IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 12.15

EXPLORING THE MEGAVOLT FRONTIER

How researchers are testing tomorrow's supergrid P. 36

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

IEEE

COMSOL MULTIPHYSICS®

5.2 NOW AVAILABLE

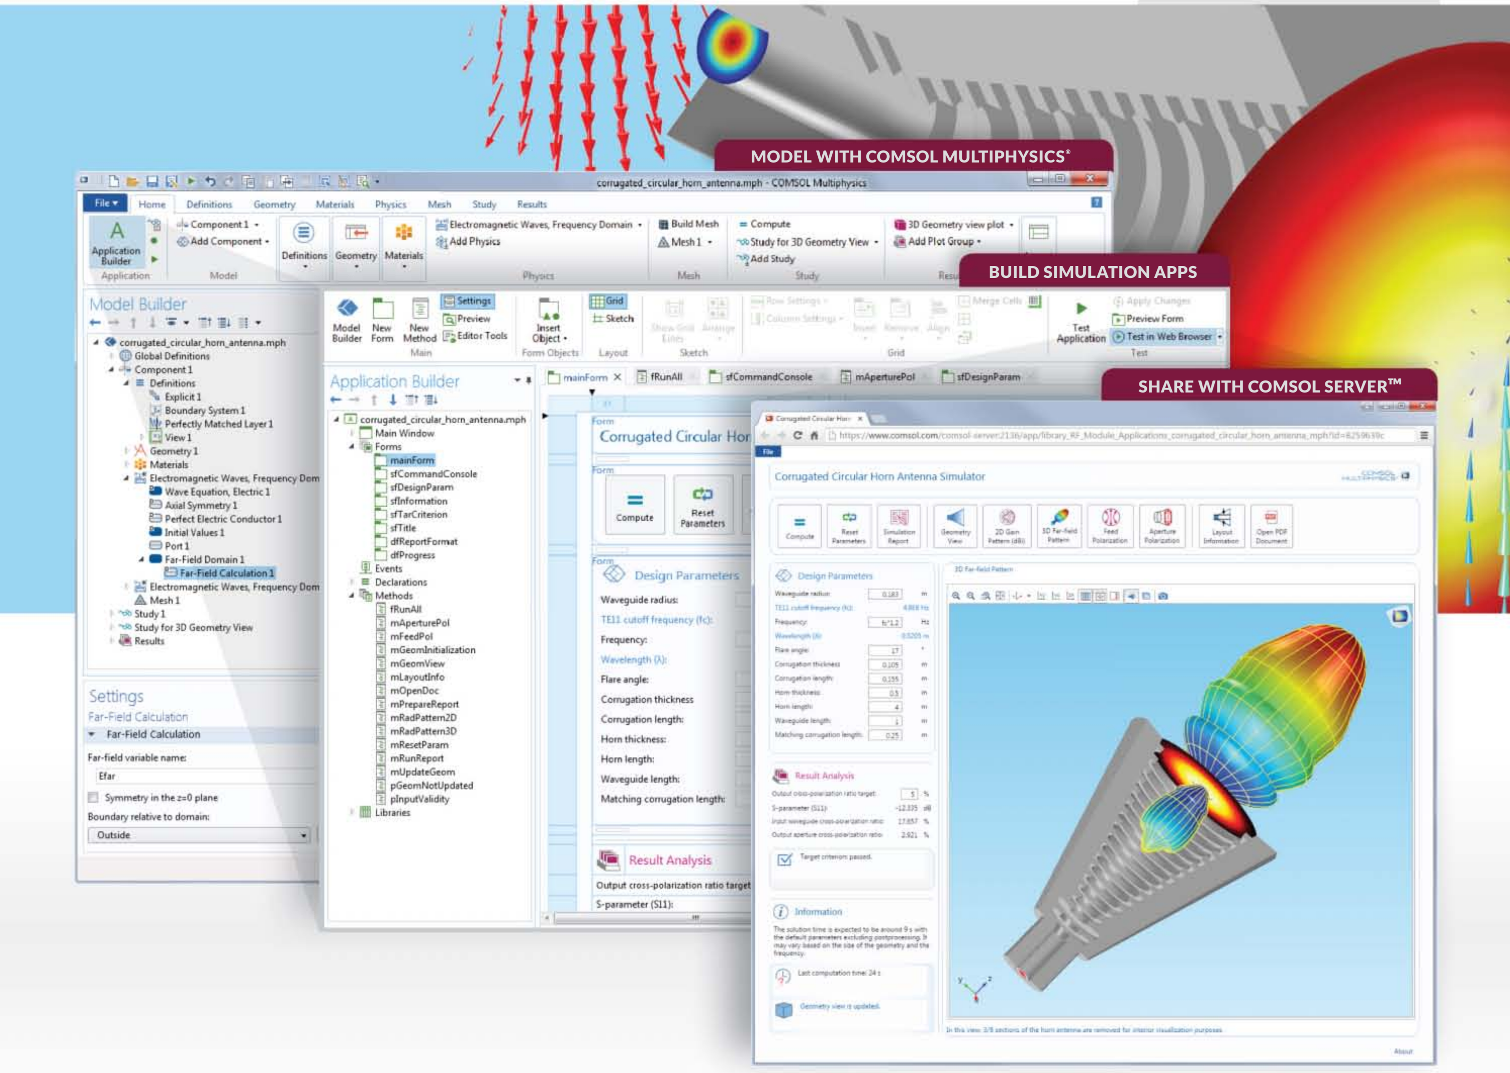

From Model to App to Server

BRING SIMULATION TO EVERYONE

COMSOL Multiphysics® and its Application Builder provide you with tools to easily design a custom interface for your multiphysics applications. Use COMSOL Server™ to distribute your apps to colleagues and customers worldwide.

To learn about simulation apps, visit **comsol.com.**

COMSOL

IEEE SPECTRUM

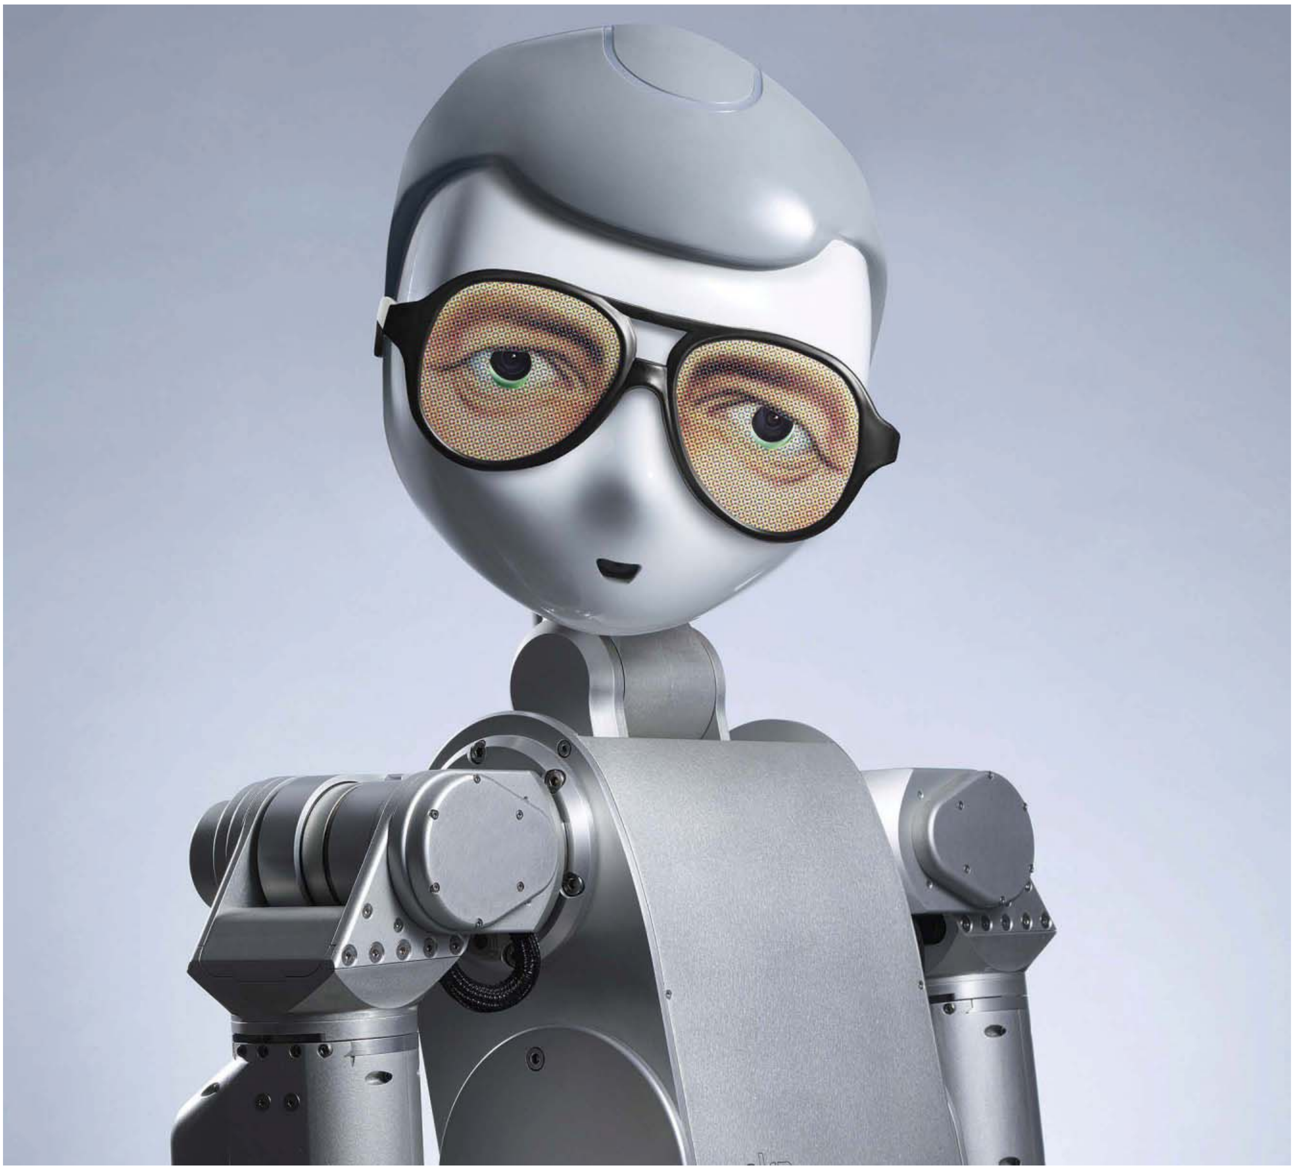

On the cover and above **Photographs for *IEEE Spectrum* by The Voorhes**

DEPARTMENTS_12.15

IEEE SPECTRUM



ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Electrically Large Aerospace Platforms—3 December
- Acoustics Simulation in Transducers, Mobile Devices, and Hearing Aids—10 December
- Modeling, Design, and Optimization of Electronic Package Designs—10 December

MASTER BOND RESOURCE LIBRARY
http://spectrum.ieee.org/static/master-bond-resource-library

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

The Institute

Available at theinstitute.ieee.org

A SMARTER HOME This month we feature the efforts of IEEE and its members to boost the IQ of today's homes. Several members in Bristol, England, retrofitted a Victorian home with sensors and other devices to keep tabs on the residents' health. IEEE Fellow Tariq Samad has spent 30 years designing gadgets to make homes more energy efficient. We also spotlight IEEE conferences, standards, and other resources that will help you get up to speed on smart-home technologies.

BUILT FOR DISASTER Students at Stevens Institute of Technology, in Hoboken, N.J., won first place in this year's U.S. Solar Decathlon for their solar-powered smart home designed to withstand flooding and hurricane-force winds.

HERZ AWARD GOES TO HASSLER Susan Hassler, *IEEE Spectrum*'s editor in chief since 1999, received this year's IEEE Eric Herz Outstanding Staff Member Award.

FROM LEFT: GETTY IMAGES; SIMPIT TECHNOLOGIES; PAUL GARNIER/GETTY IMAGES

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2015 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 52, issue No. 12, International edition. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

BACK STORY_

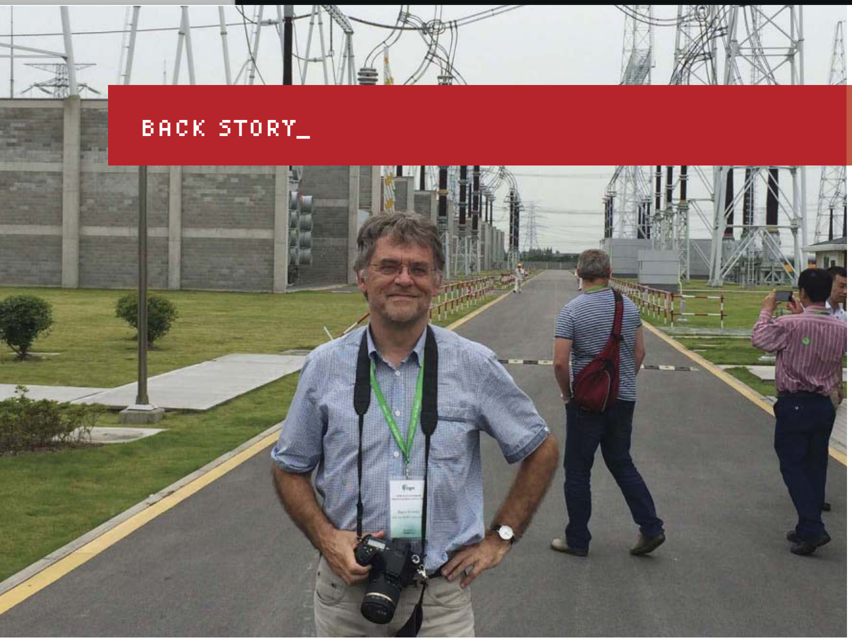

ROBERT LEROUX

Visiting Powerful Places

AS DIRECTOR OF INNOVATION AT KEMA LABORATORIES, in Arnhem, Netherlands, René Smeets is no stranger to power. The lab tests circuit breakers and transformers built for ultrahigh-voltage (UHV) transmission systems, ensuring that these components can control the titanic current flows unleashed during short circuits. KEMA's strategies to mimic those extreme conditions in the lab, which Smeets describes in the article "Safeguarding the Supergrid," give him an intimate understanding of how the components will perform when they're deployed in vast, nation-spanning transmission networks.

