phased array, which is commonly used as a diagnostic tool in medicine, uses a grid of elements to generate many small ultrasonic pulses. By using electronics to control the timing and interaction of the individual pulses, operators can create a single wave front and control the direction of the wave. Phased-array technology is now routinely used in periodic inspections of nuclear power plants, but the technology has the potential for continuous monitoring, where a single transducer is fixed in place and electronic beam steering is used to scan critical structures. This technique can check for degradation in coarse-grain materials like cast stainless steels and can also look for flaws in welded areas.

Finally, an approach drawn from seismology could be useful to monitor the formidable concrete structures in a nuclear power station. In this "diffuse field" technique, an ultrasonic pulse is introduced into a coarse-grained material such as rock, concrete, or cast stainless steel. As the ultrasonic wave propagates through the substance, the grains interfere with the initial pulse of energy and send echoes back to the transducer. The resulting signal, showing all the interactions from within the textured material, provides a distinct signature for that material. This signature changes if the substance's elastic properties vary or if a crack or other degradation is introduced. So far, diffuse ultrasonic tools are being used only for research in the nuclear industry, but their potential for inspections and long-term monitoring has been clearly demonstrated.

IN WITH THE NEW: Some components of aging nuclear plants can be replaced. This new lid for the reactor pressure vessel was installed in the Davis-Besse nuclear power station. *PHOTO: AMY SANCETTA/AP PHOTO*

IF THE UNITED STATES wants to continue relying on nuclear power to keep one out of five lightbulbs lit, the NRC needs to be assured that a sound technical basis exists to support a second round of reactor life extension. By about 2020, the NRC must decide if it needs to establish additional rules and standards for subsequent licenses to allow for operation to go from 60 to 80 years. These additional rules would give operators a clear framework for the crucial and expensive decisions before them. No decision from the NRC would, in effect, preclude extended operations because it takes many years to plan for component replacements, refurbishments, and upgrades. If utilities and other nuclear plant operators don't have an explicit framework from the NRC by 2020 that enables them to schedule their capital investments, they'll have no choice but to start planning the decommissioning of the country's nuclear reactors.

Retrofitting and upgrading nuclear reactors will not be cheap. Some plants have already reported that they will spend up to $1 billion per plant to support the 40- to 60-year license extension. It may be that economics, rather than technology, will in the end determine if it's feasible to extend the life of a given plant beyond 60 years. But there's also a strong financial case for keeping aging reactors running: The loss of the existing plants after 60 years of operations would be a crippling economic blow. In the United States, the annual electricity demand is projected to increase about 21 percent by 2030 to roughly 5000 billion kilowatt-hours. It's difficult to imagine meeting that demand without the help of most of the country's 104 existing reactors.

If the United States decides against further license extensions, massive investments of trillions of dollars will be needed to replace the more than 100 GW of base-load generating capacity represented by the country's aging nuclear reactors. Whether the investment would be in new nuclear plants, cheap natural gas plants, or renewable energy facilities, it would be a monumental national project to replace the power we'd lose. Keeping a careful eye—and ear—on our aging nuclear infrastructure may be the more attractive option. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/nuclear0812

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

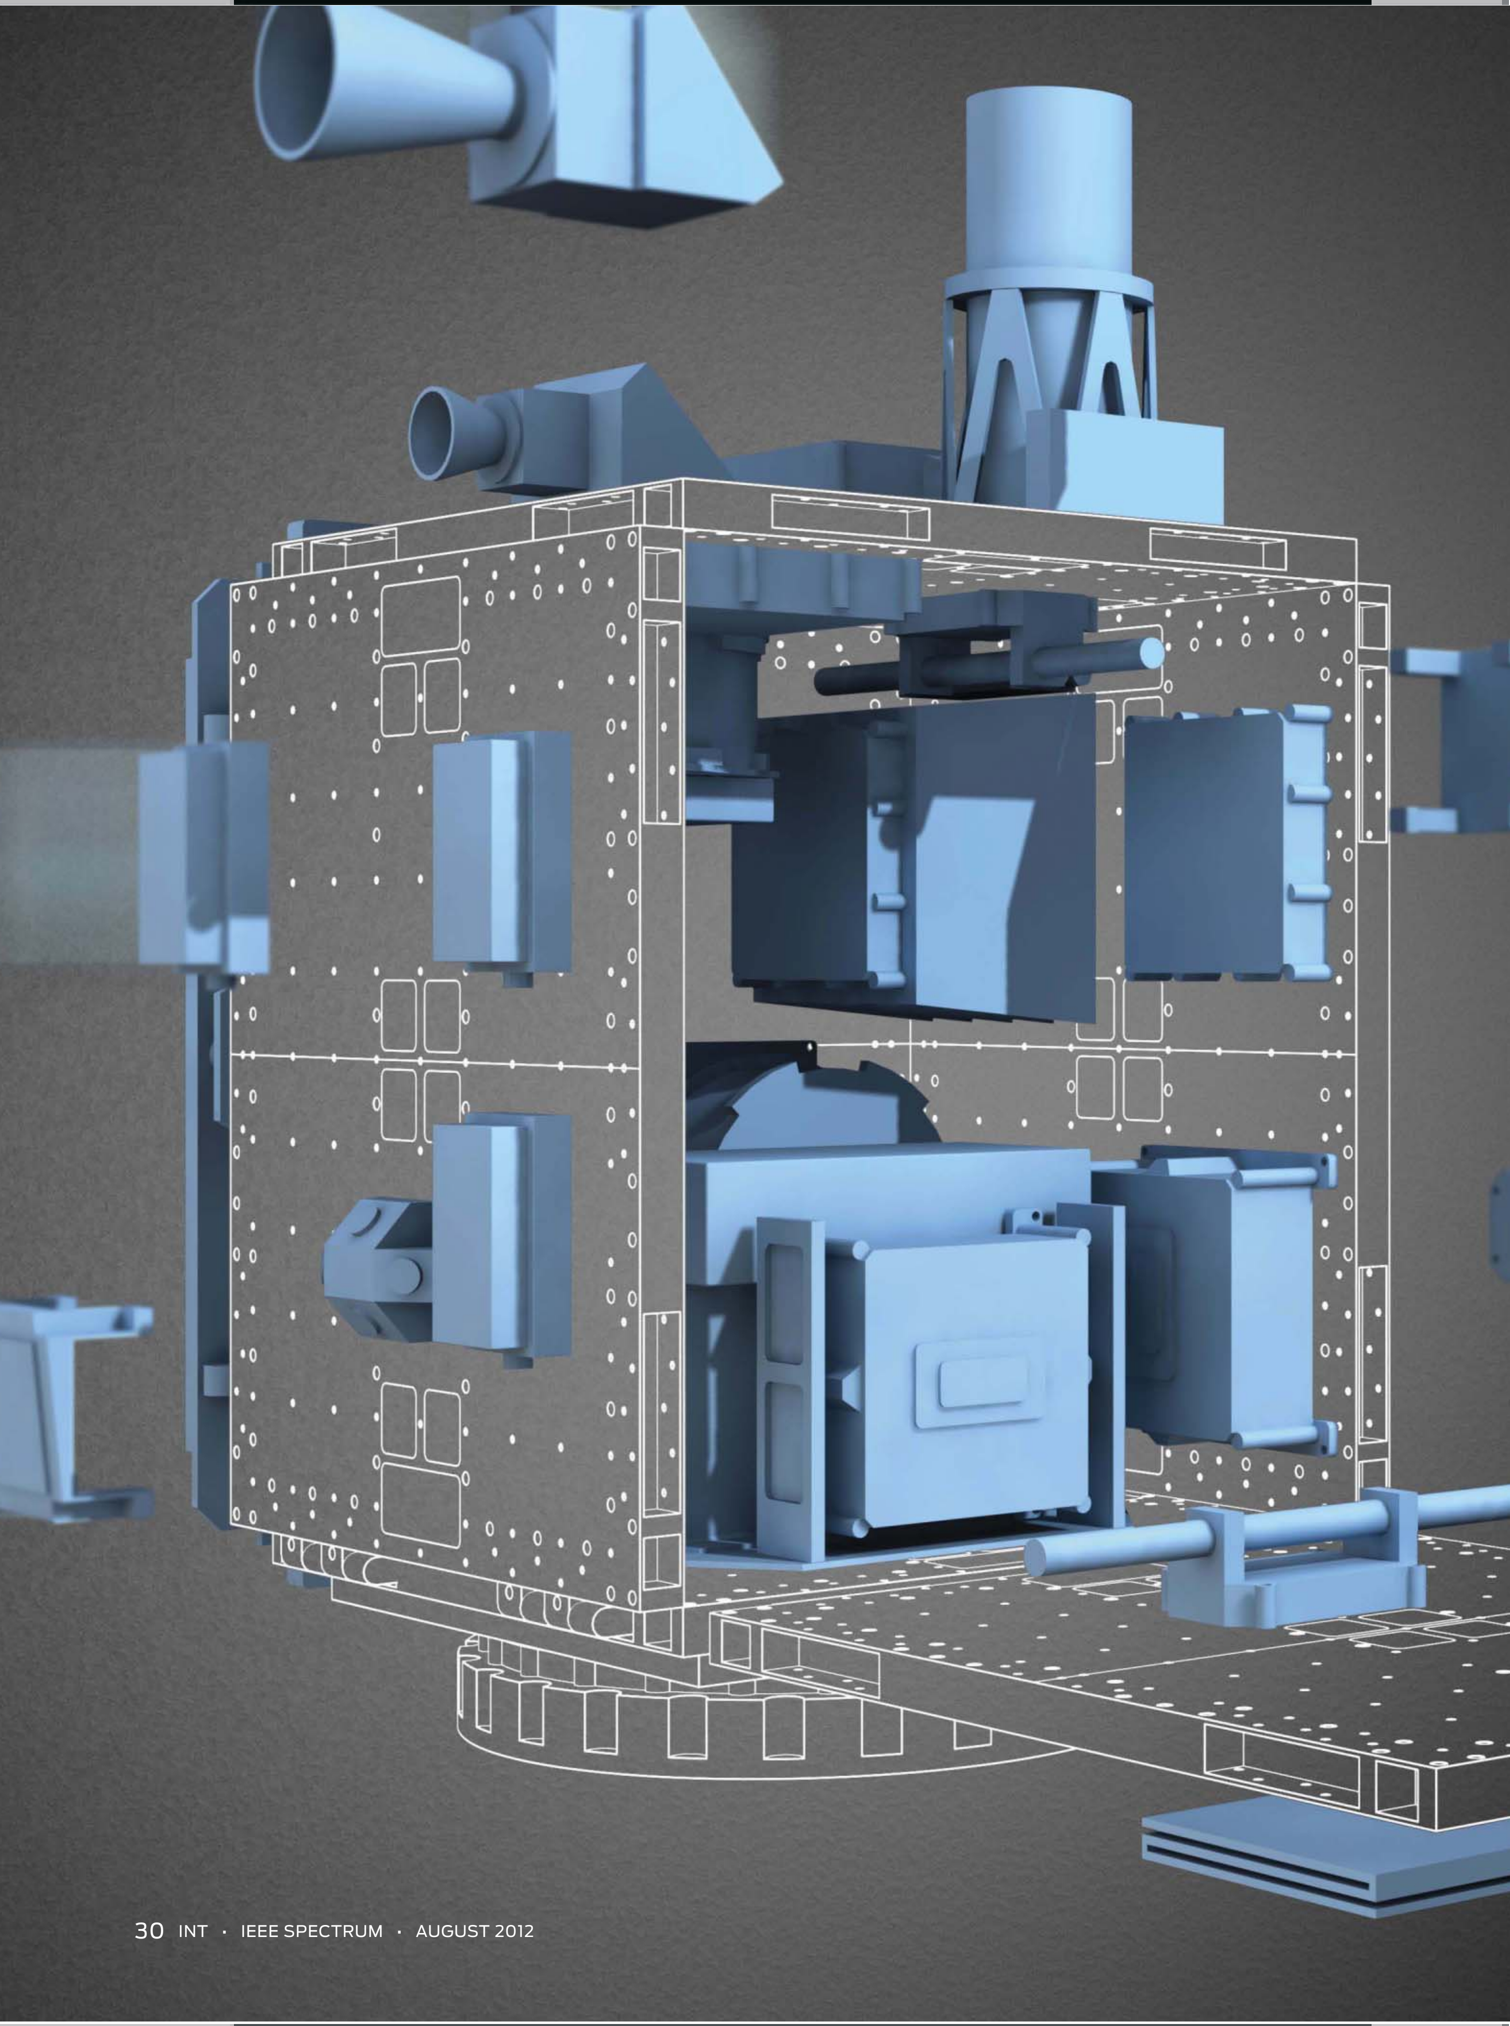

spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

Satellite design doesn't have to be rocket science

PLUG-AND-PLAY SATELLITES

WHEN YOU BUY A MOUSE FOR YOUR COMPUTER, removing the packaging is probably the hardest part of integrating it into your home system. Once you plug in the USB cable, you click on the mouse, and it just works. For it to "just work," of course, a great many things have to happen in the background: Via the USB cable, the mouse's circuitry receives power, initializes, and is recognized by the computer as a valid device. Then the driver software takes over, identifying the device as a mouse and not, say, a printer or a keyboard. Finally, a rapid succession of electrical messages traverses the cable, and these messages are translated into commands that then move the cursor smoothly across your computer screen. The fact that you don't need to know any of this to operate a mouse is by design: The mouse's computer chips and embedded software conceal the device's complexity.

