IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 04.17



THE BEST AND BRIGHTEST ON THE ROAD, FROM A THREE-WHEEL SPEEDSTER TO A HYBRID MINIVAN

# TOP TECH CARS 2017


The all-electric, back-to-the-future **Morgan EV3** goes 240 kilometers on a single charge

IEEE


## Design Smarter

CST

From initial concept to final prototype, every stage of product development means balancing multiple requirements; performance, reliability and cost on the one hand, and specifications, legal regulations, and deadlines on the other. Throughout the design cycle, CST STUDIO SUITE® can help to achieve these goals.

New features in the 2017 version of CST STUDIO SUITE focus on analysis. From individual components such as antennas and filters through to full assembled systems, devices can be simulated quickly and accurately thanks to the array of new modeling, synthesis and analysis features. Filter Designer 3D, CHT solver, EMC Interference Task and antenna-to-antenna coupling in the asymptotic solver all build on decades of research and development that have made CST STUDIO SUITE the industry-leading electromagnetic simulation software.

See the full picture and develop the potential of your ideas with CST STUDIO SUITE 2017. Smart design for smart devices.

IEEE SPECTRUM



On the cover **Photograph by Morgan Motor Co.**

PHOTOGRAPH BY **The Voorhes**

DEPARTMENTS_04.17

IEEE SPECTRUM




ADDITIONAL RESOURCES

## Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Understanding Stimulated Brillouin Scattering With a Multiphysics Approach in COMSOL—6 April

## White Papers

Available at http://spectrum.ieee.org/whitepapers

- Hardware and Measurement Abstraction Layers
- IoT Device Test Challenges
- Your Design Problems, Answered
- Accelerate R&D With Multiphysics Simulation
- Using Loadability Studies to Comply With NERC PRC-025

## The Institute

Available at theinstitute.ieee.org

- **MEET JIBO** Virtual assistants will soon become part of the family. Jibo, the first social robot for the home, can order dinner when you're running late from work and read bedtime stories to kids. We interviewed IEEE Fellow Roberto Pieraccini, director of advanced conversational technologies at the robot's manufacturer (also called Jibo), to learn more.
- **BIG BREAKS** Learn how four famous inventors, including Nikola Tesla and Grace Hopper, got their start.
- **A MODEL SCIENTIST** We interviewed Danica Kragic—a robotics engineer, fashion model, seamstress, and IEEE Fellow—who told *Vogue* magazine that she dreams of designing a robot that can sew.

**IEEE SPECTRUM**
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2017 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 54, issue No. 4, International edition. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

FROM LEFT: STINGRAY MARINE SOLUTIONS; ETYMOTIC RESEARCH; THE UNIVERSITY OF UTAH

BACK STORY_


ALEX EMMONS

# Hot on the Trail of Cold Warriors

SHARON WEINBERGER, AUTHOR OF THIS MONTH'S FEATURE "The Bunny, the Witch, and the War Room," is no stranger to weird corners in the U.S. military-industrial complex. In her first book, *Imaginary Weapons: A Journey Through the Pentagon's Scientific Underworld* (Nation Books, 2006), she described a controversial attempt to devise weapons based on a metastable form of the element hafnium. Soon afterward, she and coauthor Nathan Hodge wrote about their journeys to various atomic facilities and research centers in *A Nuclear Family Vacation: Travels in the World of Atomic Weaponry* (Bloomsbury Publishing, 2008).

Her *IEEE Spectrum* article is based on a chapter of her newly released third book, *The Imagineers of War: The Untold Story of DARPA, the Pentagon Agency That Changed the World* (Alfred A. Knopf). In it, she describes what might be the strangest undertaking by American military researchers during the Cold War: an effort to harness powers of the paranormal. You read that right. They hoped to make military use of ESP, psychic phenomena, and similar "X-Files" fodder.

The whole notion was ludicrous, of course, but many people in the intelligence and defense communities–including people with scientific training–thought otherwise at the time. "The Pentagon has been uniquely susceptible to bad science," says Weinberger. To its credit, the Advanced Research Projects Agency (as DARPA was then called) rejected the absurd idea that paranormal phenomena could be exploited. But others in government were not so quick to reach that conclusion.

We don't want to spoil your fun by revealing too much, but we will disclose that the aforementioned bunny, witch, and war room are *not* being used metaphorically in the title. You'll have to read Weinberger's article, though, to understand how that's even possible. ■

**CITING ARTICLES IN IEEE SPECTRUM** *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, Past Forward is in *IEEE Spectrum*, Vol. 54, no. 4 (INT), April 2017, p. 56, or in *IEEE Spectrum*, Vol. 54, no. 4 (NA), April 2017, p. 60.

# CONTRIBUTORS_

## Michael Dumiak

Dumiak, a Berlin-based journalist who writes for *Smithsonian Air & Space*, *Popular Science*, and *The Lancet*, often feels he cannot afford salmon. In "Lice-Hunting Underwater Drone Protects Salmon" [p. 7], he explores tech that should make the fish cheaper. "Autonomous lasers shooting parasites underwater–that got my attention," Dumiak says. "As I learned more about it, the setting, the nature, the personalities, the economics became more complex and compelling."

## Glenn Edens

In 2014, Edens was managing a team at the Palo Alto Research Center that was working on a promising new Internet architecture known as content-centric networking. His group was deeply divided over competing versions, and Edens eventually decided to scrap one project to focus on the other. "In a perfect world, you would keep all the projects going," he says. "In the real world, you don't have that luxury." Edens writes about the results in "The Packet Protector" [p. 40].

## Robert Hebner

Hebner is the director of the Center for Electromechanics at the University of Texas at Austin. Previously he served as acting director of the National Institute of Standards and Technology, among other positions in the U.S. government. An IEEE Fellow and holder of a Ph.D. in physics, he has published more than 150 research papers and reports. In this issue, Hebner tackles one of the most important topics of our time: the future of the power grid [p. 48].

## Lawrence Ulrich

Back in November, when Ulrich test-drove the Chevy Bolt EV, he concluded that the vehicle could easily surpass its estimated 238-mile (383-kilometer) range and that Chevrolet, not Tesla, had thus become the first automaker to deliver an affordable long-range electric car. That puts the Bolt at the top of *IEEE Spectrum*'s annual review of automotive tech, the seventh one Ulrich has done for us [see "Top Ten Tech Cars," p. 24].

## Robin Finlay

Finlay, a stylist, and Adam Voorhes, a photographer, are the wife-and-husband team known professionally as The Voorhes. Their assignment for this issue was to create eerie Cold War–period images for "The Bunny, the Witch, and the War Room" [p. 34]. For their photos of an era-appropriate engineer, Robin dressed up their assistant, while Adam offered suggestions based on his memories of his father, an electrical engineer.

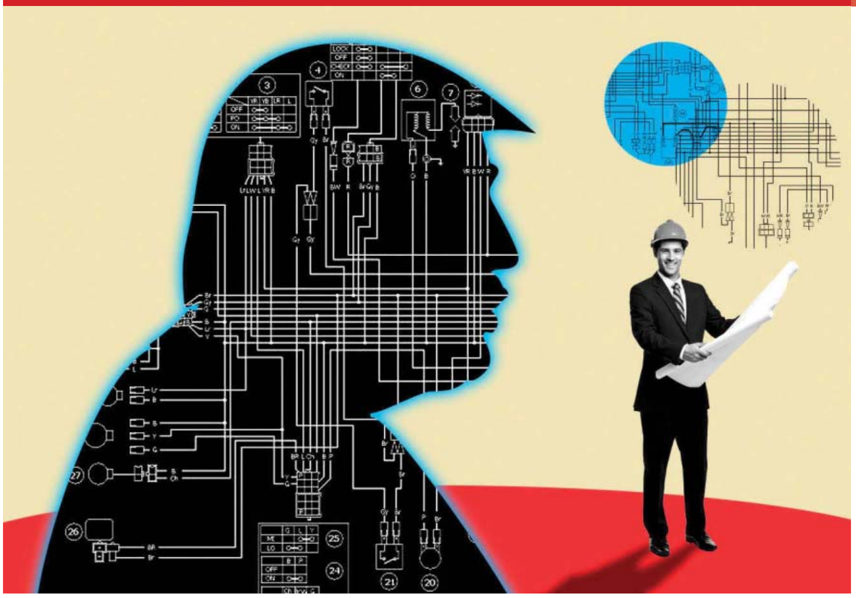

# President Trump Needs an Engineering Adviser

## The new president could use advice from a practical problem solver

Donald Trump, let an engineer help you.

We get that you find scientists—despite their amazing feats of discovery and selfless pursuit of new knowledge—too ideological, independent, and consumed by contentious claims about climate. So like the last Republican president, you will probably go a long time before selecting a science adviser (George W. Bush had been in the White House over seven months before he chose a physicist).

Here's a bold idea: Why not break with tradition and do without a science adviser? Why not choose an engineering adviser instead? Engineers make things work, and keep them working, which is exactly what you need to succeed.

While you're skeptical of elites, there's a lot to like about engineers. They build stuff, they don't just talk about it. Engineers embrace action, not theory. They possess a sense of urgency befitting your own taste for haste. And like you, they find scientists too condescending and credit grabbing.

Engineers do brag, of course, but usually only after their creations are humming. Unlike scientists, who often require unfettered freedom, engineers work within limits. And engineers have an ethos of responsibility. In their professional codes, they hold paramount the safety, health, and welfare of the public.

Engineers are also a diverse tribe—and we know you are diversity challenged. The subfields, such as electrical, civil, and mechanical engineering, share a commitment to problem solving and an enthusiasm for drawing on a broad menu of resources.

As a result, engineers often are the first responders when technological fixes are required to manage or eliminate complex emergencies. Pragmatic responses demand the high-level involvement of engineers, which will be easier and quicker to accomplish when one of their own sits among your top advisers.

For an infrastructure junkie like you, engineers are kindred spirits. Complex systems sit at the heart of the nation's critical infrastructure, the engineers' domain.

Selecting an engineer to join your inner circle of advisers would honor tradition, too. The first presidential science adviser, who served Franklin Roosevelt during World War II, was an electrical engineer. Not only did Vannevar Bush deliver technologies that helped defeat the Germans and the Japanese, he did so by embracing help from the sort of capitalist industrial titans that you value.

Herbert Hoover, the only modern American president to practice engineering, graduated Stanford with a geology degree in 1895 and worked as a mining engineer. Hoover built his reputation as a great humanitarian by defeating immense logistical barriers to the delivery of food aid to Europe during and after World War I.

Scientists have long believed they have a special capacity to help politicians—and a special claim on the public. That there will be a "March for Science" in Washington on 22 April but no "March for Engineering" speaks volumes about the different cultures of these two communities. Engineers generally eschew grand public gestures. Although they may be blunt, engineers work well in teams, where roles, lines of responsibility, and objectives are clear.

The collaborative aspects of engineering might perplex or even alienate mavericks like you, President Trump. As historian Henry Petroski has noted, engineers usually succeed in making a considerable mark on the world only when receiving help from "an army of assistants."

Along with learning to accept help from others, an engineer could teach you about the benefits of sharing credit. If great scientists personify "solitary genius," engineers value humility and working in the background, attributes that you might try out someday as president. —G. PASCAL ZACHARY

**G. Pascal Zachary** is a professor of practice at Arizona State University's School for the Future of Innovation in Society. The views expressed here are the author's own and do not represent positions of *IEEE Spectrum* or IEEE.

**CORRECTION:** In a photo caption for "Dunking the Data Center" [March] we misspelled the name of Microsoft Research's director of special projects. His name is Norm Whitaker.

PHOTO-ILLUSTRATION: Isabel Espanol

# NEWS

3.4 MILLION METRIC TONS: SALMON, TROUT, AND SMELTS FARMED IN 2014

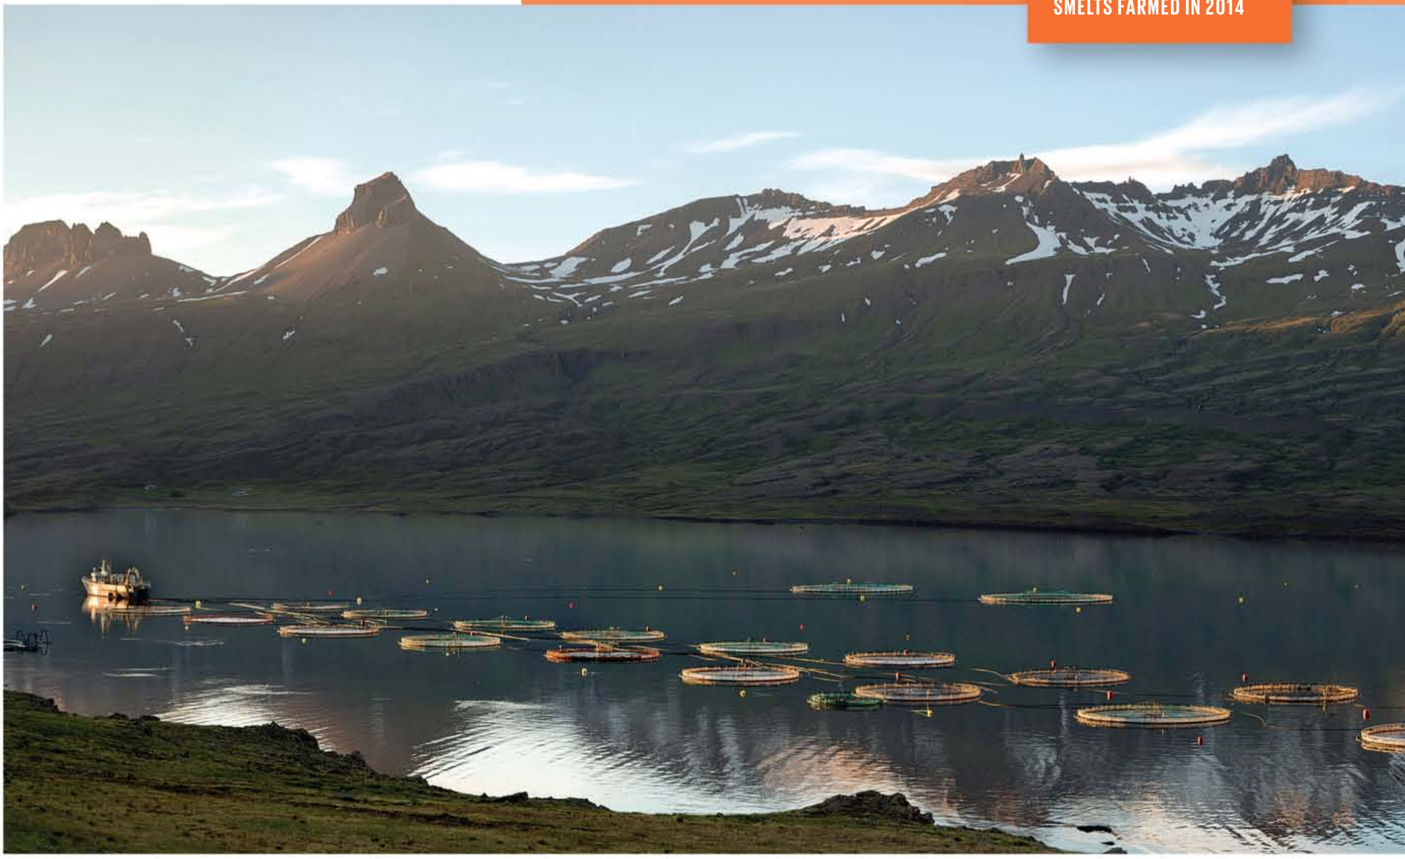

# LICE-HUNTING UNDERWATER DRONE PROTECTS SALMON

**Laser-wielding robot will rid fish farmers of costly pest**

JERZY STRZELECKI

**A FJORD'S FRIEND:** Aquaculture pens like these in Iceland, breed fish, but also the sea lice that feed off them.

**The pens holding cold-water fish** in the far North Sea fjords of Norway and quiet lochs of Scotland, whose depths are normally inky black after dark, now have an eerie night glow. There's a perfectly rational explanation: Thousands of laser pulses are doing an important job, lighting up fish stocks in the latest bid to control a pesky parasite that can injure or even kill farmed salmon, and devastate an industry estimated to reel in about US $10 billion annually.

The problem is sea lice, marine ectoparasites. Two particular sea lice species, *Lepeophtheirus salmonis* and *Caligus elongatus*, attach to salmon and sea trout, feeding off their tissues, blood, and protective external mucus membranes. Though the lice find the flesh and blood of wild and farmed fish equally palatable, the problem is particularly acute in densely populated salmon farm pens, where the parasites' food source are collected in a stationary place. Most pens typically keep 50,000 to 150,000 fish swimming inside a mesh perimeter. So, while the lice are a relatively harmless problem for fast and free-moving wild fish, the captive host in the pens find themselves unable to escape the louse onslaught.

The lice can easily cause injuries dire enough to make the fish unsuitable for bringing to market. As few as 11 lice can kill a smaller fish; »

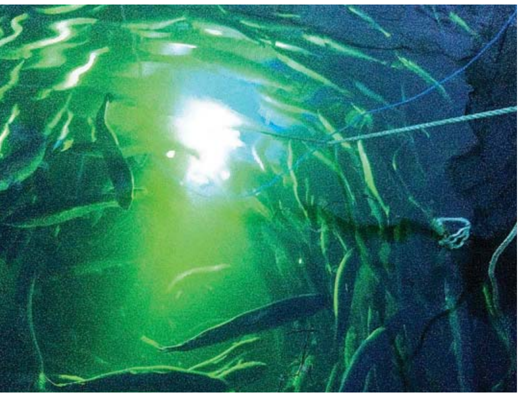

**EXTERMINATE!** The Stingray bot [above] is attached to a buoy and lowered into a fish pen [top left]. The bot's image sensors and software scan for lice attached to nearby salmon. If they spot one, the bot lets rip with a 530-nanometer (green) laser [bottom left].

STINGRAY MARINE SOLUTIONS (3)

a bigger group can leave an adult salmon wearing a grievous, bone-exposing "death crown."

The latest of many attempts to rein in sea lice involves a software-and-camera-controlled underwater laser drone. Along with colleagues at his Oslo-based company, Beck Engineering, Esben Beck, a young Norwegian designer and engineer, developed the system that's turning fish hatcheries into laser light shows. A couple of stereo cameras zero in on an individual louse attached to a fish in the pen; a thin laser beam shoots the bug, killing it but leaving the fish unharmed. This laser-beam killing machine, called the Stingray, is now being marketed by a Beck spin-off. According to the company, Stingray Marine Solutions, the device can lay waste to tens of thousands of lice each day.

Inside the Stingray's watertight aluminum package (which is about the size of a boxer's heavy punching bag) are a surgical diode laser of the sort used in dentistry, ophthalmology, and hair removal; a computer running image-matching software; small thrusters to move it through a pen; a winch for a buoy; and a 220-volt power source.

The software's lice-identifying actions are akin to face matching on a mobile-phone camera, but faster. The software triggers the laser if it registers two matching frames confirming that the cameras are pointed at a louse. The resulting 530-nanometer-wavelength beam will not hurt a highly reflective fish scale, but it will turn a small, darkish-blue louse into a floating crisp at a distance of up to 2 meters.

The Stingray node is designed to be mostly autonomous. Its custom software can consider temperature, oxygen levels, and salinity when deciding where to position itself and when to fire laser pulses.

The idea of, um, shooting fish in a barrel has captured the attention of the farmed fish industry. Some of its biggest players—Lerøy Seafood Group, Marine Harvest, and SalMar—backed the project with a little under $2 million in seed funding. The Stingray node, first made available for sale in 2014, is now zapping lice at 100 salmon farms in Norway and was introduced in Scotland at the end of 2016.

The point, says John Breivik, Stingray Marine Solutions managing director, is to cut down on the use of other treatments, such as bringing all the salmon on board a boat to be washed with hot water, flushing them through strong currents to blast off the lice, or dousing them with chemicals such as hydrogen peroxide. These treatments can be effective but expensive, and each has drawbacks—and some have quite severe critics. Breivik's claim is that two Stingray nodes have enough zapping power to decimate the lice population inside a single fish pen.

But Stingray's backers are taking a wait-and-see approach. Harald Sveier, technical manager at Bergen, Norway–based Lerøy, which together with Marine Harvest represents half of the Norwegian salmon farming industry, says his company thinks the concept's good and believes in it. With so much money at stake, though, Lerøy wants more documentation and hands-on proof.

The sea-louse issue is a nettlesome problem needing a solution. As the *New York Post* recently reported, lox prices may be headed up as a result of, among other concerns, the clouds of sea lice inflicting damage on fish. Salmon farming is controversial in some quarters, with experts and activists criticizing what they say is a self-inflicted problem. But given the world appetite, don't be surprised if more farmers decide to use underwater lasers to help shoot their way out of the situation.
—MICHAEL DUMIAK

# MEDICAL ROBOTS GO SOFT

## Three malleable machines could work safely inside the human body

**The inside of the human body is** mostly squishy (that's a technical term), and our soft innards don't always fare well when hard objects are placed inside. Not only can sharp edges damage organs and blood vessels, but the body's defense system can also surround the foreign object with scar tissue and interfere with its intended function. So researchers are working on soft robots that may be better tolerated within the body, permitting machines to make intimate contact with human tissue without jeopardizing safety.

These three experimental bots are designed for different purposes—two are implants, one is a potential surgical tool—but they all showcase a gentler kind of robotics, enabled by new materials and flexible actuators.

**1) HEART HUGGER**

Inside the chests of 41 million people around the world, failing hearts gradually become less effective at the vital task of pumping blood. Some heart failure patients get on the list for a transplant, while others receive metallic pumps called ventricular assist devices (VADs) to help their faltering organs. But VADs increase the risk of blood clots, which can form as the liquid flows over surfaces made of metal and plastic.

In search of a better pumping assistant, an international team of researchers invented a silicone sleeve that slips over the heart's exterior, thus keeping the robot from contact with flowing blood as it rhythmically squeezes the organ. The sleeve's design is inspired by the arrangement of real heart muscles, with an inner layer of material that contracts using concentric rings and an outer layer that contracts in a helical fashion. The early-stage device uses 14 pneumatic actuators (6 in the concentric layer, 8 in the helical layer), which the team can activate individually by filling them with air, allowing for experimentation with different patterns of contractions. In experiments with live pigs, the researchers demonstrated that the device can either detect and match the natural rhythm of the heart or override a faulty rhythm with a steady beat.

Soft robots have the potential to do more than supplement a failing body, says study coauthor Ellen Roche, a postdoctoral researcher at the National University of Ireland, Galway, who will move to MIT in September to start as an assistant professor of mechanical engineering. "If you can match the organ's native properties, you're going to do a better job of aiding it—and maybe you can try to rehabilitate it or help it recover function," she says. "If you just take over its function, it's not going to get better."

**2) DRUG DOSER**

When you think of Swiss clockwork, "soft" and "yielding" may not be the first adjectives that come to mind, so researchers at Columbia University get full marks for creativity. By replicating a watch mechanism called the Geneva drive in a soft hydrogel, they created a biocompatible robot that could tick along to release doses of drugs from inside the body.

Samuel Sia, a Columbia University professor of biomedical engineering, created the watch-inspired biobot with one simple hydrogel gear that's embedded with iron nanoparticles, enabling researchers to turn it with an external magnet. Each click forward brings a hollow chamber into line with an opening so a dose of liquid can flow out. Sia suggests that in cancer care such implants could enable the localized delivery of a chemotherapy drug and spare the rest of the body from the drug's toxic effects. When he tested his device in mice with bone cancer, he found that drug doses delivered by the biobot killed more tumor cells and spared more cells elsewhere in the body than a typical systemic chemo treatment. What's more, with an external controller to trigger the device's every tick, doses could be delivered only when a doctor sees fit.

The trickiest part of the design process was getting the material just right, Sia says, not too soft, but not too hard. "You don't want to lose the nice properties of hydrogels, but if your material is collaps-

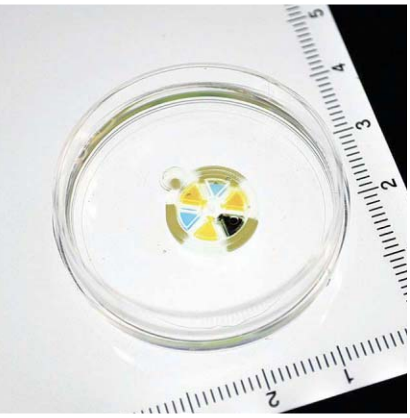

TOP: ELLEN ROCHE/HARVARD UNIVERSITY; BOTTOM: SAU YIN CHIN

NEWS

ing like Jell-O, it's hard to make robots out of it," he says. "It has to be stiff enough to work like a tiny implantable machine."