Sometimes, however, this expert likes to get a broader view and see the equipment he tests in its natural habitat. In September, while in Shanghai for a meeting convened by the State Grid Corp. of China, Smeets and other UHV specialists got a tour of State Grid's nearby Liantang substation [above]. The utility says this facility, which receives electricity from coal-fired power plants in the interior province of Anhui, handles more high-voltage power than any other substation in the world. The brand new 1,100-kilovolt transmission system routinely delivers 6,900 megawatts of power to Shanghai. State Grid already has plans to scale it up to an astounding 10,000 MW.

What impressed Smeets most in this superlative substation? "All the components used in the station were Chinese-made," he says. This fact didn't surprise him, as KEMA Labs has tested many pieces of Chinese equipment in recent years. But it did demonstrate that China has developed a world-leading UHV industry. It's an exciting time to be a UHV tourist, Smeets says: India is now building a 1,200-kV network, which will set a new mark for the highest-voltage transmission system. "It will be the pride of India," Smeets says. He's looking forward to a tour. ■

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, And Now This... is in *IEEE Spectrum*, Vol. 52, no. 12 (INT), December 2015, p. 60, or in *IEEE Spectrum*, Vol. 52, no. 12 (NA), December 2015, p. 72.

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

Mark Anderson

Anderson, an author and journalist in Northampton, Mass., has long been interested in the stranger side of nuclear energy. In this issue, he highlights the progress that some research groups are making toward alternative designs for fusion generators [p. 9]. "Conventional fusion is trying to solve the problem with a tool kit that's sometimes as old as fusion science itself," Anderson says. New materials and methods could "spark big innovation."

Carter M. Armstrong

Armstrong, an IEEE Fellow, is vice president of engineering at the Electron Devices Division of L-3 Communications, in San Carlos, Calif. His September 2012 *IEEE Spectrum* article "The Truth About Terahertz" dissected the hype surrounding T-ray devices. In this issue, he argues that, contrary to popular belief, vacuum tubes continue to outperform solid-state devices in select arenas [p. 26]. "We still depend on them like crazy," Armstrong notes. "And yet funders are letting vacuum electronics die on the vine."

Ralph Etienne-Cummings

Etienne-Cummings and coauthors Christoph Posch and Ryad Benosman discuss neuromorphic vision systems in "Giving Machines Humanlike Eyes" [p. 42]. Etienne-Cummings works at Johns Hopkins University; Posch and Benosman are at the Institute of Vision, in Paris. The trio's joint research grew out of their frequent meetings at a summer workshop on neuromorphic engineering in Telluride, Colo. "A lot of our collaboration came out of that workshop," says Etienne-Cummings. "And a lot of mountain biking."

Robin Finlay

Finlay and her husband, Adam Voorhes, are a photography team in Austin, Texas. They created three images for this issue: the knowing robots for "Giving Machines Humanlike Eyes" [p. 42] and the dramatic cover shot of an exploding transformer tower. To produce the explosion, the couple filled small balloons with propane and acetylene and then set them alight. "We found that a balloon about the size of a large grapefruit was sufficient to get a nice ball of fire without removing anyone's eyebrows," Finlay says.

Ioan Stefanovici

Stefanovici and coauthor Andy Hwang are computer science grad students at the University of Toronto, studying under third coauthor Bianca Schroeder. In "Battling Borked Bits" [p. 32], they describe their efforts to understand errors in dynamic random-access memory. "It's a fairly rare phenomenon," says Stefanovici. "You buy DRAM and assume it works." But as the authors show in their article, sometimes that's a faulty assumption, particularly for a giant data center or supercomputer.

12.15

PHOTO-ILLUSTRATIONS BY **Gluekit**

MARCEL STEGER/CORBIS

Teaching Children to Talk, Not Text

Face time promotes empathy, curiosity, and learning

MY CHILDREN, YOUNG AMERICAN TEENAGERS, don't have cellphones yet. They think I'm monstrously out of touch with the century into which they were born. I think I'm conflicted and all over the place on this issue. As the editor of *IEEE Spectrum*, I could be showering them with all the latest and greatest devices. But I'm not, entirely.

I'm uncomfortable with their technology use and look for ways to minimize it. When I visit their school and see them hunched over their laptops typing while the teacher waves her hand at the classroom smart board, I feel disconcerted, just as I do when I'm in an adult meeting and no one is making eye contact, instead looking down, or perhaps under the table, at a glowing screen.

I don't feel any need for them to be texting from the dugout during baseball practice. Watching their friends do it gives me the same unnerved feeling I get in an elevator full of people communicating with everyone but the people they are standing next to. In addition to screen time, I have rules about headphone time and about Internet usage.

Yet I'm guilty of all the same behavior, and more: checking e-mail while standing on the lip of the Grand Canyon, texting from a friend's memorial service, scanning my phone for information about the chronology of geological formations while helping my son with his homework. I'm sure you can think of your own examples.

And now I know why I experience such cognitive dissonance about the technology I also promote and cherish. In her brilliant new book, *Reclaiming Conversation: The Power of Talk in the Digital Age* (Penguin Press), Sherry Turkle has distilled and explained the ambivalence many of us feel about being, along with our families and friends and colleagues, always on.

Turkle is Abby Rockefeller Mauzé Professor of the Social Studies of Science and Technology in the Program in Science, Technology, and Society at MIT and the founder and current director of the MIT Initiative on Technology and Self. A clinical psychologist and devoted student of the impact of technology on human behavior, she has, for decades, chronicled the psychological impact our connected devices have on us as individuals and communities and cultures.

In *Reclaiming Conversation*, her thesis is that our phones and devices are changing not only what we do but also who we are, and not for the better. Based on her fieldwork, it seems that many of us are using our phones ostensibly to keep in touch but also to keep the messiness of actual human relationships and the physical world around us at a safe distance, at arm's length. Living in these self-created vacuums is zapping our ability to experience our unique capacity for empathy and intuition, for curiosity and imagination, the remarkable human characteristics that distinguish us from our very clever devices.

Turkle's solutions don't involve going all Thoreau and tossing our pocket screens away. Some involve ingenious engineering. Why not build a phone designed to do a task and release us, she says. What if companies measured device success not as a function of lots of time spent on them but as a function of time well spent?

The bulk of her correctives center on reclaiming face time from screen time, and reprioritizing the value of in-person conversation and active listening over texting and Facebooking, Snapchatting and tweeting.

If we have these conversations with our children now, in person, perhaps they'll remember that technologies wield power and change us, really change us, as augmented reality, virtual reality, maybe even brain prosthetics, push into their world.

So I think I'll read them snippets of Turkle's book. Or listen to it with them in the car. They'll roll their eyes impatiently while we try to talk about it at dinner. But *Reclaiming Conversations* contains antidotes to the tech-user dilemmas that will confront them as they're raising their own spawn. —SUSAN HASSLER

NEWS

US $2,999,999,000: APPROXIMATE COST DIFFERENCE BETWEEN SEQUENCING A HUMAN GENOME IN 2000 AND IN 2015

GETTY IMAGES

DECODING A BABY'S GENOME IN 26 HOURS

A specialty processor diagnoses critically ill infants in record time

A BETTER BABY: Quickly scanning a critically ill infant's genome could help him survive.

In intensive care units for newborn babies, genetic disorders are the leading cause of death, so diagnosing the problem quickly is paramount. Now, in a record-breaking 26 hours, pediatricians can scan and analyze the entire genome of a critically ill infant, thanks largely to a hardware system designed to handle the big data of genetics. In a recent study published in *Genome Medicine*, the unit reduced the time for the key analysis step from 15 hours to a mere 40 minutes.

The Dragen Bio-IT Processor, from the California-based company Edico Genome, plugs into a server and can be integrated seamlessly into a hospital or research facility's existing workflows, says Edico CEO Pieter van Rooyen. This specialized add-on, he argues, provides compute power that would otherwise require expensive racks of servers or slow connections to the cloud. »

ISTOCKPHOTO

In a full genome scan, machines record the sequence of the 3.2 billion "letters" that make up a person's DNA and look for the roughly 5 million variations that make that person unique. The plummeting cost of such scans is helping doctors find many new uses for them–and in the process causing a new conundrum. "Genetics will be the biggest big-data problem that ever existed," says van Rooyen. Others agree: A study in *PLoS Biology* predicted that within a decade the computation demands of genetic data will trump those of all other domains, including both astronomical research and YouTube.

By getting from genome scan to diagnosis in 26 hours, researchers at Children's Mercy Hospital in Kansas City, Mo., demonstrated how speedy genomic analysis can keep up with clinical demand. For each critically ill infant, sequencing machines first did the brute-force work of recording a genome's 3.2 billion letters. Using gold-standard Illumina HiSeq machines, researchers drove down the time this step takes from 25 hours to about 18 to 21 hours.

But lead researcher Stephen Kingsmore, the pediatrician and genomics expert who led the Children's Mercy study, says the "remarkable gains" in speed came from the Dragen, which identified all the variations in the genome of each sick baby. Then the researchers used in-house software that searched through those variants and automatically flagged those associated with a disease matching the baby's symptoms. In a prior study, Kingsmore's team showed how diagnoses based on genome scans can dramatically change treatment plans: For example, one baby with liver failure received the proper surgeries and pharmaceutical treatments based on the accurate diagnosis of a rare genetic disorder and is now a healthy 2-year-old.

The Dragen system delivers speed gains thanks to both hardware architecture and software specifically designed for genomic data, says van Rooyen. The system is a board with a reconfigurable processor chip and 32 gigabytes of memory. The sequencing machines' raw data streams onto the Dragen board without caching and is distributed to the chip's four compute engines. These components work in parallel to put all the letters of the person's genome in their proper order by matching them to a reference genome, which is stored in dedicated memory on the board itself in order to reduce demands on the server. The processor pushes data though a pipeline of about 140 operations at a blistering rate of 400 megabytes per second. "Everything moves at once," van Rooyen says.

Once this ordering step is complete, the processor is completely reconfigured to tackle the step of identifying variations in the genome. (That switchover takes just 20 seconds to complete.)

SOURCES OF BIG DATA IN 2025

By 2025, the storage needs of new genomics data will far outstrip those of any other data source, according to a study by scientists at the University of Illinois at Urbana-Champaign and Cold Spring Harbor Laboratory, in New York. By that year, they predict that 100 million to 2 billion human genomes will have been sequenced.

PROJECTED ANNUAL STORAGE IN 2025

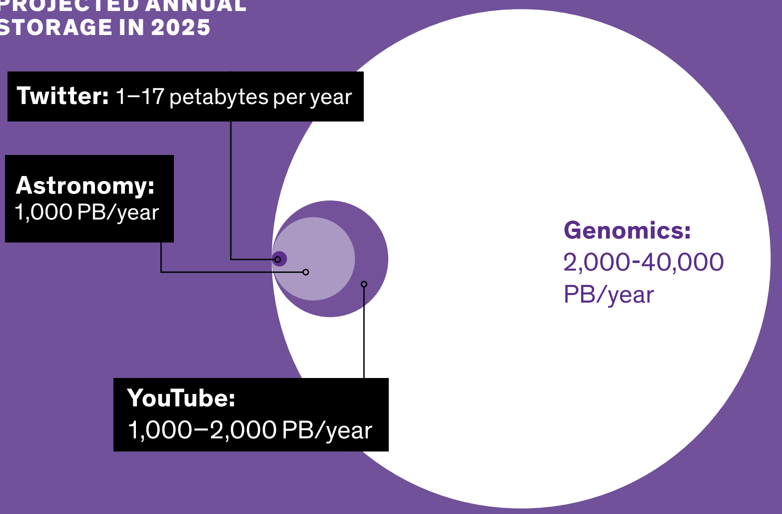

Source: "Big Data: Astronomical or Genomical?" *PLoS Biology*, 7 July 2015.

Then as the variants are identified, data describing them is compressed and sent back to the server. As most operations occur on the Dragen board, this system uses very little of the server's memory and doesn't tax its CPU, consuming about 13 to 18 percent, in Edico's estimate.

While the Dragen's statistics are impressive, not everyone is convinced that the plug-in processor offers unique benefits. One skeptic is Michael Schatz, an associate professor of quantitative biology at Cold Spring Harbor Laboratory, in New York, and coauthor of the paper that compared data sets from genomics, astronomy, and YouTube. He argues that specialized processors lock users into certain data formats and analysis methods. "They are very good at a few things," he says, "but the data keep changing and methods keep improving." Schatz thinks that doctors and researchers working in big-data genomics will be better off investing in general-purpose computer clusters, which "can easily be repurposed from one application to the next or from one data type to the next."