Without those chips and software, the challenge of getting that mouse of yours to work with your computer would be a demanding one indeed. At the least, you'd need an intimate understanding of network protocols and timing sequences of pulse trains, and you'd have to devote many painstaking hours to designing the circuitry, constructing software drivers, writing test code, and more. And that's just for your mouse. Multiply that undertaking for your keyboard, your printer, your scanner, and other such devices, each with its own unique interface and functions, and you'd quickly find yourself in a major endeavor to integrate those disparate parts into a working computer.

Now you have a good idea of what developing a typical satellite is like today. Hugely complicated and highly customized, they take many years and many billions of dollars to build and frequently come in far over budget and years past deadline. For example, the U.S. Department of Defense's Transformational Satellite Communications System was proposed in 2003 at a cost of US $12 billion, with the first of its five satellites to be launched in 2011. By 2006, the cost had grown to nearly $16 billion, and the schedule had slipped to 2014. By the time the program was canceled in 2009, estimates for its eventual cost had reached $26 billion, and further delays were anticipated, despite reductions in its original capabilities.

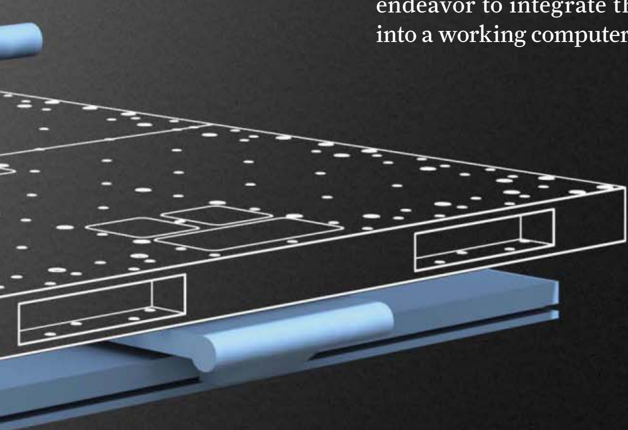

by JAMES C. LYKE

illustrations by JOHN MACNEILL

This was not an isolated case—indeed, it's universal. So eight years ago, a few of us at the Air Force Research Laboratory (AFRL) set out to find a better way. Along with a small cadre of researchers from industry, government, and academia, we have been studying the example of the personal computer and "plug-and-play" concepts from other industries in search of lessons we could apply to the task of building better spacecraft. Traditionally, satellite designers strive to increase raw performance or system capabilities by turning to faster processors or more sophisticated sensors. But we took a very different approach, concentrating instead on slashing the time it takes to go from inception to launch. Our goal was, well, lofty: to build a working satellite in just six days.

We not only succeeded, we also laid the groundwork for others. Our plug-and-play concept is now being incorporated into several satellite projects, including a tactical battlefield spacecraft being developed by Northrop Grumman Corp. for the U.S. military. The basic architecture is also a perfect fit for spacecraft that carry an eclectic mix of instruments, such as weather and environmental monitoring satellites, and those that have an inherently modular design, such as communications satellites. And although our initial efforts have been aimed mainly at small satellites (that is, those weighing less than 450 kilograms), the plug-and-play architecture should work just as well regardless of the system's size.

Meanwhile, parallel trends in satellite design and in manufacturing should make the adoption of plug-and-play satellites more palatable to the otherwise quite conservative aerospace industry. These include the enormous popularity of the miniature satellites known as CubeSats. Legions of college students and other space enthusiasts have worked on these tiny orbiters, in the process seeding the notion that spacecraft design need not be a high-overhead endeavor. The other development is 3-D printing, which allows quick production of three-dimensional objects directly from digital designs; embraced by the DIY crowd, this kind of rapid, on-demand manufacturing is also the perfect complement to plug-and-play satellite design.

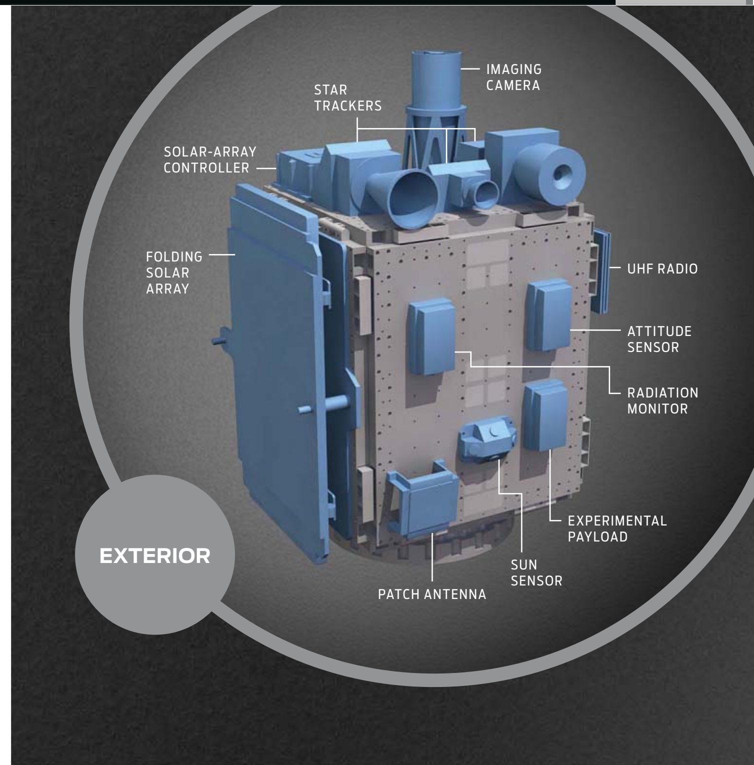

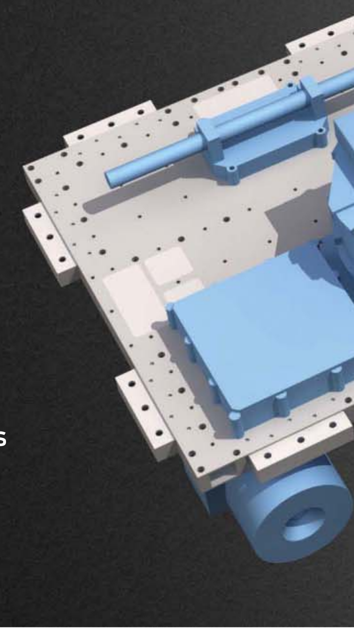

UNPLUGGED

WITH THE MONARCH approach to satellite design, all the components—including the host computer, reaction wheels, and communications module—have standard connectors. On a cube-shaped satellite like this one, up to 48 components can be connected to the inside and outside of the spacecraft's six hinged "smart" panels.

For decades, aerospace engineers and program managers have grappled with the growing complexity and expense of the systems they send into orbit. But despite many efforts to streamline and standardize, the development process has become only more protracted and expensive. In the 1990s, NASA administrator Daniel S. Goldin used the phrase "faster, better, cheaper" to herald the agency's series of Earth-observing satellites and interplanetary probes. By the end of the decade, though, a number of these stripped-down missions had failed to meet their original cost and schedule marks. In fact, three of the Mars missions launched under this banner—the Mars Observer, Mars Climate Orbiter, and Mars Polar Lander—failed altogether.

You might think that satellites are expensive because of their exotic components. Those aren't cheap, of course, but the vast majority of the cost is actually associated with labor. Spacecraft developers compulsively check, cross-check, and cross-cross-check every detail. Some teams even include people responsible for tracking down the exact date and facility where each transistor, capacitor, and integrated circuit was made. A former NASA engineer once boasted to me that every transistor on one spacecraft he helped design was backed by a file drawer's worth of documentation.

And there are a lot of components—in some satellites, the mission-critical computers are triplicated. In a process known as triple-modular redundancy, the three computers compare computations, and if one of them is at odds with the other two, the majority prevail. In some cases, the software for each computer is created by a different team, to reduce