**3) GENTLE GRABBER**

Today's surgeons use various gripping tools to wrangle with your viscera; they may use one tool to nudge aside fat tissue, for example, while using another to cut away tumors on the kidney. While surgeons are scrupulous in their attempts to avoid unnecessary harm to the tender flesh, they might have an easier time with Xuanhe Zhao's soft tools.

Zhao, an MIT associate professor in the mechanical as well as the civil and environmental engineering departments, devised a series of robotic gadgets made of hydrogel, each fabricated in the form of interlocking cubes with hollow interiors. To activate a device, a syringe pump injects water into the robot in certain combinations to make it curl up or stretch out, producing fast and forceful movements. One grabber bot [below], which resembles a five-fingered hand, demonstrated its deftness by seizing hold of a swimming goldfish—and then releasing it unharmed. Zhao's team is now collaborating with medical groups on hydrogel "hands" that could grasp organs in robotic surgeries.

Another application could be a soft robot that wraps around the intestines and contracts rhythmically, mimicking the wavelike peristalsis process that moves food through the digestive tract. For tomorrow's soft robots, it seems that everything inside the body is up for grabs. —ELIZA STRICKLAND

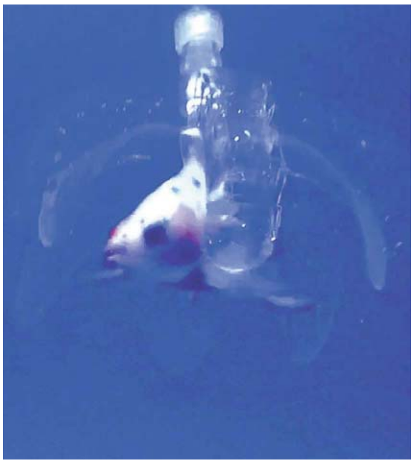

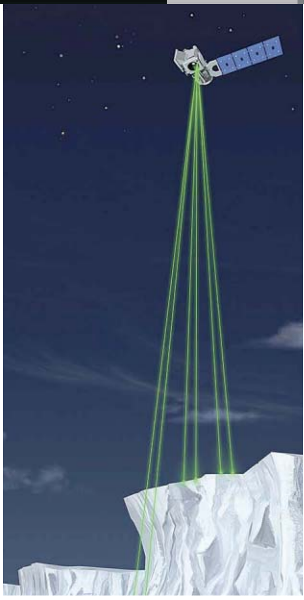

# TWILIGHT FOR THE GOLDEN AGE OF EARTH OBSERVATION?

**NASA's new geosensing satellites may be on the chopping block. The timing could hardly be worse**

**When leaders of the** congressional committees that approve NASA's missions and budgets put forth their priorities in February, only space science and deep-space exploration made the cut. Conspicuously absent was Earth science—a US $2 billion function within NASA that tracks our rapidly changing environment.

Add in White House skepticism of climate science, and what experts call today's "golden age" of monitoring Earth via satellite faces some serious challenges.

That age began in 2009, when President Barack Obama responded to a U.S. National Research Council warning that budget cuts had left the United States' Earth-observing system "at risk of collapse." NASA, the lead federal agency for satellite development, saw its Earth-science budget rise 56 percent between 2008 and 2016, and it placed eight new Earth-observing satellites with state-of-the-art sensors in orbit during that period.

The data they deliver inform a widening range of activities—crop planning and management, wildfire risk assessment, extreme air pollution warnings, and more. NASA delivered 1.42 billion data products in 2015—174 times as many as it delivered in 2000.

More missions are in the pipeline, such as NASA's second Ice, Cloud, and land Elevation Satellite (ICESat-2), whose primary objectives are tracking melting polar ice sheets and glaciers and quantifying the carbon locked up in the globe's forests.

ICESat-2, however, exemplifies both the present strength of the U.S. Earth-observation program and a less visible weakness. To understand why, you need a sense of the ambitious nature of ICESat-2's mission.

Rather than rerunning the first ICESat mission, which ended in 2009, NASA redesigned the laser altimeter to boost its impact. One laser beam firing sporadically became six beams firing 365 days a year; higher-precision digital photon counting replaced analog detection of beams bouncing back from Earth.

ICESat-2 should enable measurement of annual elevation changes in ice sheets at ±4-millimeter accuracy (and better for other targets), and at 17 times the spatial resolution of its predecessor, according to Thorsten Markus, chief of cryospheric sciences at NASA's Goddard Space Flight Center, in Maryland. Such data, he says, will elucidate some basic physical processes that elude climate models, and thus improve their predictions.

TOP: GODDARD SPACE FLIGHT CENTER/NASA; BOTTOM: HYUNWOO YUK/MIT

**ORBITAL ICE SPY:** IceSat-2 would fire six lasers earthward to measure changes to the polar ice, glaciers, and forests.

But pushing for the best has not come cheap. Instead of $300 million for an ICESat rerun, NASA's estimate for ICESat-2's development started at $559 million and has grown to $764 million. Including operations for up to seven years, the mission could cost nearly $1.1 billion, according to a NASA inspector general's report. Launch dates, meanwhile, have slipped from 2015 to 2018.

Delays and cost creep in ICESat-2 and other missions, as well as several failed launches, put a significant tarnish on NASA's Earth-observation boom. Extending existing missions to avoid gaps in data creates risk, according to the inspector general: "More than half the Agency's 16 operating missions have surpassed their designed lifespan and are increasingly prone to failures that could result in critical data loss."

Similar risks confront the National Oceanic and Atmospheric Administration, a key partner in climate and weather observation, according to a February report by Congress's watchdog agency, the Government Accountability Office. NOAA's polar weather data currently come from a dying NASA demonstration mission. Failure prior to the launch of the agencies' long-awaited Joint Polar Satellite System would degrade weather forecasts, "exposing the nation to a 15 percent chance of missing an extreme weather event forecast," writes the GAO.

If the golden age of Earth observation harbored weak spots before the 2016 election, experts say the new administration introduces new risks. One is the $54 billion in belt-tightening proposed for federal agencies by President Donald Trump.

Another is potential interference with climate science. In February, Lamar Smith, chairman of the House Committee on Science, Space, and Technology, repeated his longstanding call for "rebalancing" of NASA's portfolio. A former chairman, Robert Walker, now a lobbyist for space-related industries, built a similar plank into the space platform that he drafted for Trump's campaign. Both men question human-induced climate change—a view held by many Republicans in Congress and Trump appointees.

Walker says expanded Earth observation under Obama came at the expense of other science programs, particularly deep-space robotic missions. He also alleges that NASA science was "tainted" by a political agenda, focusing on impacts from burning fossil fuels and neglecting natural climate influences such as volcanic eruptions. "There's an extremely complex system that involves a lot more than $CO_2$," he says.

The Intergovernmental Panel on Climate Change's 2014 assessment, however, expressed "very high confidence" that volcanic eruptions caused only "a small fraction" of the warming observed since the Industrial Revolution. And it cites "robust evidence" from satellite data showing that natural factors have had "near-zero" effect since 1980.

The notion that human activities alter climate is not a political invention but a scientific judgment based on the gigabytes of data beamed daily to Earth from a gilded era's orbiting sensors. "It's not a belief," says NASA's Markus. "That's what the data show."
—PETER FAIRLEY

# SPECK-SIZE COMPUTERS: NOW WITH DEEP LEARNING

## Michigan labs' "micromotes" aim to make the IoT smarter

**Computer scientist David Blaauw** pulls a small plastic box from his bag. He carefully uses his fingernail to pick up the tiny black speck inside and place it on the hotel café table. At 1 cubic millimeter, this is one of a line of the world's smallest computers. I had to be careful not to cough or sneeze, lest it blow away and be swept into the trash.

Blaauw and his colleague Dennis Sylvester, both IEEE Fellows and computer science professors at the University of Michigan, were in San Francisco in February to present 10 papers related to these "micromote" computers at the IEEE International Solid-State Circuits Conference (ISSCC). They've been presenting different variations on the tiny devices for a few years now.

The broad goal of the Michigan Micro Mote ($M^3$) initiative is to make smarter, smaller sensors for medical devices and the Internet of Things—sensors that can do more with less energy. Many of the microphones, cameras, and other sensors that make up the eyes and ears of smart devices are always on alert, and frequently they beam personal data into the cloud because they can't analyze it themselves. Some forecasts have predicted that by 2035, there will be 1 trillion such devices. "If you've got a trillion devices producing readings constantly, we're going to drown in data,"

NEWS



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

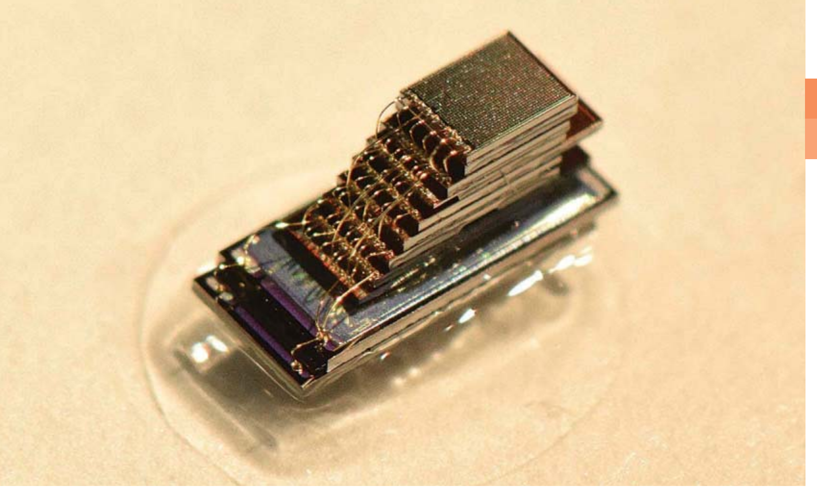

**MIGHTY AND MINI:** One of several varieties of the University of Michigan millimeter-scale computers incorporates 1 megabyte of flash memory.

UNIVERSITY OF MICHIGAN AND TSMC

says Blaauw. By developing tiny, energy-efficient computing sensors that can do analysis on board, Blaauw and Sylvester hope to make these devices more secure, while also saving energy and bandwidth.

In San Francisco, they described micromote designs that use only a few nanowatts of power to perform tasks such as distinguishing the sound of a passing car and measuring temperature and light levels. They showed off a compact radio that can send data from the small computers to receivers 20 meters away–a considerable boost compared with the 50-centimeter range they reported last year. They also described their work with TSMC (Taiwan Semiconductor Manufacturing Co.) on embedding flash memory into the devices and a project to bring on board dedicated, low-power hardware for running artificial intelligence algorithms called deep neural networks.

Blaauw and Sylvester say they take a holistic approach to adding these new features without ramping up power consumption. "There's no one answer" to how they and their engineers do it, says Sylvester. If anything, it's "smart circuit design," Blaauw adds. (They pass ideas back and forth rapidly, not finishing each other's sentences but something close to it.)

The memory research is a good example of how the right trade-offs can improve performance, says Sylvester. Previous versions of the micromotes used 8 kilobytes of static RAM, which makes for a pretty low-performance computer. To record video and sound, the tiny computers need more memory. So the group worked with TSMC to bring flash memory on board. Now they can make tiny computers with 1 megabyte of storage.

Flash can store more data in a smaller footprint than SRAM, but it takes a big burst of power to write to the memory. With TSMC, the group designed a new memory array that uses a more efficient charge pump for the writing process. The memory arrays wind up being a bit less dense than TSMC's commercial products, but still much better than SRAM. "We were able to get huge gains with small trade-offs," says Sylvester.

Another micromote they presented at ISSCC incorporates a deep-learning processor that can operate a neural network while using just 288 microwatts. Neural networks are artificial intelligence algorithms that perform well at such tasks as face and voice recognition. They typically demand both large memory banks and intense processing power, and so they're usually run on banks of servers often powered by advanced GPUs. Some researchers have been trying to lessen the size and power demands of deep-learning AI with dedicated hardware that's specially designed to run these algorithms. But even those processors still use over 50 milliwatts of power—far too much for a micromote. The Michigan group brought down the power requirements by redesigning the chip architecture. For example, they situated four processing elements within the memory (in this case, SRAM) to minimize data movement.

The idea is to bring neural networks to the Internet of Things. "A lot of motion detection cameras take pictures of branches moving in the wind—that's not very helpful," says Blaauw. Security cameras and other connected devices are not smart enough to tell the difference between a burglar and a tree, so they waste energy sending uninteresting footage to the cloud for analysis. Onboard deep-learning processors could make better decisions, but only if they don't use too much power. The Michigan group imagines that deep-learning processors could be integrated into many other Internet-connected things besides security systems. For example, an HVAC system could decide to turn the air-conditioning down if it "sees" several people putting on their coats.

After demonstrating many variations on these micromotes in an academic setting, the Michigan group hopes they will be ready for market in a few years. Blaauw and Sylvester say their startup company, CubeWorks, is currently prototyping devices and researching markets. The company was quietly incorporated in late 2013. Last October, Intel Capital announced it had invested an undisclosed amount in the tiny computer company. —KATHERINE BOURZAC



PHOTO COMPOSITE BY Jin Liwang/Xinhua/Redux

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

# A TRIAL CALLED QUESS

**THE WORLD'S FIRST** satellite built specifically with hackers in mind was launched from China in August. The Chinese and Austrian researchers behind the Quantum Experiments at Space Scale, or Quess, hope their work will help improve digital encryption by yielding hack-proof key distribution. Part of the Quess experiment will focus on mastering a phenomenon called quantum entanglement for encryption.

A pair of entangled photons share a single quantum state, no matter how much distance is put between the two particles. So, a ground station like the one pictured here can use pairs of photons, entangled on the satellite, to produce a shared random secret key that ensures a secure communications channel with another base station thousands of kilometers away. By their very nature, the particles can't be observed by a third party without the eavesdropper introducing anomalies that reveal the snooping or render the key completely useless.

THE BIG PICTURE NEWS

# RESOURCES

1958: THE YEAR JOHN C. KOSS INTRODUCES THE WORLD TO STEREO HEADPHONES

RANDI KLETT

## EARPHONE SHOW-DOWN, PART ONE

### THE ETYMOTIC ER4, AUDIO-TECHNICA ATH-IM70, AND LKER i1

RESOURCES_REVIEW

**Even for dedicated audiophiles, the musical experience nowadays is frequently** delivered by earphones. Happily, there's been an explosion in earphone ingenuity, with many models offering stunning sound reproduction and remarkable value. If you're still using the earbuds that came with your smartphone, it's time to trade up. • I recently auditioned a sample of this new earphone wave. Here I'll consider a couple of new models from Etymotic Research, a brand I've admired for years, along with a Chinese product that stunned me with its performance/price ratio. In a follow-up article, I'll consider three other surprising, audiophile-grade Chinese earphones. • Etymotic Research got its start in 1983 producing earphones for industrial, professional, medical, and scientific uses. The company had a big breakthrough in the consumer market in 1991, when it offered the ER-4S, which it claims were the first noise-isolating earphones designed to be inserted into the ear canal. The ER4 line remains Etymotic's flagship, so I was eager to check out its latest update. • Etymotic prides itself on making what it calls "reference quality" earphones—that is, earphones with a frequency response that's as flat as possible across the audio spectrum. This is in sharp contrast to almost all other earphones produced today, which either emphasize bass response (the Beats line is notorious in this regard) or have a V-shaped response that emphasizes bass and treble at the expense of midrange frequencies. Such an intentionally distorted response can make pop music, particularly hard-driving rock, hip-hop, and rap, seem more gripping and involving. • Etymotic achieves a relatively flat response in large part through the use of a ▸


**EVERYONE LIKES EXTRAS:** Etymotic ER4 earphones come with a generous assortment of accessories, including spare eartips in various sizes, a plug adapter, and a large, sturdy travel case.

ETYMOTIC RESEARCH

single balanced-armature driver in each earphone. In a balanced armature, the sound signal's current flows through a coil surrounding an armature, producing a magnetic field to vibrate the armature. The armature is connected to a diaphragm that, as it moves, changes the volume of air in an adjacent space. This changing air volume produces pressure changes that push air out of a nozzle, resulting in a sound wave. Dynamic-driver earphones, on the other hand, are similar to conventional speakers in that they produce a sound wave by vibrating an element, typically cone shaped, which directly displaces air. The advantage of a balanced-armature driver here is its efficiency over a range of frequencies.

Both balanced-armature and dynamic-driver earphones can have just one sound-producing element (driver) per ear, in which case they are called mono drivers. Or they can use multiple elements, each covering a different range of the audio spectrum. For example, the JH Audio JH16V2 Pro earphones, at US $1,499, have an arguably ridiculous 10 balanced-armature drivers *per ear.*

There are two variations of the Etymotic ER4: a "Studio Reference" version and an "Extended Response" (both $349). Each has a frequency range from 20 hertz to 16 kilohertz, and the frequency response charts provided for the two different models seem identical. So, perhaps not surprisingly, I could detect no difference in performance between these two models.

I compared the Etymotic with three other earphones: the Audio-Technica ATH-IM70, which uses dual-dynamic drivers in each earphone; the LKER i1, an inexpensive, Chinese-made earphone that's also a dual dynamic; and my own Etymotic ER-4 microPro earphones, which I bought circa 2003.

The first thing to know about the ER4s is that in order to hear bass strongly with them, it's necessary to get a good fit in your ear canal. Etymotic's models are meant to go deeper into your ear canal than most other brands, and they use a triple-flange eartip to seal out external noise. The ER4s come with eartips in several different sizes; take your time to figure out which ones are right for you, or the bass will seem anemic.

The ER4s really shine with small-ensemble music, such as classical chamber music, jazz, acoustic rock, vocal, and folk. I picked a few of my favorite recordings—including Lindi Ortega's "Ashes," the Beta Band's "Dry the Rain," and Chris Isaak's version of "Solitary Man"—and listened to them repeatedly with all five sets of earphones. I played FLAC (Free Lossless Audio Codec) music files stored on my Samsung smartphone through a DragonFly Red DAC (digital-to-analog converter) connected to the earphones.

With the ER4s, I was impressed with the detail in the guitars and other instruments; the firmness of the bass; and the clarity and realism of the vocals. By contrast, with the dual-dynamic driver units, the bass sounded boomier and the vocals more suppressed, and even slightly muddy sometimes, in relation to the bass.

But all that said, the current ER4s didn't sound significantly better to me than my 13-year-old ER-4 microPro. That's not really a bad thing; for 25 years Etymotic has steadfastly refused to join the boomy-bass bandwagon.

The big surprise came when I turned to the LKER i1 and compared it to the dual-driver Audio-Technica ATH-IM70. The ATH-IM70 has been my go-to earphone for several years now. It offers remarkable sound quality for its price of about $80. It emphasizes bass slightly, but not nearly as much as many others. However, after hours of listening with the LKER i1, I can report that it offers sound quality astonishingly close to that of the ATH-IM70—and for a price of about $25.

In a final round of testing, I listened to The Killers' song "Somebody Told Me" with all five earphones. It was the hardest-rocking song I could think of. On this one, the ATH-IM70s and the LKER i1s really came into their own.

The verdict: If your musical tastes run heavily to power chords, save yourself about $270 and get the ATH-IM70s. Or, even better, save yourself $325 and get the LKER i1s. Either one is going to be a lot better than the earbuds that came with your smartphone. But if you listen to a lot of small-ensemble music, take a listen to the Etymotic ER4s. You'll reacquaint yourself with midrange frequencies, and you just might find that you like them.

—GLENN ZORPETTE

RESOURCES_HANDS ON

# HOLD BACK THE TIDE

## BUILD AN IoT-FREE WIRELESS WATER CUTOFF TO PROTECT AGAINST LEAKS

**WHEN APPLIANCES ATTACK:** A single stuck valve led to a leak that slowly but relentlessly damaged my home so thoroughly it had to be gutted and rebuilt.

DAVID SCHNEIDER

**You've seen it in the movies—** the protagonist inexplicably blinks awake from deep sleep because some silent but menacing force threatens. Something like that happened to me early one Sunday morning not long ago. And as soon as I stepped out of bed, I knew things were going to be bad, because the floor of my second-story bedroom was covered with water.

It didn't take long to identify the source as I splashed down the hallway: the washing machine. An inlet valve had gotten stuck in the open position after we put in a load late the previous night. Eventually, the door of this front loader gave way. For several hours, water sluiced from our second-floor laundry closet onto just about every horizontal surface in the house—and quite a few of the vertical ones, too.

Fast forward six months. To fix the water damage, the interior of my house had to be taken apart and put back together. I decided that I never wanted to go through that ordeal again. But what would be the best way to protect the house from plumbing disasters? And could I do it without opening my house to another sort of threat, one from Internet of Things hackers?

## WATER WATCH

Water sensors are scattered throughout my house at strategic positions [top]. They send radio signals to a central station [middle], which is monitored by an Arduino. If the Arduino detects an alert from a sensor, it triggers a valve attached to the main water line [bottom].

DAVID SCHNEIDER (3)

One measure I took was a simple one. Scorning the cheap plastic catch pan under the washer—which did nothing to help on that fateful night—I put in a steel tray that lines the entire floor of the laundry closet.

My next bit of insurance was to install a valve with a mechanical timer on the feed line to the washing machine. This closes automatically 2 hours after being opened. So most of the time the water valve feeding the washer is closed.

But what if all that failed—or if there was a water leak from some other source? This called for active electronic measures.

Ironically, I had built a device that monitored water flow in real time not long before. But that system wouldn't have detected anything anomalous because it was designed to look for anomalously large water flows. And the flow that flooded my house wasn't abnormally high—it was just in the wrong place. A simple audible water alarm would have done the trick, and I did buy a few of those, but they don't help if nobody is home to hear them. The solution, obviously, was to install a system that would automatically block off the main supply line should water ever overflow onto the floor.

A little scouting online revealed no shortage of such systems available for purchase. But they had two drawbacks. First, many were part of network-based home-automation systems. I really don't like the thought of hackers in Eastern Europe messing with my ability to take a shower. Also, these systems were rather pricey, costing much more than what the relatively simple hardware seemed worth to me. So I set about putting together my own flood-prevention system.

It wasn't difficult because, well...I cheated. Instead of building the whole kit and caboodle from scratch, I decided to modify an inexpensive wireless alarm (US $29). This accepts radio signals from various kinds of sensors that transmit on 433 megahertz. For my house, I deployed four water sensors ($20 each) at strategic locations around my house.

The only other component I needed to purchase was a motorized ball valve ($67), which I had a plumber install when he was visiting to fix the water heater. This valve is actually quite nifty. It has a manual override knob and stores energy in a capacitor, allowing it to operate when power is removed. When you apply power, it shifts from open to closed, and when you remove power it goes back to open. Internal limit switches ensure that the motor is powered only when changing positions.

So all I needed to do was to make the wireless alarm trigger the motorized valve. I combed through my boxes of leftovers from various other projects, where I found an Arduino Duemilanove looking for a new home, a 12-volt wall wart power adapter, an LM317 voltage regulator, and a 5-volt relay.

The 12-V wall wart powers the wireless alarm and the LM317, which I configured to output about 9 V, which then supplies the Arduino. I used the LM317 to drop the supply voltage, because while it's acceptable to feed 12 V directly to an Arduino, the onboard regulator gets uncomfortably hot. The wall wart also powers the ball valve when the 5-V relay commands it to close.

The wireless alarm has four red LEDs, each of which light up when a signal is detected. I connected one side of each of these LEDs to four of the Arduino's analog-input pins. The voltage on the anodes drops to less than 3 V when an alarm is triggered. So the Arduino simply polls each of the four inputs constantly to see if any of them is reading under 3 V. If so, the Arduino energizes the 5-V relay.

Skeptics will point out that my fail-safe will itself fail if my house loses power. Yes, sure, but I'm not too concerned about that. The bigger worry I had when I put this system together involved where to place the sensors. If I put them too close to sources of water, say, a tub or kitchen sink, they might trigger from a simple splash. So I tucked them in places where only a significant spill would appear. But I suppose a true water leak could develop with the water flowing in such a way that it missed a sensor. Even so, I do sleep better at night knowing that an Arduino is watching over me.