One way or another, doctors must start scaling up. Van Rooyen predicts that in a few years, every infant born in the developed world will have his or her genome sequenced in the hospital. "It's just a matter of time before clinical genomics will be with us everywhere," he says. "It's prudent to have the infrastructure ready." He envisions the Dragen processor outputting its analysis directly into a patient's electronic medical record, where actionable intelligence would be flagged for the physician.

Automating medicine to this degree, from genome sequencing to diagnosis, will be necessary if we want to make use of today's best genetic technologies, says Kingsmore, the pediatrician. "If we're going to scale this, it has to be on the backs of smart machines," he says. "We want to take humans out of the equation, because we're the bottleneck."

—ELIZA STRICKLAND

ALTERNATIVE FUSION PROJECTS WARM UP

Physicists pin their hopes on new magnets and fast lasers

The past few months have proved that hope for nuclear fusion as the ultimate clean and nonpolluting energy source springs eternal. One reactor plan projects a tantalizing gigawatts-per-year net energy out of its still-on-the-drawing-board idea. Another scheme uses the same reaction as the first but seeks smaller-scale reactors. A third uses the familiar "heavy hydrogen" reactions of decades past—deuterium and tritium hydrogen isotopes combining to create helium, neutrons, and energy—but relies on possibly transformative design changes enabled by using the latest superconducting magnets.

To be clear, unexpected errors or oversights could still ground one or all of these efforts. But when so much of the research world is depending on the overdue and overbudget US $20 billion ITER project, each of these efforts counteracts the monoculture mind-set in fusion research that has been the subject of some industry questioning and criticism.

"There's inertia in having the established magnetic [fusion reactors], and it's a mature technology that's being used," says Dennis Whyte, director of MIT's Plasma Science and Fusion Center. "But the new superconducting technology has improved over even the last three to four years. Even since we started the project, the capability of the technology has improved."

The new reactor design Whyte now touts is a familiar torus-shaped tokamak—much like the current generation of doughnut-shaped plasma fusion reactors such as NSTX in Princeton, N.J.; ST25 in Milton Park, England; and the projected ITER project in France. But unlike the existing tokamaks, this one is smaller and has more powerful magnets. The magnets' power comes from a second-generation breed of superconductor made with rare-earth barium copper oxides. REBCOs, as they're called, can continue functioning at higher

CLEAN LINES: The pure tungsten electrodes inside LPP's dense-plasma-focus device resist the vaporization that would add unwanted atoms to the fusion reactions.

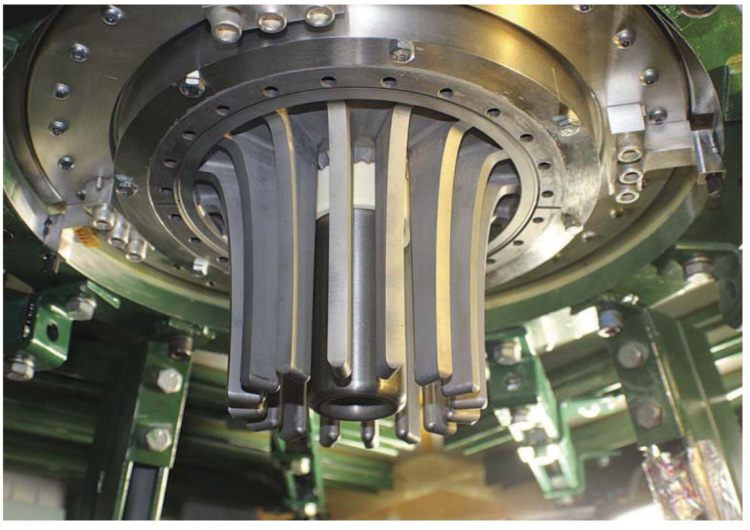

FOCUS FUSION

temperatures and magnetic-field strengths than most other superconductors today. For instance, compared with ITER, Whyte's "affordable, robust, compact" (ARC) reactor nearly doubles the magnetic field strength, from ITER's peak strength of 13 teslas to an ARC peak magnetic field of 23 T. That's important because the potential power that can be extracted from a reactor scales as the magnetic field to the fourth power.

ARC began as a class project at MIT, Whyte says. And now it's seeking partners to bring its drawing-board design to life. Plus, Whyte says, students are still adding new innovations—like high-temperature alloys, modular 3-D-printed core designs, and clever thermal cooling fins that could offer twice the cooling capability. "This is technology that is just ready to blossom," Whyte says. "This is an incredibly exciting time to be in engineering."

Lawrenceville Plasma Physics' Focus Fusion device is in the advanced research phase, says LPP president and chief scientist, Eric Lerner. It has now had 12 test runs using a heavy-hydrogen mix like the plasma soup that fuels tokamak reactors. But, says Lerner, these preliminary reactions are just to test the system. Ultimately the "dense plasma focus" device is designed to combine protons with boron atoms to generate three helium atoms plus energy—and crucially no neutrons that can contaminate other nearby atoms in the fuel or reactor and turn them radioactive.

According to the LPP website, the device consists of two coaxial cylindrical metal electrodes sandwiching a low-pressure gas. A mega-ampere-scale capacitor bank discharges millionth-of-a-second bursts of power that induce a current and a megagauss magnetic field. The device, LPP says, can get the gas up to billions of degrees Celsius for billionths of a second. The company's calculations suggest these conditions could ultimately be sufficient to spark proton-boron fusion.

Lerner says his group, based in Middlesex, N.J., is now trying to get atomic impurities out of the plasma, because fusion reactions are extremely sensitive to foreign atoms in the mix. But he says he expects to be testing his device on proton-boron mixtures within the next year.

"I find their results interesting but preliminary," says Robert Hirsch, who from 1972 through 1976 directed fusion research for the U.S. Atomic Energy Commission and Energy Research and Development Administration. Hirsch also chaired an independent review panel of the dense-plasma-focus device last year. "There is simply too much happening in these experiments to draw conclusions early on."

Another promising unorthodox fusion reactor design was announced this year. It too proposes to fuse protons and borons to create energy plus innocuous helium as its ash. But instead of large, high-strength magnets, it would use some of the shortest and most powerful laser pulses in the world today to produce short bursts of high-strength magnetic field.

Heinrich Hora, emeritus professor of theoretical physics at the University of New South Wales, in Australia, says a spate of recent experiments in Europe have led his group to conclude that an "avalanche" fusion reaction could be triggered at the trillionth-of-a-second timescales of a petawatt-scale laser pulse—whose brief bursts pack a quadrillion watts of power. If scientists could exploit the avalanche correctly, he says, it could lead to a breakthrough in proton-boron fusion.

"You put in 30 kilojoules and get 1 billion joules out," Hora says of his team's preliminary calculations. "The whole avalanche process is a state of plasma that's very much unexplored. More experiments are now necessary."

His group's conceptual reactor design was published in July in the journal *Laser and Particle Beams*. And Hora along with collaborators in the United States and Europe presented separate papers outlining their research in September at the Inertial Fusion Sciences and Applications conference in Seattle.

"How can we afford not to be doing this?" Whyte says of the smaller and cheaper unorthodox approaches to finding nuclear energy's holy grail. "Fusion is hard. But it's worthwhile."
—MARK ANDERSON

FOUR NEW WAYS TO CHILL CHIPS

Engineers look beyond traditional chip design for cooling fixes

Things are getting a bit too hot in the microprocessor world. Again.

Moore's Law has always come with the caveat that more transistors, switched at a higher frequency, means more heat. Over the years, chipmakers have used tricks like throttling back clock speeds and putting multiple microprocessor cores on a chip to spread out the heat.

But heat continues to stifle chip performance. Hot spots on today's processors can reach power densities of 1 kilowatt per square centimeter, much higher than the heat inside a rocket nozzle. A growing fraction of transistors on advanced microprocessors are not even operated at any one time because they would generate too much heat, says Avram Bar-Cohen, a program manager at the microsystems technology office of DARPA. "As we put more and more transistors on them, this 'dark silicon' fraction has gone from 10 to 20 percent, in some cases more," he says.

There's only so much that processor designers can do to keep chips from generating too much heat, and it's time for some new ways to get that heat out.

The conventional approach to dissipating the heat is to attach the silicon die to a copper or aluminum block carved with elaborate fins and ridges. A plastic fan blows air across the metal. As you can guess, these systems can be bulky, noisy, and power hungry.

Plus, heat sinks and fans won't cut it for future processors, which will be 3-D stacks of ICs. Such layering can trap heat between chips, making getting rid of it even harder.

Researchers are exploring better ways to cool chips and gadgets, either by redesigning tried-and-tested methods or drastically overhauling them. The challenge for all of these technologies is whether the semiconductor industry will be ready to integrate them, says Ravi Prasher, a director in the energy technologies division at Lawrence Berkeley National Laboratory, in California, who has studied various chip-cooling technologies.

Here's a look at some methods that could beat the heat. —PRACHI PATEL

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

KEEPING IT COOL: Fujitsu's superslim refrigerator is designed to suck heat out of a smartphone [top]. CoolChip has combined the heat sink and fan for added efficiency [center]. Water flows through microchannels on a chip [bottom].

FROM TOP: FUJITSU; EVAN ACKERMAN; ROB FELT/GEORGIA TECH

TINY WATER PIPES

Instead of transferring heat through many millimeters of heat-conducting layers and then into the air, one idea is to put something cold mere micrometers away from the transistors generating heat. This can be done by carving microfluidic channels into a chip or substrate and pumping a fluid coolant through them.

In September 2015, engineers at Georgia Tech demonstrated the first on-chip embedded liquid cooling system on a commercial field-programmable gate array. (FPGAs contain logic blocks that a customer can configure using software to form any computational circuit, with such varied applications as defense, astronomy, and medical imaging.) Electrical and computer engineering professor Muhannad Bakir and his colleagues used conventional microfabrication techniques to etch 200-micrometer-tall channels into the silicon on the back of the chip and connected the channels to external water tubes; the pump and reservoir will eventually be integrated with the chip's packaging. The system slashed the FPGA's operating temperature by more than 60 percent compared with that of an air-cooled device. "The technology is compatible with all silicon chips and packaging technologies," Bakir says.

Another key advantage of liquid cooling: It could be embedded between 3-D-stacked high-power chips. To show this potential, Bakir's team embedded interconnects into the cooling channels that could be used to link stacked chips.

THE FRIDGE

If water isn't enough to do the job, some research teams are looking at evaporating refrigerants inside microfluidic channels. By turning into vapor when heated, a coolant can suck out even more heat than a circulating liquid, says Bruno Michel, who leads the advanced thermal packaging group at IBM Zurich Research Laboratory. The vapor turns back to liquid in a condenser outside the chip and returns to repeat the process. Michel says that this two-phase cooling approach, which utilizes low pressure to boost vaporization, should allow them to reach DARPA's target of cooling chip hot spots by 1 kW/cm^2. The tricky part is to regulate the pressure of the system so that the coolant never fully vaporizes. "You have to control the amount of boiling and the vapor quality," he says.

Simple versions of such two-phase cooling are making their way into smartphones. In March 2015, Fujitsu said it had developed a refrigeration system less than a millimeter thick.

THE FAN OF FANS

CoolChip Technologies, a startup spun out of MIT, has improved the fan-cooled heat-sink design of old by rolling the two into one. The combination device, which the company calls a kinetic cooling engine, is a flat, circular piece of aluminum carved with fins that spins like a fan. It is half the size of conventional cooling fans, emits much less noise, and can remove 50 percent more heat, according to the company. Microsoft has hired CoolChip to make a system for the Xbox One game console. By early 2016, CoolChip's technology should be in tens of thousands of gaming-unit cooling systems made by the Taiwanese company Cooler Master.

CARBON NANOTUBES

Thermal greases and polymers are used right now as interface materials for transferring heat from the silicon die to the heat sink. Carbon nanotubes would be excellent substitutes. They have some of the highest-known heat conductivity and are flexible; a rigid material would expand and potentially damage the chip or its package.

The problem is that nanotubes are very stable and don't easily form chemical bonds that connect them to other parts of the chip package. So "it's difficult to have good heat transfer between them and heat sources or sinks on either side," says D. Frank Ogletree, a physicist at Lawrence Berkeley National Laboratory. Ogletree has used organic molecules to form covalent bonds between nanotubes and metal, improving heat flow sixfold. A big remaining hurdle is that only about 5 percent of nanotubes, which are grown in vertical arrays on silicon, touch the metal surface. "If someone could solve this challenge, then it will really be a compelling technology," says Berkeley Lab's Prasher.

4G FOR ALL

Software-defined radio will let communities build open-source 4G networks

Conventional cellphone networks aren't for everybody. In the state of Oaxaca, Mexico, rural communities high in the Sierra Juárez mountains have been building their own 2G cellular networks using mostly open-source and low-cost software and hardware. These communities have liberated themselves from dependence on large commercial networks that had neglected their sparsely populated region.

They can do this in part thanks to software able to take over functions that once required hardware and to a parallel movement to make such software open source and almost free. So-called software-defined radios make it easier for tinkerers and researchers to prototype their own components for a 4G LTE technology cellular network. While a full open-source version of 4G LTE is not yet complete, a proposed agreement between some academic and industrial partners would ensure that open versions in the works will be interoperable with commercial systems.

It has taken decades to get this far. Telecommunications engineer Raymond Knopp of Eurecom, a research institute at Sophia Antipolis, in the South of France, has been experimenting with open versions of cellular radio since 1997, when he and others began playing with an open version of today's third-generation standard. "It was something that looked a lot like 3G," he recalls, "but it couldn't interconnect to it."