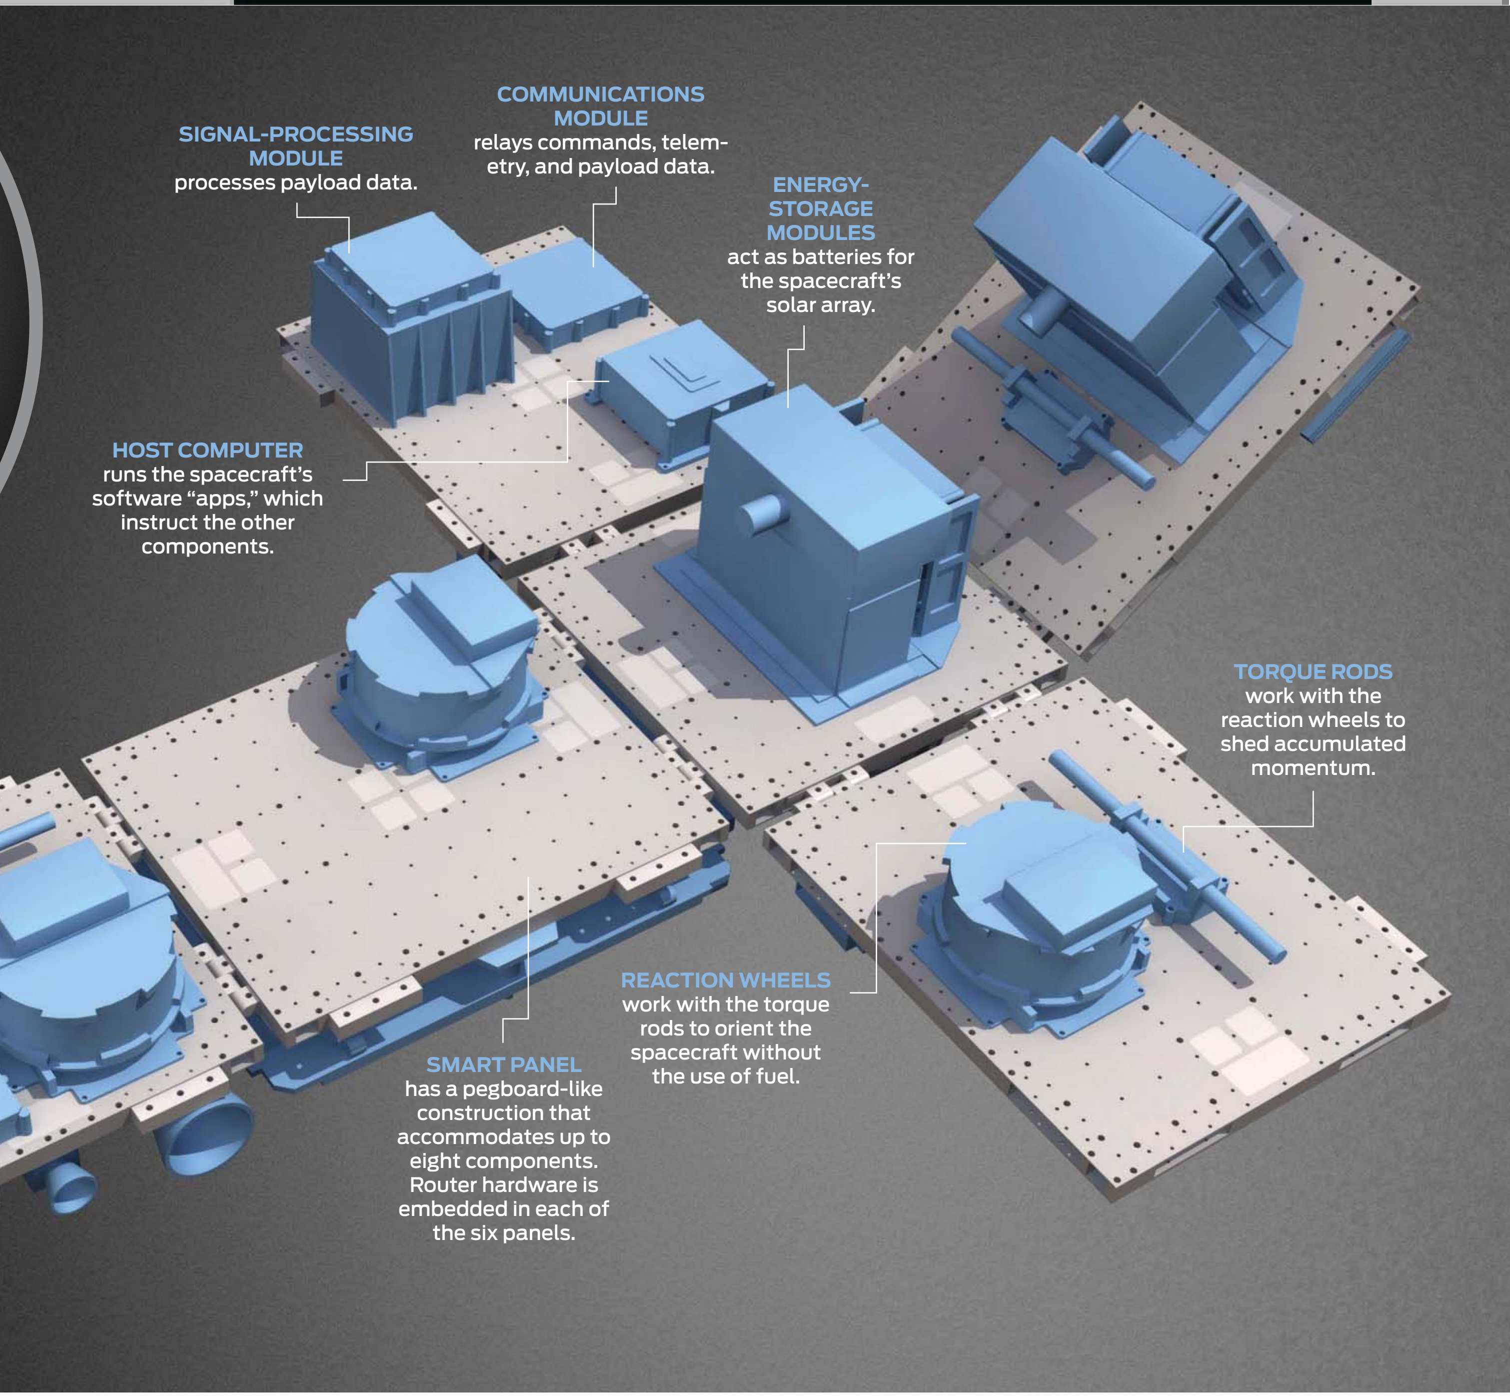

the chance that all three computers will fail at the same time while performing the same task.

A couple of other labor-intensive strategies also add to the cost. One is the use of legacy components. The thinking is that components that have been flown on existing spacecraft pose a lower risk. But reusing component designs from prior missions often requires significant reworking of hardware interfaces and flight software. It's kind of like trying to construct a jigsaw puzzle from the pieces of old puzzles: To make the pieces fit, you have to fashion additional unique pieces to fill in the gaps and hold the whole thing together.

The planning process is also quite inefficient, laborwise. The typical spacecraft is designed and built as though there will be no significant changes, even though requirements inevitably evolve, especially in protracted developments. Imagine what your architecture bills would be like if you were building a house and every time you decided to change a major appliance or a piece of furniture, you had to redesign the entire building.

We have a better idea, which we call Monarch, for "modular open network architecture." It harnesses the basic idea of plug-and-play systems, in which even complex systems can be formed quickly and reliably by arranging a number of existing building blocks. The architecture is open, not proprietary, and uses publicly available interfaces, akin to the USB and Ethernet standards found in computers.

USB and Ethernet allow you to build personal computers modularly and then combine them into networks. A computer's components are basically black boxes: Usually you only need to know a component's function—whether it's a mouse, keyboard, Web camera, or whatever—and what connector it requires, but not how it's internally constructed.

With Monarch, a spacecraft's components can similarly be thought of as black boxes—only instead of mice or keyboards, they may be gyroscopes or scientific instruments. These components, all of which have standard connectors, can be easily combined to form a Monarch system. A bigger, more complex spacecraft would simply have more black boxes than a smaller, simpler one.

Monarch allows for just three classes of black boxes. There are *endpoints*, which are components that perform a function, such as thermometers, cameras, and radios; *routers*, which connect two or more components; and *hosts*, which are comparable to the central processing unit of a PC. The host runs a variety of open-source "apps" that let it control the components. These apps are more complicated than typical iPad or Android apps, but like them, each is a relatively small program that carries out a specific function. For example, a spacecraft guidance app enables the satellite to track its position and motion and uses control algorithms to make any necessary adjustments to the thrusters. Another app gathers telemetry, by which the satellite constantly monitors its health and status in a form suitable for downlinking to ground control. That's basically all there is to Monarch: Components are combined to form a satellite system, and the system's apps use those components to carry out the satellite's mission.

When our group first came together in 2004, we devoted a number of months to carefully considering the best way to connect the Monarch components. When, say, a new communications device is added to the satellite, how would the host and router know what it is and what to do with it? Personal computers use software drivers to do this job; the driver acts as a sort of bridge between the operating system and the component. But we didn't want to use drivers, which need to be updated every time a component changes. If you've ever had to track down a driver for your new printer only to discover that it doesn't yet exist for the operating system you use, you know how annoying that can be. So instead we embedded a bit of code in each Monarch component that tells the host and router whether what's just been connected is a high-frequency radio receiver, an antenna, or a gyroscope.

In rethinking our components, we channeled the American inventor Eli Whitney. By championing the concept of interchangeable parts in guns, Whitney revolutionized the firearms business. Of course, a modern satellite can have a far wider assortment of components than did the guns of Whitney's day, but the general idea still holds: When the parts are standardized and share a common interface, they can be quickly assembled; any defective or outdated part can be easily removed and replaced without having to start from scratch. And the labor required—and therefore the overall cost—is kept low.

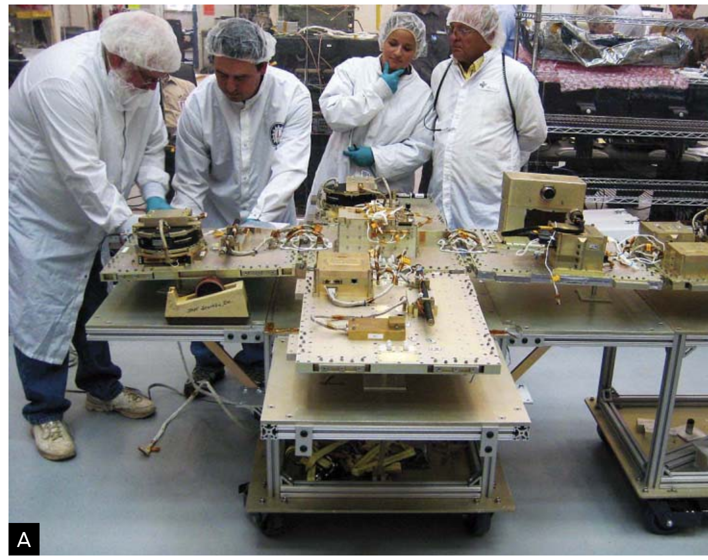
A


C


D

After laying the basic groundwork for the Monarch architecture, the next step was obvious: to see if we could use it to do what we'd set out to do—that is, design and build an actual, working satellite in six days. We created a test bed within the space electronics branch of the AFRL, located at Kirtland Air Force Base in Albuquerque. Here we cobbled together an inventory of "plug-and-play" components from spare parts left over from old space missions. The assortment included reaction wheels, used to rotate the spacecraft; torque rods, used to stabilize the spacecraft; and sun sensors, which measure the spacecraft's orientation toward the sun. We reworked these hand-me-downs by adding circuitry to them to match our Monarch connectors and programmed the interfaces to communicate using the Monarch protocols.

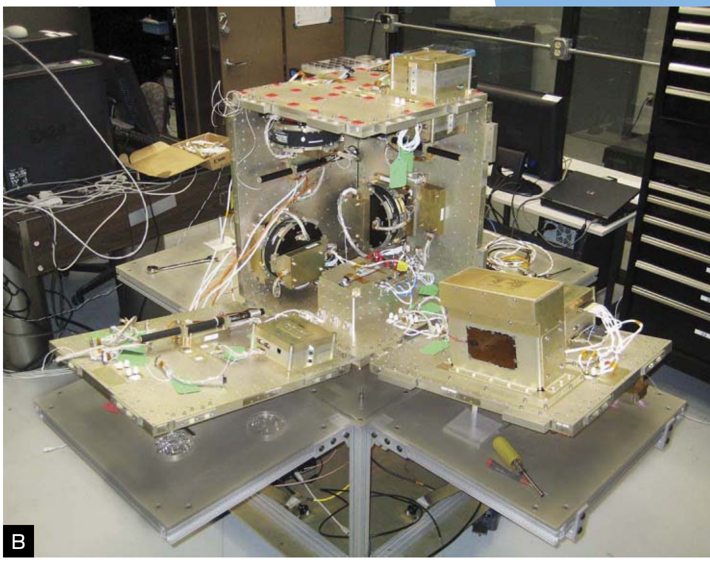

B

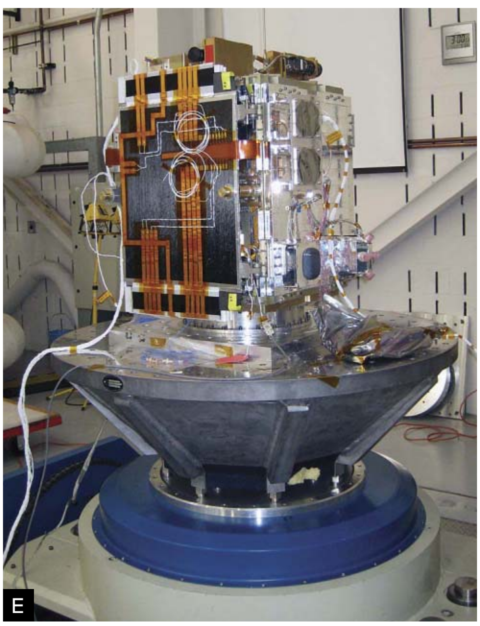

E

SOME ASSEMBLY REQUIRED: In experiments at the Air Force Research Lab, workers can build an entire Monarch satellite in 4 hours. Following a careful script, they start by connecting the spacecraft's six panels [a], which are hinged for easy access. Once the interior components are installed, the spacecraft is closed up [b] and additional components are connected on the exterior. The satellite then undergoes testing in a thermal vacuum chamber [c]; unlike a traditional highly customized satellite, however, it does not require weeks or months of tests. An adapter ring [d] is added to the bottom to attach the spacecraft to the launch vehicle [e]. After launch, a carefully controlled separation would release the spacecraft from the launch vehicle.
PHOTOS: AIR FORCE RESEARCH LABORATORY

We also fashioned standard walls for our six-day spacecraft: hinged panels that look like pegboards (with rows and columns of holes every 5 centimeters) and into which up to eight components can be quickly and easily connected and networked. The final satellite would consist of six of those panels, studded with two or three dozen components.

Our inaugural system was called the Plug-and-Play Satellite-1, or PnPSat-1. During a series of exercises, we quickly put together a variety of plans representing prospective satellite designs; these were based loosely on the architectures of existing traditional satellites, such as TacSat-2, a small imaging satellite developed by AFRL and launched in 2006. A typical PnPSat session would start with a mock design activity, using a fairly crude approximation of the design tools that a future plug-and-play spacecraft team would presumably have at its disposal. From this design activity would emerge the satellite's flight software and bill of materials—that is, the list of Monarch components needed. We would also come up with a set of instructions, such as the sequence in which to connect components to panels and panels to each other.

After we hammered out the design, we had to actually assemble the bird. Our team functioned kind of like a NASCAR pit crew, with the carefully scripted instructions guiding our moves. In our initial trials, we mainly used mock-ups of the various components, and the assembly took around 4 hours. Eventually, though, we were working with actual hardware aimed at a flight-capable orbiter, and we got the assembly time down to just over an hour. (To get a sense of how we did it, you can view the time-lapse video from one of our trials, at http://spectrum.ieee.org/satellite0812.) The sessions were exhausting, but we knew the clock was ticking, and that made the work exciting and a lot of fun.

Our PnPSat-1 was built to demonstrate plug-and-play principles, including whether or not they result in a space-worthy craft. To that end, the satellite was scheduled for launch in August 2008, but it got bumped from the roster several months before liftoff—maybe not such a bad outcome, considering that the SpaceX Falcon 1 rocket it would have ridden on failed just after launch.