**—DAVID SCHNEIDER**

# RETHINKING CS101
## RUSS MILLER WANTS STUDENTS TO EMBRACE PARALLEL ARCHITECTURES FROM DAY ONE

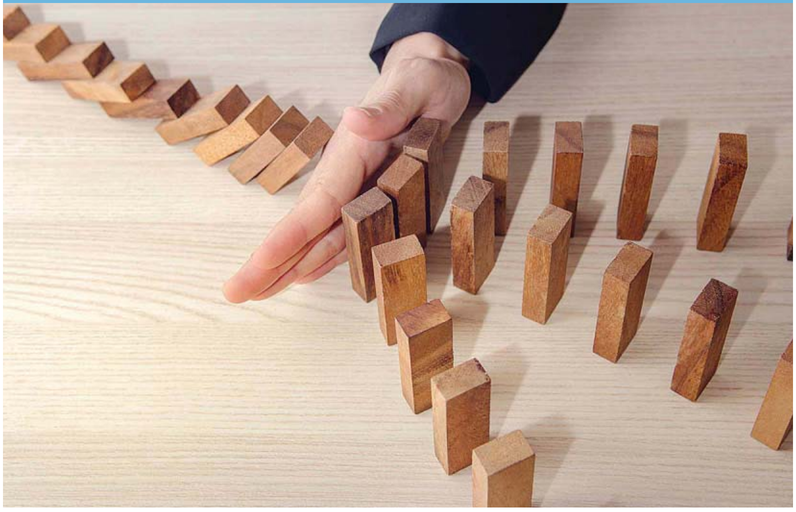

Over 350 computer science students take Russ Miller's Discrete Structures course every fall semester. About 90 percent are freshmen. By week five, they are breaking problems into small chunks and learning ways to solve each chunk at the same time—in parallel.

Today, multicore processors power our laptops and cellphones. Distributed cloud servers or supercomputer clusters process large data sets to improve Facebook news feeds or predict the weather. To take full advantage of these systems, you need parallel algorithms. "It's a parallel world," says Miller, a computer scientist at the State University of New York at Buffalo. "Why is no one teaching a course in parallel algorithms to freshmen?"

Currently, most introductory computer science courses start with sequential programming, in which the computer performs just one instruction at a time. Universities that integrate parallel thinking into their undergraduate curricula tend to offer only an upper-level elective. Others that do spread parallelism throughout the curriculum start no sooner than the second or third course. University needs can vary, but Miller believes that teaching parallel thinking "becomes harder the longer you wait," whereas it can become "second nature" if you do it early enough. So in 2013, Miller changed the State University at Buffalo discrete mathematics course to teach parallel algorithms. It has no prerequisites.

He says that most required discrete mathematics courses teach some material students won't need until their junior or senior year. "That's just the most ridiculous waste of everybody's time," he says. "They just check out." After spending three or four weeks covering the basics of standard logical thinking and divide-and-conquer strategies, Miller dives into parallelism. He gives context by first explaining 1960s- to 2000s-era parallel computing architectures. The rest of the semester he dives into general, hardware-agnostic parallel algorithms for tasks such as searching and sorting. Students learn about such topics as image-segmentation applications.

After each lesson about a new algorithm, students mathematically analyze its theoretical performance on old parallel hardware architectures. They learn some ways to change the algorithm to work on modern, real-world architecture, such as a cloud or grid.

Because it's not about programming per se, Miller can skip thorny implementation details, such as syntax or debugging methods, and have plenty of time to teach students a parallel-first mind set.

"Something like that could work," says Mehran Sahami, a computer scientist at Stanford who cochairs the Association for Computing Machinery steering committee on computing curricula. The ACM and the IEEE jointly introduce new guidelines roughly every 10 years: The latest, issued in 2013, recommend integrating parallel education throughout the curriculum.

But some educators find their ability to deeply embrace parallelism is constrained by other demands. Some instructors of introductory computer science courses, such as Steven Bogaerts at DePauw University in Greencastle, Ind., spend about a week on threads (subsections of a program that can run in parallel) and how to stop them from accessing the same resources at the same time. But to Bogaerts, "it's just already a very full course," so it's hard to do much more.

And some point out that there's more to enhancing code performance than just parallelism. "Of all the ways of getting performance, parallelism is among the hardest," says Charles E. Leiserson, a computer scientist at MIT who teaches a junior- and senior-level course on performance engineering. Leiserson says that parallelizing algorithms doesn't guarantee they'll run faster than sequential algorithms when you implement them on real hardware. He says other factors are important to manage and understand, such as the memory hierarchy or compiler.

Miller, who does not cover the memory hierarchy outside of small examples, is optimistic about his students' futures. He says they have three-and-a-half years to learn how to write efficient code. In his course, they're learning to solve problems at a high level. "I want them to think that the world's open to them," he says. —ANDREW SILVER

ISTOCKPHOTO

# PLANET OF THE COWS

**FOR YEARS I HAVE TRIED TO IMAGINE** how Earth would appear to a comprehensive and discerning probe dispatched by wonderfully sapient extraterrestrials. Of course, the probe would immediately conclude, after counting all organisms, that most individuals are either microscopic (bacteria, archaea, protists, fungi, algae) or very small (insects) but also that their aggregate weight dominates the planetary biomass. ● That would not be really surprising. What these tiny creatures lack in size they more than make up in numbers. Microbes occupy every conceivable niche of the biosphere, including many extreme environments. Bacteria account for about 90 percent of the human body's living cells and as much as 3 percent of its total weight. What would be surprising, however, is the picture the probe would paint of the macroscopic forms of animal life, which is dominated by just two vertebrates–cattle (*Bos taurus*) and humans (*Homo sapiens*), in that order. ● Unlike the extraterrestrial scientists, we do not get an instant readout. Even so, we can quantify cattle zoomass and human biomass (anthropomass) with a fair degree of accuracy. The numbers of large, domesticated ruminants in all high-income countries are known, and they can be reasonably estimated for all low-income and even pastoral societies. The Food and Agricultural Organization of the United Nations puts the global cattle count at about 1.5 billion head in 2015. ● To convert these numbers into living ruminant zoomass requires adjustments for age and sex distribution. Large bulls weigh more than 1,000 kilograms (about 2,200 pounds); American beef cows are slaughtered when they reach nearly 600 kg, but Brazilian cattle go to market at less than 230 kg; and India's famous Gir milk breed weighs less than 350 kg when mature. A good approximation is to assume an average sex- and age-weighted body mass of 400 kg; that implies a total live cattle zoomass of some 600 million metric tons.

Similarly, when calculating the total mass of humanity it is necessary to consider the age and body weights of populations. Low-income countries have much higher shares of children than affluent nations (in 2015 about 40 percent in Africa compared to about 15 percent in Europe). At the same time, the rates of overweight and obese people range from the negligible (in Africa) to 70 percent of the adult population (United States). That is why I am using specific means for different continents derived from available population age and sex structures and from anthropometric studies and growth curves for representative countries. That complex adjustment produces a weighted mean of about 50 kg per capita, which, given a total of 7.4 billion people, implies a global anthropomass of about 370 million metric tons in 2015.

This means that the cattle zoomass is now at least 60 percent larger than the anthropomass and that the live weight of the two species together is about one billion metric tons. Even the largest wild mammals add up to only a small fraction of those masses: The 350,000 savanna elephants in Africa, with an average body weight of 2,800 kg, have an aggregate zoomass of less than one million metric tons, less than 0.2 percent of cattle zoomass.

The aggregate mass of cattle and humans is crushingly larger than the total mass of all wild vertebrates, and it clearly leaves too little space for the multitude of other species. Cows and men occupy much of the available land, consume much of its photosynthetic product, and generate an increasing amount of greenhouse gas.

No wonder we are in the midst of mass-scale species extinction, with no readily acceptable and effective relief in sight. By 2050 there will be 9 billion people and, most likely, 2 billion cattle, together augmenting their already crushing dominance of Earth. ■

ILLUSTRATION BY Stuart Bradford


# ESCAPING THE POWERPOINT PRISON

We have met the Devil of Information Overload and his impish underlings, the computer virus, the busy signal, the dead link, and the PowerPoint presentation. *—James Gleick,* The Information

**PITY POOR POWERPOINT.** Has there ever been a piece of software that elicited such negative feelings? Meeting goers have long been bored senseless by slides that were transition-heavy and content-light, living in fear of both **ant fonts** (impossible-to-read small type used by someone trying to cram as much text on a slide as possible) and the **triple delivery** (having the same text displayed onscreen, on a handout, and spoken aloud). Business leaders have raged against the productivity suck of employees spending hours making trivial tweaks of their presentations. Per Walter Isaacson in his book *Steve Jobs* (Simon & Schuster, 2011), one of the first things Jobs did on his return to Apple was prohibit PowerPoints. "People who know what they're talking about don't need PowerPoint," Jobs said. People who did *not* know what they were talking about were said to have knowledge that is only **PowerPoint deep**. The U.S. military became top-heavy with **PowerPoint Rangers**: military officers who excel at creating presentations.

PowerPoint carried on, of course, but a change is in the wind. The past few years have seen a welcome flowering of innovative ways to present information.

It began early in the new millennium with the **PechaKucha** (Japanese for *chit-chat* or *chatter*), a presentation consisting of 20 slides shown for 20 seconds each, without pause. (Not surprisingly, this format has led to some people calling this style of presentation the **PechaKucha 20×20**.) That's a mere 6 minutes and 40 seconds, if you're keeping score. A similar idea is **the lightning talk**: a short and strictly timed presentation, usually a maximum of 5 minutes long and sometimes with a maximum number of slides as well. An example is the **Ignite talk**, which is restricted to 5 minutes and 20 slides, each of which advances automatically after just 15 seconds. Another 5-minute presentation format is the **data blitz**, which is most often used to present academic research or scientific findings.

And some are blazing completely new presentation trails. An example is the **Takahashi method**, created by Japanese programmer Masayoshi Takahashi, where each slide consists of a single word in an extremely large font (a kind of *anti* ant font, if you will). The **Lessig method**, named after its principal proponent, Harvard law professor Lawrence Lessig, is similar in that it places on each slide nothing but a word, phrase, or short quotation.

Another novel strategy is **speed geeking**. This is an event in which the presenters are stationed throughout a large room or hall. The audience is divided into groups, and each group is assigned to a different presenter. After a short talk, the groups rotate to the next speaker, and this continues until each group has heard each presentation.

What's a speaker to do if she has an hour's worth of material to present in a world conditioned by the 20-minute TED talk? Easy. Break up her material into **talklets**, which are relatively short, stand-alone presentations. She could also construct **à la carte content**, which involves presenting only a subset of the presentation based on audience interests or questions.

But is all that nothing but **attention deficit theater**, as some critics suggest? Are we so conditioned by 140-character tweets that we can't listen to someone speak for more than a few minutes? That might be true, but I prefer to see these new presentation formats as a positive trend. Albert Einstein once said that if you can't explain something to a 6-year-old, then you don't understand it yourself. In that spirit, perhaps we can say that if a modern-day presenter can't explain a concept in the space of, say, a PechaKucha, then he needs to refine his ideas before foisting them on an audience. ■

ILLUSTRATION BY Edmon de Haro

**BOLT UPRIGHT: A Chevy Bolt takes shape on an assembly line, where economies of scale help make this super-long-range EV superaffordable.**

# 2017 Top Ten Tech Cars

The all-electric Chevy Bolt, with a range of 383 kilometers, tops our list

By Lawrence Ulrich

**You can show people an Italian supersedan** that set a lap record on the famed Nürburgring circuit. You can point them to a sleek, affordable electric runabout that can go from Paris to Luxembourg on a single charge. But they'll probably just want to know one thing: When are cars going to drive themselves?

The specter of the fully autonomous, virtually self-aware car has captured imaginations and dominated headlines like no other automotive story in decades. And for at least one very good reason: Such vehicles promise to dramatically reduce or even eliminate deaths and injuries in automobiles.

Yet as we spotlight the year's 10 most technically innovative cars, there's not a fully autonomous car available in any showroom in the world. My own view is that you won't see one for 5 to 10 years–especially if you're referring to (trumpets, please) Level 4 autonomy, the point at which human drivers can snooze or read a novel while the robotic chauffeur does all the work.

To be sure, advances are coming at a startling speed. Companies as established as Ford and as new as Google's spin-off, called Waymo, are racking up countless autonomous kilometers in test vehicles, many of them on public roads. Seeing that, the hype-driven tech media continues to proclaim that we're oh-so-close to showroom cars that can go it alone entirely.

It's wishful thinking. Even Tesla's vaunted (and misleadingly named) Autopilot doesn't come close to it. One unfortunate Tesla driver learned that lesson the hard way, milliseconds before perishing when his Autopilot failed to detect and brake for a semi truck directly in his path.

And in 2017, every major global automaker is playing it safe and emphatically holding the line at Level 2 autonomy, which demands that a driver be ready to retake control at a moment's notice. They're focused on ever-improved lane keeping, automated braking, and other systems that serve as an electronic copilot–one that never outranks the driver. Humans must remain vigilant and (mostly) keep their hands on the wheel.

Whether that makes you cheer or jeer, you'll find plenty of other technology triumphs that are in automobiles available now. Here we've assembled the best of these, including breakthroughs in performance, materials, electric power trains, and energy efficiency–and a couple of advances in *semi*-autonomous systems as well. Dig in, enjoy the bounty, and keep your eyes on the road.

GENERAL MOTORS CO.

TOP 10 TECH CARS 2017

# Chevrolet Bolt

## This EV is fast, long range, and affordable

GENERAL MOTORS CO.

THE CHEVROLET BOLT hums off of Mulholland Highway and swings east toward Malibu, a view of frothy Pacific surf filling its steeply raked windshield. After 3 hours of driving, from suburban L.A. snarls to wilderness canyons, the all-electric Bolt is well on its way toward topping its official, EPA-rated range of 383 kilometers (238 miles).

It's my first drive in the Bolt, whose unassuming shape masks seriously ambitious thinking and technology. And I don't mind saying I'm having a blast. But there's more testing to do. Now, ahead of me, a Toyota full of youngsters that's adorned with a Pepperdine University bumper sticker is ambling down the Pacific Coast Highway, part of the famous Highway One, which cuts a spectacular swath from the redwoods of Northern California nearly all the way to San Diego.

Easing up behind the Toyota, its occupants completely unaware of my stealthy approach, I nudge the Bolt left, hammer the gas–er, accelerator–and feel the signature, instant-on shove of 361 newton meters (266 foot-pounds) of electric torque. The front-mounted electric motor, its magnets made from dysprosium (it's a rare earth element–you can look it up), cranks out 149 kilowatts (200 horsepower).

And that's not the only wizardry under the hood. Its motor maxes out at 8,810 rpm, almost twice the revs of the Chevy Spark EV's motor. As a result, the Bolt's single-speed transmission can adopt a much shorter gear ratio to boost efficiency.

The upshot? The Bolt is no Ludicrous-mode Tesla, but it's still seriously quick, able to scoot from 0 to 100 kilometers per hour (62 miles per hour) in about 6.3 seconds. The Bolt flies past the Toyota, and the only clue to the accelerative violence taking place is a mellow electric hum.

The Bolt has already won a more-important race: With all eyes on Tesla's long-delayed Model 3, it's Chevrolet that has delivered the world's first affordable long-range electric car.

It remains to be seen how many people are yearning to go electric but have been deterred by, for example, the US $120,000 price of the Tesla Model S P90D. At $29,995, after figuring in the generous $7,500 U.S. federal tax credit, the Bolt is poised to usher in a new era of pure electrics. The Chevy likely renders short-range EVs obsolete, along with the "range anxiety" that turned off buyers who shied away from a car that couldn't quite get to Grandma's house and back.

Chevy made waves when it disclosed it had achieved an industry-low cost of $145 per kilowatt-hour for the Bolt's battery, even angering its LG Chem supplier, which feared that other customers would demand a similar deal. Bill Wallace, GM's chief battery engineer, said it was less about arm-twisting its supplier and more about applying low-cost methodologies—from raw materials to manufacturing—while ensuring that the Bolt battery wouldn't flame out like a Galaxy Note 7.

To carve out enough range for about 4 hours of freeway travel—or perhaps a week's worth of around-town operation—the Bolt packs a remarkably compact and power-dense battery that spans its entire floor. Its 60-kWh capacity matches the battery in the base-model Tesla Model S.

Peaking at 160 kW, the Bolt's 288 lithium-ion cells have a "nickel-rich" chemistry, which boosts the ratio of nickel beyond the typical 1:1:1 spread with manganese and cobalt of this type of battery. The new nickel-boosting technology not only stores more energy but can also run hotter, by about 5 °C, with no loss in performance and durability.

Rather than the air cooling employed by EVs like the Nissan Leaf, the Bolt adopts the more costly and sophisticated liquid cooling of its cousin, the plug-in hybrid Chevrolet Volt. That system should make for longer battery life, especially for owners in hotter climes.

That underfloor battery does double duty as a structural unit, boosting chassis rigidity by 28 percent. I sensed the car's impressive solidity and quiet in my first minutes behind the wheel, including the first time I flung the Bolt through a sharp turn. This little squirt is legitimately capable and fun to drive. The Bolt goes about its business beautifully, with that spunky surge of electric torque, confident handling, and the angelic, distant-choir hum of an EV.

The Bolt's econobox styling and interior are my only real complaints, a seeming lost opportunity for GM to steal some of Tesla's high-design cachet and buyers. But in GM's defense, 60-kWh batteries aren't cheap, and the Bolt met its daunting price target while delivering big on the technology, engineering, and performance under the skin.

And there's still good stuff inside. I especially liked the cool, Cadillac-derived rearview mirror that uses a camera to dramatically expand the view. Wireless phone charging, a 4G LTE Wi-Fi hotspot, automated pedestrian braking, and forward-collision and lane-departure monitors are among the snazzy options.

A clever driver's display brings welcome transparency to that pressing EV question: How far can I really go in this thing? In my EV testing, including of the Tesla models, predictions of remaining mileage or battery level have been highly unpredictable. Crank the heater or play with that zippy Ludicrous mode and the mileage display suddenly drains away like sand in an hourglass. In contrast, the Chevy sandwiches a large numeric readout of remaining distance between a best-case and worst-case scenario. Another digital bar shows which direction you're trending, based on driving style, terrain, climate-control usage, and ambient temperature.

Outstanding efficiency is another upside. Even with a few energy-sucking romps to the Bolt's electronically limited top speed of 148 km/h (92 mph), in my test I saw the electrical equivalent of 1.96 liters per 100 km (120 mpge). Officially, the Bolt earns a rating of 119 mpge from the U.S. Environmental Protection Agency. That figure breaks down to 1.84 L/100 km (128 mpge) in the city and 2.14 L/100 km (110 mpge) on the highway. Tesla's rear-wheel-drive Model S is larger and faster, but it goes only 75 percent as far on a given amount of juice.

The Chevy also offers the "one-pedal" driving that EV fans quickly fall in love with. Four adjustable levels of regenerative braking return progressively higher amounts of energy to the battery, allowing you to step off the gas or squeeze a steering-wheel paddle and smoothly decelerate without touching the foot brake. It's especially enjoyable in city traffic or on twisty downhill roads, and Chevy says the max-level regen can boost the Bolt's real-world driving range by up to 5 percent.

As for charging, the Bolt's 7.2-kW onboard charger will stuff an empty battery in less than 10 hours on your basic 240-volt, 32-ampere charger—perfect for overnight home charging, or at workplaces and retailers that offer this Level 2 charging. For an extra $750, time-pressed owners can buy a DC fast-charging coupler that adds 145 km (90 miles) of range in 30 minutes, 257 km (160 miles) in 60 minutes, and a full charge in 2 hours. The network of megapowered DC chargers is in its infancy, with 868 U.S. locations and about 1,800 plugs at last count, but it's growing.

Buyers of EVs tend toward optimism, anyway. Fortunately, the Bolt offers plenty of other reasons to feel hopeful about the electric future. ■

▽

ENGINE
**149 kilowatts**

0-100 km/h
**6.3 seconds**

MAX SPEED
**148 km/h**

RANGE
**383 km**

EV EFFICIENCY
**120 mpge**

PRICE
**US $29,995**

# Ford Focus RS

## Born for the rally race

**Back in the 1980s,** when Volkswagen birthed the hot hatchback with the GTI, owners like me thought 110 horsepower was a big deal. Things being what they are today, we now have the Ford Focus RS, which spins up a borderline-ridiculous 261 kilowatts (350 horsepower) and 475 newton meters (350 foot-pounds) of turbocharged torque from a dinky 2.3-liter four-cylinder engine. That's more force than you get from many V-8s.

Ford's little beastie is designed to handle the dirt and snow of rally racing, or your best simulation—including a 447-kW (600-hp) version that superstar racer Ken Block will drive in the FIA Rallycross series.

But my first taste of it comes at the Monticello Motor Club, the devilish 6.6-kilometer road course in New York's Catskills region. My test pairs me with another driving star, Ben Collins, formerly The Stig from the BBC show "Top Gear." Collins proceeds to drive the wheels off the thing while bringing me up to speed on the Ford's technical calling card: a trick torque-vectoring all-wheel-drive system by the Michigan-based company GKN.

After a few laps, Collins and I trade places, and I'm soon tearing up my home course at speeds more in line with a pure sports car than a hatchback. Before the front-driven wheels can spin, the system will preemptively send 100 percent of the torque to either or both rear wheels, delivering uncanny control and explosive acceleration out of corners.

As a driver toggles through Normal, Sport, and Track modes, the GKN unit activates subsystems, each referencing new multiaxis calibration tables, that steadily trade traditional all-wheel-drive balance for take-no-prisoners performance.

The coup de grâce is Drift mode, which enables the tire-smoking, spinning-top drifting that has spawned a million Internet videos. You'll want to set aside money for periodic replacement of the Michelin Pilot Sport Cup 2 tires. They're going to need it. ■

**TAKE A SPIN:** The Ford Focus RS's engine spins up turbocharged torque that tests the tires when the car is in Drift Mode, essentially a controlled, curving skid.

TOP 10 TECH CARS 2017

# Alfa Romeo Giulia

## Digital smarts with an old-school feel

WE HERE AT *IEEE SPECTRUM* cheerfully accept our mandate of extolling the wonders of digital this and electrified that. But sometimes we do long for the quaint pleasures of simpler cars of days gone by. In this regard, the Giulia jogs our memories like a 45-rpm record, even though its bravura analog driving feel is backed by modern advances.

This spicy Italian peaks with the Quadrifoglio version, its 377 kilowatts (505 horsepower) achieved by turbocharging the blazes out of a V-6 with a displacement of a mere 2.9 liters.

This Giulia slays the 0-to-97-kilometer-per-hour (60-mile-per-hour) run in just 3.8 seconds, and peaks at a class-topping 307 km/h (191 mph). Numbers like those help explain how it briefly claimed the record as history's fastest production sedan on Germany's Nürburgring circuit.

But speed and statistics don't tell the full Giulia story, as I learned while lapping the Alfa at Sonoma Raceway in the California wine country. We're talking heart-stopping, high-wire handling and steering that's quicker than that of many sports cars, with an insane 11.8-to-1 steering ratio: Just touch the wheel and the car turns. The rear-drive Quadrifoglio gets an ultrastiff carbon-fiber driveshaft to trim weight, cutting back on the parasitic energy losses en route from the engine to its final wheeled destination. Carbon fiber forms the hood and roof. There's a twin-clutch, torque-vectoring rear differential to boost cornering. Zoom beyond 120 km/h (75 mph) and an active front splitter adjusts its angle by 10 degrees to boost front-end downforce.

ENGINE
**377 kilowatts**

0-97 km/h
**3.8 seconds**

MAX SPEED
**307 km/h**

FROM LEFT: FORD MOTOR CO.; FIAT CHRYSLER AUTOMOBILES; MORGAN MOTOR CO.

The Alfa is also the first car to adopt the brake-by-wire technology of Frankfurt-based Continental, which severs the physical link between the driver's pedal and the brakes themselves. Yes, the Integrated Braking System saves up to 4 kilograms (9 pounds) compared with conventional stoppers. But it also lays the groundwork for autonomous brakes that don't require a pedal at all and react to threats faster than any human could.

The pedal communicates with a digital control module that melds an antilock braking system with stability control while simulating the feedback your foot would get from conventional brakes. Alfa still has work to do to make the brakes feel more natural and progressive. But the company claims they can slow you down from 97 km/h (60 mph) to a full stop in 31 meters (102 feet)—a mighty achievement indeed.

One other traditional driving link is broken: You can't get a manual transmission. The consolation is an excellent eight-speed ZF automatic transmission. Humans are invited to thwack its (literally) cool, rabbit-eared aluminum shift paddles to their heart's content. Switch Alfa's console DNA Pro Drive to its zestier performance modes—Alfa claims shifts take less than 100 milliseconds. Considering the Alfa's warp-speed approach to everything, we'll take the company's word on that one. ■


**BRAVURA ANALOG RESPONSE:** The Alfa Romeo Giulia gives you that old-time performance by means of digital magic, such as brake-by-wire.