Starting around 2003, Knopp and his colleagues signed a series of contracts with Airbus, Alcatel-Lucent, National Instruments, and the French space agency to develop components of LTE and WiMax networks for hard-to-reach parts of France, or for quick deployment in emergency situations. To simplify the intellectual property issues, they developed the software using open-source licenses. "Open source was a way for us not having to hide anything, to avoid having people sign nondisclosure agreements left and right," he says. By 2010, they realized that they had most of the components for a 4G LTE network, the fastest version of 4G. They call it OpenAirInterface.

IN THE OPEN: Once cellular networks dotted Oaxaca's mountain towns, residents soon began asking for mobile broadband service.

FAIRWAVES

Network engineer Paul Sutton of Trinity College Dublin says that OpenAirInterface is probably the most complete open 4G project so far. "LTE is a serious step change in computational complexity," he says. But because programmers outside industry now have access to high-end computing power, it is more realistic for them to write their own experimental code for managing cellular networks. A Google software engineer, Ben Wojtowicz, almost single-handedly developed openLTE, an open-source LTE software implementation, in his spare time. Similarly, developer Fabrice Bellard created his own closed-source, commercial 4G LTE base-station software.

Cellular hardware is also growing more accessible to smaller organizations. Fairwaves, in Somerville, Mass., makes the software-radio-based mobile networks used in Oaxaca. The cost is about an order of magnitude less than that of conventional mobile network operators' base stations, in part because it does away with the need for multimillion-dollar core network servers by moving the network software onto the tower. The community network operates on GSM bands, with some data transmission comparable to that of the second generation of cellular networks, or 2G. Yet the company's latest base station is "theoretically compatible" with 4G, says Fairwaves CEO Alexander Chemeris, and "there is a lot of interest in doing open-source 4G."

So far, Fairwaves and OpenAirInterface have demonstrated their work at recent mobile network conferences. And Sutton's startup, Software Radio Systems, in Cork, Ireland, will soon complete handset software and is developing its own base station. Knopp has formed an industry group called the OpenAirInterface Software Alliance, which is in discussions with the European Telecommunications Standards Institute, the organization that defines cellular network standards. The goal is to ensure that future standards–including 5G–will be compatible with open-source software. The alliance's industry partners recently agreed to a new license for collaborative, experimental coding that would allow them to claim royalties on commercial applications of their contributions to new technology platforms. The license and open software could enable small outfits such as Software Radio Systems to innovate on top of a shared infrastructure and participate in the latest telecom generation. "The number-one big challenge is to get a very large community around what we have," Knopp says.
—LUCAS LAURSEN

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

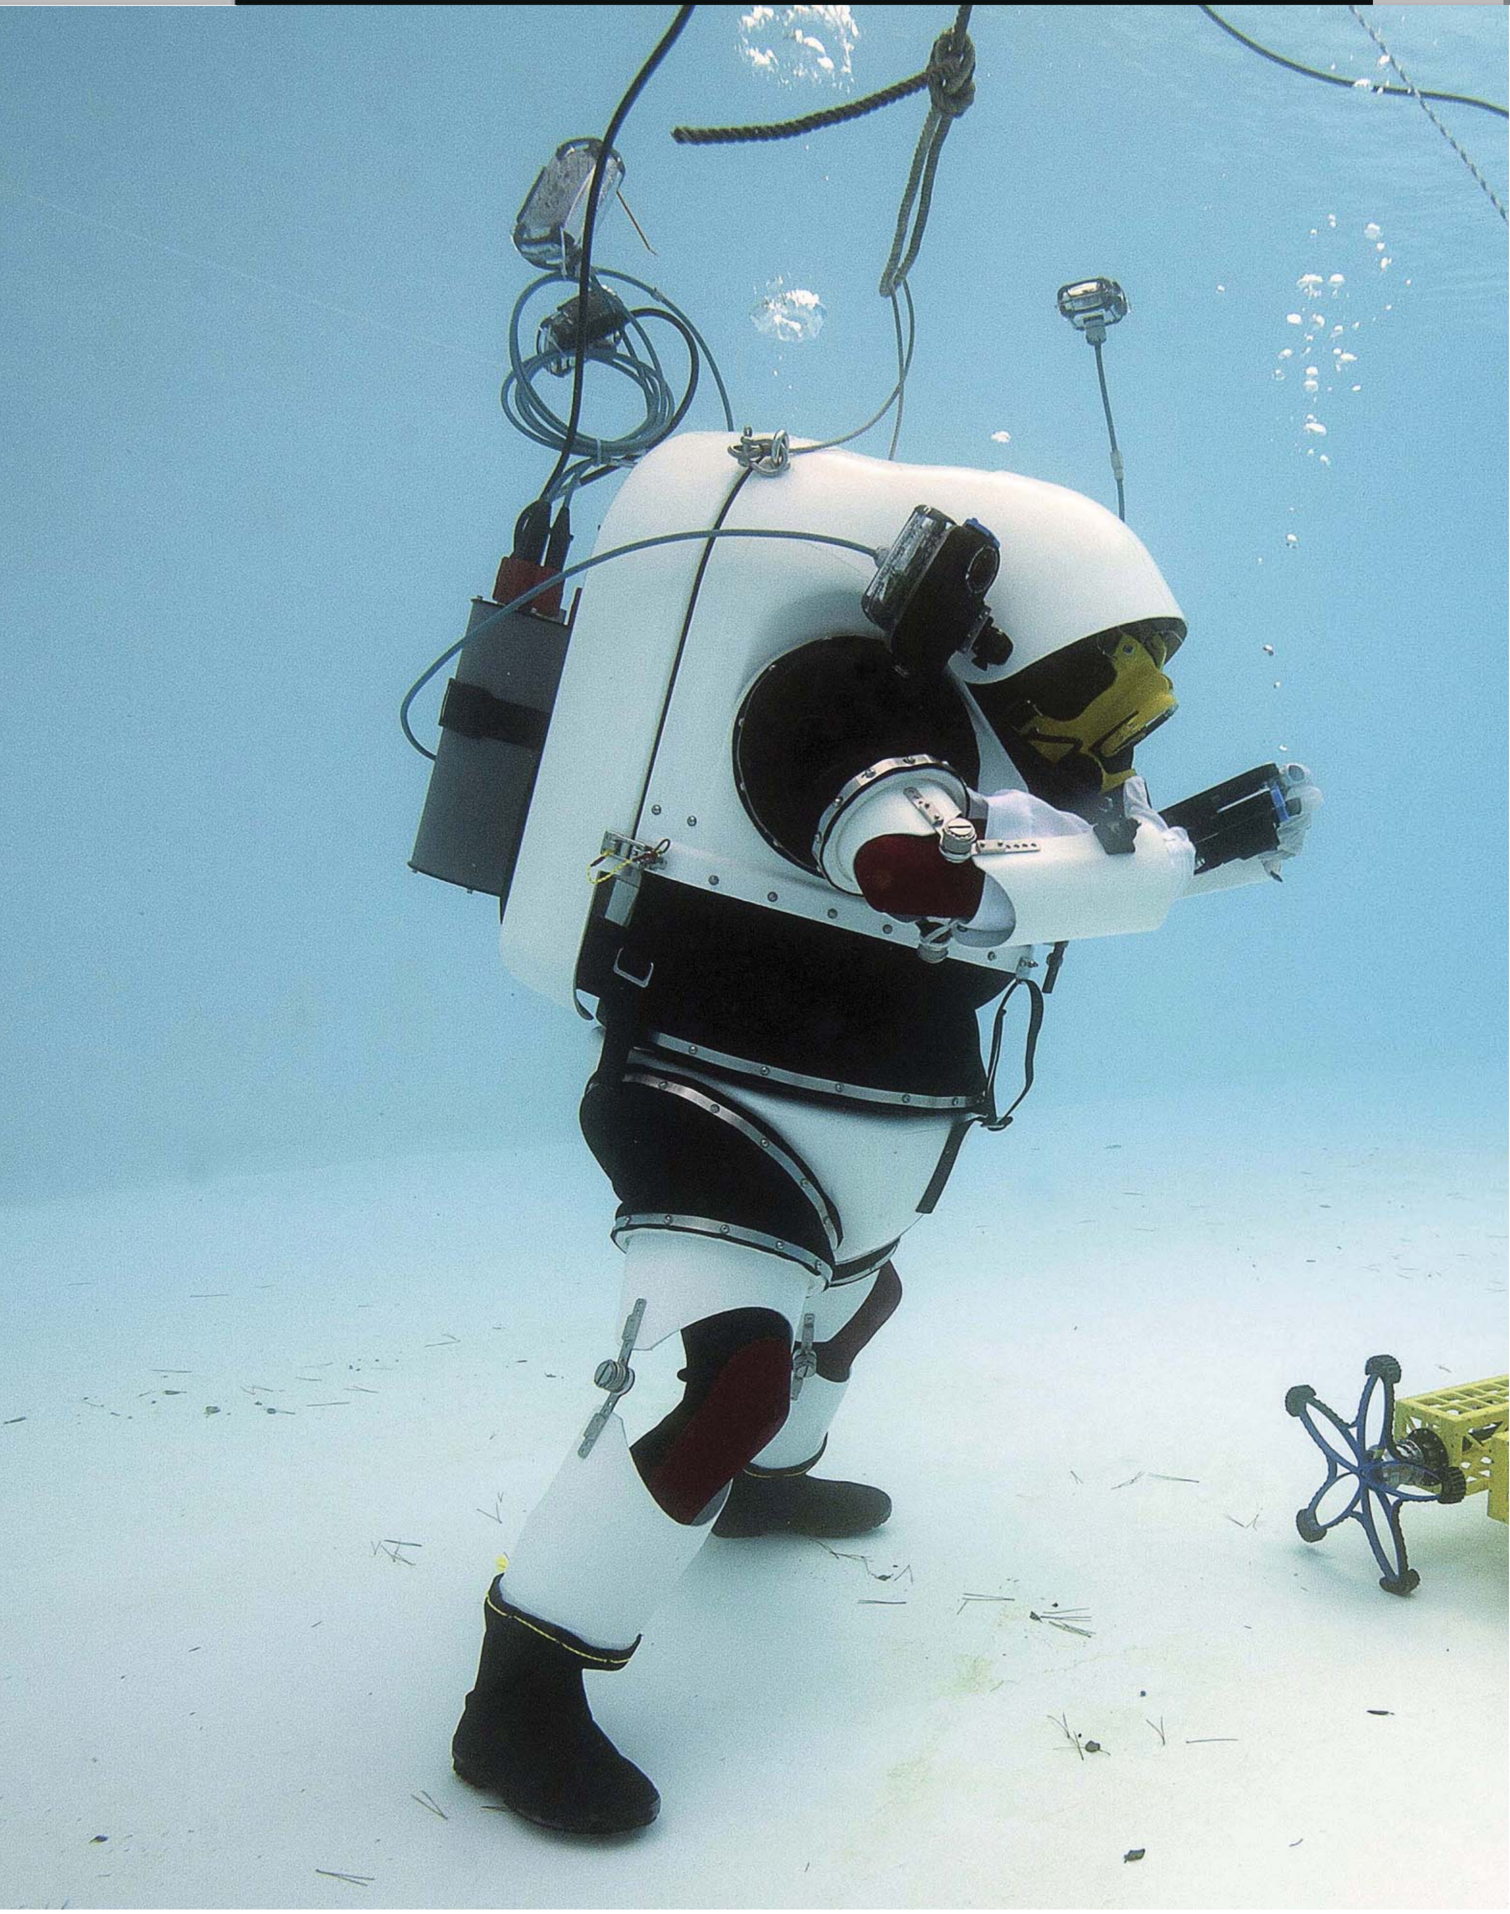

PHOTOGRAPH BY Boris Horvat/AFP/Getty Images

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

DEEP SPACE DIVING

TO PREPARE for the rigors of space, astronauts just as often take to the waters below as they do to the skies above. This particular practice session took place in September on a faux seafloor at the facilities of Comex (Compagnie Maritime d'Expertises), a leading engineering and deep-sea diving company, based in Marseille, France. The aim is to let European Space Agency (ESA) astronauts get used to moving around—and in this case, using a handheld computer to drive a robotic vehicle—despite the bulkiness and limited freedom of motion of the Gandolfi test suit developed by Comex and ESA.