Monarch is now being used in the design and construction of several satellites. The 400-kg tactical battlefield spacecraft being developed by Northrop Grumman is called the Modular Space Vehicle (MSV). It will have a variety of swappable radio and electro-optical payloads, including communications modules and tactical intelligence, surveillance, and reconnaissance units. The 24-month effort to develop, build, and test the MSV should yield a launch-ready vehicle by the end of 2013. If it proves successful, the plan is to turn over future manufacture of the MSV to a satellite "factory," called the Rapid Response Space Works, in Albuquerque. One day, the plan goes, a battlefield commander will be able to order up a dedicated spacecraft, and the factory will have it ready in a matter of months or even weeks, rather than years.

OUR BIGGEST accomplishment may have been roiling the satellite design community. Some of our colleagues have been excited by what we've done, while others have questioned its significance and indeed the whole notion that you can design an entire satellite around a plug-and-play architecture. Clearly, we've touched a nerve, and that's good, because it's helping raise fundamental questions about how spacecraft should be built. It's bolstering the efforts of our group and of other spacecraft designers striving to make greater use of standardization to reduce the costs of satellites and space launches.

One of those workers is Bob Twiggs, a former Stanford professor who's now at Morehead State University, in Kentucky. Years ago, to teach students about satellites, Twiggs came up with the idea of CubeSats. These diminutive spacecraft have since become a viral phenomenon, with hundreds of groups worldwide (including high schools and Kickstarter-funded teams) building them.

The typical CubeSat is a simple 10-cm-per-side cube weighing no more than 1.33 kg and yet containing all the essential elements of an autonomous satellite, including power, communications, and a useful payload. For those not content with just one cube, double- and triple-high CubeSats are available. You can go the other way as well: Twiggs's latest project, the PocketQub, breaks up a single CubeSat into eight bite-size chunks.

One of the main reasons for the growing popularity of the CubeSat is the dispenser that carries the tiny spacecraft into orbit and spits it out into space. Officially known as a Poly-Picosatellite Orbital Dispenser—but more commonly referred to by its acronym P-POD (pronounced "peapod," of course)—it looks like a surplus munitions container. Developed by Jordi Puig-Suari at California Polytechnic State University, the P-POD can house up to three CubeSats, and the ruggedized container can be readily added to practically any launch vehicle, taking advantage of the fact that most rockets have unused carrying capacity. To date, at least 40 CubeSats have been successfully launched, with another hundred or so planned in the next five years. The CubeSats now in orbit are accomplishing such tasks as studying radio-wave propagation through the ionosphere and measuring cosmic-radiation flux at low earth orbit.

The CubeSat and Monarch concepts share the common vision of a simplified, standardized architecture that allows spacecraft to be designed quickly and at low cost. But most CubeSats are not plug-and-play, so they are subject to scaled-down versions of the same problems that afflict big satellites, such as the need to integrate customized components that lack a standard interface. With CubeSats, these problems are mitigated by the satellites' simplicity and the fact that they're typically built by small, tight-knit teams. There's no reason, though, that a Monarch approach can't be applied to CubeSats, and in fact our group has created several dozen briefcase-size plug-and-play versions of CubeSats, which we use for training purposes. We can pull apart the pieces and switch them out in minutes. Just as Eli Whitney wowed members of the U.S. Congress by quickly disassembling several rifles and then reassembling them from the pile of parts, our demonstrations typically leave observers amazed.

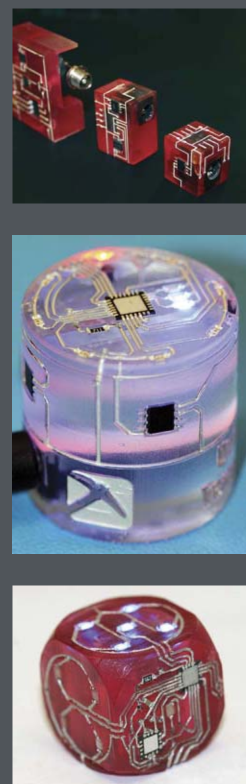

DIFFERENT DIMENSIONS: The top and center photos show 3-D printed magnetometers. At bottom is a 3-D printed gaming die. When the die is rolled, the embedded microcontroller and accelerometer identify the top surface and illuminate the appropriate LEDs.

PHOTOS: W.M. KECK CENTER FOR 3D INNOVATION

THREE-DIMENSIONAL printing has great potential for both Monarch and for CubeSats. This rapid manufacturing method, being developed largely outside the aerospace industry, allows essentially any shape to be produced quickly out of a variety of materials, including plastics, metals, ceramics, cement, even sugar. A group of academics and enthusiasts led by Gil Moore, a retired U.S. Air Force Academy professor, is now working on two 3-D-printed CubeSats. For the first, called Rampart, they're printing the satellite's primary structures, deployable solar cell arrays, and even a propulsion tank directly from computer-aided-design drawings. The material they're using is Windform XT, a nylon with carbon microfibers, which is very strong and doesn't emit much gas—both desirable qualities for space applications. The group's other CubeSat, called PrintSat, was created mainly to measure Windform XT's performance in space and has already been selected by NASA for launch as early as next year.

Of course, many satellite parts, such as motors and complex microcircuits, can't be produced using current 3-D printers. Most can print in only one material, and building something complicated like a satellite or a smartphone requires hundreds of materials. Within about five years, though, that may be possible. We're now collaborating with researchers at the W.M. Keck Center for 3D Innovation at the University of Texas at El Paso, which has amassed one of the world's most impressive armadas of 3-D printing equipment. Some of the center's most interesting efforts involve printing in mixed materials, to construct things like electronic blocks. By printing the dielectrics and conductors and then using pick-and-place machinery to position the tiny electronic components, it should be possible to form customized electronic assemblies on demand.

A company in Orlando, Fla., called nScrypt is pushing that concept even further. Its engineers believe that in five years or so they will have built a single machine capable of printing and assembling an entire smartphone, including the electronics, the wireless components, and the display. Once that happens, we will be very close indeed to being able to quickly and conveniently print all the components for an entire Monarch spacecraft.

So can Monarch, CubeSats, and 3-D printing usher in a new era of satellite design, one that slashes the billion-dollar costs and colossal inefficiencies that plague today's projects? Disruptive concepts are always a tough sell, particularly in the aerospace community, where technology adoption strategies are described with phrases like "crawl, walk, run." But the business-as-usual alternative—big spacecraft that vastly exceed their budgets and take years to build—is unsustainable. If there's a way to build well-engineered, capable spacecraft at a fraction of the cost and time, we need to give it a try. We have nothing to lose—and billions of dollars to save. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/satellite0812

TO BUILD A BRAIN

A MINIATURE, MASSIVELY PARALLEL COMPUTER, POWERED BY A MILLION ARM PROCESSORS, COULD PRODUCE THE BEST BRAIN SIMULATIONS YET / BY STEVE FURBER

PHOTOGRAPH BY DAN SAELINGER; PROP STYLIST: ARIANA SALVATO

For all their progress, computers are still pretty unimpressive. Sure, they can pilot aircraft and simulate nuclear reactors. But even our best machines struggle with tasks that we humans find easy, like controlling limbs and parsing the meaning of this paragraph. ¶ It's a little sobering, actually. The average human brain packs a hundred billion or so neurons—connected by a quadrillion (10^{15}) constantly changing synapses—into a space the size of a cantaloupe. It consumes a paltry 20 watts, much less than a typical incandescent lightbulb. But simulating this mess of wetware with traditional digital circuits would require a supercomputer that's a good 1000 times as powerful as the best ones we have available today. And we'd need the output of an entire nuclear power plant to run it.

Closing this computational gap is important for a couple of reasons. First, it can help us understand how the brain works and how it breaks down. There is only so much to learn on the coarse level, from imagers that show how the brain lights up when we remember a joke or tell a lie, and on the fine level, from laboratory studies of the basic biology of neurons and their wirelike dendrites and axons. All the real action happens at the intermediate level, where millions of networked neurons work in concert to produce behaviors you couldn't possibly predict by watching a handful of neurons fire. To make progress in this area you need computational muscle.

And second, it's quite likely that finding ways to mimic the brain could pave the way to a host of ultraspeedy, energy-efficient chips. By solving this grandest of all computational challenges, we may well learn how to handle many other difficult tasks, such as pattern recognition and robot autonomy.

Fortunately, we don't have to rely on traditional, power-hungry computers to get us there. Scattered around the world are at least half a dozen projects dedicated to building brain models using specialized analog circuits. Unlike the digital circuits in traditional computers, which could take weeks or even months to model a single second of brain operation, these analog circuits can model brain activity as fast as or even faster than it really occurs, and they consume a fraction of the power. But analog chips do have one serious drawback—they aren't very programmable. The equations used to model the brain in an analog circuit are physically hardwired in a way that affects every detail of the design, right down to the placement of every analog adder and multiplier. This makes it hard to overhaul the model, something we'd have to do again and again because we still don't know what level of biological detail we'll need in order to mimic the way brains behave.

To help things along, my colleagues and I are building something a bit different: the first low-power, large-scale digital model of the brain. Dubbed SpiNNaker, for Spiking Neural Network Architecture, our machine looks a lot like a conventional parallel computer, but it boasts some significant changes to the way chips communicate. We expect it will let us model brain activity with speeds matching those of biological systems but with all the flexibility of a supercomputer.

Over the next year and half, we will create SpiNNaker by connecting more than a million ARM processors, the same kind of basic, energy-efficient chips that ship in most of today's mobile phones. When it's finished, SpiNNaker will be able to simulate the behavior of 1 billion neurons. That's just 1 percent as many as are in a human brain but more than 10 times as many as are in the brain of one of neuroscience's most popular test subjects, the mouse. With any luck, the machine will help show how our brains do all the incredible things that they do, providing insights into brain diseases and ideas for how to treat them. It should also accelerate progress toward a promising new way of computing.

TRADITIONAL CMOS chips were not invented with parallelism in mind, so it shouldn't come as a big surprise that they have trouble mimicking mammalian brains, the best parallel machines on Earth. A few comparisons show why brain modeling is such a thorny problem. The logic gate in an integrated circuit is typically connected to just a few neighboring devices, but the neurons in the brain receive signals from thousands—sometimes even hundreds of thousands—of other neurons, some clear on the other side of the brain. Also, neurons are always at the ready, responding as soon as they receive a signal. Silicon chips, by contrast, rely on global clocks to advance computation in discrete time steps, an approach that consumes a lot of power. To top it all off, while the connections between CMOS-based processors are fixed, the synapses that link neurons are always in flux. Connections are constantly being forged or reinforced or phased out.

Given all these differences, it's a wonder we can even begin to tackle the problem of simulating brain activity. But there have actually been some pretty impressive supercomputer models that have managed to reproduce neuron operation with great fidelity. The ongoing Blue Brain Project, led by Henry Markram at the École Polytechnique Fédérale de Lausanne, in Switzerland, is a prime example. The simulation, which began in 2005, now uses a 16 384-processor IBM BlueGene/P supercomputer and data collected from very detailed studies of brain tissue to simulate 10 000-neuron sections of the rat brain, each section no larger than the head of a pin.