## Morgan EV3

### Three wheels, now electric

**LIKE A TRIKE:** The Morgan EV3 is the environmentally safe therapy for midlife crises.

**Dropping** an electric power train into a Morgan is like putting a nuclear reactor in a clipper ship. This romantic, inimitable three-wheeler still employs the ash-wood frame and hand-beaten sheet metal that the originals did. *In 1909*. That's when H.F.S. Morgan founded his little shop in Malvern, England.

But instead of depending on that old V-twin motorcycle engine hammering away up front, Morgan fans can race into the 21st century—in stealthy, silent fashion—with a rear-mounted electric motor. As incongruous as it may seem, electric drive suits the Morgan well because it exploits the roadster's compact design. Consider the stats: a curb weight below 500 kilograms (1,102 pounds), much of it in the form of a 20-kilowatt-hour lithium-ion battery housed in its tubular space frame, and a 46-kilowatt liquid-cooled AC motor driving the single rear wheel. Morgan figures you'll manage 240 kilometers (150 miles) on a charge. To trim weight, the EV3 is the first Morgan with carbon composite panels, used for the bonnet, tonneau cover, and side pods. The face is inspired by the aero-engine race cars of the 1930s, with burnished brass bars up front actually serving as conductive cooling fins for the battery.

Morgan expects the EV3 to scoot from 0 to 100 kilometers per hour (62 miles per hour) in under 9 seconds, on par with the internal-combustion version, and to top 145 km/h (90 mph). Yes, your mom's Toyota Camry is faster. But if you've ever been fortunate enough to drive a Morgan or even experience it from the passenger seat, you'll know that it feels like it's traveling 200 km/h when it's going only 100, plunging you into a hair-mussing euphoria of sensory stimulation and hovering danger.

Looking for a means to enjoy a satisfying midlife crisis with a low carbon footprint, no legal bills, and no secret child-support payments? Look no further, friends. ■



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

## ▷ Cadillac CT6

### A lightweight land yacht

**IT'S TWO-THIRDS ALUMINUM:** **The Cadillac CT6 is long yet light, and it's packed with energy-saving technology.**

**American luxury** sedans were once dismissed, often rightly, as lumbering land yachts. But Cadillac has honed its reputation with assassins like the 477-kilowatt (640-horsepower) CTS-V, and now it has applied the same thinking to its flagship sedan. Yes, the CT6 has a street presence that stretches to nearly 5.2 meters (17 feet), which is longer than many SUVs. But remarkably, this Caddy is lighter than a Mercedes S550 by some 450 kilograms (1,000 pounds).

The advanced Omega chassis uses 11 different materials to lose those pounds, including aluminum castings that, viewed in a cutaway model, look like computer-modeled works of art. The 13 castings dramatically reduce the number of connection points in the ultrarigid chassis, each a potential stress point: Instead of 35 stamped parts in the structure, now there are only two. Throw in aluminum body panels and nearly two-thirds of the Caddy, by weight, is formed from the lightweight metal.

The Cadillac feels so sprightly that it gets along fine with a thrifty four-cylinder engine, in turn the lightest CTS at 1,659 kg (3,657 lb). The pace quickens with an optional 301-kW (404-hp) twin-turbocharged V-6. The lusty V-6 saves fuel by deactivating unneeded cylinders and by stopping and starting the engine at every pause. It also conserves energy with the turbocharger's featherweight titanium-aluminide turbine wheels.

Cadillac pioneered the magnetic suspension now used by such carmakers as Ferrari and Audi, and the latest version of its Magnetic Ride Control reacts and adjusts to bumps and cornering forces in the blink of an eye. Strike that: *Ten times as fast* as the blink of an eye, Cadillac says. ■

TOP 10 TECH CARS 2017

## ▷ Hyundai Ioniq

### An EV that wastes no watts

**THRIFTY WITH ENERGY:** **The Hyundai Ioniq wrings more range out of a watt-hour than any other EV around.**

**The most energy-**efficient car in America isn't a Tesla, Toyota, or even the remarkable Chevrolet Bolt. It's a Hyundai. With a range of 200 kilometers (124 miles), the all-electric Ioniq comes nowhere near the Bolt's astounding 383 km (238 miles). But the Hyundai's gasoline equivalent of 1.73 liters per 100 kilometers (136 mpge) tops all current EVs.

Let's do some quick math. Americans will pay about US $3.36 to charge the Ioniq's 28 kilowatt-hour battery pack, at the average electricity rate of 12 cents per kilowatt hour. So, for less than the price of 1.5 gallons of unleaded, an Ioniq can cover 124 miles, the distance from Philadelphia to Washington, D.C.

Not ready to make the leap to a pure EV? The Ioniq offers hybrid and plug-in hybrid versions, with the standard hybrid delivering up to 4.1 L/100 km (58 mpg). Less gawky-looking than a Prius, these Ioniqs match the slippery 0.24 drag coefficient of the Tesla Model S, making them the most aerodynamic models on U.S. roads. With automakers and environmentalists now focused on automobiles' full carbon footprint, from manufacturing to disposal, the Ioniq's recycled interior plastics are combined with volcanic stone and powdered wood, reducing weight by 20 percent and allowing easy recycling.

Hyundai is experimenting in California with a way to lure buyers from conventional cars: a cellphone-style subscription plan. It bundles unlimited mileage, charging costs, scheduled maintenance, and replacement of worn-out items into a single monthly payment. Don't be surprised to see other EV makers follow suit. ■

RANGE
**200 km**

EV EFFICIENCY
**136 mpge**

HYBRID EFFICIENCY
**58 mpg**

**QUICK, CLEAN DIESEL:** **The Porsche Panamera comes in many tech-packed versions, but the diesel is really a standout.**

FROM LEFT: GENERAL MOTORS CO.; HYUNDAI; PORSCHE


# Porsche Panamera

**The world's fastest diesel**

**THE ORIGINAL PANAMERA** sedan, in 2009, dazzled everyone with its performance but drew brickbats for its Hunchback-of-Stuttgart styling. The 2017 edition looks like a proper Porsche while packing even more technology and performance under that sleeker skin.

The new Panamera is the first of myriad models built on the VW Group's clever MSB architecture, a modular layout that will allow many vehicles—including an upcoming Bentley Continental GT—to share power trains, steering, and much else, regardless of the size and shape of the car. A pair of twin-turbocharged engines includes a 2.9-liter V-6 and a 4.0-L V-8, the latter delivering a monstrous 410 kilowatts (550 horsepower), a 0-to-97-kilometer-per-hour (60-mile-per-hour) blast in 3.7 seconds, and a 305-km/h peak. Europeans will get a 4S Diesel model whose V-8 turbodiesel spools up 310 kW (418 hp) and a titanic 850 newton meters (627 foot-pounds) of torque. Porsche calls it the world's fastest production diesel, combining a 285-km/h (177-mph) top speed with up to 900 miles of range on a single tank. That's enough to go from Paris to Rome.

ENGINE
**410 kilowatts**

0-97 km/h
**3.7 seconds**

MAX SPEED (DIESEL)
**285 km/h**

And, unlike the Volkswagens that failed emissions testing (and covered it up with a nefarious digital "defeat device"), this Porsche carries a tank of urea-based DEF fluid to neutralize nitrous oxides.

The technical goodies are impressive. Electronic shocks, rear steering, and torque vectoring are all overseen by the new 4D Chassis Control, which analyzes driving situations in three spatial dimensions—pitch, roll, and yaw—and optimizes response in real time (hence the "fourth dimension"). The system thus networks various electronic suspension systems that once worked independently. It obsessively monitors the vehicle's trajectory and driver inputs, and coordinates every system to maximize agility, stability, and safety. Models equipped with a three-chambered air suspension can include optional active stabilizer bars, controlled via 48-volt electromechanical actuators. These counteract cornering loads to keep the body flatter than a Marine's haircut.

An electronically controlled steering rack mounts at the rear subframe to steer rear wheels opposite to the fronts at lower speeds, for snappier moves in tight corners. Or it can steer them in tandem with the front wheels at high speeds to keep you on course.

Also new is InnoDrive, which analyzes GPS data from up to 3 km away to autonomously steer, shift, and control the Porsche's speed. That radar- and camera-based system also reads road signs to automatically keep the Panamera at the speed limit. And a new thermal-imaging night vision system, in tandem with LED Matrix headlamps, will alert drivers to animals or pedestrians in your path, far beyond high-beam range—if they're not already fleeing your prodigiously fast approach. ■

TOP 10 TECH CARS 2017

# Chrysler Pacifica Hybrid

**Energy efficiency today, robotic driving tomorrow**

SURPRISINGLY, THERE'S NEVER BEEN A hybrid minivan of any kind on the U.S. market. So, when Chrysler rolled out not only a hybrid minivan but a frugal plug-in model, it was a big deal.

This year's Pacifica sheds 113 kilograms (250 pounds) thanks to a lightened platform, which includes aluminum sliding doors and an aluminum-and-magnesium lift gate, both featuring hands-free operation. The minivan's 3.6-liter V-6 runs on the fuel-saving Atkinson cycle—as in hybrids like the Toyota Prius—and links to an ingenious twin electric-motor arrangement that can send as much as 194 kilowatts (260 horsepower) to the front wheels. A 16-kilowatt-hour lithium-ion battery pack with 96 cells is designed and built in Michigan and, together with related electronics, weighs in at 295 kg (650 lb) and stretches 75 centimeters (2.5 feet). It fits right in the receptacle where the Stow 'n Go seats disappear into the floor in the nonhybrid version. Many current hybrids use one electric motor-generator exclusively to recapture energy, another to provide propulsion, but the Chrysler innovates again: Its dual electric motors integrate an electronic clutch and planetary gearset. That allows both motors to act as traction motors, boosting the efficiency and scalability of the system.

How efficient? During all-electric operation it reaches the equivalent of 2.8 liters per 100 kilometers, or 84 mpge. And where, say, a Toyota Prius plug-in can manage just 18 kilometers (11 miles) on batteries alone, the Pacifica's battery will take you about three times that distance before the gasoline engine springs to life—at which point this minivan is still good for 7.4 L/100 km (32 mpg), versus just 10.7 L/100 km (22 mpg) for the nonhybrid model. To ease the load on the battery and heat the big seven-passenger cabin, an engine-powered 7-kW electric heater steps in when necessary.

The automatic transmission can constantly vary gear ratios by varying the speeds of three devices—the engine, motor-generator, and main electric trac-

ENGINE
**194 kilowatts**

EV EFFICIENCY
**84 mpge**

GASOLINE EFFICIENCY
**32 mpg**

## Infiniti QX50

**The world's first variable compression-ratio engine**

**We tend** to eschew concept cars for our Top Ten list on the grounds that anyone can build a prototype and float pie-in-the-sky claims. But the QX50 is an exception, and for more than its energetic crossover design. First, this Infiniti is real, and its technology is coming to showrooms soon, including the first variable compression-ratio engine ever. The new 2.0-liter, four-cylinder engine can adjust its piston stroke to vary compression from about 8:1 to 14:1, heightening efficiency at the lower range and performance up top.

Infiniti pledges to have the engine in at least one new Infiniti—likely this QX50—by 2018, with a target of 198 kilowatts (266 horsepower) and 390 newton meters (288 foot-pounds). The QX50 itself will replace the current, compact EX35 crossover. And with Nissan, Infiniti's parent, targeting 2020 for its first fully autonomous car, the QX50 also previews new ProPilot technology, able to negotiate stop-and-go highway traffic and also recognize and track surrounding vehicles. ■

FROM LEFT: INFINITI; FIAT CHRYSLER AUTOMOBILES; DAIMLER



POST YOUR COMMENTS at http://spectrum.ieee.org/toptechcars0417

tion motor–through the planetary gearset. That provides all the flexibility and efficiency of a belt-driven continuously variable transmission (CVT), but without the power lag or rubber-band feel you get from a CVT when you punch the accelerator. The engine and transmission can be a bit whinier than those of some hybrids, but the Pacifica drives well and is stuffed with the latest digital goodies.

I dazzled a full load of children and adults by demonstrating the Pacifica's fully automated parking tech, in both parallel and perpendicular modes. The van's theater system proved a boon with four children wedged into the back. They all enjoyed the built-in games on dual 10-inch touch screens. In addition, the Pacifica benefits from the ongoing trickle-down of technology from the luxury segment–automated collision warnings and braking, automated high beams, and adaptive cruise control.

Fire up the Uconnect phone app to remotely charge the Chrysler in just 2 hours, when it's connected to a 240-volt Level 2 charger, or program it to juice at off-peak hours. You can also monitor its charge status or preheat and cool the cabin using electricity rather than gasoline.

At this year's Detroit Auto Show, this minivan sparked the kind of buzz that's typically reserved for a Tesla. The Google car spin-off Waymo announced a fleet of 100 fully autonomous Pacificas–co-engineered by Fiat Chrysler and Waymo–that have already gone into testing. They're scheduled to roll out in a ride-sharing service as early as 2018.

An exciting minivan? I *never* saw that coming. ■

**AN EXCITING MINIVAN:** Even in all-electric mode, the Chrysler Pacifica Hybrid gives superb performance and range while using energy efficiently.

# Mercedes-Benz AMG E63 S

## It avoids crashing—but if it does, it helps you survive


**DANGER, WILL ROBINSON:** When the AMG E63 S senses an imminent crash, it braces every part of you, even your inner ear.

ENGINE
**450 kilowatts**

0-100 km/h
**3.3 seconds**

TORQUE
**800 newton meters**

**Most owners**, I suspect, will prefer to drive Mercedes's new missile by their flawed human selves. I can't blame them. The supersedan's 4.0-liter biturbo V-8, shared with the AMG GT supercar, pours out up to 450 kilowatts (603 horsepower). This Benz will hit 100 kilometers per hour (62 miles per hour) in a supercar-like 3.3 seconds, which is ridiculous for a sybaritic, all-wheel-drive midsize sedan.

Yet when it's time to take a break from perspiring action, the new Drive Pilot system surrounds the Benz in radar beams that peer 250 meters ahead and 80 meters behind the car. They're supplemented by separate 3D and 2D cameras that read lane markers and traffic ahead.

The car can still get flustered by tighter curves or vague road markings, as I found on a test drive in New York. But it will just about drive itself on the highway at up to 160 km/h (100 mph), chiding you with flashing alerts if you take your hands off the wheel for too long. Hit your turn signal and the system scans the surroundings before executing an automated lane change.

If all else fails, Mercedes's latest suite of Pre-Safe features add cool new safeguards: Active seat bolsters thoughtfully shove occupants 8 centimeters toward the vehicle's center just before a side impact to mitigate injuries by, among other things, getting bodies moving in the proper direction before a strike.

As you realize a crash is inevitable, you will be pleased to know that among the safeguards is Pre-Safe Sound. It's an interference signal that comes through the audio system in the milliseconds before impact, triggering the ear's stapedius muscle to contract, protecting the eardrum from injuries due to high acoustic pressures.