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

RESOURCES

HOLIDAY GIFT GUIDE 2015

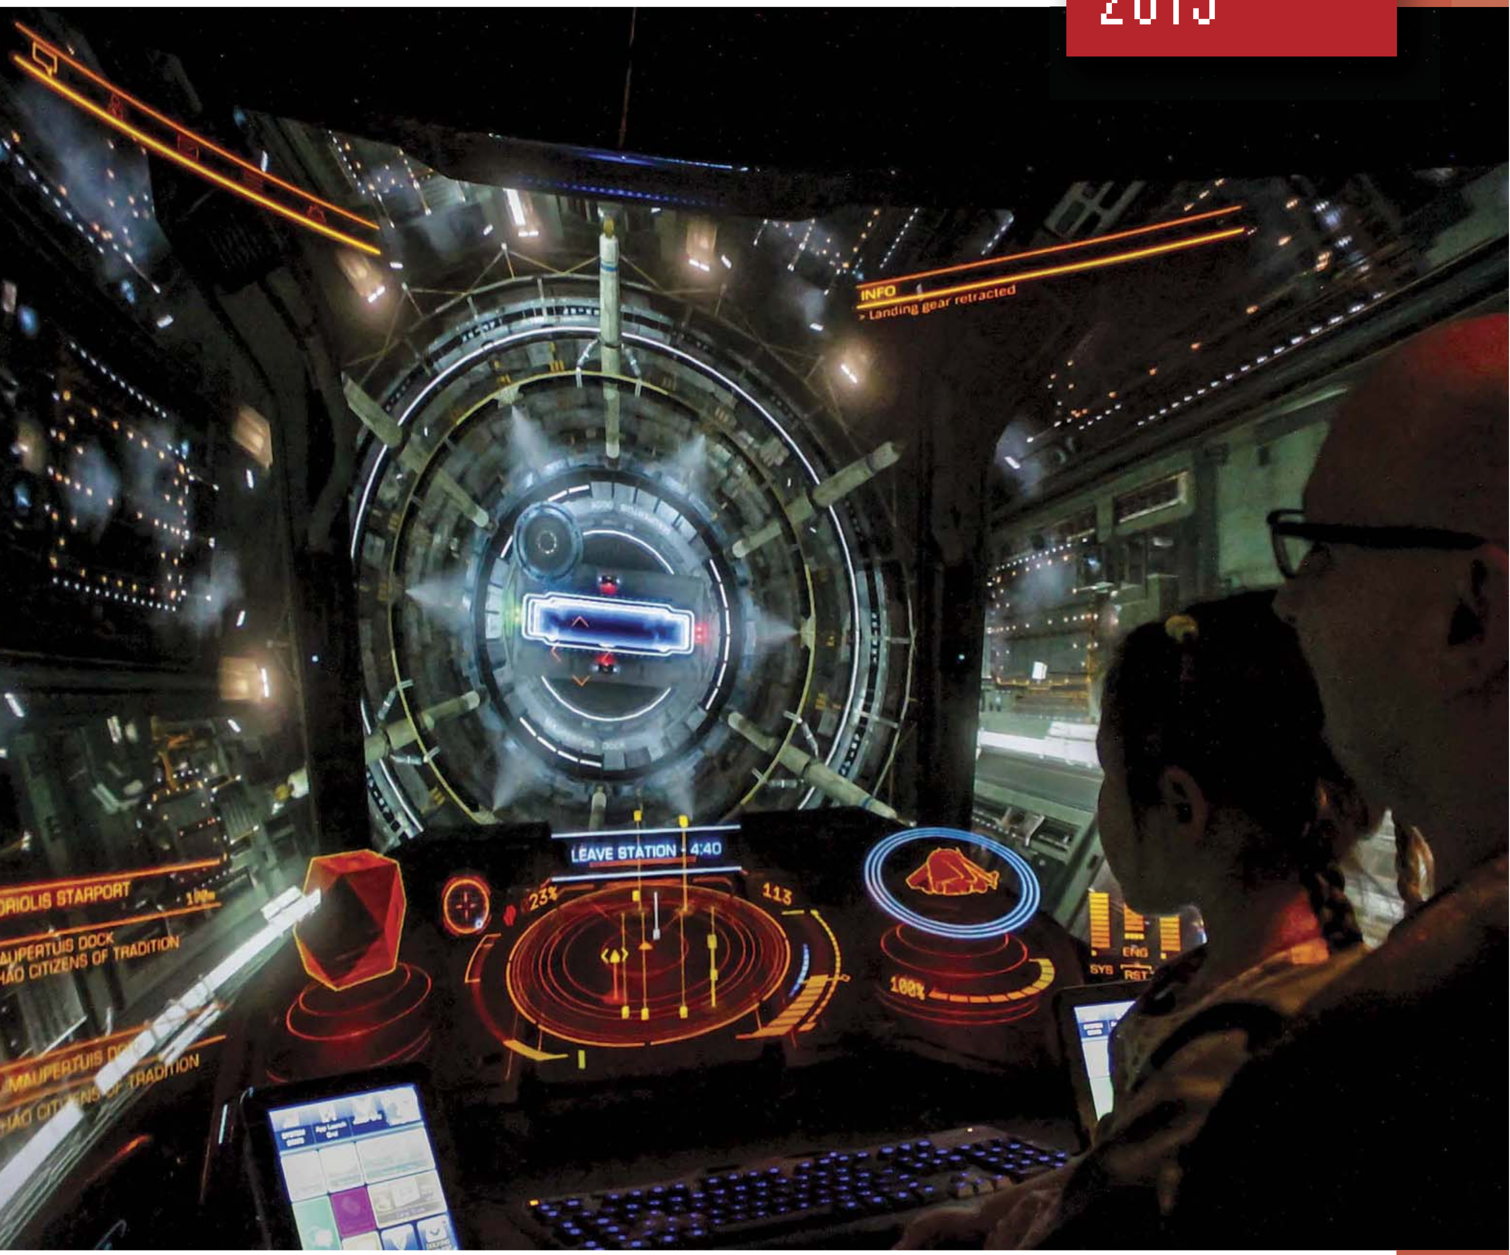

SIMPIT TECHNOLOGIES

IEEE SPECTRUM'S ANNUAL SELECTION OF GADGETS FOR TECHIES

NINEBOT ONE

I will admit that when I first saw an early version of the US $950 Ninebot One motorized wheel demonstrated at CES in January 2014, I thought it was a stunt product. When I wobbled around on the show floor the following year, I wondered what lunatics were buying them. But sales in parts of Europe have apparently been brisk, and I've even seen one or two zooming around in the hipper areas of New York City, so I guess I should put aside my old-fogeyness and admit that—with practiced riders—it *does* look kind of fun. **—STEPHEN CASS**

ICARUS 180–AVENGER

Virtual-reality goggles promise great things, but they do have an inherent drawback. Because they completely block the wearer's view, all but the simplest hand or gesture controllers are impractical. If you enjoy zipping through simulated realities in even a moderately complicated vehicle—whether a race car, airplane, or spaceship—you often need at least an entire keyboard's worth of buttons. The $3,125 Icarus 180–Avenger "sim pit" may be the solution, with a semicircular screen 1.7 meters in diameter that's fed by two HD projectors (the price quoted includes the pit, projectors, and software). You'll need to add your own PC gaming rig, but most software should run without modification, letting you turn a patch of your home into the Grand Canyon or the Orion Nebula. **—S.C.**

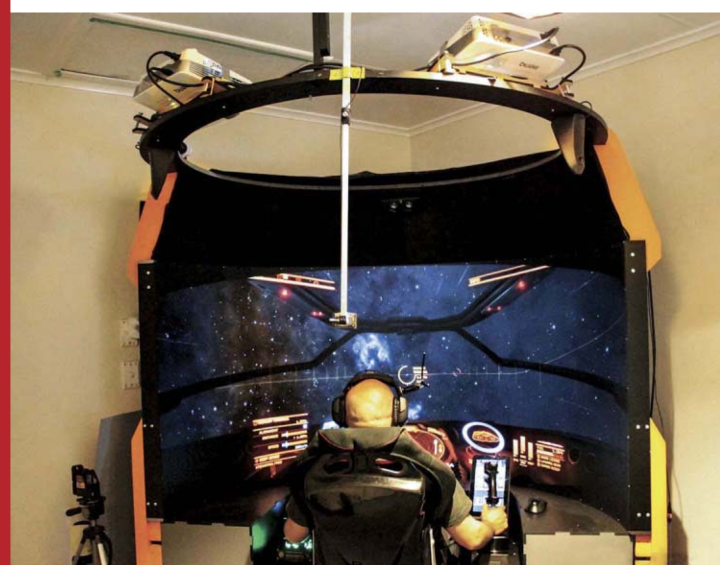

JAMSTIK+

The $300 Jamstik+ is intended to be a teaching instrument. Sensors under the Jamstik's frets detect finger placements; an iPad app displays where your fingers are and where they should be for a given chord. Users can play along to various tunes *Guitar Hero*–style. And, of interest to circuit benders and other electronic musicians, the company says it is open to sharing the software-development kit for custom projects on request. **—S.C.**

CLOCKWISE FROM TOP LEFT: NINEBOT (2); SIMPIT TECHNOLOGIES; ZIVIX

RESOURCES_GIFT GUIDE 2015

GEAR VR

For those not willing to splurge on a sim pit—and straddling the line between the cheap-and-cheerful $25 Google Cardboard virtual-reality goggles and the sophisticated-but-pricey $350 developer's version of the Oculus Rift headset—there is the Gear VR (most places are selling it for $100). Also designed by Oculus—in a clear move to stave off having its lunch eaten by low-cost competitors—the Gear VR holds a Samsung Galaxy Note 4. Like Cardboard, it uses the smartphone's screen to generate graphics. However, in addition to having better optics than the Cardboard, the Gear VR also features a touch pad and other controls to make interacting with VR applications easier. **—S.C.**

DIGITSOLE

With memories still fresh of how last year's polar vortex gripped the offices of *IEEE Spectrum* (along with most of the eastern United States), electrically heated insoles seem like a very attractive proposition, even with a price tag of $200. This being 2015, the rechargeable insoles are (what else?) controlled by a smartphone app, which allows you to adjust the temperature for each foot individually and track the number of steps walked. **—S.C.**

HYDROFOIL DRONE

Parrot, which helped create the consumer drone market with the 2010 launch of its AR.Drone, has bridged water and air with its $160 Hydrofoil drone. The drone is actually a two-in-one deal, with an aerial minidrone that can fly independently or be seated in a frame to power a hydrofoil hull. —S.C.

BB-8 BY SPHERO

This month sees the release of the highly anticipated *Star Wars: The Force Awakens*. The movie's trailers have already created a fan-favorite character: the BB-8 droid, seen rolling around landscapes with a head that seems to float on top of a spherical body. To the surprise of most, BB-8 turned out to be not a computer-generated special effect but a physical prop. And to the satisfaction of many, you can now own a (small) BB-8 for $150. How do they do it? Magnets! **—S.C.**

CLOCKWISE FROM TOP LEFT: SAMSUNG; DIGITSOLE; PARROT; LUCASFILM LTD.

BOOKS

THE ART OF ELECTRONICS, 3RD EDITION

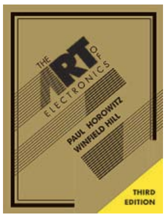

It might seem odd to recommend a $120 textbook as a gift, but there are probably few working EEs around today who don't have a copy of the cherished but somewhat outdated 1989 second edition of *The Art of Electronics* (Cambridge University Press). The latest edition takes in 25 years of advances, particularly in regard to microcontrollers and programmable logic devices, but it is still written in the same signature style as earlier editions. The book is littered with invaluable practical tips and even has frequent humor (EE humor, but humor nonetheless!). **—S.C.**

THE SPARKFUN GUIDE TO PROCESSING

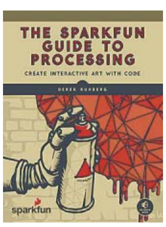

While more comprehensive books about processing exist, *The SparkFun Guide to Processing* (No Starch Press, $30) is an excellent primer on this Javaesque language for multimedia art, especially for casual or younger programmers. A quick read through the relevant chapter will have you—within minutes—visualizing data downloaded from online repositories or communicating with an Arduino microcontroller. **—S.C.**

HOLLYWEIRD SCIENCE

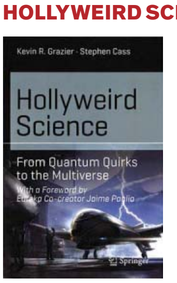

IEEE Spectrum's own senior editor Stephen Cass pairs up with Kevin R. Grazier, a physicist and science advisor to productions such as *Battlestar Galactica* and *Gravity*, on this highly entertaining yet authoritative account of science and technology as viewed through the lens of Hollywood (Springer, $25). With good humor and abundant enthusiasm, the pair explore the plausibility of dozens of sci-fi scenarios presented in TV shows and film.
—JEAN KUMAGAI

RESOURCES_GIFT GUIDE

RESOURCES_HANDS ON

WATER STATS ON TAP

USE AN ARDUINO AND A RASPBERRY PI TO CREATE A REAL-TIME FLOW METER

CALIFORNIA, WHERE I ONCE LIVED, IS IN THE MIDST of a severe drought. Thankfully, I now reside in North Carolina, a much wetter state. But even here there have been years when water was in short supply and people were asked to conserve. ● Most folks are happy to comply with such requests. But that's not easy to do when information about your usage comes only once a month on your water bill. You could, of course, track usage more finely by hoofing it out to your water meter periodically and noting the readings. But who has the discipline for that? So I cobbled together a monitor that can show how much water my household is using in real time. ● My system doesn't require plumbing alterations, and it transmits the information via my home wireless network, so I can view water use on a connected device anywhere in the house. Now when my teenage kids take profligate showers, I can berate them with quantitative measurements! ● To measure water flowing from the municipal supply, I employed a digital compass—that is, a 2-axis magnetometer—that I originally purchased some time ago from SparkFun Electronics to use as a direction finder. SparkFun no longer sells that module, but for even less money (US $15), you can now purchase a 3-axis magnetometer, which I suspect would work just as well or even better.

BOOKS, FROM TOP: CAMBRIDGE UNIVERSITY PRESS; SPARKFUN ELECTRONICS; SPRINGER; RIGHT: DAVID SCHNEIDER

LIQUID DATA

A magnetometer is valuable here because of the way mechanical water meters work. To isolate the gears of the actual meter from water exposure, a rotor is typically attached to a magnet in a separate chamber through which the water flows. As the water spins the rotor, this magnet turns another magnet in the meter's gearbox. Consequently, these coupled magnets produce a magnetic field around the meter that varies depending on how fast the water is flowing.