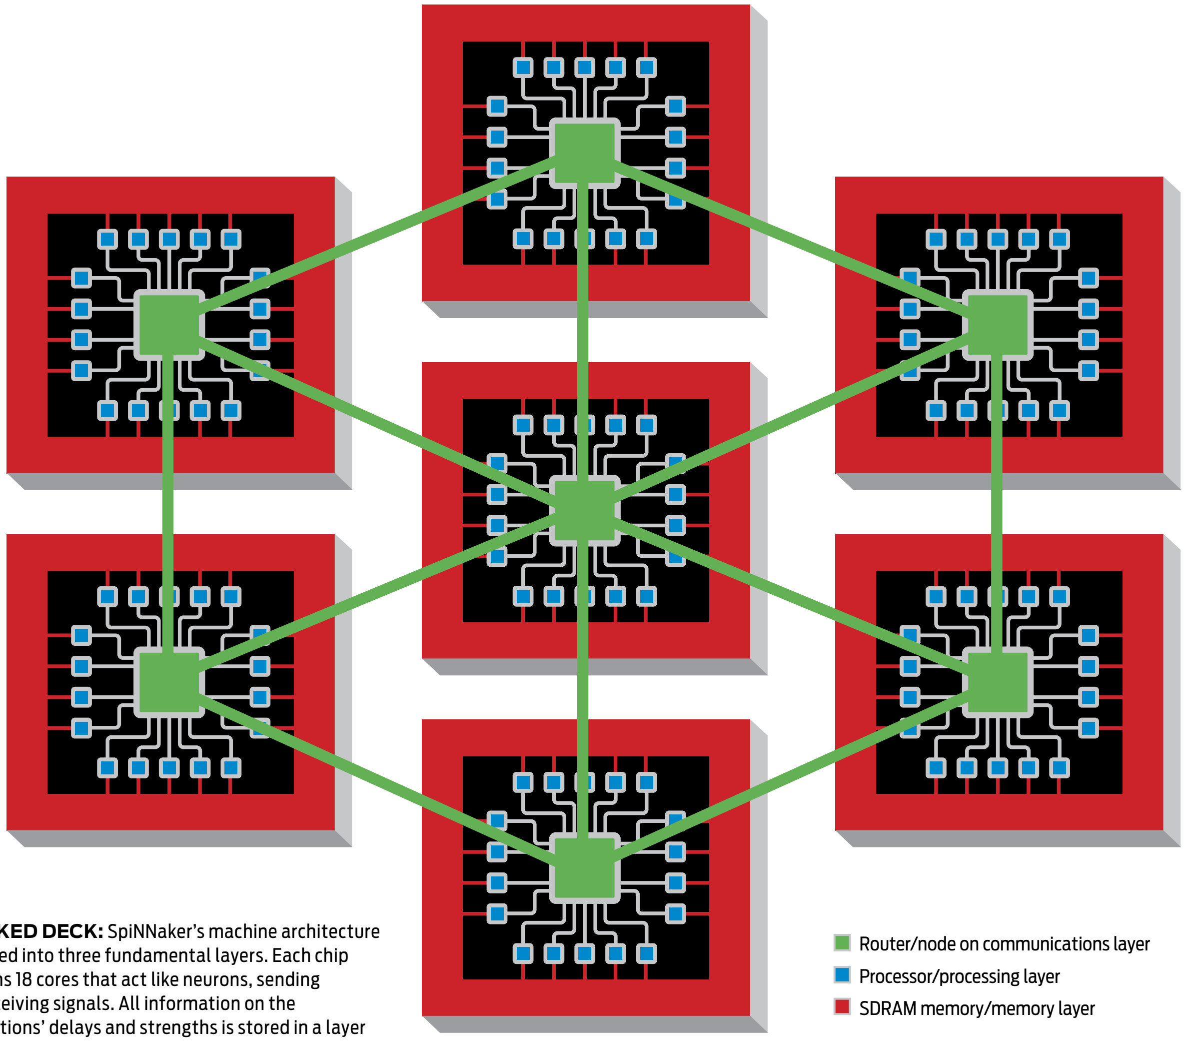

STACKED DECK: SpiNNaker's machine architecture is divided into three fundamental layers. Each chip contains 18 cores that act like neurons, sending and receiving signals. All information on the connections' delays and strengths is stored in a layer of synchronous dynamic RAM (SDRAM) on each chip, and all signals pass through a separate router layer.

Another team, led by Dharmendra Modha at IBM Almaden Research Center, in San Jose, Calif., works on supercomputer models of the cortex, the outer, information-processing layer of the brain, using simpler neuron models. In 2009, team members at IBM and Lawrence Livermore National Laboratory showed they could simulate the activity of 900 million neurons connected by 9 trillion synapses, more than are in a cat's cortex. But as has been the case for all such models, its simulations were quite slow. The computer needed many minutes to model a second's worth of brain activity.

One way to speed things up is by using custom-made analog circuits that directly mimic the operation of the brain. Traditional analog circuits—like the chips being developed by the BrainScaleS project at the Kirchhoff Institute for Physics, in Heidelberg, Germany—can run 10 000 times as fast as the corresponding parts of the brain. They're also fabulously energy efficient. A digital logic circuit may need thousands of transistors to perform a multiplication, but analog circuits need only a few. When you break it down to the level of modeling the transmission of a single neural signal, these circuits consume about 0.001 percent as much energy as a supercomputer would need to perform the same task. Considering you'd need to perform that operation 10 quadrillion times a second, that translates into some significant energy savings. While a whole brain model built using today's digital technology could easily consume more than US $10 billion a year in electricity, the power bill for a similar-scale analog system would likely come to less than $1 million.

Speed could actually be a disadvantage in some cases. If, for example, you want to develop a robot brain that can handle visual or audio inputs, it helps to have a neural model that works at about the same speed as a brain does so that it behaves at natural speeds. There are some ways around that problem. The Neurogrid project at Stanford, for example, builds brain models by operating analog circuits below their transistors' threshold voltage, which slows operations down to a biologically realistic rate. This approach comes with a caveat, however: It relies on fairly large circuits, which are hard to scale up in an economical way.

But as speedy and efficient as analog circuits are, they're not very flexible; their basic behavior is pretty much baked right into them. And that's unfortunate, because neuroscientists still don't know for sure which biological details

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

are crucial to the brain's ability to process information and which can safely be abstracted away. That's certainly the case for dendritic trees—the branching network of neuron inputs that collect the signals arriving from all incoming connections. There is biological evidence that suggests that two such signals coming from the same branch affect a neuron differently than they would if they came from different branches, but it's still not clear whether this has a big impact on brain operation and needs to be included in models.

In 2005, my colleagues and I set out to find a good compromise between the shortcomings of the traditional digital and analog approaches to brain modeling. We wanted to come up with a system that would be capable of modeling brain activity in real time, as analog circuits do, yet be as programmable as a general-purpose digital computer.

We ended up with SpiNNaker, which received £5 million ($8 million) from the United Kingdom's Engineering and Physical Sciences Research Council in 2006. Four U.K. universities—Cambridge, Manchester, Sheffield, and Southampton—are involved in the project, along with three industry partners, ARM, Silistix, and Thales, which contributed the processor and interconnect technologies.

The basic idea behind SpiNNaker is pretty simple. The machine will consist of 57 600 custom-designed chips, each of which contains 18 low-power ARM9 processor cores. Such chips are, of course, eminently programmable. At the center of each chip, we place a specially designed router that receives and directs all the packets coming from the cores and forms links with neighboring chips. We stack 128 megabytes of synchronous dynamic RAM, or SDRAM, on top of each chip to hold the connectivity information for up to 16 million synaptic connections.

As with most other brain models, SpiNNaker's operation is centered on the "spike"—an idealization of the electrical impulse sent out by firing neurons. The information needed to model a spike is tiny: You can condense it down to a single packet containing just 40 bits. But things get complicated when you set out to pass around as many of those packets as the brain does. To model even 1 percent of the human brain could involve wrangling 10 billion packets a second, each of which might need to be sent along to dozens of other chips containing hundreds of processors.

Such traffic is tough for even the best parallel computers to handle. Their architectures are optimized for quickly passing around big chunks of data from one point to another, but they perform very badly when juggling a great number of very small packets. The problem lies in the way the communications system is organized. Because they're designed to be very flexible and decentralized, conventional supercomputers off-load much of the routing information to the data packets they ferry. Each packet carries all its routing information in a header, and this header can't really be scaled down to fit the size of a packet. This approach works fine if you're transporting data in large chunks, but it becomes a burden when it comes to small packets. A tiny 40-bit spike may need to carry 10 times that amount of data in order to be routed properly. Setting up such headers on many tiny packets wastes a lot of energy and drastically reduces speed by clogging bandwidth.

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/brain0812

We've eliminated this problem by taking those routing responsibilities away from the processors. In SpiNNaker, a processor modeling a spiking neuron sends a small packet that uniquely identifies the neuron to the router at the center of the chip. When a router receives a packet, it looks up the packet's

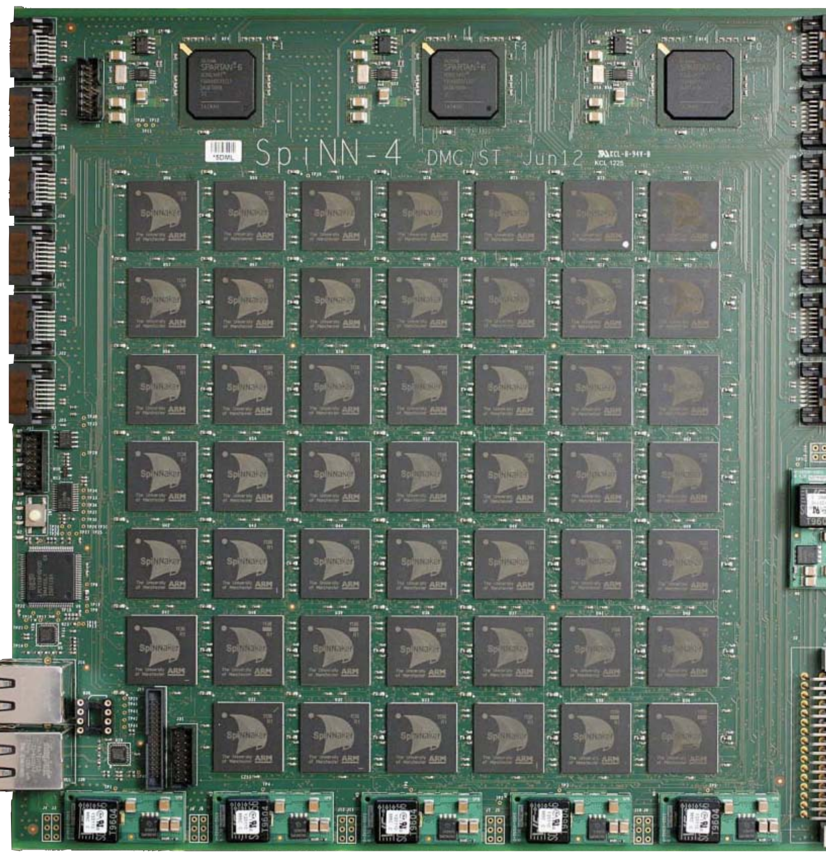

unique identifier in a precomputed table that lists all the connections between neurons. Then the router passes copies of the packet out to other processors on the same chip or to routers on six adjacent chips. All the processors do is receive spikes and, if the total spike input is strong enough, generate new spikes.

In SpiNNaker we cannot implement anything like the hundreds of thousands of physical connections that are sometimes found among individual neurons. However, we can make up for that weakness by exploiting the computational power of the cores as well as the millionfold speed advantage that signals moving along metal wires have over biological ones. Because modeling a single spike requires only a fraction of the core's time, we can save on space and power by packing about a thousand simulated neurons onto every processor. The output signals generated by the interaction of those thousand neurons come in the form of spikes, which we send out using only the wires that connect each of the processors and routers. We can keep all these overlapping signals in order by using careful multiplexing.

LEFT: NORCOTT TECHNOLOGIES LIMITED; RIGHT: UNIVERSITY OF MANCHESTER

By designing our machine in this way, we've thrown out some of the central axioms of parallel processing. Two of the key ones are the need for synchronization among the many programs running on the processors in the machine, and the expectation of deterministic operation—the idea that if you run the same program twice, you get exactly the same result both times. All the processors in SpiNNaker run in real time, and no attempt is made to impose global synchrony using a central clock. This approach mimics the asynchronous way the brain works. Communications between processors are initiated whenever a sender wants to send, and signals arrive at the receiver unheralded and must be handled, ready or not. This means that, just as in the brain, the precise ordering of signals is unknown, and the results can differ in minor ways from one run to the next.

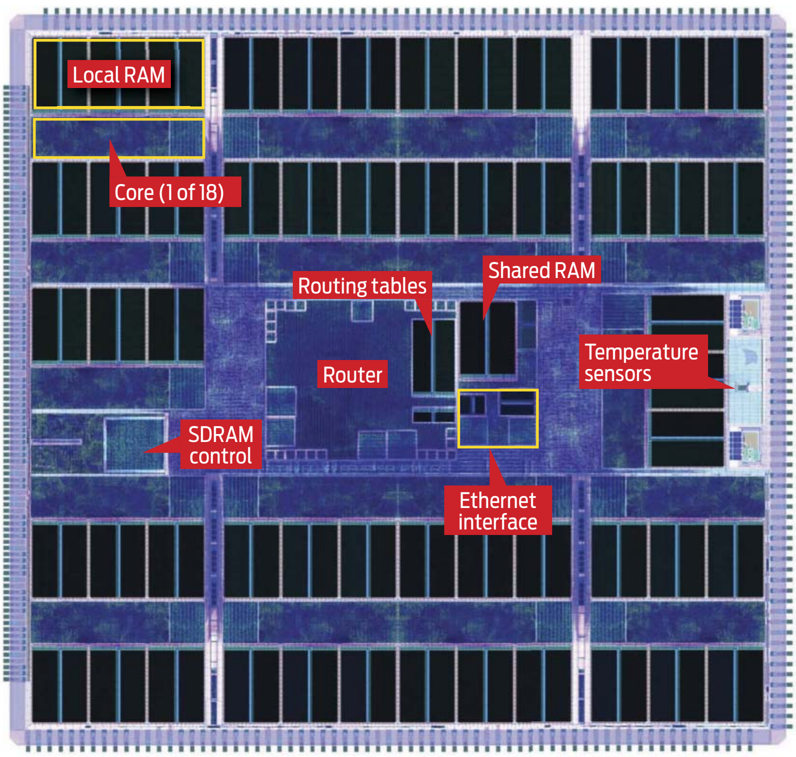

CHIPS AHOY: To keep SpiNNaker as compact as possible, the machine's chips are packed together in sets of 48 onto 23-centimeter-square boards [left]. A SpiNNaker chip contains 18 ARM9 cores [above], each with local RAM. Cores communicate with one another and with more-distant cores via a router at the center of each chip. All the information on the connectivity of the system is uploaded to these routers.