Shhh... can you hear that? It's the sound of German engineers thinking. ■

# THE BUNNY, THE WITCH, AND THE WAR ROOM

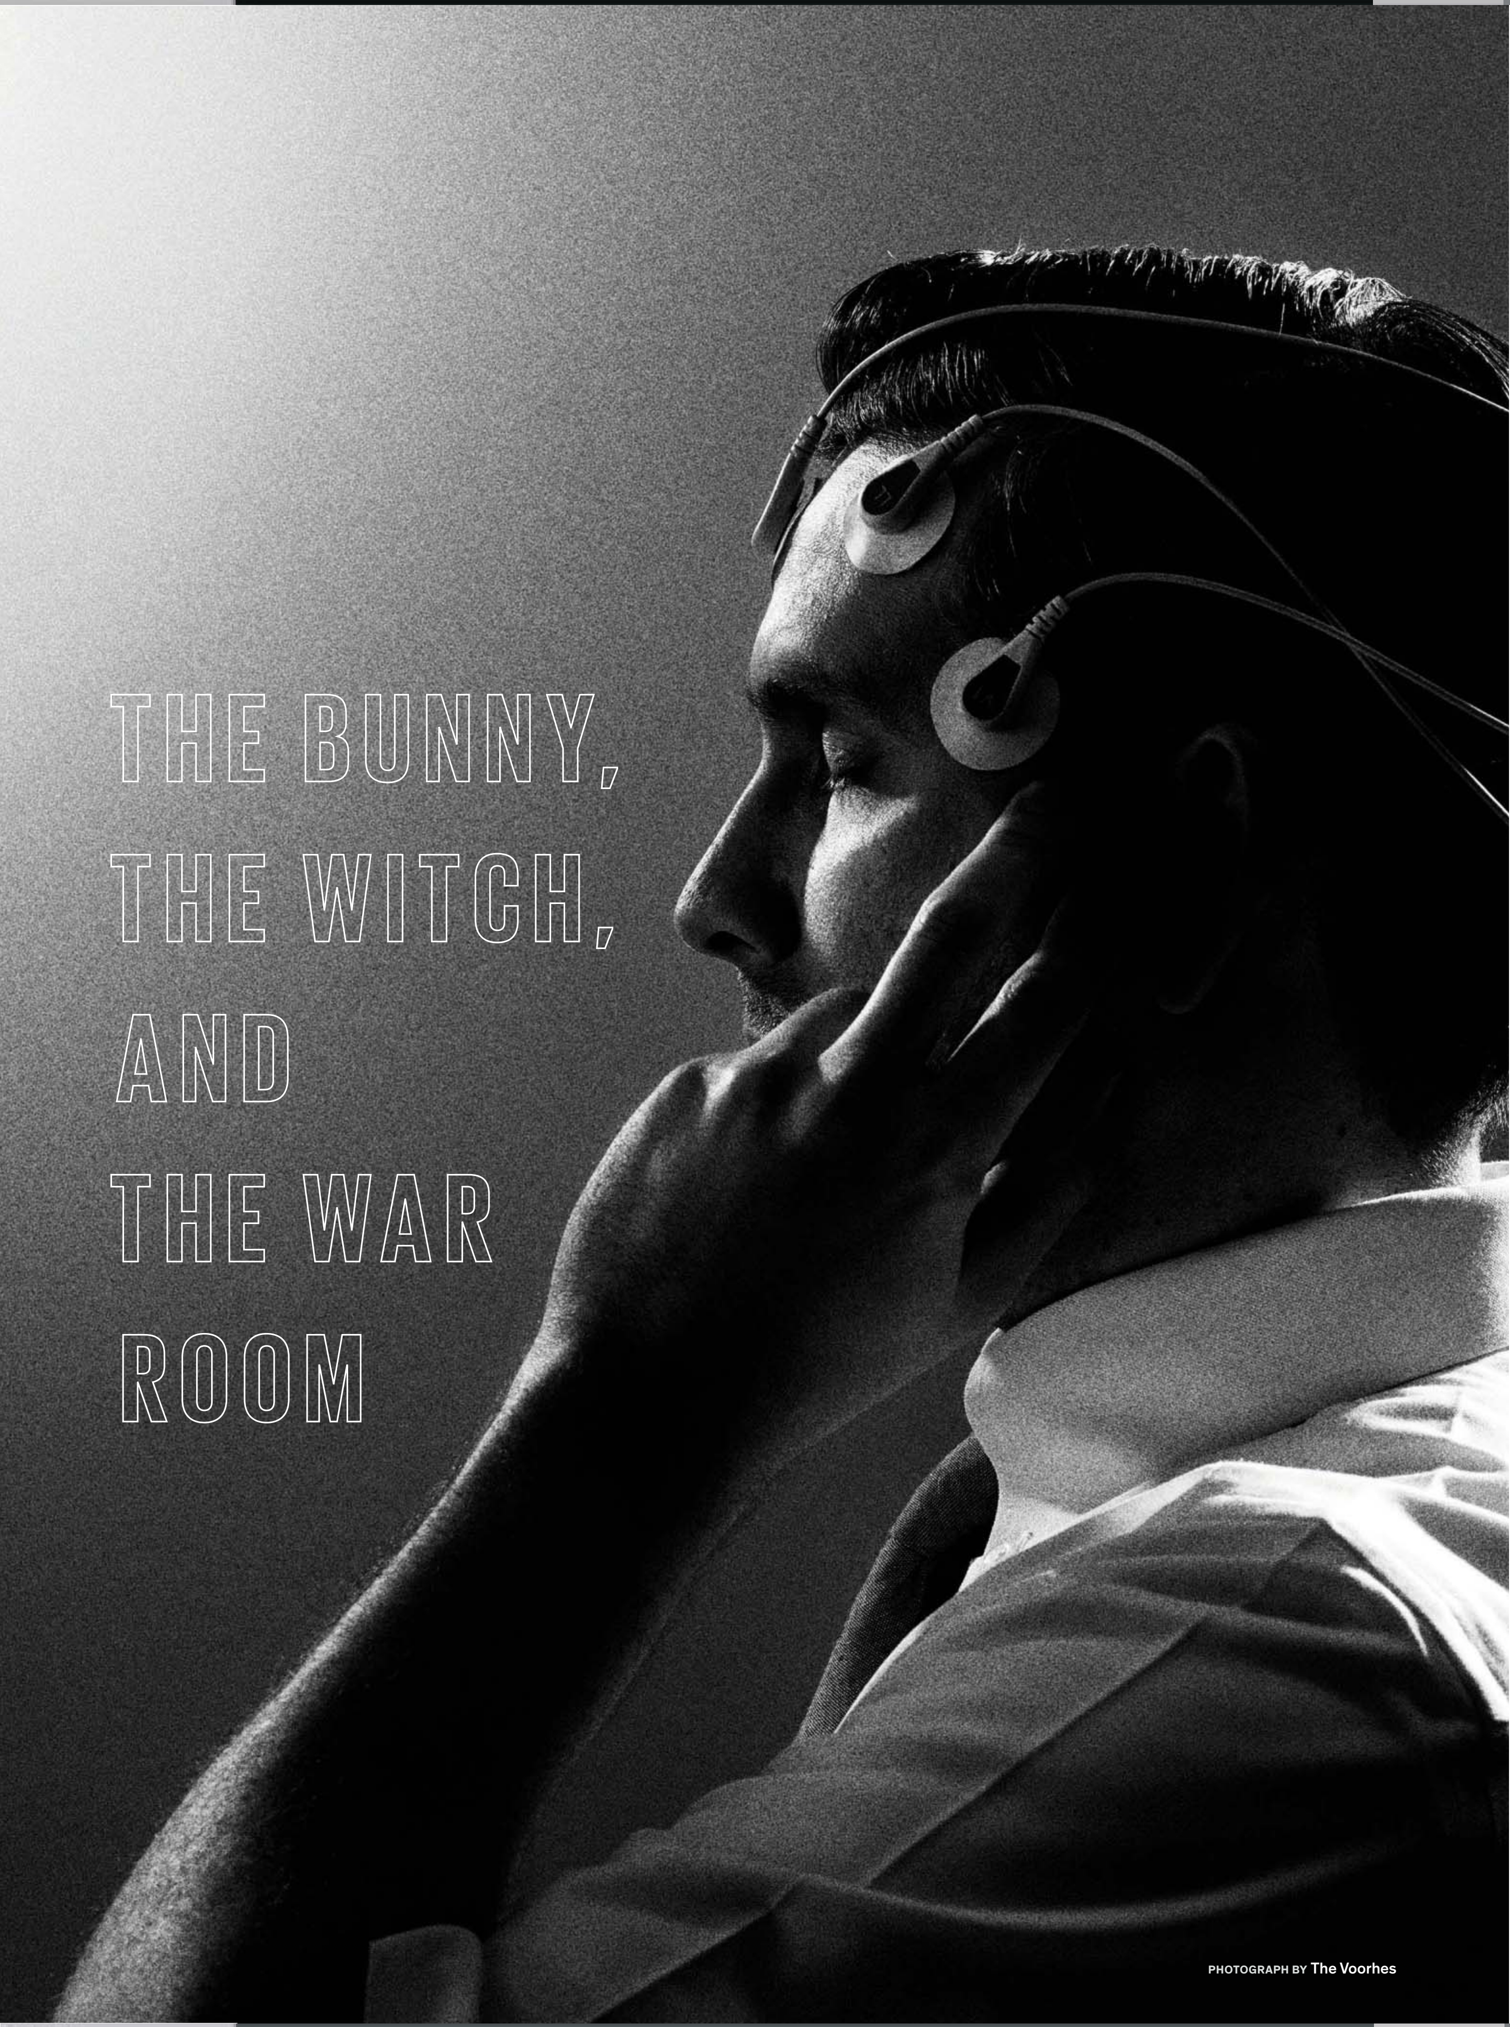

PHOTOGRAPH BY The Voorhes

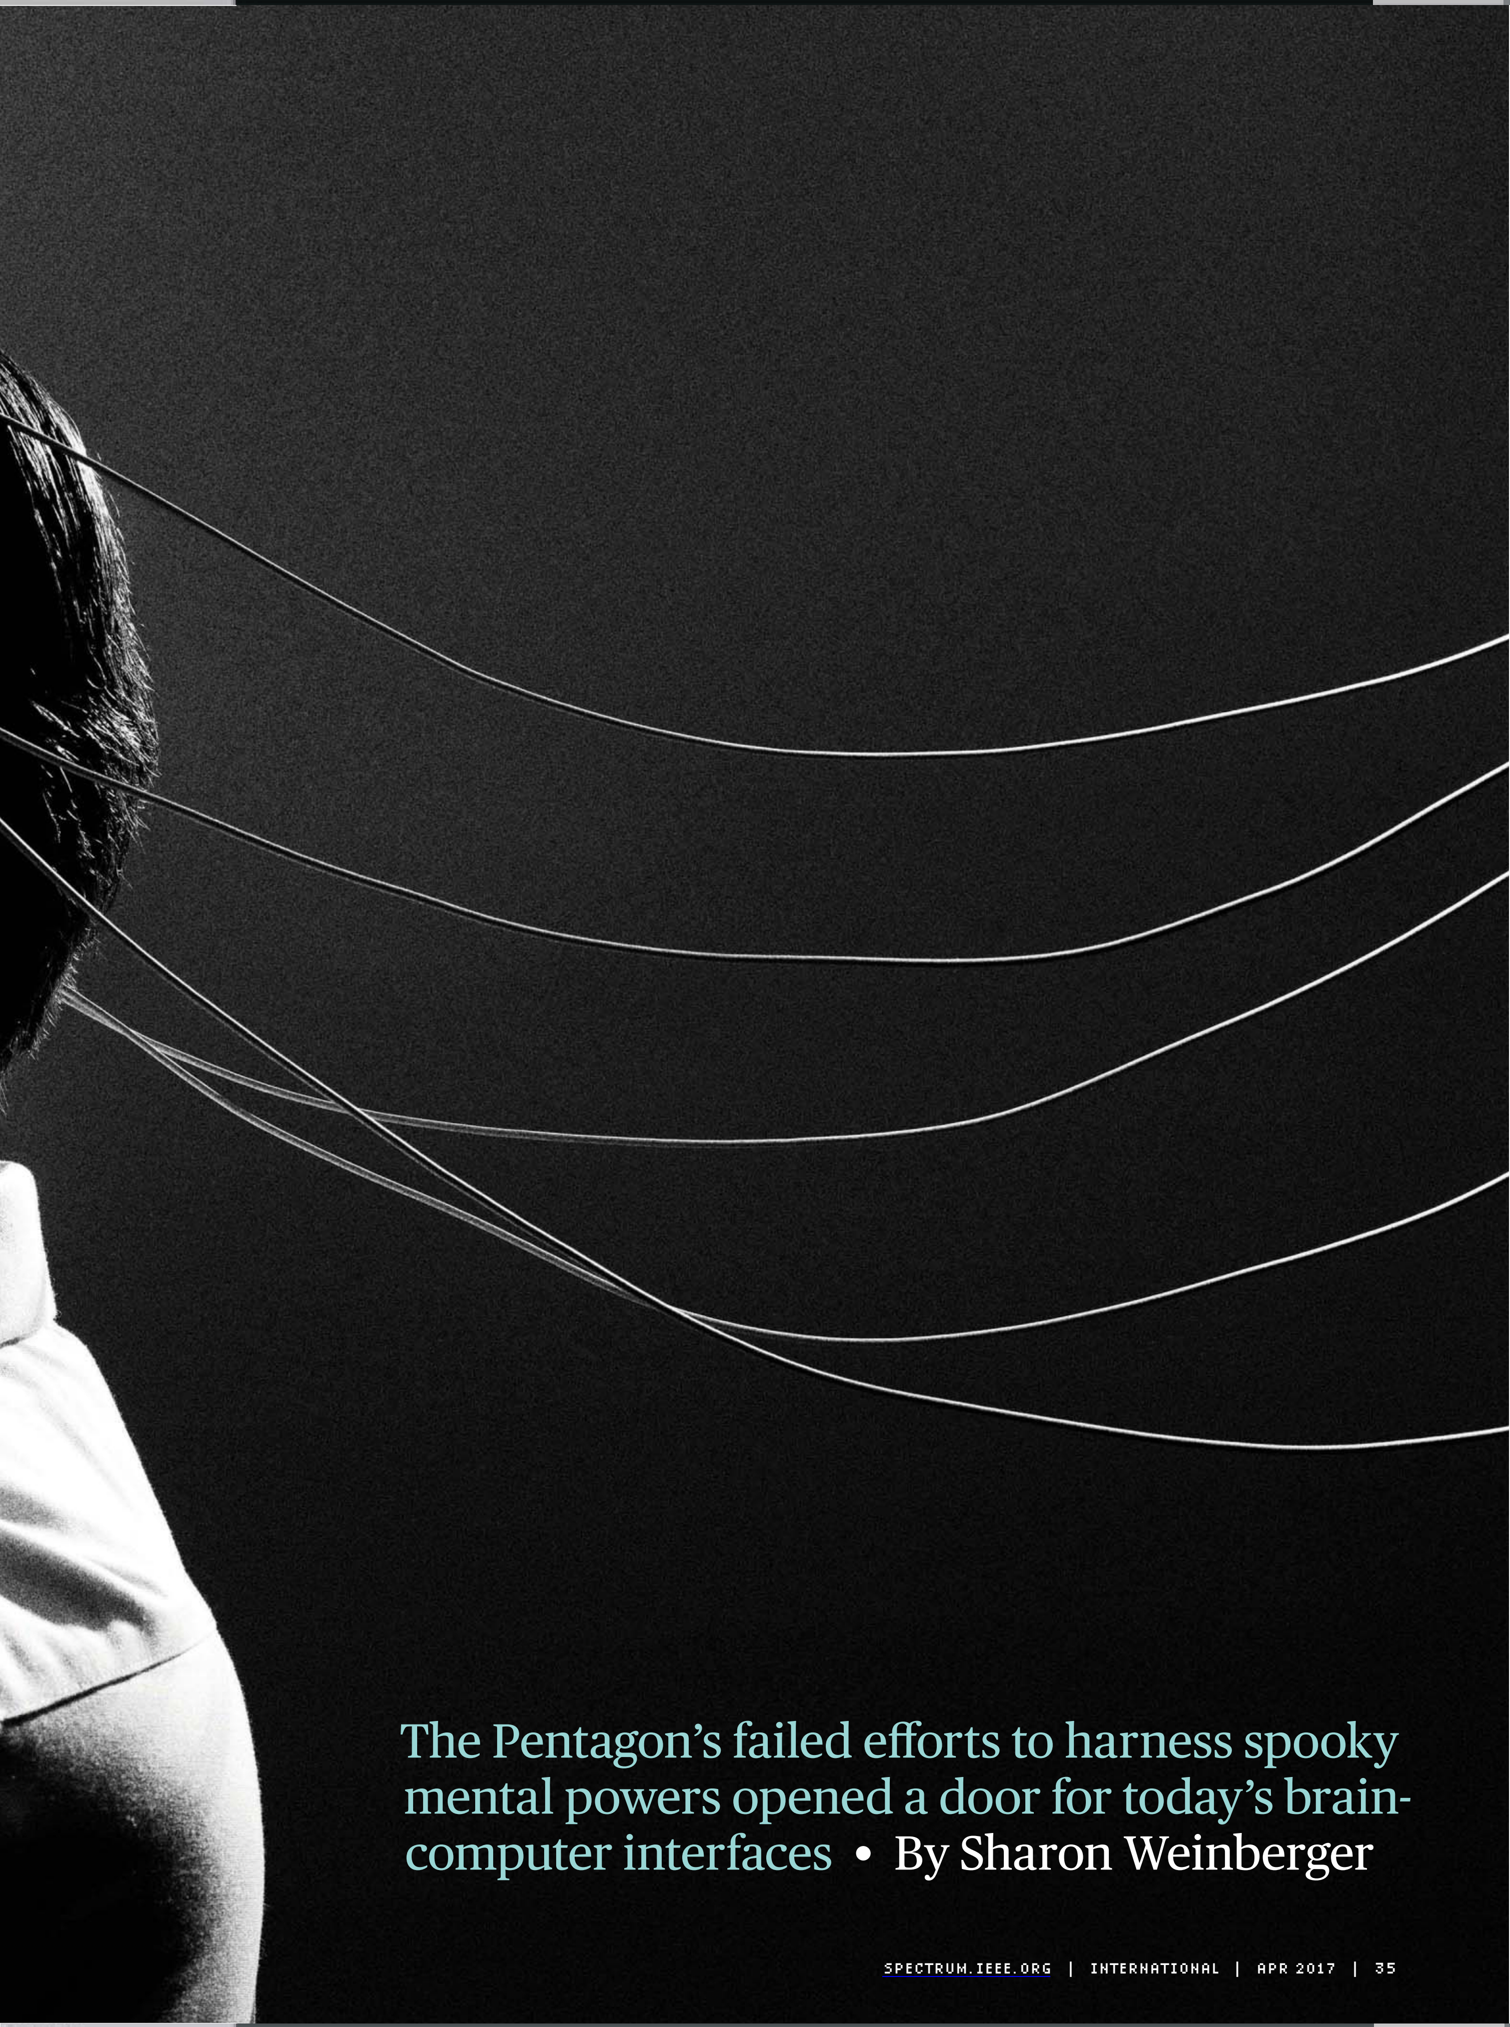

The Pentagon's failed efforts to harness spooky mental powers opened a door for today's brain-computer interfaces • By Sharon Weinberger

AMONG THE STANFORD RESEARCH INSTITUTE'S many classified research projects in the early 1970s was a contract supported by the Central Intelligence Agency's Office of Technical Service, a division headed by Sidney Gottlieb, perhaps the most notorious scientist ever to work for the spy agency. The secret program was testing different forms of parapsychology, such as whether humans had the ability to use their minds to visualize or even influence remote objects. Believing the work was showing promise, Gottlieb one day invited the director of the Advanced Research Projects Agency (ARPA), Stephen Lukasik, over to his CIA office to discuss it.

Gottlieb, a chemist by training, was both an unconventional thinker and an unwavering patriot, who believed his work served the good of the nation. "Friends and enemies alike say Mr. Gottlieb was a kind of genius, striving to explore the frontiers of the human mind for his country," read the 1999 *New York Times* obituary of Gottlieb, "while searching for religious and spiritual meaning in his life." In the end, however, Gottlieb would be remembered most for what looked like a willful contempt of common decency.

As the head of the Office of Technical Service, Gottlieb led a wing of the CIA whose failed innovations to assassinate the Cuban leader Fidel Castro included poison pens and exploding seashells. He also worked on one of the agency's most notorious projects: the use of LSD as a mind-control drug. Under Gottlieb's supervision, LSD was tested on unwitting human guinea pigs, including, among other unfortunate victims, the mentally ill, prostitutes, and even one army scientist who later committed suicide. When the program was first exposed in 1975 by the Rockefeller Commission, and then detailed by the congressional Church Committee, Gottlieb's public legacy as some sort of mad scientist was all but assured.

The day Lukasik went to visit Gottlieb—in 1971, as Lukasik recalls it—the CIA scientist was in fine form. What Gottlieb wanted to discuss was bunny rabbits and nuclear Armageddon.

In the early 1970s, the Soviet Union and the United States were locked in a cat-and-mouse game involving nuclear submarines. Submarines equipped with nuclear missiles were difficult to spot when prowling the deep seas, making them a potent weapon. But there was no good way to tell submarines deep underwater that they needed to launch their missiles. And coming to the surface periodically to receive communications would make them vulnerable to detection and attack.

That was where Gottlieb's new pet project came into play. In 1970, the best-selling book *Psychic Discoveries Behind the Iron Curtain* (Prentice-Hall) described the enthusiasm of the Soviet Union and other Eastern Bloc countries for psychic phenomena of all sorts. "Major impetus behind the Soviet drive to harness ESP [extrasensory perception] was said to come from the Soviet Military and the Soviet Secret Police," the authors, Sheila Ostrander and Lynn Schroeder, asserted. The book detailed dozens of investigations into psychic phenomena conducted behind the Iron Curtain, ranging from Kirlian photography, which sought to capture the "aura" of living things, to telepathic projection of emotions. The idea that the Soviets were investing money in parapsychology quickly became a self-reinforcing justification for the Americans to do the same.

According to *Psychic Discoveries*, one theory of parapsychology the Soviets were testing involved a projected emotional link between a newborn and its mother, which allowed the mother to "sense" her offspring's death even over long distances. Because actually killing a newborn human child was not really an option, they resorted to experimenting with baby rabbits and their mothers. The experiment was as ghastly as it sounded: A baby rabbit would be killed out of sight and sound of its mother, while scientists in a separate lab room observed the mother for a reaction.

The Soviets claimed it worked and could be used for communicating with submarines, even if they never quite laid out the protocol for how this would be done. Presumably, a mother rabbit would be kept aboard the submarine, with a submariner assigned to monitor it for signs of distress. The idea was not that an overly excited mother rabbit would prompt a nuclear exchange, but such a sign could be used, as Lukasik put it, as a "bell ringer for Soviet boomers." It would be a signal for the submarine to surface and get a more detailed message, such as an order to launch its nuclear missiles.

The very absurdity of the scenario did not dissuade Gottlieb. The CIA had begun funding the



PHOTOGRAPHS BY The Voorhes

Stanford Research Institute to conduct a "quiet, low-profile classified investigation" into parapsychology. Gottlieb was interested in having ARPA look at that work and possibly support it.

Although the purported Soviet experiments sounded dubious, antisubmarine warfare was an area that ARPA was pursuing. Perhaps more important, the late 1960s and early 1970s had sparked widespread interest in parapsychology, even among some members of Congress, who were pressuring agencies like ARPA to support it.

"I thought this was a lot of bullshit," Lukasik admitted, but he figured that at the least the agency could make a good faith effort to see if there was anything worth funding.

The scientist selected to lead the parapsychology investigation was ARPA's resident expert in counterculture, George Lawrence. The 39-year-old cut a distinctive figure at ARPA. Even in the 1960s and 1970s, the agency was almost as straitlaced as any other part of the Pentagon. It was home to intellectual freethinkers, but they were largely drawn from the hard-science faculties of universities, the defense industry, and the military, not exactly the hotbeds of 1960s counterculture.

Lawrence was an exception. He had adopted at least the trappings of a bohemian, favoring bell-bottoms and wide-collared shirts. Lawrence had a penchant for research that captured the cultural zeitgeist of the late 1960s, when explorations of mind-body interactions and consciousness research combined science and spiritualism. He was also part of a small but growing number of psychologists fascinated by computers.

His first major program, started in 1970, was in biofeedback, a relatively new area of investigation that involved training people to control physiological functions, such as breathing and heart rate, by providing subjects with real-time information from sensors. The idea was that a person could essentially will his or her way to a different physical state. It melded biology with Eastern philosophy and evoked comparisons to Timothy Leary's promotion of LSD. Scientists thought biofeedback might enable people in stressful situations to slow their heart rates or lower their blood pressure purely through mental concentration.

The justification for ARPA's interest in this field was to help troops in combat. Biofeedback could, in theory, allow soldiers to shoot more accurately, or even to slow their bleeding after being shot, by letting them control their heart rate. Researchers hypothesized that pilots of damaged aircraft could be taught to lower their heart rate and blood pressure, in order to carry out emergency procedures without panicking. There was little in the way of documented experiments, however, and Lawrence considered biofeedback an area ripe for examination.

**Editor's note:** This article is based on a chapter of the author's newly released book, *The Imagineers of War: The Untold Story of DARPA, the Pentagon Agency That Changed the World* (Alfred A. Knopf). In this excerpt, Weinberger recounts one of the more bizarre episodes in the history of that agency, revealing some real-life government X-Files.

His ARPA program was the first systematic exploration of the field, bringing the scientific method to an area dominated by anecdotes. But when Lawrence wanted the researchers to travel to Vietnam to test biofeedback in the field, no one wanted to go. "Someone wrote to a congressman and said I was trying to coerce university professors into going into the jungles in Vietnam," he said. "I thought they would look at it as an interesting adventure, the way that I did."

It probably did not matter in the end, because Lawrence concluded that the more ambitious applications for biofeedback, such as soldiers' ability to slow down their heart rate enough to prevent them from bleeding out, were probably unattainable. On the flip side, Lawrence wrote, at least no one was going to die by consciously willing his heart to stop.

While the biofeedback program was not necessarily successful, it cemented Lawrence's reputation at ARPA as the go-to guy for counterculture ideas. So it was not a total surprise when Lawrence was assigned to look at the CIA's parapsychology research to see if it was something that ARPA should fund.

**WHEN ARPA LAUNCHED ITS INVESTIGATION** into parapsychology, it was hard to tell how seriously anyone, even Lawrence, really took it. At least at face value, Lawrence embraced the assignment. He played with Kirlian photography to see if it could really capture auras, attended a parapsychology conference in Scotland, and traveled around the country meeting witches, psychics, and other purveyors of the paranormal. He liked the witches most of all.

But Lawrence's most famous psychic investigation—and one that ended up attracting national attention when it was picked up by the press—involved a trip he made in December 1972 to the Stanford Research Institute, where the physicists Russell Targ and Hal Puthoff were being funded by Gottlieb's CIA office to investigate psychic phenomena.

At the time, that work was focused largely on testing the skills of Uri Geller, a charismatic Israeli entertainer turned paranormalist. Geller's most well-known spectacle was bending spoons, purportedly with his mind. He also claimed a host of other psychic abilities, such as thought

PHOTOGRAPH BY The Voorhes

projection and "remote viewing," the term given to the ability to describe objects in far-off, or at least unseen, locations. Remote viewing was of particular interest to the national security community, because it would in theory enable spying on foreign bases and technology.

Puthoff and Targ, who were eager to get mainstream recognition, agreed to host Lawrence for an informal demonstration but told him he could not observe their controlled experiments. Lawrence invited two other scientists to accompany him: Ray Hyman, an amateur magician and university psychologist, and Robert Van de Castle, a professor of sleep studies who believed in psychic premonitions, including his own. Van de Castle, whom Lawrence knew from graduate school, studied the ability of people to predict the future and receive thoughts while dreaming. "He and Hyman and I made this trip to the Stanford Research Institute, where Geller was going to convince me that his stuff was valid, and I was going to pump a lot of money into it," Lawrence said.

Hyman and Van de Castle showed up to meet with Puthoff, Targ, and Geller. Lawrence, who'd been drinking heavily the previous night, swaggered in late, looking a mess, according to Van de Castle. And so the day started with one hungover military scientist, one amateur magician turned psychologist, a professor who studied psychic dreams, two seemingly credulous physicists, and Uri Geller, the would-be psychic superweapon. It went downhill from there.

Geller began his repertoire by demonstrating his ability to mind-read numbers. The Israeli performer dramatically covered his eyes with his hand and had Lawrence write down a number on a piece of paper. Hyman, sitting to the side, later recalled that he could clearly see that Geller was peeking, watching the motion of Lawrence's hand.

In another demonstration, Geller wanted to show his psychic ability to receive someone's thoughts, so he took Van de Castle aside into a separate room. Geller asked Van de Castle to choose a cartoon from a magazine and draw it by hand, because magazine pictures were harder to "receive." Both pictures–the original and the hand-drawn duplicate–were placed in separate envelopes. Van de Castle placed the envelope with the original picture in his breast pocket and the one with the hand-drawn image under his elbow. Geller then instructed the professor to close his eyes, stood directly behind him–close enough to touch him–and prepared to receive Van de Castle's thoughts. Geller soon emerged triumphant: He had drawn a stick figure facsimile of the image, a feat no one observed because only Van de Castle was in the room and he had his eyes closed the entire time.

Hyman was perplexed: What were the conditions of the experiment? Why did no one observe Geller drawing the image? The answers were evasive at best. And so it went with the rest of the demonstrations. Either Geller could not or would not perform under close scrutiny, or when he did seem to get results, there was little credible examination. "Targ and Puthoff, from the way I have encountered them by day in their laboratory, seem to emerge as bumbling idiots rather than as respected, accomplished physicists," Hyman wrote. He believed Geller's work had all the classic hallmarks of a trained magician: befriend, distract, and dazzle.

If Hyman was doubtful of Geller's psychic capabilities, Lawrence was outraged. In one demonstration, Geller moved a compass needle by 5 degrees. Lawrence, stomping his foot to imitate what he believed Geller had done, moved the needle 45 degrees. It was clear that Puthoff and Targ were not going to get ARPA funding.

**ALTHOUGH GELLER'S DEMONSTRATION** had failed to impress Lawrence, the idea of reading people's minds captured the ARPA scientist's imagination. The same year he visited the Stanford Research Institute, Lawrence launched a different sort of mind-reading project: Instead of relying on the paranormal, researchers would use measurable brain signals to control a computer.

The brain-driven computer dreamed up by Lawrence, using what he called "biocybernetic communication," was outright audacious. The

OPPOSITE PAGE: CLOCKWISE FROM LEFT, HELLA HAMMID; THE LAWRENCE FAMILY; HELLA HAMMID; STEPHEN LUKASIK; BERNARD GOTFRYD/GETTY IMAGES machine would not just be controlled through inputs provided by a keyboard or joystick; it would interact directly with the human mind, using sensors that monitor brain activity.

ARPA's biocybernetics program funded a raft of researchers tapping brain signals, such as Jacques Vidal, a UCLA researcher who coined the term *brain-computer interface*. "Can these observable electrical brain signals be put to work as carriers of information in man-computer communication or for the purpose of controlling such external apparatus as prosthetic devices or spaceships?" Vidal wrote in a seminal paper in 1973. Within a few years, Vidal's research yielded promising results: In one experiment, test subjects were able to move an electronic object through a maze on a computer screen just by thinking.

Those were fantastic times, according to Emanuel Donchin, then a professor at the University of Illinois who was funded by Lawrence for research on detecting and interpreting brain signals. ARPA wasn't the only agency supporting such work, but it was the most important.

On the other hand, ARPA programs like biocybernetics were often outrageously optimistic about their military applications. "Soon, for example, a computer monitoring electrical brain activity of an aircraft pilot...should be able to determine whether a warning signal not only had been seen but that the pilot understood its significance and intended to respond appropriately," one early program description read. "I made it up," Lawrence recalled of some of the more fantastical applications. The challenge of biocybernetics was weighing the fantastical applications it offered–brain-driven computers and mind-controlled aircraft–with the reality that such work was decades away.

**LAWRENCE'S BRAIN-DRIVEN COMPUTERS** were on the edge, but so was killing bunny rabbits to communicate with submarines or funding an Israeli magician to remotely view Soviet bases. ARPA was a place in the early 1970s that tolerated and even encouraged exploring such outlandish ideas, but unlike some other agencies, it required good science.

In one final meeting to discuss ARPA's possible funding of parapsychology, Lawrence sat with Lukasik and CIA officials who had been funding such work. At the end, one of the CIA officials turned to Lawrence and said, "Dr. Lawrence, what do you think about all this?"

At that point, Lawrence's investigation of psychic phenomena had introduced him to a colorful array of mystics and frauds. "You have been wasting your money," he exploded in frustration. "Every damn dime of this is nonsense."

There was dead silence. Lukasik quickly changed the subject, and no one ever asked Lawrence to look at parapsychology again. Nor did ARPA ever fund a psychic program. "I worked so long, and so hard, and dealt with so many fools and charlatans," Lawrence later recalled. "There is no question in my mind that all of it is bunk."

Geller's advocates, who believed the magician could help the United States spot Soviet submarines, looked on Lawrence's conclusion with great disappointment. But Lawrence helped save ARPA from the embarrassment that befell the intelligence community when it was revealed the nation's spies had spent tens of millions of dollars on psychics. And for those who questioned whether ARPA's open-ended investigation of parapsychology was a good idea at all, the reality is that the same attitude that allowed Lawrence to meet witches and psychics also enabled him to pursue the brain-driven computer.