I figured that the field changes produced by the magnetic coupler mechanism would provide a booming signal outside the meter. But in fact, the magnetic signal is quite subtle. There are probably a few reasons for this: The steel housing of the meter likely shields the meter's internal magnetic field, and the moving parts inside the meter probably include various steel components that generate spurious magnetic fields at higher frequencies. At least that's what I suspected was going on. In any event, the signal I had to work with was a mess.

So my first challenge was to program an Arduino attached to the magnetometer to transform that noisy magnetic signal into a flow rate. I toyed with the idea of using a Fourier transform to pick out the dominant frequency corresponding to the flow rate, but instead I plumped for autocorrelation. That is, the program multiplies a short sample of the signal by a time-lagged version of itself and sums up the results. To find the dominant frequency, the Arduino code increments the lag between the two samples and looks for a peak in the summed results. That requires much less processing and seems pretty robust with respect to noise and harmonics.

Some testing with the garden hose showed that this arrangement could discern flow rates between gentle and gushing. The low end is limited by how long you let the Arduino collect measurements

To remotely monitor water flowing through my home's meter [top] I attached a magnetometer and used an Arduino to read and process the raw data. The Arduino sends the data by cable to a Raspberry Pi, which generates a simple Web page showing usage that I can access over my home's Wi-Fi network.

DAVID SCHNEIDER (3)

before processing them. I settled on about 3 seconds, which lets me measure down to a near dribble. But at the very highest flow rates (as when I used an outdoor spigot with no hose attached), my system faltered.

You might expect that high rates would produce a signal that was easier to measure than at low rates, but that's not the case. When the flow is really high, the magnets inside the water meter rotate at a clip that exceeds the magnetometer's fastest possible sampling rate, leading to a phenomenon called aliasing. (It's what makes wagon wheels in those old Westerns look like they're spinning backward instead of forward once a certain speed is reached.) As a result, the data simply weren't valid above a certain flow rate.

Short of experiencing a burst pipe, though, nobody in my house would be running water at or above that threshold, so I decided not to worry about this limitation and proceeded to work on how to communicate the information. There was also the issue of how to power the Arduino and magnetometer I had placed at the water meter, which is well separated from my house.

The solution was a 15-meter length of four-conductor telephone cable. Two conductors brought 9-volt DC power to the Arduino; the other two carried the digital output (transformed to RS-232 voltage levels using another board from SparkFun) to a Raspberry Pi, which I placed on my porch next to a regular AC power socket. From there, it was easy to have two "wall warts" power the Pi and the Arduino, and the Pi was close enough to my router for it to get onto my home wireless network.

I programmed the Pi in Python to serve up a single Web page, using the Flask framework for creating Web apps. That page shows flow rate and total water usage, updated every 10 seconds, and has a button to reset the water-usage counter to zero.

I have not yet attempted to calibrate my setup so that the rate is shown in sensible units, say liters per minute. So the number for total usage is also given in arbitrary units at the moment. But the Arduino and Pi in combination do an admirable job of informing me about relative flow rates. The whole family can now compete to see who can linger the least in the shower.

—DAVID SCHNEIDER

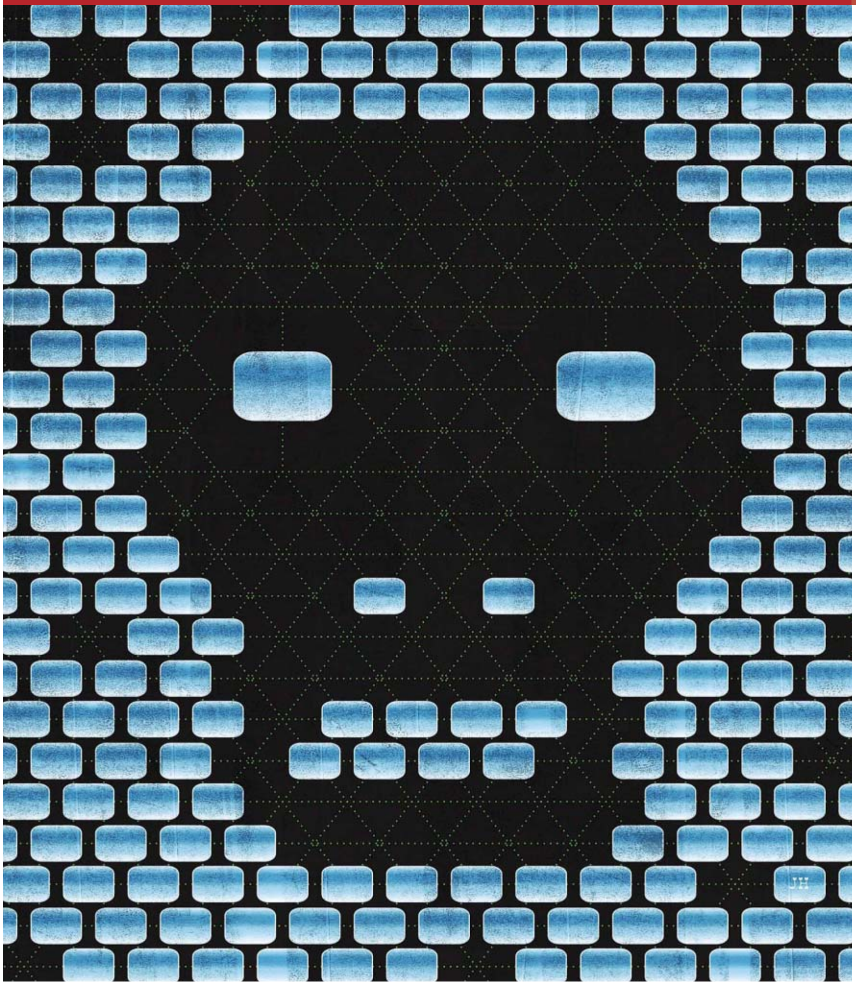

NETTING THE WEB

SOME INNOVATIONS COME VERY CLOSE to being purely good, with downsides too minimal to dwell on. Many health-care measures are in this category: inoculation against polio, fortification of flour with iron, electronic monitoring of vital functions in an intensive care unit. But for other innovations it is impossible to offer a clear net appraisal. • Consider cars: Is having the "freedom of the road" (or the freedom to sit in gridlock) worth the eye-irritating, lung-searing, life-shortening effects of photochemical smog? Does the convenience of on-demand mobility outweigh 1.3 million annual deaths on the roads? • But perhaps the best, perhaps unsurpassable, case of this category of innovation is the World Wide Web. For every undoubtedly good thing that the Web has brought us, I can easily cite a bad one. Here is just a sample of such contrasts. • The wealth of available information in a multitude of languages needs superlatives to describe, but I am absolutely certain that most of that information is either wrong, questionable, irrelevant to anybody but its providers and a few close targets, or supremely ephemeral. How could it be otherwise when images of latte art are among the most posted visuals?

As a lifelong interdisciplinarian I used to run up and down the stairs of many libraries every day, a tiring but excellent exercise; now I just search for electronic copies from the comfort of my house (not quite so healthy, even if I do most of my Web searches while standing up). Now the range of what is accessible is stunning: A few weeks ago I was reading a facsimile of Charles-Augustin de Coulomb's pioneering study of ergonomics, "Résultat de plusieurs expériences..." in *Mémoires de l'Institut national des sciences et arts* for the year 1799. But what is disappearing is just as stunning. When I checked the first 30 websites cited in my 2005 Oxford book *Creating the Twentieth Century,* I discovered that 19 of them (63.3 percent) were no longer up. The sites were defunct.

Of course, brick-and-mortar libraries also used to retire some of their old holdings, the first step usually being to place old, crumbling volumes into underground storage. But that loss of old books was nowhere near so rapid as the Web's loss of pages. Who ever heard of a library losing nearly two-thirds of a collection in a single decade!

Every reader of *IEEE Spectrum* is a near-constant Web user and could extend such contrasts almost ad infinitum. The Web lets you order the strangest consumer items directly from another continent–and get your Visa card information hijacked in the process by a teenager in Latvia. The Web contains practically every great poem ever written in scores of languages–but also the outpourings of semi-educated, self-appointed, and endlessly opinionated, logorrheic bloggers. Then again, what else can we expect when pressing a few keys is all that's required to publish a rant? The descent to the lowest common denominator was thus inevitable, blessing us with selfies and kittens, and providing an addictive alternative to thinking, working, or exercising with more than 70 billion hours of YouTube videos every year.

Are these positive or negative considerations when trying to net the value of the Web? Or are they the very essence of its existence? ■

ILLUSTRATION BY John Hersey

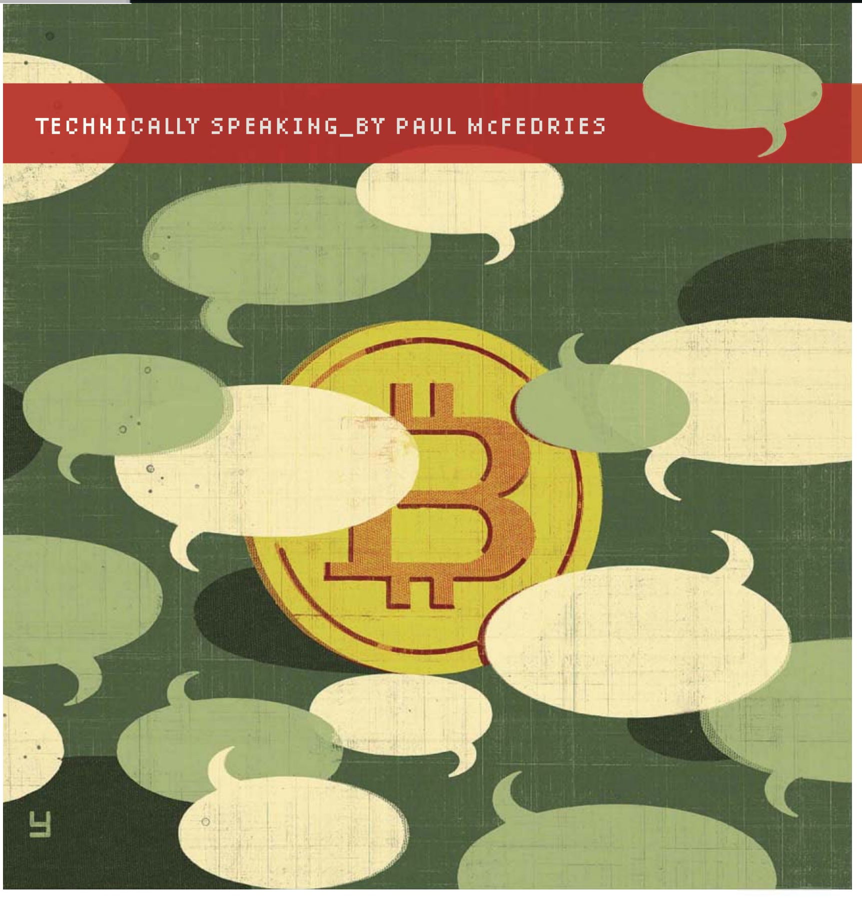

BITCOINAGES

What is needed is an electronic payment system... allowing any two willing parties to transact directly with each other without the need for a trusted third party. *—Satoshi Yakamoto, inventor of Bitcoin*

IS BITCOIN DESTINED TO BECOME THE FACEBOOK of money, or the Second Life? That is, will Bitcoin become universal, a medium of exchange used by just about everyone, or will it remain a curiosity, a tool used only by a dedicated fringe (in Bitcoin's case, criminals and privacy extremists)? As a mere language columnist, I haven't the foggiest, but I *can* tell you that, no matter what its fate, Bitcoin is generating new words and phrases at an inflationary rate. Is this profusion of new terms undermining (that's a pun, as you'll see) Bitcoin's quest for universality? • The confusion begins at the beginning, with the simple act of categorization: What is Bitcoin, exactly? The person known as Satoshi Yakamoto, who invented Bitcoin back in 2008, called it **electronic cash**. Other terms in use are **cryptocurrency**, **virtual currency**, **digital currency**, **digital cash**, and **math-based currency**. (This last term reminds me of something the entrepreneur Chris Dixon posted to his Tumblr in 2013: "Three eras of currency: Commodity based, e.g. Gold; Politically based, e.g. Dollar; Math based, e.g. Bitcoin.") Others are calling it **MoIP** (Money over Internet Protocol), a play on VoIP (Voice over Internet Protocol). The fact that many people are speculating in bitcoins has led some to categorize it as a **crypto-commodity**, and even a **moneylike informational commodity**. • We find more solid ground when we look at the Bitcoin system itself, which is best described as **trustless**. That is, instead of a trusted third party such as a bank brokering financial transactions, transactions occur peer-to-peer—directly between buyers and sellers—and are verified using well-known and well-tested methods of public-key cryptography, with some extra twists thrown in. The basic unit is the **electronic coin**, or bitcoin. (For the record, **bitcoinage** also comes in several smaller units, including the **millibitcoin**—0.001 bitcoins—the **microbitcoin**—0.000001 bitcoins—and the **satoshi**—0.000000001 bitcoins.) The important point is that all transactions are stored in a distributed database called the **blockchain**, an electronic ledger—sometimes called a **consensus ledger** or **cryptoledger**—that contains the complete history for all existing bitcoins (which are committed to the ledger periodically in groups called **blocks**, hence the name).

Everyone running the Bitcoin client software has a copy of the blockchain, and some clever cryptographic sleight of hand ensures that the ledger cannot be falsified or manipulated. This process involves generating a *hash* value for each block, which requires a nontrivial amount of computational work. Doing the required computations is called **mining** because once a user discovers a hash for a transaction block, she is given some bitcoins—called a **block reward**—as compensation for helping to maintain the system's integrity and security.