The basic operation of SpiNNaker involves mapping a problem onto the machine—setting up the connectivity graphs in the machine's routing hardware—and then letting the model run with the spikes flying where and when they may.

Building a digital computer in this way comes with a lot of flexibility advantages. With SpiNNaker, there is effectively no difference between communicating with a nearby processor and one that's many chips away. We can upload any neural network we'd like, and the exact way that processors are connected should have no bearing on how fast that neural network can be modeled. In a sense, the SpiNNaker machine could be considered a rewirable computer—an enormous version of the field-programmable gate array chip, or FPGA, specialized for neurons. With appropriate tweaking, it should be able to model any part of the brain we choose.

OUR FULL 57 600-chip machine won't be finished until the end of 2013, but we've already made some progress. Since we accepted delivery of the first SpiNNaker test chip in May 2011, we've built circuit boards containing four such chips, for a total of 72 processor cores. We've mounted this prototype system onto a simple wheeled robot and shown that the robot can perform real-time processing of basic visual information, like following the path of a white line of tape. It's certainly not a difficult task for a modern computer, but it shows that SpiNNaker chips can be connected to form a real-time neural network and can interact with the world through real-world sensors and actuators. We recently received the first 48-node boards, which will be used to build the upcoming system.

When complete, the full million-processor SpiNNaker machine will occupy 10 or so standard 19-inch racks and consume 50 to 100 kilowatts of power. That's still about a hundred times as much as a comparable analog model would need, but then again it's only about a hundredth the power you'd need for an equivalent supercomputer. We also have room to improve. To save money, our processors were built using a decade-old, 130-nanometer chip manufacturing process. If the project produces good results, we could move to a much smaller feature size for our integrated circuits and potentially drop power consumption by a factor of 10.

Although we've carefully designed and simulated the machine, there are still quite a few engineering questions we'll have to sort out once the machine is built. We'll need to figure out the best way to divide large networks into pieces that are small enough to be mapped onto a single processor, to cope with run-time faults, and to package everything into software that our neuroscientist and psychologist collaborators will find easy to use.

Even partial progress toward understanding how the brain works could yield dramatic benefits. We work with psychologists who use neural networks to model and test treatments for reading disorders caused by strokes or similar brain damage. These networks are trained to read—translate text to speech—and are then selectively damaged to reproduce the clinical pathology. SpiNNaker will allow these models to become more detailed and sophisticated, which could help psychologists select better therapies.

We also expect that our computer architecture could help fields outside of brain modeling that can also benefit from computers capable of dividing problems into a very large number of small processes. Some areas that could benefit include computer graphics, circuit modeling, and drug discovery.

SpiNNaker won't get us all the way to full-scale simulations of the human brain. But the machine's communications architecture could help pave the way for better-networked analog chips that could get us there. It will also help show us what information we need to make good models. Then we can really put our brains to use. ❏

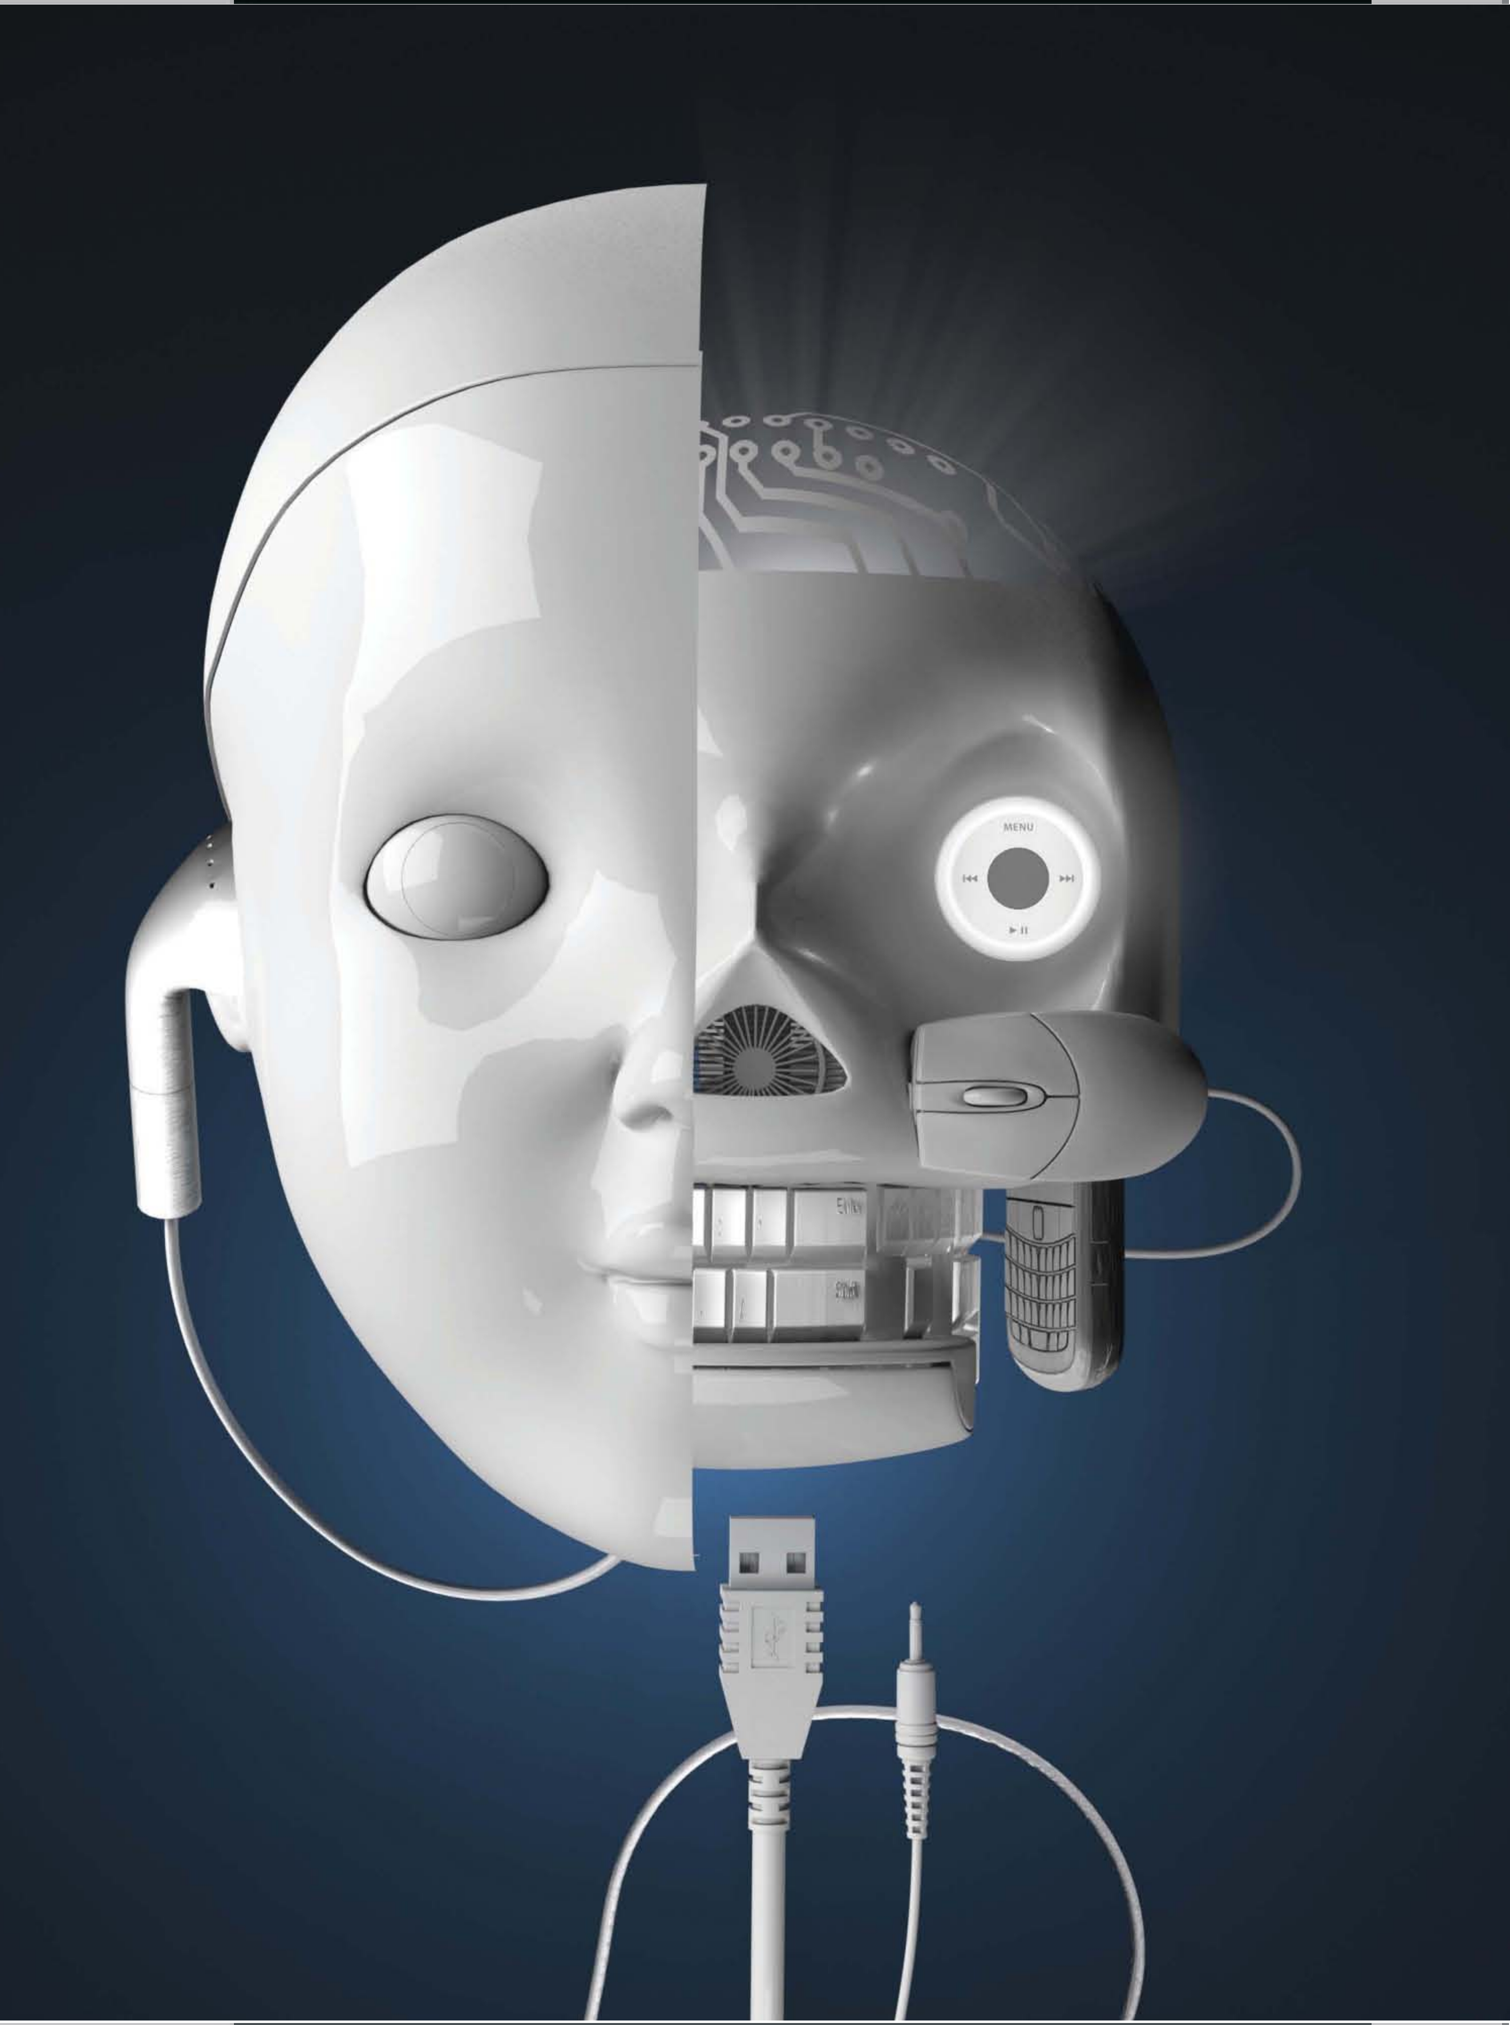
MENU

OUR TOOLS ARE USING US

Human brains can't cope with today's technology

BY WILLIAM H. DAVIDOW

ILLUSTRATIONS BY EDDIE GUY

COMMENTARY

In my days as an engineer, I ran the microprocessor division at Intel Corp. I then became a venture capitalist, investing in companies that built semiconductors, computers, networking systems, and Internet-related services. I focused on products that helped businesses run more effectively and gave little thought to how they might affect our minds, social interactions, and governance.