The intelligence community's support of psychics continued through 1995, producing claims of successful results but little in the way of scientific evidence. Biocybernetics, on the other hand, blossomed. It was an audacious idea in the early 1970s, when the ability to read brain signals was crude at best. By 2013, however, biocybernetics had spawned an entire industry of brain-computer interfaces used for such diverse applications as commercial video games, car sensors, and tools that allow "locked in" patients, those with no way to communicate with the external world, to type messages and control external devices. Applications that were once decades away are now being built, and Lawrence's early vision is becoming a reality.

As for parapsychology, Lawrence joked years later that maybe he should not have been so forthright with his criticism, instead playing it out even longer. "At the very least," he said, "I could have met some more witches." ■

**THE X-TEAM:** Brought together in this photo collage, these men each played a part in ARPA's dabbling with the paranomal in the 1970s. Pictured here are [clockwise from left] physicist Russell Targ, ARPA psychologist George Lawrence, physicist Hal Puthoff, ARPA director Stephen Lukasik, and magician Uri Geller, all shown as they appeared during that era.

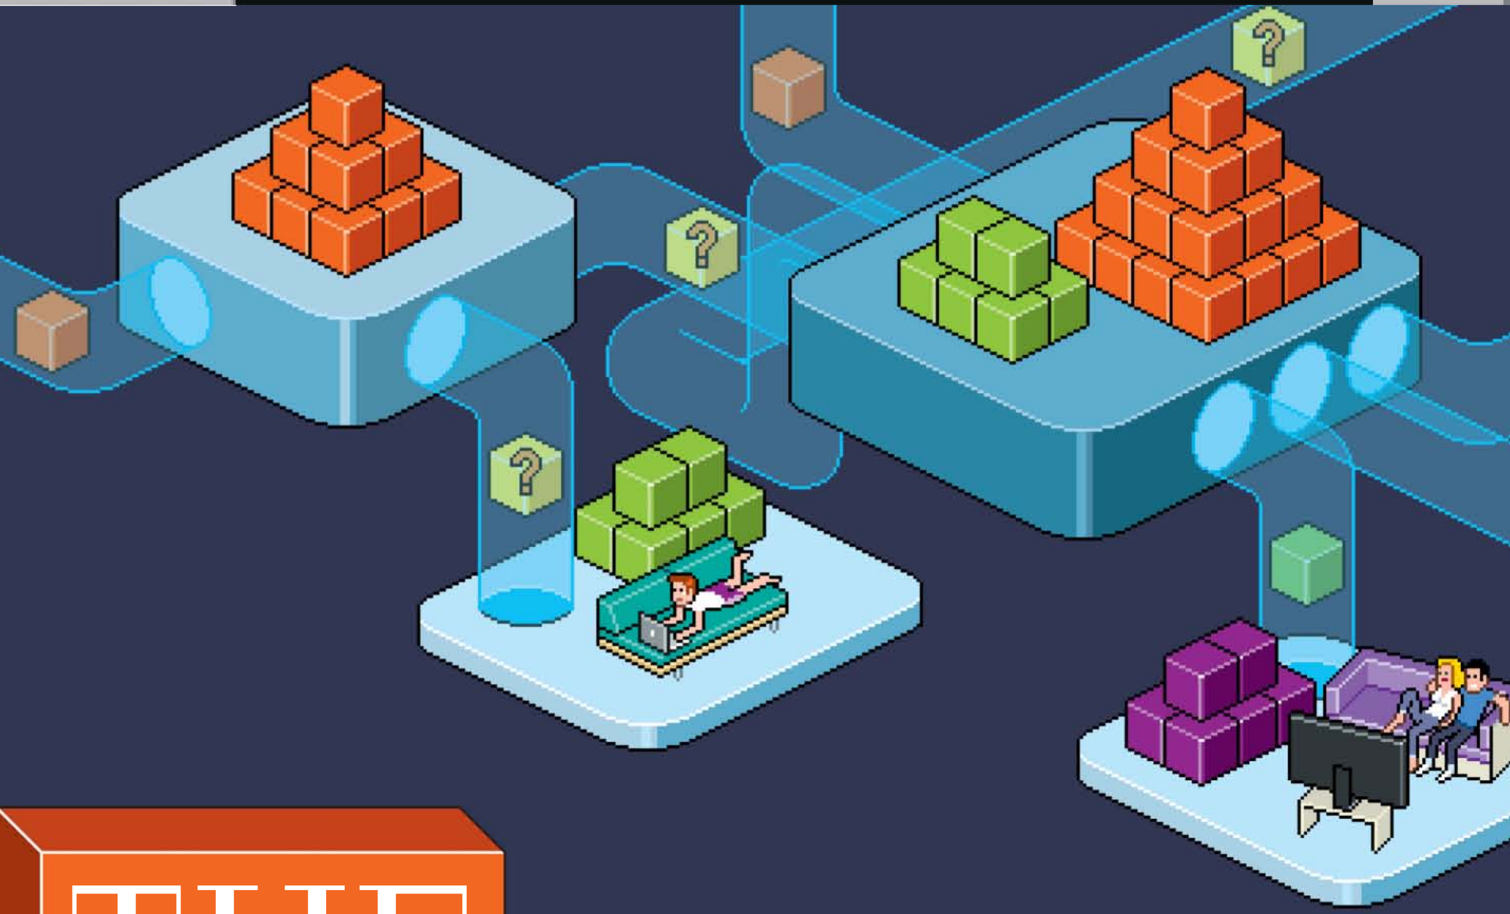

# THE PACKET PROTECTOR

By GLENN EDENS & GLENN SCOTT

**Content-centric networking** is safer and more reliable than traditional Internet protocols



ILLUSTRATION BY QuickHoney

# THE INTERNET IS MORE THAN 45 YEARS OLD,

and it's starting to show its age. To be sure, it has served us wonderfully well. Its underlying technologies delivered the World Wide Web (still a young adult at around 28 years old) and our global communications network. Even as its user base has swelled to 3.4 billion, these technologies have scaled admirably. ■ Today, however, all of those users demand a level of performance that the Internet was never designed to deliver. The authors of the original Internet protocols, who began their pioneering work in the late 1960s, designed them for a network to be used mainly for sending electronic mail from one computer to another. Now, though, people spend far more time streaming Netflix movies. Oftentimes, one piece of content must be distributed to hundreds of thousands or millions of users simultaneously–and in real time. ■ With its growth and shifts in usage, the Internet is being severely strained. That's why you're often stuck watching a "buffering" message when you're trying to watch a viral video. ■ Increasingly, network engineers find themselves scrambling for patches to improve performance or searching for ways to squeeze slightly more capacity from this creaky infrastructure. Looking ahead to what it hopes will be a golden era for the Internet, Cisco Systems expects global traffic to grow by 22 percent per year through 2020. It's hard to imagine that happening, however, if the original framework remains in place.

What we really need is an Internet that can provide more bandwidth and lower latency to many users at once–and do it securely. Along with our colleagues at the Palo Alto Research Center (PARC), in California, we've developed a better Internet architecture. We call it content-centric networking, or CCN. Our approach fundamentally changes the way information is organized and retrieved and improves network reliability, scalability, and security.

After a decade of development, we're now testing our concept: In January 2016, PARC released the open-source code for CCN software. Since then, more than 1,000 copies have been downloaded by individuals, universities, and industrial-research organizations. Companies including Alcatel-Lucent (now part of Nokia), Huawei, Intel, Panasonic, and Samsung have also had substantial R&D efforts focused on one or more aspects of CCN in recent years. In February, Cisco announced that it had acquired the CCN platform that we originally developed at PARC.

As the Internet stretches to its next billion users, all of whom will want to stream videos and upload content to their heart's content, it's time for us to rethink the way the Internet was built. Although we don't expect CCN to completely replace the Internet's protocols, we are convinced that an alternative architecture can offer better performance and security in many cases.

**THE ORIGINS OF** the modern Internet were largely patterned on technologies that support the public telephony system. Just like the telephone system, the early Internet needed a set of addresses to identify users and instructions to spell out how information should be routed and protected throughout the network. Over time, many of the same methods that had been used to make telephone calls reliable and secure, and which enabled the telephone system to scale up, were reproduced in the Internet.

That strategy wasn't perfect, though, because early telephones relied on circuit-switched networks. In such a network, users sent a stream of information over a single connection established at the start of a call. The Internet is a packet-switched network, which means it separates the many bits that make up a piece of content into smaller pieces (packets) of data. On the Internet, packets can be sent over different paths on the network and are reassembled into the original piece of content at its destination.

Each piece of online content is stored on a user's server (known as a host), often that of the content's creator. To retrieve the content, other users must navigate to that server with their requests. Then the packets must be sent back to the requestor.

Along the way, routers must constantly juggle these packets. To help routers direct packets, the Internet's

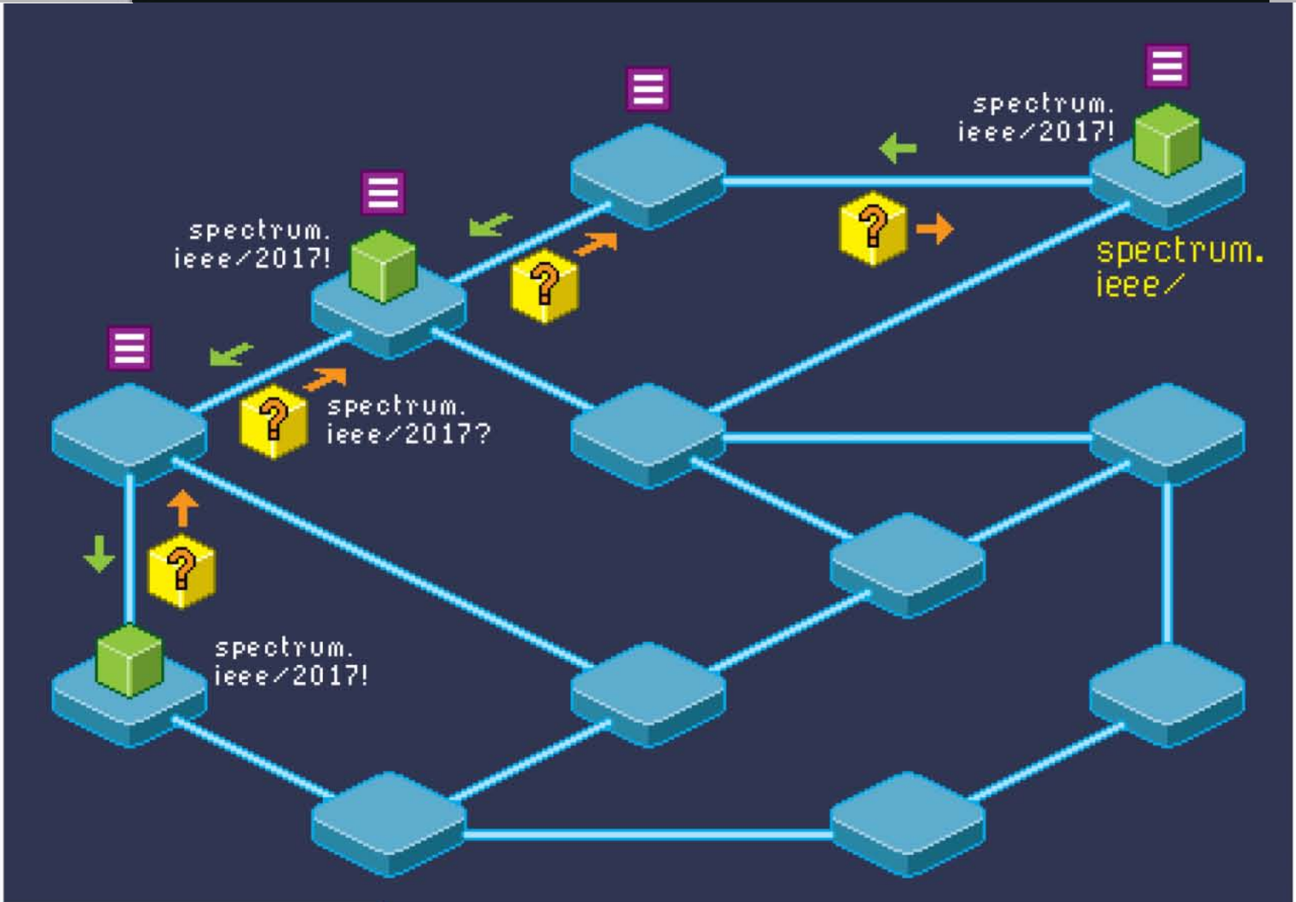


# A NEW WAY TO ROUTE

**TO RETRIEVE AN *IEEE SPECTRUM* ARTICLE** in a content-centric network, a node issues an interest packet [yellow] for content labeled with the routable prefix of spectrum.ieee/2017. Nearby nodes [blue] forward the request [orange] until it finds a content packet [green box] with a matching label. Then, the nodes return that content to the requester by retracing the same path [green arrows]. Along the way, any node may copy and store a popular piece of content in its content store [purple] for future requests.

founders came up with a clever system of assigning a unique IP address to every computer or server. The IP address is similar to a phone number. Routers move information along by reading bits on an incoming content packet that indicate its intended destination, looking up the recipient's IP address in routing tables, and forwarding the packet in the recipient's direction.

Whenever an Internet user types in a URL or email address, a Domain Name System translates those characters into the matching IP address. The DNS system is basically a phone book for the Internet. Through these servers, a URL, such as http://spectrum.ieee.org/, becomes an IP address, such as 23.197.245.16.

To enable packets to travel between hosts once a user has the correct IP address, the Internet's founders also built a shared network architecture. It has four basic layers, each with a distinct function. The first layer is the physical link over which information is sent, such as copper wires, fiber-optic cables, cellular towers, and home routers. The second and third layers are governed by common rules known as protocols that define how information is named and routed. Here, the main protocols are the Internet Protocol, which manages addresses for packets and hosts, and the Transmission Control Protocol, which describes how information is transferred. The upper layer has to do with specific applications, with more protocols that translate information for display within different Internet browsers and other programs.

These basic protocols have supported the Internet for more than four decades of growth. But they do have some shortcomings. For example, they don't always organize content in the most efficient way. Nor do they incorporate security prescriptions such as encryption by default. While these functions can be accomplished by adding more protocols, doing so can increase latency and further burden the network with additional traffic.

With CCN, we've developed a new architecture based on how information is organized within the network, rather than the IP addresses of hosts. That's why it's called content-centric networking—it's based on how content is named and stored, instead of where it is located. We've designed new protocols that can find and retrieve content from wherever it happens to be in the network at a given time and also perform many

ILLUSTRATION BY QuickHoney

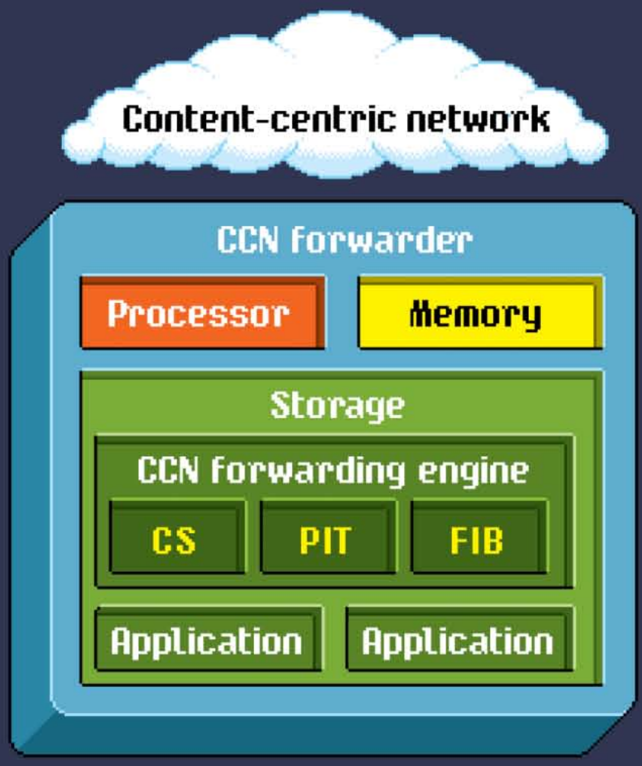


## POWER TO THE NODE

**SERVERS AND ROUTERS** can find content packets located anywhere on a CCN network by consulting two specialized tables: the pending interest table and the forwarding information base. The FIB lists where content is currently stored, while the PIT traces how past requests were forwarded. Nodes can also pluck content packets they've cached in their own content store (CS) to satisfy requests.

additional tasks that could make networks faster, more resilient, and more secure.

To understand how all of this works, we'll walk you through how a CCN network finds and retrieves a content packet for a user who is interested in reading an article or watching a video on the Web.

**IN TODAY'S INTERNET,** there is only one kind of data packet–one that carries both content and requests for content between users. But in a CCN network, there are two types: content packets and interest packets. They work together to bring information to users. Content packets are most like traditional data packets. The bits in a content packet may specify part of an ad on a Web page, a piece of a photo in an article, or the first few seconds of a video. Interest packets, on the other hand, are like golden retrievers that a user sends out onto the network to find a specific content packet and bring it back.

When you visit a Web page, your computer needs to fetch about 100 pieces of content on average. A piece of content could be a block of text, a photo, or a headline. With CCN, when you navigate to a website or click on a link, you automatically send out interest packets to specify the content you would like to retrieve. Typing in a single URL, or Web address, can trigger a user's browser to automatically send out hundreds of interest packets to search for the individual components that make up that page.

Both interest and content packets have labels, each of which is a series of bits that indicate which type of packet it is, the time it was generated, and other information. The label on a content packet also includes a name that designates what bits of content it holds, while the label on an interest packet indicates which content it wishes to find. When a user clicks on a link, for example, and generates a flurry of interest packets, the network searches for content packets with matching names to satisfy that request.

The name on a packet's label is called a uniform resource identifier (URI), and it has three main parts. The first part is a prefix that routers use to look up the general destination for a piece of content, and the second part describes the specific content the packet holds or wishes to find. The third part lists any additional information, such as when the content was created or in what order it should appear in a series.

Suppose a Web surfer's browser is using CCN to navigate to this article on *IEEE Spectrum*'s website. The network must find and deliver all the content packets that make up the complete article. To make that process easier, URIs use a hierarchical naming system to indicate which packets are needed for the page, and in what order. For example, one content packet might be named spectrum.ieee/2017/April/ver=2/chunk=9:540. In this example, spectrum.ieee is the routable prefix for the second version of the article, and the specific packet in question is the ninth packet of 540 that make up the complete article.

Once a CCN user has clicked that link or typed it in as a Web address, the user's machine dispatches an interest packet into the network in search of that content, along with other interest packets to search for packets 10 and 11. As the interest packet for number 9 travels along, each router or server it encounters must evaluate that interest packet and determine whether



ILLUSTRATION BY QuickHoney

it holds the content packet that can satisfy its request. If not, that node must figure out where in the network to forward the interest packet next.

To do all of this, every node relies on a system known as a CCN forwarder. The forwarder operates on components that are similar to what you'd find in a traditional router. A CCN forwarder requires a processor, memory, and storage to manage requests. The forwarder also runs a common software program called a forwarding engine. The forwarding engine decides where to store content, how to balance loads when traffic is heavy, and which route between two hosts is best.

The forwarding engine in a CCN network has three major components: the content store, the pending interest table, and the forwarding information base. Broadly speaking, CCN works like this: A node's forwarding engine receives interest packets and then checks to see if they are in its content store. If not, the engine next consults the pending interest table and, as a last resort, searches its forwarding information base. While it's routing information, the engine also uses algorithms to decide which content to store, or cache, for the future, and how best to deliver content to users.

To understand how that system improves on our existing Internet protocols, consider what happens when a new interest packet arrives at a node. The forwarding engine first looks for the content in the content store, which is a database that can hold thousands of content packets in its memory for quick and easy access, like the cache memory in a conventional router. But CCN has a key difference. Unlike the traditional Internet protocols, which permit content to be stored only with the original host or on a limited number of dedicated servers, CCN permits any node to copy and store any content anywhere in the network.

To build its own content store, a node can grab any packet that travels through it, keep a copy of it, and add that copy to its store to fill future requests. This ability means that content isn't stuck on the server where it was originally created. Content can move throughout the network and be stored where it's needed most, which could potentially enable faster delivery.

Currently, large companies such as Netflix pay a lot of money to store extra copies of their most popular content on content delivery networks built from regional data centers. With CCN, the entire Internet could act like one big content delivery network. Any server with available memory—not just the servers that Netflix manages—could store the first 3 seconds of a popular Netflix film. Later, we'll explain how special security precautions built into the most basic layer of CCN make it possible to securely copy and store content in this way.

To return to our example, if the forwarder finds the content it's looking for in the node's content store, the system sends that content packet back to the user through the same "face," or gateway, by which the interest packet entered the system. However, when an interest packet arrives, that node might not hold a copy of the needed content in its content store. So for its next step, the forwarding engine consults the pending interest table, a logbook that keeps a running tally of all the interest packets that have recently traveled through the node and what content they were seeking. It also notes the gateway through which each interest packet arrived and the gateway it used to forward that content along.

By checking the pending interest table (PIT) whenever a new interest packet arrives, the forwarding engine can see whether it has recently received any other interest packets for the same—or similar—content. If so, it can choose to forward the new interest packet along the exact same route. Or it can wait for that content to travel back on its return trip, make a copy, and then send it to all users who expressed interest in it.

The idea here is that these PIT records create a trail of bread crumbs for each interest packet, tracing its route through the network from node to node until it finds the content it's seeking. This is very different from conventional networks, where routers immediately "forget" information they've forwarded. Then, that forwarder consults the PIT at each node to follow the reverse path back to the original requester.

Suppose, though, that an interest packet arrives at a node and the forwarding engine can't find a copy of the requested content in its content store, nor any entry for it in the pending interest table. At this point, the node turns to the forwarding information base—its last resort when trying to satisfy a new request.

Ideally, the forwarding information base (FIB) is an index of all the URI prefixes, or routable destinations, in the entire network. When an interest packet arrives, the forwarding engine checks this index to find the requested content's general whereabouts. Then it sends the interest packet through whatever gateway will move it closer to that location and adds a new entry to the pending interest table for future reference. In reality, the FIB for the entire Internet would be too large to store at every node, so just like today's routing tables, it is distributed throughout the network.

In a traditional network, routers perform a similar search to find the IP address of the server that holds the bits of information a user wishes to retrieve and figure out which gateway to send the request through. The difference here is that with CCN, the forwarding information base finds the current location of the information itself on the network rather than the address of the server where it's stored.

**BY FOCUSING ON** the location of content rather than tracking down the address of its original host, a CCN network can be more nimble and responsive than today's networks. In fact, our studies indicate that the CCN model will outperform traditional IP-based networks in three key aspects: reliability, scalability, and security.

CCN improves reliability by allowing any content to be stored anywhere in the network. This feature is particularly useful in wireless networks at points where bit-error rates tend to be high, such as when data is transmitted from a smartphone to a cell tower, or broadcast from a Wi-Fi access point. Current Internet protocols leave error recovery to higher levels of the protocol stack. By keeping a copy of a content packet for a short while after sending it along, a CCN node reduces the upstream traffic for packets that need to be retransmitted. If a packet fails to transmit to the next node, the previous node does not need to request it again from the original host because it has its own copy on hand to retransmit.

The pending interest table can also make it easier for networks to scale. By grouping similar interest packets together, it can reduce the bandwidth needed to satisfy each request. Instead of sending a new request back to the original host for each identical interest packet that arrives, a node could satisfy all those requests for interest packets with identical copies of the content it has stored locally. If the record shows that there has been a lot of demand for a viral cat video, the algorithms within that node may prompt it to keep an extra copy of all those packets in its content store to more quickly satisfy future requests.

Boosting reliability and making it easier to scale networks are two important benefits. But to us, the most important advantage of CCN is the extra security it offers. In traditional networks, most security mechanisms focus on protecting routes over which information travels (similar to the strategies used in early circuit-switched telephone networks). In contrast, CCN protects individual packets of information, no matter where they flow.

Currently, two users can establish a secure connection through established Internet protocols. The two most common of these are HTTPS and Transport Layer Security. With HTTPS, a user's system examines a digital certificate issued by a third party, such as Symantec Corp., to verify that the other user is who she claims to be. Through TLS, users negotiate a set of cryptographic keys and encryption algorithms at the start of each session that they both use to transfer information securely to each other.

With CCN, every content packet is encrypted by default, because each content packet also comes with a digital signature to link it back to its original creator. Users can specify in their interest packets which creator they would like to retrieve content from (for example, Netflix). Once they find a content packet with that creator's matching signature, they can check that signature against a record maintained by a third party to verify that it is the correct signature for that piece of content.