As Bitcoin slouches toward the mainstream, it has generated a lot of new words, but will it generate enough users to supplant our current system of so-called **fiat money**? A few years ago, when the venture capitalist Paul Graham was asked if he thought Bitcoin would have a lasting impact, he said that it "has all the signs. Paradigm shift, hackers love it, yet it's derided as a toy. Just like microcomputers." The PC succeeded because it became indispensable, just as in the 21st century Facebook became indispensable. Second Life hasn't passed that test, so we'll see if Bitcoin manages the feat. It's a long shot, but in an age characterized by the **Snowden effect**—the increased awareness of, and the desire to be protected from, government and corporate surveillance—Bitcoin's hour may have come round at last. ■

ILLUSTRATION BY Brad Yeo

THE QUEST *for the* ULTIMATE VACUUM TUBE

THE COLD-CATHODE TRAVELING-WAVE TUBE, AN ULTRACOMPACT, ULTRAEFFICIENT SOURCE OF RF WAVES, MAY FINALLY BE WITHIN REACH

BY CARTER M. ARMSTRONG

PHOTOGRAPH BY Sue Tallon

IN JULY 1962, the Telstar 1 satellite took an enormous leap toward the globally connected world we now take for granted. It relayed from space, for the first time ever, live television images and telephone calls between continents: specifically, a ground station in Andover, Maine, and other stations in England and France. It accomplished this feat thanks to a microwave repeater that had at its heart a slight but powerful vacuum device known as a traveling-wave tube. The 30-centimeter-long, glass-walled electron tube was at the time the only device capable of boosting a broadband television signal with enough power to cross an ocean. Solid-state devices just weren't up to the task.

More than a half century later, traveling-wave-tube amplifiers *still* dominate satellite communication. That's right—your ultrahigh-definition satellite TV and satellite radio come to you courtesy of vacuum tubes in space.

Of course, there's a huge difference between Telstar's 3.5-watt, 4-gigahertz amplifier and one of the dozens of highly efficient microwave amplifiers on, say, the DirecTV-15 satellite, launched earlier this year. The latest generation of traveling-wave tubes can provide up to 180 W at frequencies up to 22 GHz, with efficiencies approaching 70 percent and rated lifetimes exceeding 15 years. Though their basic function is the same—amplifying RF signals—just about everything else has changed: the design, the testing, the materials, and the fabrication.

That's my point. In the six decades since vacuum tubes lost out to solid-state devices in computers, receivers, and power supplies, vacuum technology has continued to evolve and branch out into new terrain, sustaining a small but skilled corps of engineers and scientists around the world, as well as a multibillion-dollar industry. That's because the traveling-wave tube and other vacuum devices continue to serve one purpose extremely well: as powerful sources of microwave, millimeter-wave, and submillimeter-wave radiation. (Vacuum tubes are also used in amplifiers for musical instruments and high-end audio, but the tubes I'm talking about are for generating radio-frequency waves, not audio waves.) What's more, they do it efficiently and over broad bandwidths, and compactly and reliably, too. And by virtue of their construction and the metals and ceramics from which they're fashioned, traveling-wave tubes are inherently hardened against radiation (unlike solid-state devices) and fairly impervious to temperature and mechanical extremes. Besides satellite communication, traveling-wave tubes are widely used in radar, electronic warfare, and other military systems.

VACUUM TECHNOLOGY CONTINUES TO SMASH THROUGH PERFORMANCE BARRIERS AND OPEN UP NEW APPLICATIONS. IT REFUSES TO FADE INTO OBSCURITY

And now, ongoing research into a new and potentially revolutionary kind of traveling-wave tube—the ultracompact and ultraefficient cold-cathode TWT—looks poised to deliver the first practical device by the end of this decade. These are exciting times for vacuum tubes. Here's why.

"IF THE TRANSISTOR HAD BEEN INVENTED FIRST, the vacuum tube would have been invented immediately afterwards." That remark, by vacuum electronics great Robert S. Symons, was his standard comeback against frequent attacks that his was an antiquated field of legacy devices soon to be eclipsed by semiconductors. Those of us who followed in Symons's footsteps of course agree that the tube's compelling attributes make it worthy of pursuit, but unfortunately we still contend with that same prejudice against the technology.

Not that we have a, er, chip on our shoulder, but in what ways are tubes superior? Let's review. At its heart, an RF source, whether solid-state or vacuum electronics, is a power converter, changing some fraction of the incoming electricity into RF power. The input power is a product of the incoming current times the incoming voltage. Traveling-wave tubes operate at high voltage—typically from a few kilovolts to tens of kilovolts—which means they need only a little current to produce a very powerful electron beam. True, you need a fabulous power supply to feed the tube, but advances in compact high-voltage electronics have solved that problem.

Solid-state devices used in RF sources, by contrast, operate at low voltage—typically 8 to 10 volts for sources based on gallium arsenide and up to about 70 V for gallium nitride sources. So to get a high output power, the current also needs to be high, which can then create excessive heat that stresses the chip's connections. The upshot is that to equal the output of, say, a single 300-W traveling-wave tube, you'll need to gang a bunch of solid-state devices together, which adds to the complexity and cost of your design; even then, the efficiency of your RF source will be at best only about half that of a modern TWT.

To appreciate how a tube does what it does, let's consider its basic parts: an *electron gun* (consisting of a cathode that's heated to generate a beam of electrons, an anode to accelerate the beam to high velocity, and in most cases a beam-controlling gate electrode near the cathode to switch the cathode current on and off); an *electromagnetic circuit* that interacts with the electron beam; a *collector* for capturing the electron beam as it exits the interaction circuit; and a *focusing magnet*, for confining the electron beam as it propagates through the tube. The first three parts—the electron gun, interaction circuit, and collector—reside in a hard vacuum (about 10^{-8} pascals), while the magnet sits outside the vacuum. We create the vacuum by heating the tube to over 500 °C for more than 24 hours, which cleans the inner surfaces of the tube and drives out residual gas. The walls of a modern

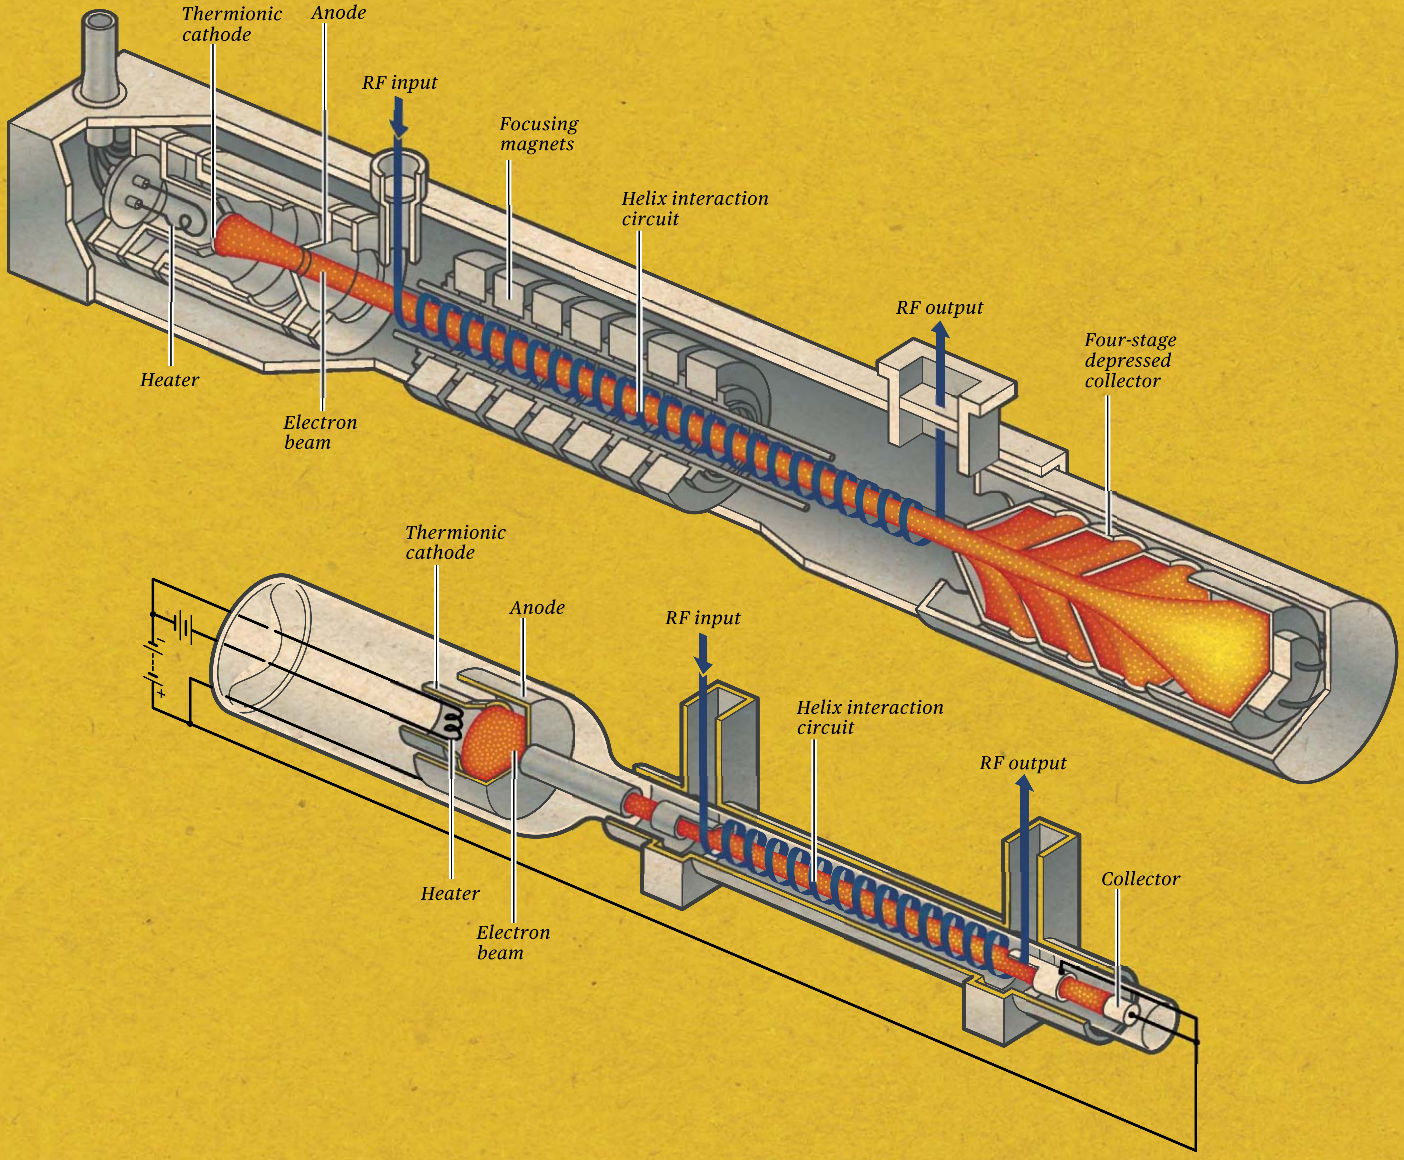

THE TRAVELING-WAVE TUBE: THEN AND NOW

For amplifying radio-frequency signals, nothing beats a traveling-wave tube. Here, we see a modern TWT [top] and one from the 1950s [bottom]. In both cases, the cathode is heated to emit an electron beam, which is injected into a helix-shaped circuit.

The RF signal to be amplified is injected onto the helix, with the resulting electromagnetic wave extending away from the helix and exerting a force on the incoming beam. Unmodulated as it entered the helix, the beam current now has an RF component. The modulation induces electromagnetic fields on the helix, which then act on the electrons, and the RF signal grows exponentially.