That lapse now comes home to me as I see people walking down the street, eyes fixed on the screens of their mobile phones, ears plugged into their iPods, oblivious to their surroundings...to reality itself. They are not managing their tools; their tools are managing them. Tools now make the rules, and we struggle to keep up.

I've spent my career developing and financing the companies that supply these profoundly powerful tools. For the most part, I thought of them as harmless, and I believed my job was simply to make the tools better so that others would use them to improve the world. Only in recent years have I become aware of and concerned about their serious side effects. And so I have decided to study them and do my best to explain those effects to the world. Here's what I've learned.

First, it wasn't always this way. Our relationship with tools dates back millions of years, and anthropologists still debate whether it was the intelligence of human-apes that enabled them to create tools or the creation of tools that enabled them to become intelligent. In any case, everyone agrees that after those first tools had been created, our ancestors' intelligence coevolved with the tools. In the process our forebears' jaws became weaker, their digestive systems slighter, and their brains heavier. Chimpanzees, genetically close to us though they are, have bodies two to five times as strong as ours on a relative basis and brains about a quarter as big. In humans, energy that would have gone into other organs instead is used to run energy-hungry brains. And those brains, augmented by tools, more than make up for any diminishment in guts and muscle. Indeed, it's been a great evolutionary trade-off: There are 7 billion people but only a few hundred thousand chimpanzees.

In the distant past our tools improved slowly enough to allow our minds, our bodies, our family structures, and our political organizations to keep up. The earliest stone tools are about 2.6 million years old. As those and other tools became more refined and sophisticated, our bodies and minds changed to take advantage of their power. This adaptation was spread over more than a hundred thousand generations.

Our social structures evolved along with the tools. Some 10 000 years ago, tribes of roaming hunter-gatherers began to stay in one place to raise crops. Agriculture made cities possible, and with cities came the arts and commerce. As transportation improved and cities grew, it became important to control distant places that supplied food and raw materials. About 6000 years ago, ancient city-states such as Uruk emerged in Mesopotamia and governed the surrounding countryside. Millennia later, Athens, the largest of the Greek city-states, controlled about 2500 square kilometers and most of the Aegean Sea. So far, so good.

But with the invention of movable type in medieval times, followed by other improved ways of connecting—better ships and roads, trains, planes, automobiles, and the Internet—technology raced ahead of us. In our own time the most striking example of the acceleration of technology has been Moore's Law, under which the density of transistors on chips has doubled every 18 months.

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

The performance of single processors long rose in tandem with the transistor count, and even after that relationship stalled in the mid-2000s, the switch to devices using many processors kept the performance curve pointing upward.

This exponential rise in capability has greatly augmented the pooling of knowledge from different sources to achieve the creative synthesis described by the 19th-century mathematician and philosopher Henri Poincaré: "To create consists precisely in not making useless combinations, and in making those which are useful and which are only in a small minority.... Among chosen combinations, the most fertile will often be those formed of elements drawn from domains which are far apart."

Drawn in part from the Internet, the newly created knowledge gets deposited back on the Internet, increasing its scope and accelerating the development of technology. Burgeoning knowledge in turn drives rapid change—it advances technology, transforms business institutions, and changes how markets work and how people interact. Governments, social institutions, and our brains struggle to keep up.

Even the nature of change itself has changed. Living creatures started out evolving in one dimension defined by the physical world and another defined by the biological world. Then came humans, who added a third dimension—the artificial one engendered by tools and technology. Now, with the widespread use of the Internet, a fourth dimension has been added—that of virtual reality, or cyberspace. It is indeed appropriate to consider this last dimension as real and distinct from the tools and technologies of the past, because however fast those things may have changed, the rate of change in the virtual space is much faster. It took a lot of time to build physical infrastructure—railroads, highways, bridges, skyscrapers, and so forth. But in virtual space, entire new infrastructures can arise overnight, as Google and Facebook have proved.

I now believe that our minds, bodies, businesses, governments, and social institutions are no longer capable of coping with the rapid rate of change. And it is obvious that this change is indeed more rapid than any comparable change that came before.

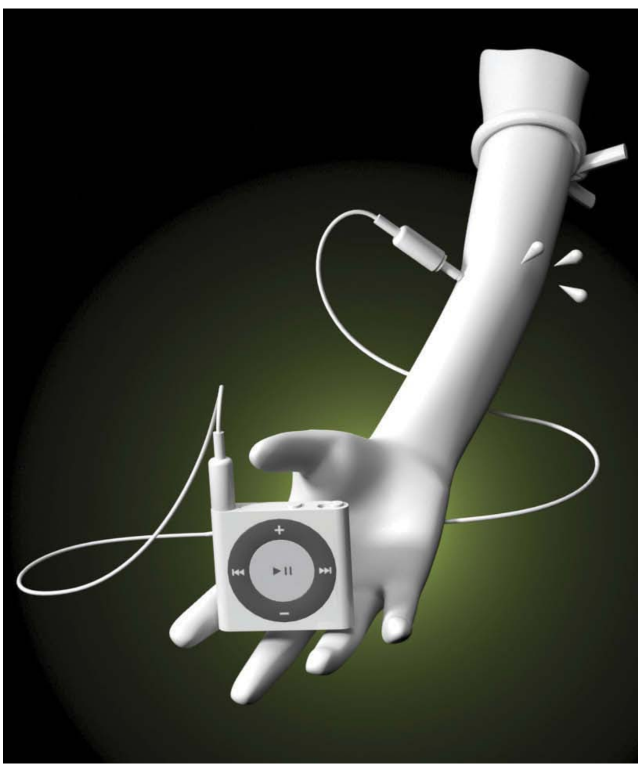

Think of the many years it took Barnes & Noble to build its retail chain of U.S. bookstores. The company set up its first bookstore in 1917, and by 2010 it was operating 717 stores. It took time for the company to find the proper locations, lease them, and stock them with inventory, and still more time to build them into viable businesses. The company was limited in what it could do because only certain physical locations were suitable for retail stores. Compare that long history to the rise of Amazon.com, which started in 1994 and was operating in virtual space throughout the United States by the next year, putting a bookstore in every home that had an Internet connection.

Barnes & Noble responded with an online strategy of its own, one that now gives pride of place to sales of e-books, and the company continues to fight on; meanwhile, competitor Borders was liquidated. Both constitute sterling examples of the "creative destruction" of capitalism, as the economist Joseph Schumpeter put it. But the fact that entire business models can come and go that fast is extraordinary. It also indicates the challenges that rapid change presents to other institutions.

An example is our financial institutions, transformed in the past 20 years by radical innovations, such as the introduction of high-speed trading—in which computers trade securities with other computers—and the immensely complex new investment vehicles known as derivatives. These innovations, enabled by improvements in computing power and telecommunications, have made markets more volatile and played leading roles in the recent worldwide economic meltdown. And our regulatory framework has failed to keep up.

In 2005 about 80 percent of the shares of companies listed on the New York Stock Exchange were traded on its floor or on its proprietary digital system; today only 20 percent are traded there. Most of the other shares are traded on alternative systems that have cropped up, systems in which many of the old rules and regulations no longer apply. High-frequency traders have exploited these lightly regulated trading systems. In the United States, computer-driven algorithms now execute 60 percent of all stock trades. One result was the Flash Crash on 6 May 2010, when the Dow Jones Industrial Average fell 1000 points in a matter of minutes. Although it recovered 600 points of the loss minutes later, the episode shows us just how little we understand about high finance and how vulnerable we are to its vagaries.

European stock exchanges have become vulnerable as well. The Better Alternative Trading System, for example, has taken market share away from the London Stock Exchange. Xavier Rolet, CEO of the London Stock Exchange, has tried to counter this threat by diversifying the exchange's business. Whether this strategy can slow the decline of the London exchange is unclear.

The Internet has also made it easier to participate in the over-the-counter market in derivatives—which are essentially side bets on the value of assets. In 2000, the notional value of these OTC derivatives was US $60 trillion; by 2007 it had risen to $600 trillion. Losses on these derivatives played a major role in the 2008 financial crisis and are still causing problems today. In late 2009, a little-known company called Markit Group created the iTraxx SovX index, which made it easier to use derivatives to place bets on the possibility of a Greek default. As the cost of insuring Greek debt based on the iTraxx SovX rose, investors shunned Greek bonds, making it harder for the country to borrow. Of course, the Greek economy was in dire straits anyway, but those problems were aggravated by the new strategies of speculation that technology has made possible.

Regulatory reform has been a case of too little, too late. Basel III, an agreement concluded in 2011 by the banking supervisors of 10 major industrialized countries, is too weak. And in the United States, the 2319-page Dodd-Frank legislation will probably prove to be too complex to achieve its goal of averting another meltdown. Ultimately, the only way to deal with financial innovation in a virtual world is through international regulatory systems based on commonly accepted principles. But as a former head of the Bank for International Settlements told me, good luck with that. The tools are making the rules.

Evidence is also beginning to pile up that our brains can't deal effectively with virtual environments. This makes perfect sense: The challenge they pose is barely half a generation old, yet our minds have been shaped by other challenges that date back thousands of generations.

And as we take bodies and brains adapted to physical space and immerse them in virtual worlds, they are not only unable to cope, they respond in unanticipated ways. As Nora Volkow, the director of the National Institute on Drug Abuse, has observed, "the technology is rewiring our brains."

We already know that physical stimuli can cause profound changes in the brain. Studies of combat veterans afflicted with post-traumatic stress disorder, for example, reveal that it produces a persistent and worrying increase in levels of cortisol, a hormone associated with stress. Increasingly we are seeing evidence of similar changes from the stress we experience when moving from the real to virtual worlds.

Linda Stone, formerly a developer of interactive media for Microsoft, has studied the multitasking of those of us who sit in front of computer screens for long stretches of time. She finds some of the cortisol-triggered responses seen in combat veterans, including elevations in heart rate and blood pressure. And in a 2010 paper for the National Bureau of Economic Research, Decio Coviello of HEC Montréal and his colleagues concluded that multitasking, by spreading attention among many functions, hurts overall performance. Russell A. Poldrack, a neurobiologist at the University of Texas at Austin, has similarly shown that multitasking retards learning.

Worse, though, is that our brains seem to crave the virtual world, with repeated exposure producing changes that resemble drug addiction. According to Gary Small, a professor of psychiatry at the University of California, Los Angeles, the excitement of getting an e-mail alert causes a release of dopamine, a neurotransmitter that reinforces the behavior and thus drives us to crave more such stimulation. Before long, it becomes impossible for people to put down their iPhones and BlackBerries. Dopamine's effects were shaped by natural selection: It helped to focus our attention so that we wouldn't be eaten by tigers. These days, it is facilitating our consumption by e-mail and text messages.

Many experts believe our Internet addiction is similar to that associated with gambling. In both cases, people find it difficult to function normally, have stable family lives, or be effective at work.

It will be years before we fully understand the lasting effects of living in virtual worlds. But until we do, it is best to approach the situation with caution. The main challenge we face is to recognize that we are designed to reside in a slower-paced physical world. This is extremely difficult to accept. We want our news instantaneously. We want to be in touch with everyone at all times. Our careers depend on our being constantly available.

But we have to make a choice: We can design our lives so that we stay in control, or we can cede the control of our lives to our tools.

If we choose the former path, then we are going to have to rethink the way we regulate financial markets, and in this and many other areas, we are going to need new and stronger laws. For example, we will need to guarantee property rights in virtual space to protect our privacy. And we will need to recognize that the nature of our institutions has changed. For many of us, a single corporation—Apple—now dominates our virtual existence: Our virtual lives are bounded by MacBook Airs to the north, iPads to the south, iPhones to the west, and iPods to the east. This company has more power over me than the regulated telephone companies ever did. Should companies like Apple be regulated?

We also have decisions to make. For my part, I have decided to stay in control. I use the miraculous tools of the information age to augment my life in physical space, rather than live in the virtual space and use the physical world to supplement it.