With this system in place, creators can allow other users to copy and store their content, because packets will always remain encrypted and verifiable. As long as users can verify the signature, they know that the content packet originated with the creator and that users can securely access the content–a motion picture, say–from anywhere it happens to be.

This security feature brings another bit of good news: Distributed denial-of-service attacks–in which hackers send a large volume of requests to a website or server in order to crash it–are more difficult to execute in CCN. Unusual traffic patterns are easier to discern in a CCN network and can be shut down quickly. On the other hand, clever attackers may just try to figure out a way to flood the network with interest packets instead. This security challenge would have to be solved before CCN could be widely adopted.

Another significant challenge is figuring out how to integrate CCN's protocols into routers running at the speeds used on current networks. Analysts are especially concerned that routers in a CCN system would have to store rather large FIB and PIT tables to track the many moving content objects on the network, which will present major computational and memory-related challenges. However, researchers are now working on this problem at Cisco, Huawei, PARC, and Washington University in St. Louis, which have all demonstrated prototype routers supporting various elements of the CCN protocols.

Meanwhile, researchers elsewhere are drawing up their own protocols for alternative versions of CCN. These groups have built similar architectures called CCN-lite and Named Data Networking, with the support of the University of Basel, in Switzerland, and the National Science Foundation. All of these projects, along with the PARC software, are part of the broader field known as information-centric networking (ICN), which includes other research not directly related to CCN about redrawing the Internet's architecture.

One project worth mentioning is the GreenICN research program, led by Xiaoming Fu of the University of Göttingen, in Germany, in conjunction with more than a dozen universities and companies. That program explores the use of new technologies, including PARC's open-source CCN software, to create more robust networks to be rapidly deployed following a natural disaster. Fu and his colleagues took advantage of CCN's ability to operate independently of any private network and better manage limited energy resources. The group demonstrated a prototype network in early 2016.

**FROM ITS EARLIEST DAYS,** no one could have predicted what the Internet would become. Now that we have more than 45 years of knowledge about how people behave online, we want to build the next generation of the Internet to be even better than the last. We believe that early commercial deployments of CCN will be carried out, probably next year, on private networks meant for video distribution or financial transactions, where the numbers of users, forwarders, and caches can be carefully controlled.

It is hard to predict the future of today's Internet at its middle age. But consider this: In 1876, Western Union passed up the chance to purchase all of Alexander Graham Bell's telephone patents and, in an internal memo, stated that "the 'telephone' has too many shortcomings to be seriously considered as a means of communication."

The pace of technological change has accelerated dramatically since the early days of the telephone, so it's hard to imagine a 125-year run for our current Internet technologies. With CCN, the Internet could begin its evolution into a faster and more secure service that billions more users can rely on for decades to come. ■



↗ POST YOUR COMMENTS at http://spectrum.ieee.org/ccn0417

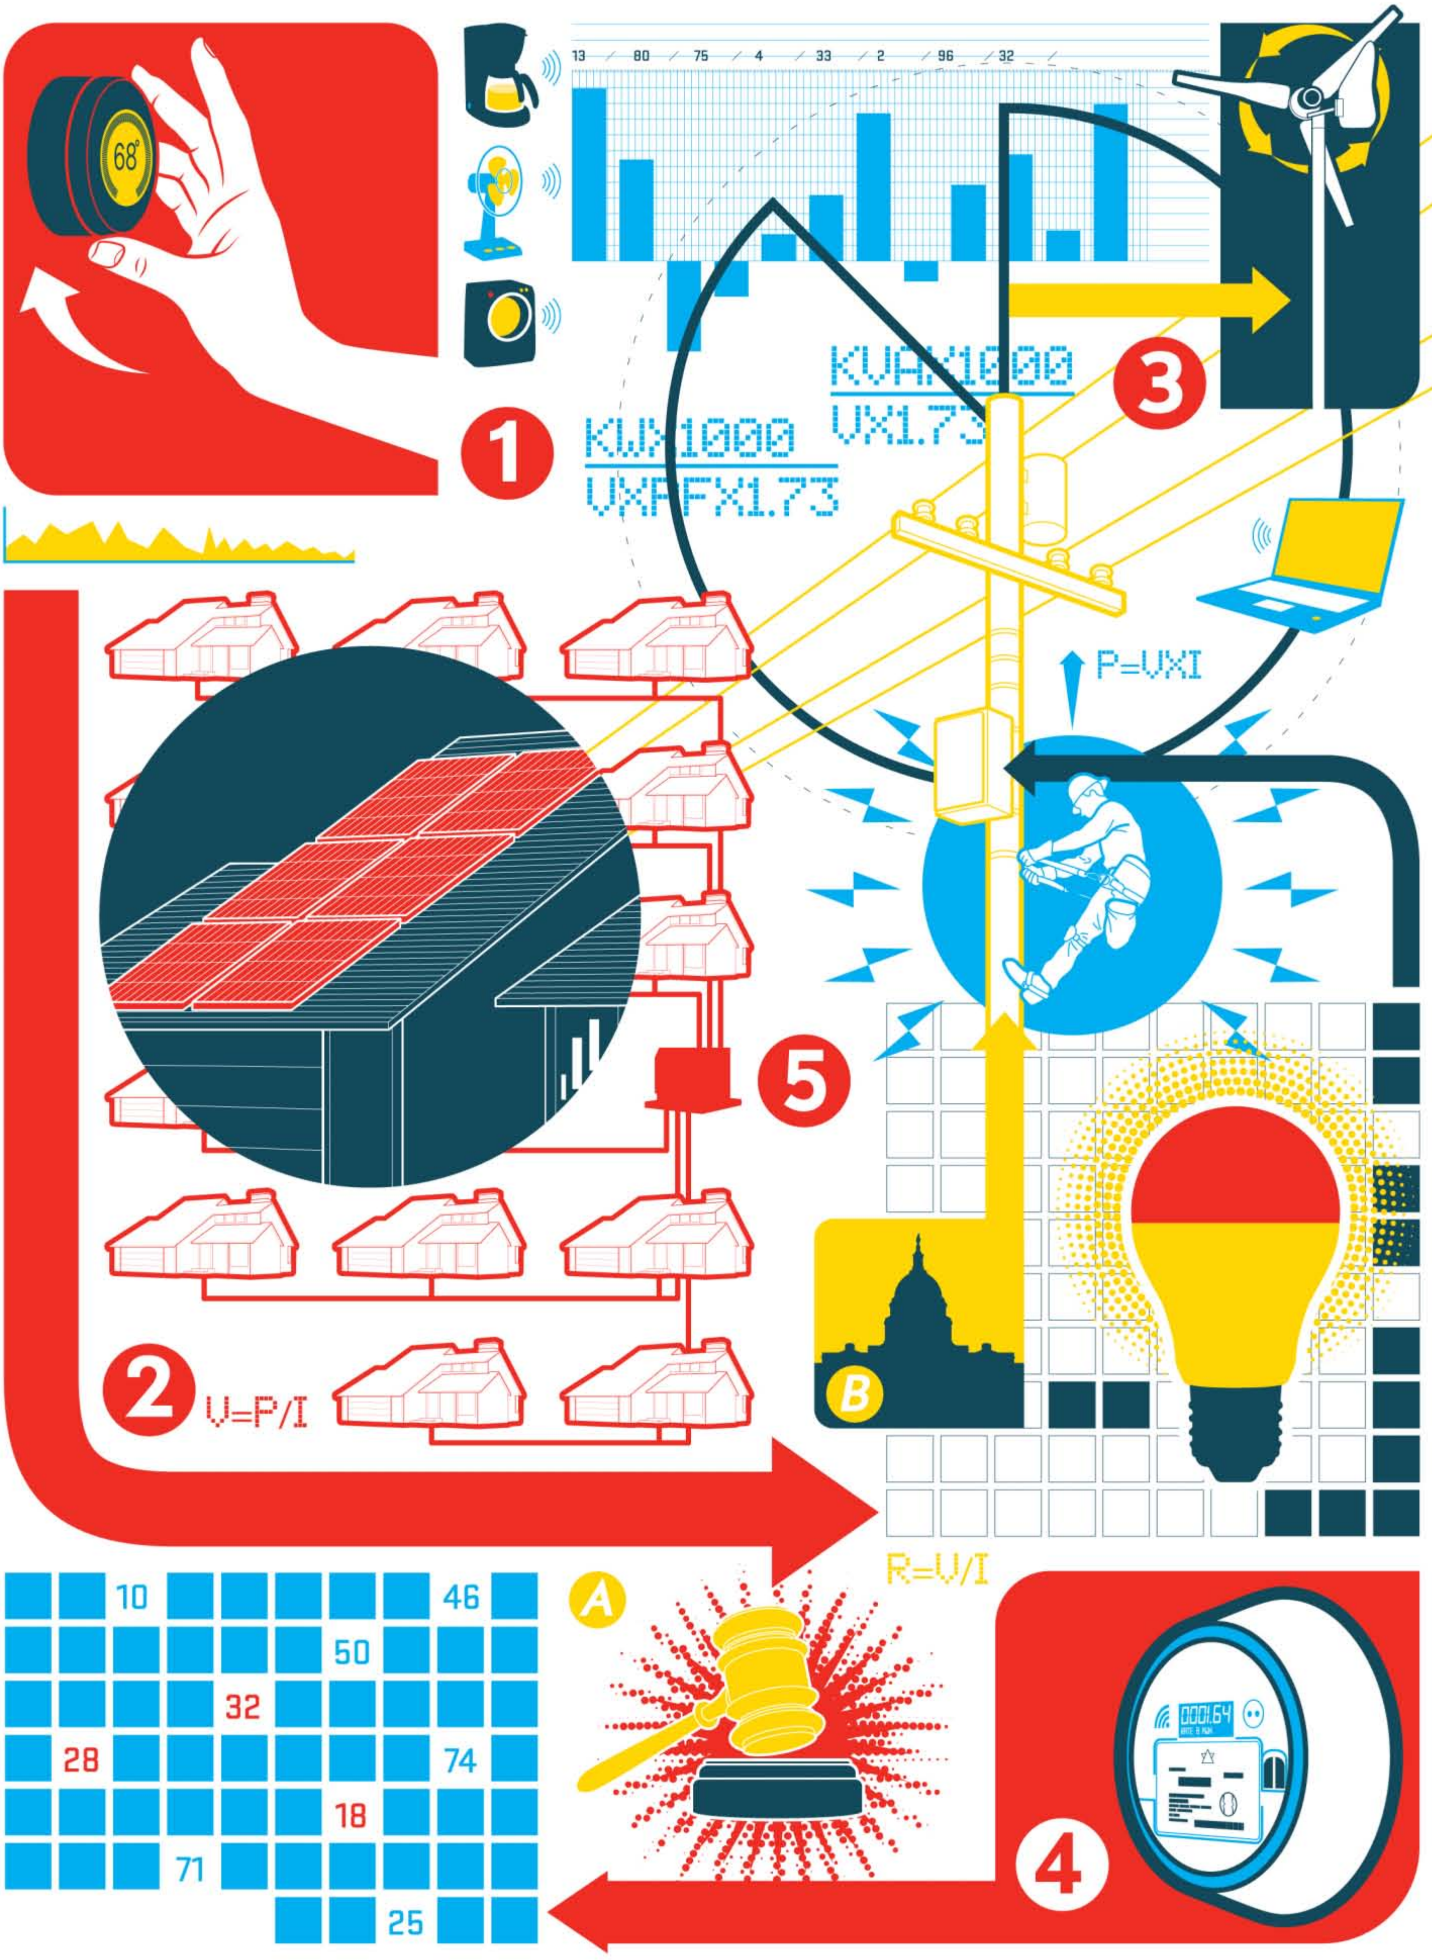
68°
13 80 75 4 33 2 96 32
KWX1000
VXPFX1.73
P=VXI
1
2
V=P/I
3
4
5
A
B
R=V/I
10
46
50
32
28
74
18
71
25
0001.64

# *The* POWER GRID *in* 2030

Rooftop solar, microgrids, and big data will revamp how we produce and consume electricity

*By* **Robert Hebner**

*Illustrations by* **MCKIBILLO**

**DEVELOPING TECHNOLOGY IS LIKE DRIVING A RACE CAR:** You push the machinery as fast as it'll go, and if you can avoid a crash, a prize awaits you at the finish line. For engineers, the reward is sometimes monetary, but more often it's the satisfaction of seeing the world become a better place. • Thanks to many such engineers driving many such race cars, a lot of progress is about to happen in an unexpected spot: the electricity sector. The power grid's interlocking technological, economic, and regulatory underpinnings were established about a century ago and have undergone only minimal disruption in the decades since. But now the industry is facing massive change. • Most observers are only vaguely aware of the magnitude of this overhaul, perhaps because it's a hard story to tell. It doesn't translate well to a set of tweets. Many people have come to think of the electric-utility business in much the same way they think of their taxes—boring, tedious, and somehow, always costing more money.

What's happening in this industry stems from technology improvements, economic forces, and evolving public priorities. As the changes dig away at the very foundation of the electricity sector, the results are likely to be anything but boring. Yet they may well cost you more money.

For about a century, affordable electrification has been based on economies of scale, with large generating plants producing hundreds or thousands of megawatts of power, which is sent to distant users through a transmission and distribution grid. Today, many developments are complicating that simple model.

At the top of the list is the availability of low-cost natural gas and solar power. Generators based on these resources can be built much closer to customers. So we are now in the early stages of an expansion of distributed generation, which is already lessening the need for costly long-distance transmission. That, in turn, is making those new sources cost competitive with giant legacy power plants.

Distributed generation has long been technically possible. What's new now is that we are nearing a tipping point, beyond which, for many applications, distributed generation will be the least costly way to provide electricity.

While it certainly helps, the declining cost of renewables and gas-fired electricity is not all that's spurring this change. To be competitive, the entire distributed system will have to work well as a whole. Quite a few technological advances are coming together to make that possible: advanced control systems; more compact, smarter, and efficient electrical inverters; smart electricity meters and the burgeoning Internet of Things; and the ever-growing ability to extract actionable information from big data.

Amid this changing scene, a picture is beginning to emerge of what a typical electrical grid may well look like in 10 or 20 years in most of the developed world. Yes, generation will be much more decentralized, and renewables such as solar and wind will proliferate. But other aspects are also shifting. For example, the distribution network—the part of the grid to which your home and business connect—will likely become more of a negotiating platform than a system that just carries electricity from place to place.

Getting to this more sophisticated grid won't be easy. Nevertheless, it's coming. What will it look like? Here is my best guess, based on my decades of experience as a government official charged with helping electric utilities get access to emerging technologies. It is the future I'm now working to help realize as an academic researcher.

**The first thing to understand** is that decentralization is going to be neither simple nor universal. In some places, decentralization will prevail, with most customers generating much of their own power, typically from solar photovoltaics. Others might use small-scale wind turbines. In regions where sunlight and wind are less plentiful, natural gas will probably predominate. Intertwined among all of those, a continuously improving version of the legacy grid will survive for decades to come.

According to the U.S. Energy Information Administration (EIA), in the first 11 months of 2016, some 48.82 million megawatt-hours of distributed solar energy were produced in the country, up 46 percent from the year before. That's still a tiny proportion, though. In 2016, about 1.4 percent of electricity in the United States came from the sun via solar panels, including both utility-scale plants and distributed ones, according to the EIA. But solar is growing fast because of its increasingly favorable economics. For example, in Chile's most recent power auction, 120 MW of solar power was the lowest-cost option, at US $29.10 per megawatt-hour.

Many analysts expect that grid-connected, distributed solar power will be fully cost competitive with conventional forms of generation by the end of this decade. In the meantime, a dizzying array of government incentives, which vary from region to region (even within one country) are helping the technology to take off.

Ultimately, the lowest-cost form of generation will dominate. But figuring out what the lowest-cost option actually is will be tricky because it will depend on both local conditions and local decisions.

For example, regulators are increasingly convinced that the burning of fossil fuels leads to significant societal costs, both from the direct exposure of those living near some power plants to their noxious emissions and from greenhouse-gas-induced climate change. Historically, these costs were difficult to quantify. So they were typically borne not by the producers or consumers of the electricity but by the victims—for example, farmers whose crops were damaged.

There is growing public interest in understanding the true cost of pollution and possibly shifting more of it to electricity producers and possibly consumers as well. Fortunately, we now have the modeling and computational capabilities to begin to put a reasonable lower limit on those costs, which gives us a defensible way to reallocate them.

Although the best strategies for reallocating those costs are still being debated, the benefits of distributed renewable generation are already very apparent—as is the feasibility. Data collected during the Pecan Street Project, funded by the U.S. Department of Energy, indicates that a house in Austin, Texas, outfitted with solar panels typically generates 4 or 5 kilowatts during the midday hours of a sunny day in summer, which exceeds the amount of power the home typically uses during such a period.

Whether or not rooftop solar makes sense for a particular homeowner, however, depends on the initial cost, maintenance costs, subsidies, the cost of grid power, and the selling price of the excess electricity generated.

The U.S. Department of Energy's SunShot initiative has as its goal making solar power cost competitive—without subsidies—by 2030. (A Chinese government agency has a similar agenda.) Specifically, SunShot's goal is to reduce the cost of distributed, residential solar power to 5 U.S. cents per kilowatt-hour by 2030; it costs about 18 cents today. Today, a 6-kW rooftop residential solar system in the United States typically costs between $15,000 and $20,000; the exact figure depends on where you live. According to data from the EIA, the average retail cost of electricity delivered by the grid in the United States is 12.5 cents per kilowatt-hour. So

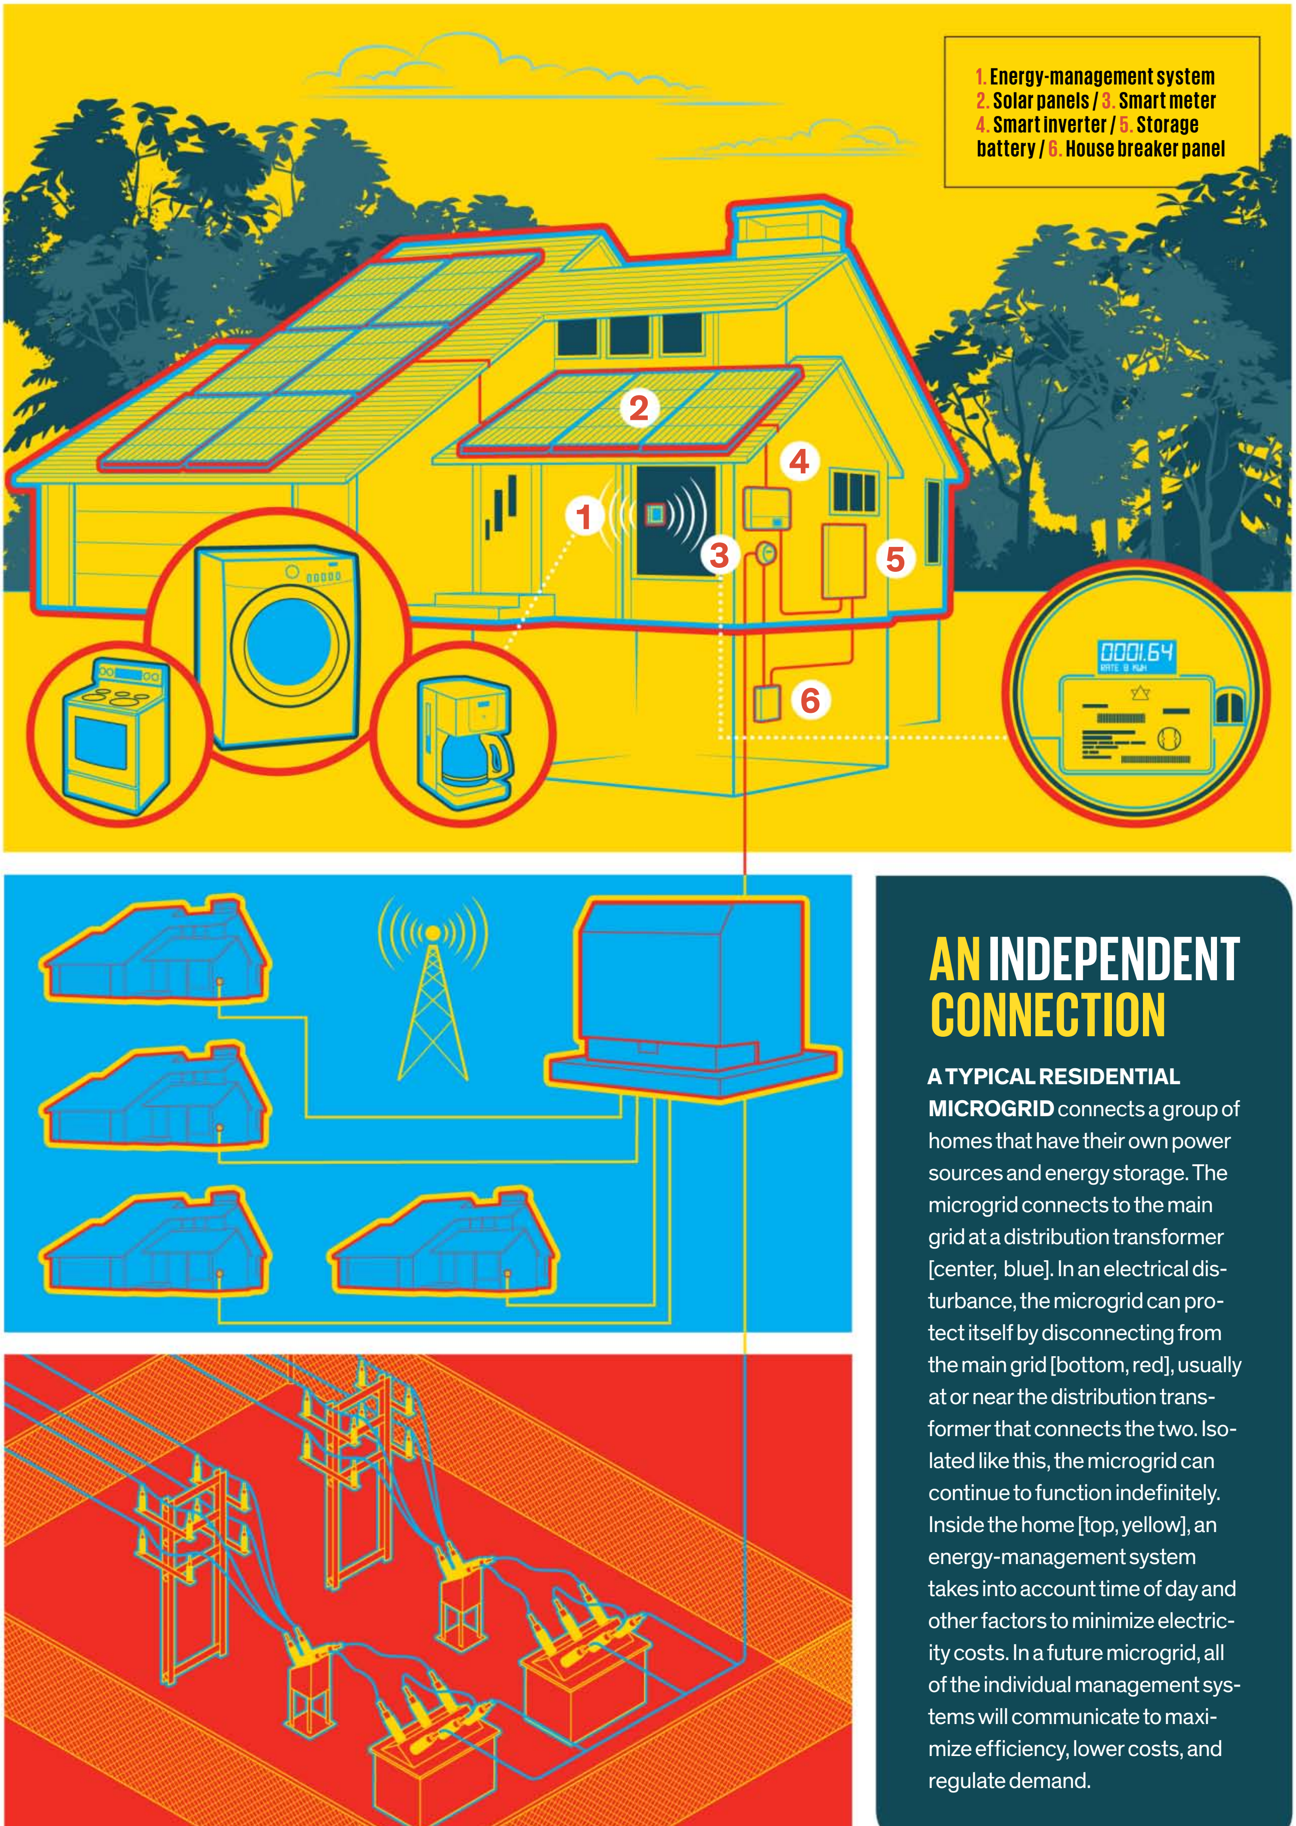


# AN INDEPENDENT CONNECTION

**A TYPICAL RESIDENTIAL MICROGRID** connects a group of homes that have their own power sources and energy storage. The microgrid connects to the main grid at a distribution transformer [center, blue]. In an electrical disturbance, the microgrid can protect itself by disconnecting from the main grid [bottom, red], usually at or near the distribution transformer that connects the two. Isolated like this, the microgrid can continue to function indefinitely. Inside the home [top, yellow], an energy-management system takes into account time of day and other factors to minimize electricity costs. In a future microgrid, all of the individual management systems will communicate to maximize efficiency, lower costs, and regulate demand.

ILLUSTRATION BY MCKIBILLO



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

at 18 cents, rooftop-generated solar is not yet, on average, competitive with grid-delivered electricity. But many governments, for example U.S. state governments, subsidize the purchase of solar-power systems to make them competitive.

Meanwhile, many utilities are experimenting with alternative-ownership options. One is community solar, in which individual consumers buy a small number of panels in a relatively large, utility-scale system. They then get monthly credits for the electricity generated without having panels on their roofs. Another experiment, being run by CPS Energy, in San Antonio, uses rooftop solar, but CPS Energy owns the equipment and pays the homeowner for the use of the roof.

One challenge with distributed solar is storage. Most solar-panel owners are using the grid as the functional equivalent of storage: They sell excess power to the grid when they can and buy back from the grid to compensate for shortfalls. This is usually the simplest and cheapest way to even out differences in production and consumption. Nevertheless, many people—most notably, Elon Musk—are betting the economics will soon favor batteries. Musk's electric-car company, Tesla, sells a battery for home use called Powerwall 2, which costs $5,500 and offers 14 kWh of storage, enough to run an average home overnight. However, adding the costs of battery storage to a solar installation to go off grid makes the costs of power significantly higher than those of ordinary electricity from the grid.

Comparing the options for expanding the use of solar power is not straightforward, however, because much depends on how the grid will evolve. For example, right now, the grid could not handle a changeover to 100 percent solar (even in areas where it would make sense, like the southwestern United States or the North African desert). The grid we have today was designed around sources whose output generally varies little from day to day. But the U.S. DOE, under its ENERGISE program, is striving to develop, by 2030, the control, protection, and other technologies needed to enable an entirely solar-powered grid.

**GLOBAL TOTAL INSTALLED WIND AND SOLAR CAPACITY**

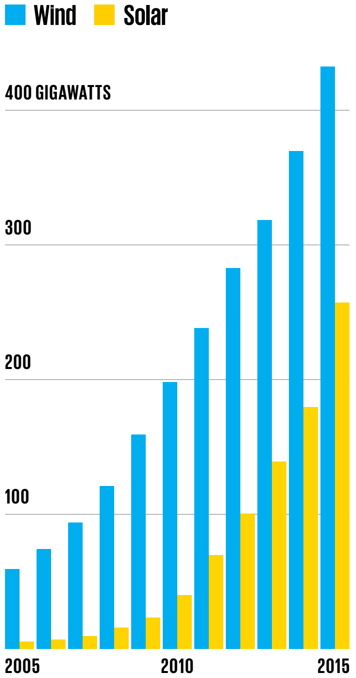


**GROWTH IN WIND AND SOLAR** has been brisk. The two together represent roughly 10 percent of the world's installed capacity but contribute only about 4 percent of production.

The grid will evolve in other ways, too, and quickly. One of the most important trends, already well under way, is the increasing use of microgrids. A microgrid is a group of connected power sources and loads. It can be as small as an individual house (often dubbed a nanogrid) or as large as a military base or college campus. Microgrids can operate indefinitely on their own and can quickly isolate themselves if a disturbance destabilizes the larger grids to which they are normally connected.

This is an important feature during both natural and man-made disasters. Consider what happened when Hurricane Ike hit the Houston-Galveston area of Texas in 2008: Blackouts were widespread, but 95 percent of the outages were caused by damage to less than 5 percent of the grid. The grid effectively distributed the effects of what was only modest equipment damage.

This isolating capability of microgrids also promises enhanced cybersecurity. That's because microgrids can help keep localized intrusions local, making the grid a much less appealing target for hackers.

When disaster strikes, whatever its cause, microgrids can limit the consequences. If it is not physically damaged, a microgrid can operate as long as it has access to a source of power, whether that's natural gas, the sun, or wind.

In the long term, with the timing depending as much on economics and regulation as technology, it is quite possible that the grid will evolve into a series of adjoining microgrids. Utilities have proposed to build such microgrid "clusters" in, among other places, Chicago, Pittsburgh, and Taiwan, a tropical island where grids are prone to storm damage. These adjoining microgrids would share power with one another and with the legacy grid to minimize energy cost and to maximize availability.

In an era of adjoining microgrids that are privately owned and operated, what will become of the utility company? There are at least two possibilities. It might simply supply power to the microgrids that need it, rather than doing that for individual customers. Or it might manage microgrids and their connections with one another and to the legacy grid. Across the United States, the concept of a utility is already being reinvented in some places as more competition is introduced. Microgrids are going to accelerate that trend.

The spread of distributed generation and the rise of microgrids will also be shaped by two other factors: the expansion of the Internet of Things and the growing influence of big data.

The Internet of Things is a boon for distributed generation because it is giving rise to industries that are mass-producing sensors, microcontrollers, software, and other gear that will be easily and cheaply adaptable for use in future, data-driven grids and microgrids. How will these things be used? Imagine a residential solar-power system of the near future. It will have "customer equipment"—solar panels, a smart inverter, and a storage battery and systems to manage loads dynamically. From time to time, the power output of that installation will be lower than usual, because of, say, a heavily overcast day.

But it would be easy to design a control system, based on readily available IoT components, that could communicate with similar systems in surrounding houses. These systems would work together, for example, to turn air conditioners on or off ahead of or behind schedule, or alter their thermostats by half a degree, to accommodate intermittent, unexpected shortfalls in capacity. What would enable this plan to work is the fact that most modern homes are well insulated, so it takes time before the internal temperature changes enough to trigger the HVAC system. The reason why homes would be grouped together in this scheme is that it would make the task easier: In the group, some homeowners would be willing to sacrifice a lot of comfort, some less. But the power needs of the group of houses would be relatively predictable and manageable, from the utility's standpoint.

Most consumers do not want to make frequent and detailed decisions on energy use. So imagine a device—let's call it an energy thermostat—that permits you to set a range of comfortable temperatures, rather than entering a single one. The wider you set the range, the less you'll pay for power. The grid or microgrid operator would use the range—yours and everybody else's—to dynamically match supply and demand on a minute-by-minute basis. On a hot afternoon, with demand at its peak, the temperature in your home would be at the top of the range.

Electric utilities will also begin making greater and much more effective use of big data. Utilities have been using data since the very beginning: When Thomas Edison opened the Pearl Street power station, in New York City in 1882, it had indicator lights to show when the load had increased or decreased enough to warrant adjustments to the dynamo producing the DC power. But that system clearly was not scalable. If a utility had to readjust its generators every time a customer came online, the industry would have died out long ago.

Having a large number of loads makes the aggregate demand predictable—and manageable. This happy condition obviously depends on there being little correlation of usage from house to house and business to business. But just suppose that at 3:00 p.m. on a hot summer day, everyone in a medium-size city turned off their air conditioners at the very same second, waited 15 minutes, and then turned them all back on again at exactly the same time. That would almost certainly cause a massive blackout.

With big-data tools, it may no longer be necessary to depend on consumers' actions being only loosely similar. It should be possible to understand how to adjust production and consumption to enhance system behavior. For example, with the energy-thermostat concept outlined above, the system operator needs to have not only the appropriate controllers but also access to real-time data to determine the risk of system failure when load-management actions are taken.

Utilities in many areas have embarked on this path using various customer incentives to permit, say, time-of-day pricing or some other form of load management by the utility rather than by the consumer. But we are now taking just baby steps. Big-data tools will soon let us take larger strides and may well one day let us run. It may be possible to use real-time operational data to optimize the performance of large sections of the grid and to predict future performance.

**Although my main goal** is to describe a hopeful vision that many of us in the utility business have for the electric grid, I would be remiss if I did not point out some of the challenges. These include financial ones, regulatory ones, and technical ones. And they come in all shapes and sizes.

One of the most fundamental is slow growth. To pay for costly system upgrades, utilities in the past would have relied heavily on growth in demand, and therefore sales. But improvements in efficiency, which consumers seek (and rightly so), have slowed growth in demand to the extent that it is now increasing at a rate lower than that of the growth in gross domestic product. And the figures are sobering: In 2014, the U.S. DOE predicted that in the period from 2012 to 2040, the demand for electricity will grow by only 0.9 percent per year. So, utilities cannot expect to fund the required system changes in the same ways as they have in the past, through growth.

Other shifts in the industry will only exacerbate these money woes. For example, in the past utilities could count on key pieces of equipment lasting a long time. But smart grids depend on electronic components, such as smart meters, controlled by software, which have shorter lifetimes and require much more frequent upgrades.

The biggest unknown is how swiftly the regulatory process can adapt. If it can't move quickly enough to keep up with the technology, expect agonizingly slow change. And what if governments try to prop up outmoded technologies with subsidies? That could drag out the process further. On the other hand, some would argue that regulators *should* slow the rate of change. Though the arguments for that are worthy of political discussion, I'm certainly not in that camp.

Historically, regulations have been driven mainly by legal and economic considerations rather than by technical ones. But now, with the pace of technology outrunning other factors, regulators in the United States and Europe are reacting to this new state of affairs in many different ways. My view is that the staffing of regulatory agencies will need to become more technically savvy if we are to navigate these turbulent waters while continuing to provide electric power with the lowest cost and highest reliability.

I'm confident that in the end, we'll have electrical grids that are less costly, more sustainable, and more user friendly than the ones that came before. The United States' National Academy of Engineering recently selected electrification as the top engineering accomplishment of the 20th century. But electrification now needs to be reengineered to meet the needs and opportunities of the 21st century. This is our chance to show that we are as good as our forebears of two, three, or four generations ago at technology, regulation, public policy, finance, and the management of change in general. And to leave to posterity a legacy as fine and enduring as the one that was left to us. ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/futuregrid0417

香港中文大學
The Chinese University of Hong Kong

Applications are invited for:-

## Faculty of Engineering
## Professors / Associate Professors / Assistant Professors (Ref. 1700004N)

The Faculty of Engineering is seeking several faculty posts at Professor / Associate Professor / Assistant Professor levels with prospect for substantiation. The professors will play a significant role in the Cyber Security Center, which will be established by the Faculty of Engineering.

Cyber security is identified as one of the Faculty's strategic research areas, to be developed by both the Department of Computer Science & Engineering and Department of Information Engineering. Talented candidates are sought to complement existing efforts and create new synergies. Candidates in the following areas are encouraged to apply:

- cryptography and computational theory in security
- network, system and software security
- data security and privacy
- computer forensic
- hardware and IoT security

Applicants should have a relevant PhD degree and a good scholarly record demonstrating potential for teaching and research excellence.

Appointments will normally be made on contract basis for up to three years initially commencing August 2017, which, subject to performance and mutual agreement, may lead to longer-term appointment or substantiation later. The exact start date can be worked out with the successful applicants.

Applications will be accepted until the posts are filled.

**Application Procedure**

Applicants please upload the full resume with a cover letter, copies of academic credentials, publication list with abstracts of selected published papers, a research plan, a teaching statement, together with names and e-mails addresses of three to five referees to whom the applicant's consent has been given for their providing reference (unless otherwise specified).

The University only accepts and considers applications submitted online for the posts above. For more information and to apply online, please visit http://career.cuhk.edu.hk.

上海科技大学
ShanghaiTech University

## TENURE-TRACK AND TENURED POSITIONS

ShanghaiTech University invites highly qualified candidates to fill multiple tenure-track/tenured faculty positions as its core founding team in the School of Information Science and Technology (SIST). We seek candidates with exceptional academic records or demonstrated strong potentials in all cutting-edge research areas of information science and technology. They must be fluent in English. English-based overseas academic training or background is highly desired. ShanghaiTech is founded as a world-class research university for training future generations of scientists, entrepreneurs, and technical leaders. Boasting a new modern campus in Zhangjiang Hightech Park of cosmopolitan Shanghai, ShanghaiTech shall trail-blaze a new education system in China. Besides establishing and maintaining a world-class research profile, faculty candidates are also expected to contribute substantially to both graduate and undergraduate educations.

**Academic Disciplines:** Candidates in all areas of information science and technology shall be considered. Our recruitment focus includes, but is not limited to: computer architecture, software engineering, database, computer security, VLSI, solid state and nano electronics, RF electronics, information and signal processing, networking, security, computational foundations, big data analytics, data mining, visualization, computer vision, bio-inspired computing systems, power electronics, power systems, machine and motor drive, power management IC as well as inter-disciplinary areas involving information science and technology.

**Compensation and Benefits:** Salary and startup funds are highly competitive, commensurate with experience and academic accomplishment. We also offer a comprehensive benefit package to employees and eligible dependents, including on-campus housing. All regular ShanghaiTech faculty members will join its new tenure-track system in accordance with international practice for progress evaluation and promotion.

**Qualifications:**

- Strong research productivity and demonstrated potentials;
- Ph.D. (Electrical Engineering, Computer Engineering, Computer Science, Statistics, Applied Math, or related field);
- A minimum relevant (including PhD) research experience of 4 years.

**Applications:** Submit (in English, PDF version) a cover letter, a 2-page research plan, a CV plus copies of 3 most significant publications, and names of three referees to: sist@shanghaitech.edu.cn. ***For more information, visit*** http://sist.shanghaitech.edu.cn/NewsDetail.asp?id=373

**Deadline:** The positions will be open until they are filled by appropriate candidates.

SOUTHERN UNIVERSITY OF SCIENCE AND TECHNOLOGY 南方科技大学

**Professor/Associate Professor/Assistant Professorship in the Department of Electrical and Electronic Engineering**

**The University**

Established in 2012, the Southern University of Science and Technology (SUSTech) is a public institution funded by the municipal of Shenzhen, a special economic zone city in China. Shenzhen is a major city located in Southern China, situated immediately north of Hong Kong Special Administrative Region. As one of China's major gateways to the world, Shenzhen is the country's fast-growing city in the past two decades. The city is the high-tech and manufacturing hub of southern China, home to the world's third-busiest container port, and the fourth-busiest airport on the Chinese mainland. A picturesque coastal city, Shenzhen is also a popular tourist destination and was named one of the world's 31 must-see tourist destinations in 2010 by The New York Times. The Southern University of Science and Technology is a pioneer in higher education reform in China. The mission of the University is to become a globally recognized institution which emphasizes academic excellence and promotes innovation, creativity and entrepreneurship. The teaching language at SUSTech is bilingual, either English or Putonghua. Set on five hundred acres of wooded landscape in the picturesque Nanshan (South Mountain) area, the new campus offers an ideal environment suitable for learning and research.

**Call for Application**

The Southern University of Science and Technology now invites applications for the faculty position in the Department of Electrical and Electronic Engineering. It is seeking to appoint a number of tenured or tenure track positions in all ranks. Candidates with research interests in all mainstream fields of electrical and electronic engineering will be considered, including but not limited to IC Design, Embedded Systems, Internet of Things, VR/AR, Signal and Information Processing, Control and Robotics, Big Data, AI, Communication/Networking, Microelectronics, and Photonics. SUSTech adopts the tenure track system, which offers the recruited faculty members a clearly defined career path. Candidates should have demonstrated excellence in research and a strong commitment to teaching. A doctoral degree is required at the time of appointment. Candidates for senior positions must have an established record of research, and a track-record in securing external funding as PI. As a State-level innovative city, Shenzhen has chosen independent innovation as the dominant strategy for its development. It is home to some of China's most successful high-tech companies, such as Huawei and Tencent. As a result, SUSTech considers entrepreneurship is one of the main directions of the university, and good starting supports will be provided for possible initiatives. SUSTech encourages candidates with intention and experience on entrepreneurship to apply.

**Terms & Applications**

To apply, please send curriculum vitae, description of research interests and statement on teaching to eehire@sustc.edu.cn. SUSTech offers competitive salaries, fringe benefits including medical insurance, retirement and housing subsidy, which are among the best in China. Salary and rank will commensurate with qualifications and experience. Candidates should also arrange for at least three letters of recommendation sending directly to the above email account. The search will continue until the position is filled. For informal discussion about the above posts, please contact Professor Xiaowei SUN, Head of Department of Electrical and Electronic Engineering, by phone 86-755-88018558 or email: **sunxw@sustc.edu.cn.**

PAST FORWARD_BY EVAN ACKERMAN

# BUILT FOR SPEED

**1963: The IBM 1460 data-processing system was an upgraded version of the** IBM 1401, one of the first transistorized computers ever sold commercially. The 1460 was twice as fast as the original, with a 6-microsecond cycle time, and it came with a high-speed printer that could output 1,100 lines per minute. Or an infinite number, if you forgot to start the stopwatch. ■

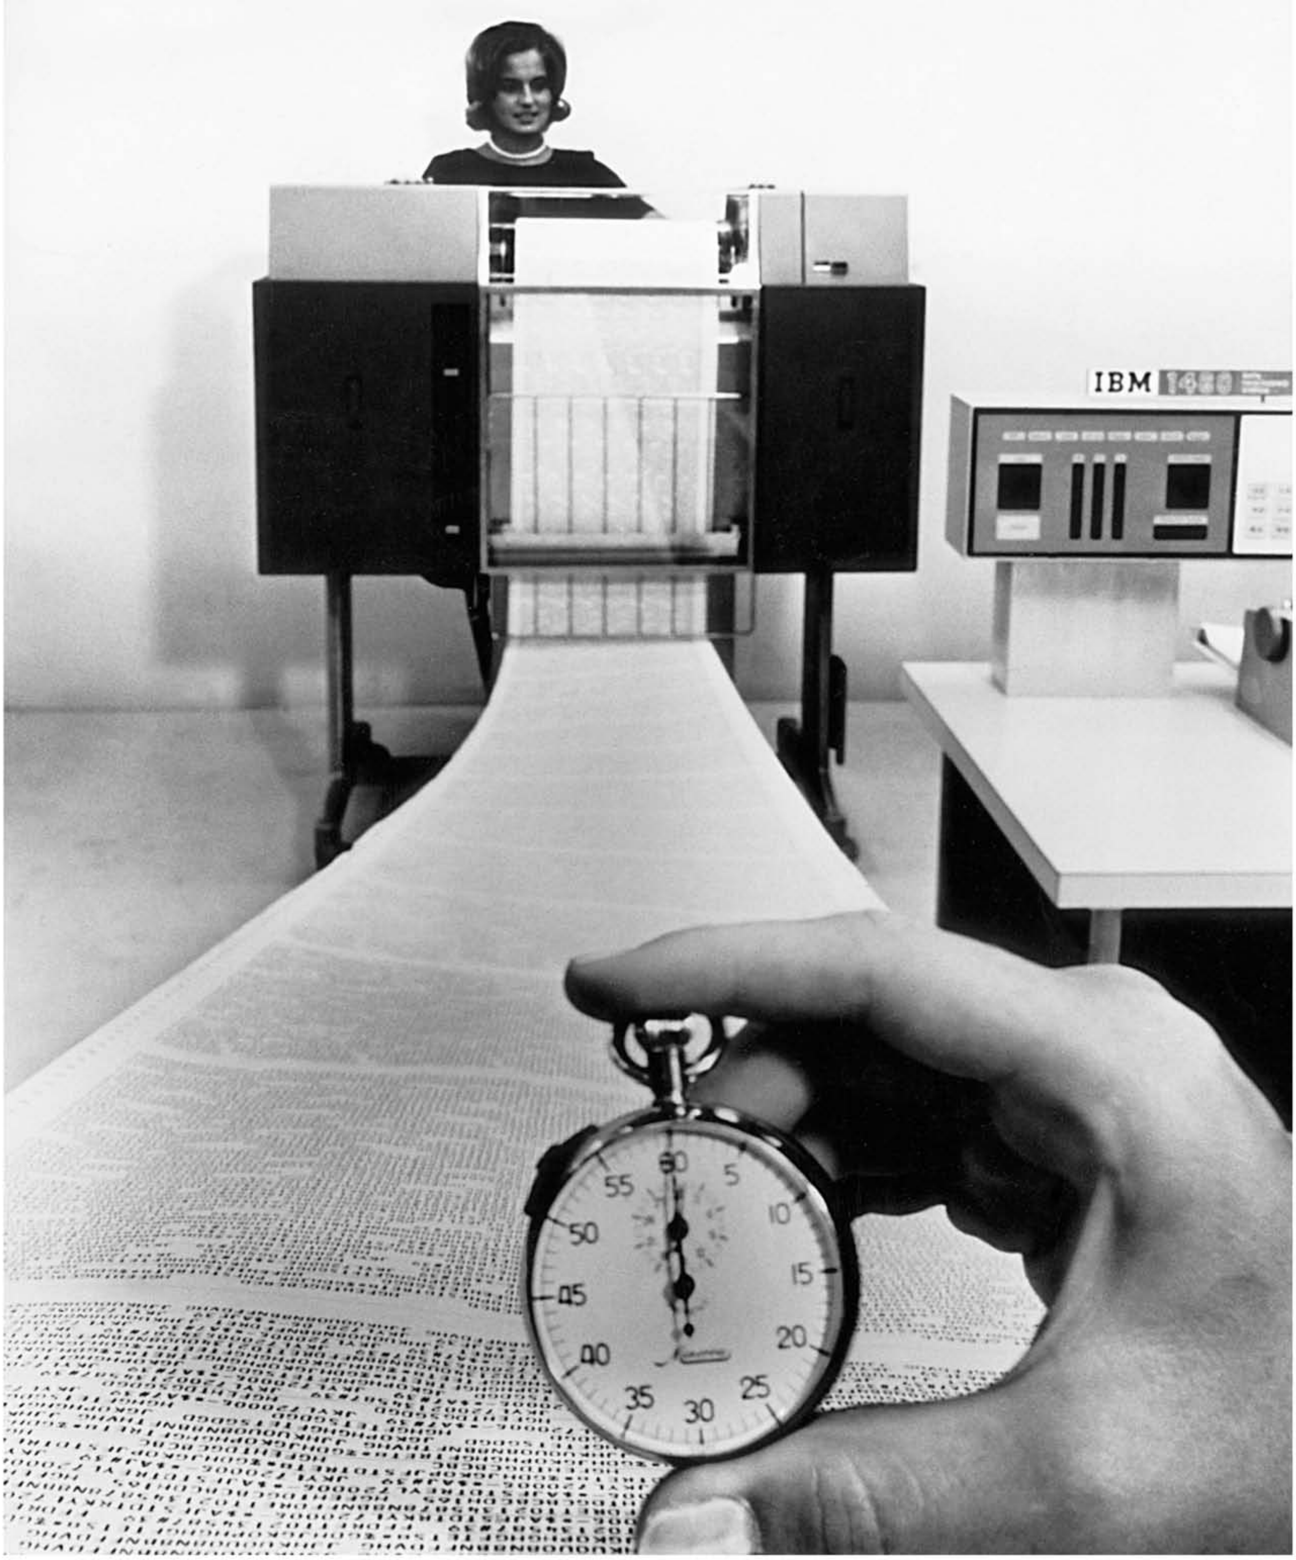


MONDADORI PORTFOLIO/GETTY IMAGES



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


MathWorks®
Accelerating the pace of engineering and science
Satellites that tango,
using automatically generated code.
That's Model-Based Design.
To achieve a breakthrough in two satellites performing autonomous formation flying and rendezvous, engineers at OHB Sweden automatically generated verification tests and the final flight code from system models. Discover Model-Based Design with MATLAB and Simulink at mathworks.com/mbd
MATLAB® & SIMULINK®
©2012 The MathWorks, Inc.
Image of satellite Tango transmitted by Mango.
© OHB Sweden