I'm reminded of an observation made to me a while back by John M. Staudenmaier, a historian of technology who is also a Jesuit priest. He pointed out that the quickest way to end a deep and meaningful conversation was to glance at your watch. What would he say today about our ever more tempting smartphones?

I have shut off most alerts and reminders on my computer and smartphone. I check for e-mail on my own schedule, just a few times a day. At home, I have built a physical wall around the virtual world. I let myself read news on my iPad anywhere in my home, but I answer e-mails and conduct business only in my office. I heed Staudenmaier's advice and never end important conversations by glancing at my smartphone. My iPhone is never present when I am out with my wife, listening to the challenges my kids are facing, or playing and laughing with my grandchildren.

My advice to you is to take control of your tools. I promise your life will be better if you aren't constantly checking to see if you've got mail. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/coevolving0812

Sun Yat-sen University / Carnegie Mellon University

Joint Institute of Engineering

Faculty Positions Available *in Electrical/Computer Engineering*

Pursue research, innovation and education in China.

Sun Yat-sen University (SYSU) & Carnegie Mellon University (CMU) are partnering to establish the Joint Institute of Engineering (JI). SYSU and CMU are embarking on an exciting opportunity to transform engineering education in China. JI will provide world-class education and cutting-edge research in China's Pearl-River Delta region, which provides rapidly growing opportunities for future technology innovation.

JI is seeking full-time faculty, who have an interest in pursuing innovative, interdisciplinary education programs and in leading research efforts in all areas of Electrical/Computer Engineering . Candidates should possess a Ph.D. in ECE or related disciplines, with a demonstrated record and potential for research, teaching and leadership. The position includes an initial year on the Pittsburgh campus of CMU to establish educational and research collaborations before locating to Guangzhou, China.

This is a worldwide search open to qualified candidates from all countries and of all nationalities, with an internationally competitive compensation package for all qualified candidates.

Please visit **http://sysucmuji.sysu.edu.cn** *or* **http://sysucmuji.cmu.edu** *for details.*

Technical University of Denmark DTU

PROFESSOR

Modelling of Electricity Markets

DTU Electrical Engineering invites applications for a position as Professor with special responsibilities (5-year term that can be extended) within Modelling of Electricity Markets. The successful candidate is expected to lead the research and head a research group that will be part of our newly formed Centre for Electric Power and Energy. A key area for the new centre will be development and analysis of electricity systems with high shares of variable and renewable power production.

Application deadline: **1 September 2012**

DTU is a technical university providing internationally leading research, education, innovation and public service. Our staff of 5,000 advance science and technology to create innovative solutions that meet the demands of society; and our 7,000 students are educated to address the technological challenges of the future. DTU is an independent academic university collaborating globally with business, industry, government, and public agencies.

SCIENCE AND TECHNOLOGY AT A GLOBAL SCALE
– SET THE STANDARDS FOR THE FUTURE
See our PhD-programmes at dtu.dk/phd

Full Professor in Internet Technology

KTH School of Information and Communication Technology, ICT seeks a Full Professor in Internet Technology.

Royal Institute of Technology (KTH) in Stockholm is the largest and oldest technical university in Sweden. Education and research spans from natural sciences to all branches of engineering and includes architecture, industrial management and urban planning. There are a total of just over 15,000 first and second level students and more than 1,600 doctoral students. The school of Information and Communication Technology, ICT, conducts research and education at leading international level in material physics, photonics/optics, electronic systems, integrated devices and circuits, as well as communication systems, **www.kth.se/en/ict**.

We expect the new professor to build and lead a world-leading research group and an educational program in the subject area of Internet Technology. The subject area includes design and analysis of internet based communication networks. Special emphasis is put on architectural and system issues in large-scale networks, as well as communication protocols for such networks.

The application deadline is September 3, 2012. For complete and official information about the position and for application procedure visit **www.kth.se/en/om/work-at-kth/vacancies**.

Available Position

Postdoctoral Position for Gas Research Center (GRC) funded project

Department **Electrical Engineering**
Location **Petroleum Institute – Abu Dhabi**

Description
A **Postdoctoral** position is currently available for the GRC project– GRC-002 **Black Powder in Gas Flow lines: Sources, Prediction and Mitigation** at the Petroleum Institute in Abu Dhabi. The project is an applied research that includes both experimental and computational modeling studies in transport and deposition of particles in flow lines. There are tremendous opportunities for research collaboration with the PI's industrial stakeholders, as well as the affiliated universities.

Qualifications: Interested candidates should have a Ph.D. in Electrical engineering with expertise in embedded systems design and confidence in the design, development, and use of electronic instrumentation are required in addition to good practical skills.

We are particularly interested in candidates with the following qualities:

1– *Strong background in analog and digital electronic design*
2– *Ability to effectively build and troubleshoot embedded hardware boards*
3– *Self-independent and self-motivated*
4– *Very good mathematical background. Experience in Modeling is a great advantage*
5– *Very good written and oral communication skills*
6– *Work experience in relevant research projects is preferable.*

Salary/Benefits: Salary and benefits will be commensurate with qualifications and experience. The total compensation package includes a tax-free 12-month base salary, and a benefits allowance that covers relocation, housing, initial furnishings, utilities; transportation (automobile purchase loan), health insurance, children education, end-of-service benefit and annual leave travel. Applicants must be in excellent health and will be required to pass a pre-employment physical examination.

To Apply: Interested candidates should submit all materials online ***https://career.pi.ac.ae***. This is a fixed term position for three years and is available immediately. The deadline for applications is **15TH September, 2012**. Only shortlisted applicants will be notified.

Centre for Nano Science and Engineering

Indian Institute of Science, Bangalore

http://www.cense.iisc.ernet.in/

We invite applications for tenure track Assistant Professor positions. The candidates should have an excellent academic and research record with at least 3 years post-PhD research experience.

We pursue interdisciplinary research across several disciplines with a focus on nanoscale systems. In particular we seek candidates with interest and expertise in (a) Nanoelectronics and photonics, (b) Lab-on-chip and biomedical devices, (c) Photovoltaics and energy storage systems, (d) Sensor systems, integration and packaging and (e) Computational nanoengineering.

The applications with a detailed CV, research and teaching plan and list of references should be sent to **fsc@cense.iisc.ernet.in**, with a copy to **fsc.cense@gmail.com**

the data

Biggest New Oil Areas

Estimated reserves, in billions of barrels [blue areas]

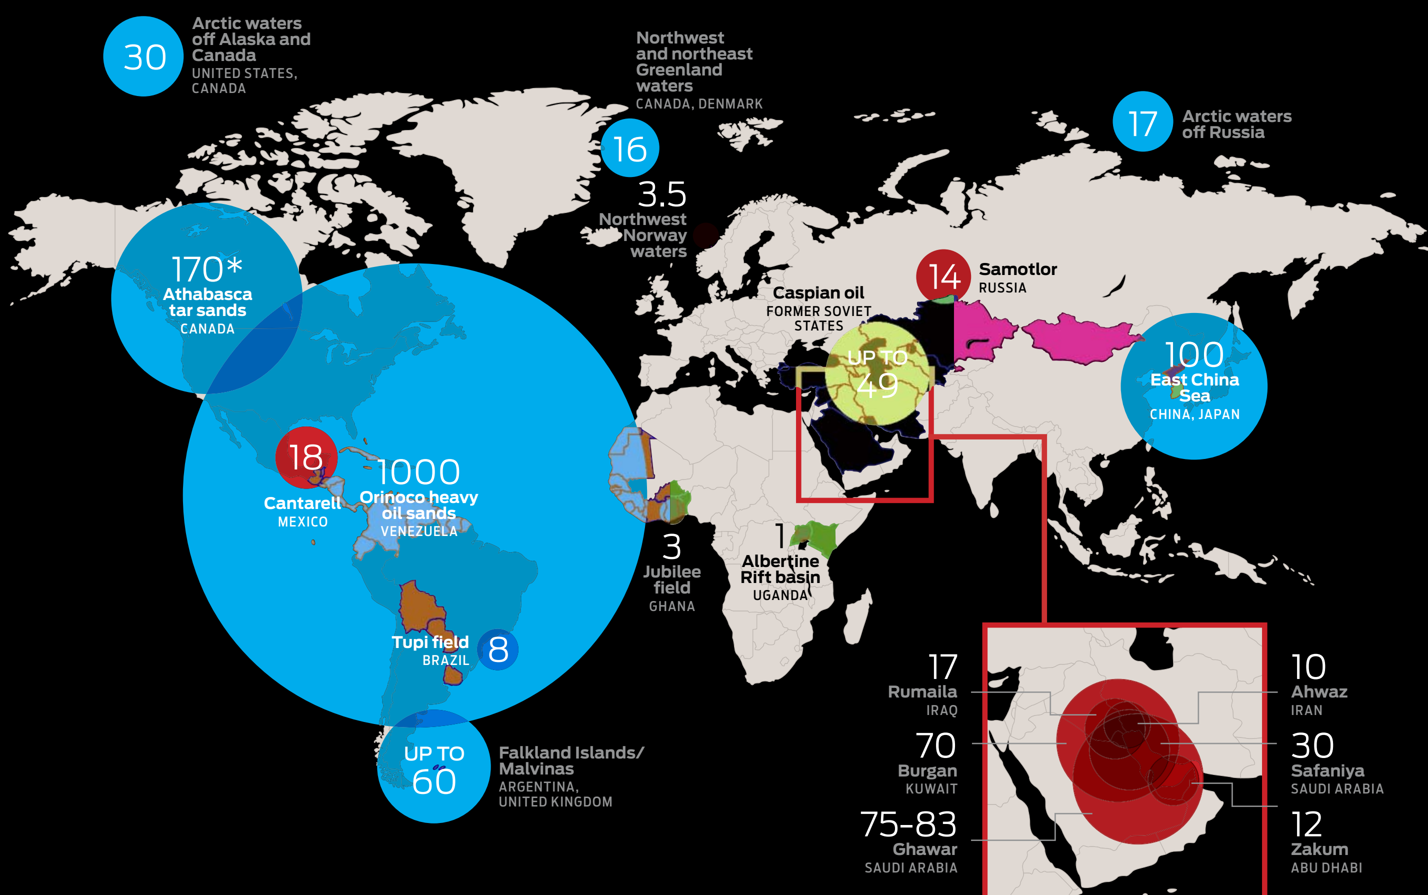

The Age of Hard-to-Get Oil

THE WORLD'S "proven" oil reserves, those that are readily recoverable with present-day technology, are estimated to be about 1.4 trillion barrels of oil. With daily global consumption around 85 million barrels and edging toward 100 million, we're on schedule to run out of easily extracted oil in a generation.

So what will the world energy economy look like 30 to 40 years from now?

IEEE Spectrum spoke to Michael T. Klare, the author of *The Race for What's Left: The Global Scramble for the World's Last Resources* (Metropolitan Books, 2012). He sees two major trends shaping the future. First, as the world keeps getting warmer, there will be more pressure to alter behavior. So global demand for oil a generation from now will not be a linear projection from the patterns of past decades.

Second, with production already declining sharply at most of the world's major existing oil fields, more and more of the oil will come from harder-to-get-at sources—ultradeep-water deposits, Arctic reserves, the oil sands of Canada's Alberta province, and the extraheavy crude of Venezuela's Orinoco Belt. If exploited to the fullest potential, Venezuela's reserves alone could satisfy world demand for another generation. But that oil will not be brought to the surface with the oil rigs we saw in the film *There Will Be Blood* (2007), notes Klare. The new technology will be more sophisticated, more resource demanding, more environmentally threatening, and of course, much more expensive. —*William Sweet*

Biggest Current Oil Fields

Total capacity, in billions of barrels** [red areas on map]

Bars below represent peak production, in millions of barrels per day. The lighter region indicates decline from peak.

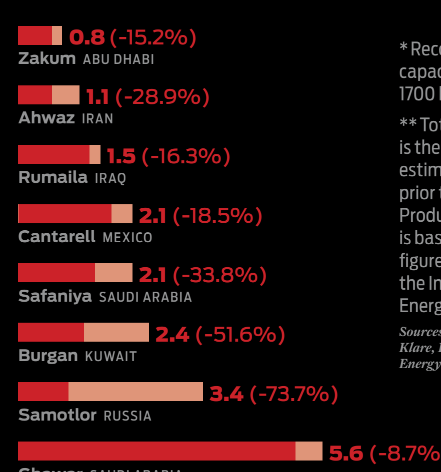

*Recoverable; total capacity is up to 1700 billion barrels.

**Total capacity is the most recent estimate of reserves prior to extraction. Production decline is based on 2007 figures reported by the International Energy Agency.

Sources: Michael T. Klare, International Energy Agency

COLUMN FIVE MEDIA, BRANDON PALACIO

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