In a modern traveling-wave tube, the beam power that is not converted to RF is recovered by the depressed collector, which acts like the regenerative brakes on a hybrid car and boosts the overall efficiency to 65 percent or more. That feature plus other advances in design, testing, materials, and fabrication have allowed TWTs to continue to outperform solid-state RF sources. And recent research on a cold-cathode TWT promises to revolutionize vacuum electronics. *—C.M.A.*

ILLUSTRATION BY James Provost

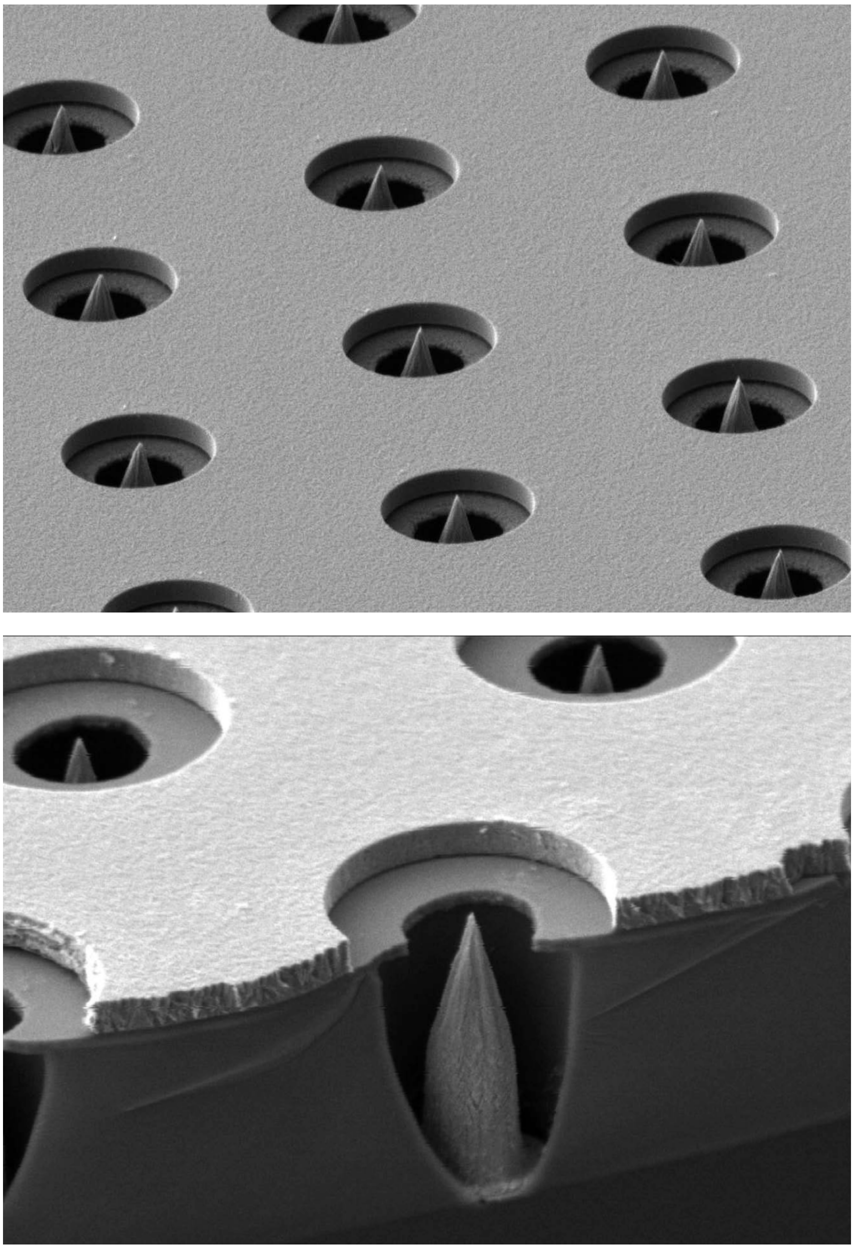

CONES IN WELLS: The cold-cathode traveling-wave tube developed by researchers at SRI International and L-3 Electron Devices relies on up to 50,000 micrometer-size molybdenum cones, each of which can emit a tiny beam of electrons at room temperature. The tube itself is shown in the opening photo of this article.

SRI INTERNATIONAL (2)

tube are made not of glass but of heat- and corrosion-resistant metals like tungsten and molybdenum, as well as iron, high-purity copper, and high-temperature ceramics.

So here's how the traveling-wave tube works: First, the electron gun injects an electron beam through the center of the interaction circuit; the most common type of circuit used in traveling-wave tubes is shaped in a helix. The RF signal to be amplified is then injected onto the helix. The resulting electromagnetic wave extends away from the helix and so exerts a force on the incoming beam of electrons. The first electrons entering the helix will accelerate in the direction of the beam; a half cycle later, the entering electrons will slow down. As the beam travels through the circuit, the faster electrons will overtake the slower ones so that bunches of electrons form.

The result is that the beam current, which was unmodulated as it entered the helix, now has an RF component at the signal frequency. The modulation in turn induces electromagnetic fields on the helix, which then act on the electrons; as a result, the signal on the helix grows exponentially–that is, we get gain. But because power must be conserved, this increase in the RF signal's strength comes at the expense of the beam's kinetic energy, and so the beam slows down.

Even so, on its pass through the circuit, only about a quarter to a third of the beam power is typically converted to RF, which means the beam exiting the circuit still retains a lot of power. To recover some of that kinetic energy, we slow the beam down further using something called a depressed collector, which is basically a series of electrodes that operate like the regenerative brakes on a hybrid car. A modern space-based traveling-wave tube has a depressed collector with four or five stages, which can recover well over 80 percent of the exiting beam energy, resulting in overall efficiencies of 65 percent or more. The tube also has fins attached to the collector to radiate heat from the tube and the satellite into the cold surroundings of space. A solid-state device has no analogue to the depressed collector.

Because of the tube's exquisitely low vacuum and the beam's near-perfect focus, very little energy is lost due to collisions as the beam passes through the circuit. The electrons in solid-state devices, by contrast, suffer a multitude of collisions as they drift through the semiconductor, and so produce waste heat right at the transistor junction, inside the device itself. It's not hard to see why heat removal is one of the biggest challenges for high-power solid-state devices. And in contrast to TWTs, their performance is much more sensitive to variations in operating temperature.

SO WHAT'S RESPONSIBLE for the major leaps in performance and efficiency from the Telstar era to today? In a way, solid-state and vacuum devices have followed parallel paths over the past 50 years, benefiting from advances in computational design, engineered materials, and microfabrication.

Computer modeling and simulation, for instance, allow engineers to model a tube from the cathode to the collector and everything in-between, with the result that it's commonplace to have prototypes work on the first pass, something unthinkable just 10 years ago. And tube designers can now create previously unimaginable multiple-beam vacuum devices, like Northrop Grumman Corp.'s multiple-beam traveling-wave tube, which provides an astounding 50 W of pulsed power at 214 GHz. The tube uses five separate electron beams rather than a single beam, so it can operate at lower voltage and requires less magnetic focusing. Understanding the complex three-dimensional physics involved in producing, focusing, and capturing those beams and modeling their conversion to RF power could be done only using modern computational tools.

Meanwhile, advances in microfabrication are making it possible to actually build these new designs. Researchers at the University of California, Davis, have used high-speed micromachining—with tool bits barely visible to the naked eye and spinning at rates exceeding 50,000 revolutions per minute—to create metallic structures with dimensional accuracies in the micrometers. Such fine structures are what you need to produce the extremely high frequencies typical of today's most sophisticated traveling-wave tubes. Borrowing a page from the semiconductor industry, engineers at Northrop Grumman and Teledyne Scientific Co. have etched traveling-wave-tube circuits onto silicon wafers. And in an extension of a MEMS technique, researchers at the U.S. Naval Research Laboratory are using a process that combines UV lithography and electroplating to make extremely accurate, all-metal structures for use in millimeter-wave traveling-wave tubes. Given the pace of advances in additive manufacturing, 3-D-printed components for vacuum tubes can't be that far off.

RESEARCH INTO THE POTENTIALLY REVOLUTIONARY COLD-CATHODE TRAVELING-WAVE TUBE COULD DELIVER THE FIRST PRACTICAL DEVICE WITHIN FIVE YEARS

ENGINEERED MATERIALS, TOO, are fomenting a minor revolution in vacuum-device manufacture. I'm guessing you've never heard of nanocrystalline scandium oxide-doped tungsten powder, but it is behind the biggest advance in tube cathodes in decades.

As mentioned, the cathode is the source, or emitter, of the electron beam that flows through a tube. Most tubes use thermionic emitters, which emit a copious stream of electrons when heated to around 1,000 °C. (You can sear a steak at 400 °C.) The most common type of cathode consists of a barium mixture dispersed within a porous tungsten matrix. Heated, the barium continuously diffuses to the surface, forming a thin oxide layer. This layer is said to have a "work function," which is just a measure (expressed in electron volts) of how difficult it is to extract an electron from it. The lower the work function, the lower the temperature you need to emit electrons, the lower the evaporation rate of the barium, and the longer the cathode lasts. The cathode life basically triples for every 0.2 electron volt's reduction in the work function. Today's cathodes have work functions of about 2.0 eV; applying a thin metal coating lowers the work function to 1.8 eV.

Now, the cathode's operating temperature is determined by its current density—that is, the amount of current emitted per unit area. Careful control of the current density in space-based traveling-wave tubes means the cathodes can last up to 20 years; terrestrial tubes, which typically operate at higher current densities, have shorter lives, of about seven years. But tube designers are always on the lookout for materials that will reduce the work function and thereby extend the tube's life.

Thus the interest in scandium oxide. Researchers have known for years that scandate cathodes, made from tungsten and scandium oxide, have low work functions. But attempts to turn them into practical, commercial components always fell short. Recently, however, Yiman Wang and her colleagues at the Beijing University of Technology demonstrated that cathodes made from nanocrystalline powders of scandium oxide and tungsten resulted in a robust emitter with a work function of 1.43 eV, more than tripling the life of a standard traveling-wave-tube cathode. This advance also allows you to operate such a cathode at a higher current density, which would vastly simplify the focusing of the fine electron beams needed in higher-frequency millimeter- and submillimeter-wave tubes.

AS GREAT AS THE SCANDATE CATHODE sounds, imagine if you replaced the thermionic emitter with a cathode that didn't need to be heated at all. That would be truly revolutionary. Vacuum engineers even have a name for such a device: the cold-cathode traveling-wave tube. A cathode that operates at ambient temperature would have all sorts of advantages: You wouldn't need a heater; the cathode wouldn't wear out; there'd be no warm-up time for the tube—it would essentially be instant on and off; and the current density could be much higher because its emission wouldn't be limited by the operating temperature. What's more, you could modulate the beam's current right at the cathode, as is done in the high-power tubes used for UHF TV transmitters, only now it would be at microwave frequencies. Cold-cathode vacuum devices would usher in a new generation of ultracompact, high-fidelity, high-efficiency amplifiers, which could be used, for instance, in powerful yet tiny RF transmitters on small satellites.

No surprise, then, that researchers have sought the cold-cathode traveling-wave tube for decades. Charles "Capp" Spindt and Kenneth Shoulders of SRI International, in Menlo Park, Calif., were among the first to articulate its virtues, in a 1966 paper entitled "Research in Micron-size Field Emission Tubes."

While there are various ways of emitting a beam of electrons without heating, the one that Spindt | CONTINUED ON PAGE 48


:(
DRAMMIT!

BATTLING BORKED BITS

DRAM GLITCHES CAN CRASH MACHINES. BETTER ERROR-AVOIDANCE SCHEMES CAN HELP

By IOAN STEFANOVICI, ANDY HWANG & BIANCA SCHROEDER

NOT LONG after the first personal computers started entering people's homes, Intel fell victim to a nasty kind of memory error. The company, which had commercialized the very first dynamic random-access memory (DRAM) chip in 1971 with a 1,024-bit device, was continuing to increase data densities. A few years later, Intel's then cutting-edge 16-kilobit DRAM chips were sometimes storing bits differently from the way they were written. Indeed, they were making these mistakes at an alarmingly high rate. The cause was ultimately traced to the ceramic packaging for these DRAM devices. Trace amounts of radioactive material that had gotten into the chip packaging were emitting alpha particles and corrupting the data.

An error has occurred. Turn the page to continue__

Once uncovered, this problem was easy enough to fix. But DRAM errors haven't disappeared. As a computer user, you're probably familiar with what can result: the infamous blue screen of death. In the middle of an important project, your machine crashes or applications grind to a halt. While there can be many reasons for such annoying glitches–including program bugs, clashing software packages, and malware–DRAM errors can also be the culprit.

For personal-computer users, such episodes are mostly just an annoyance. But for large-scale commercial operators, reliability issues are becoming the limiting factor in the creation and design of their systems.

Big Internet companies like Amazon, Facebook, and Google keep up with the growing demand for their services through massive parallelism, with their data centers routinely housing tens of thousands of individual computers, many of which might be working to serve just one end user. Supercomputer facilities are about as big and, if anything, run their equipment even more intensively.

In computing systems built on such huge scales, even low-probability failures take place relatively frequently. If an individual computer can be expected to crash, say, three times a year, in a data center with 10,000 computers, there will be nearly 100 crashes a day.

Our group at the University of Toronto has been investigating ways to prevent that. We started with the simple premise that before we could hope to make these computers work more reliably, we needed to fully understand how real systems fail. While it didn't surprise us that DRAM errors are a big part of the problem, exactly how those memory chips were malfunctioning proved a great surprise.

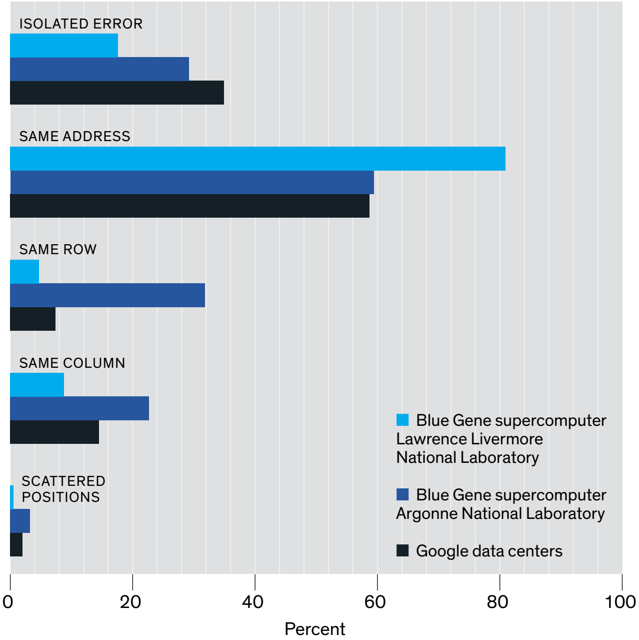

FLAVORS OF GOOFS: Most often DRAM errors repeatedly affect the same address—or perhaps just the same row or column—rather than being isolated occurrences.

To probe how these devices sometimes fail, we collected field data from a variety of systems, including a dozen high-performance computing clusters at Argonne National Laboratory, Lawrence Livermore National Laboratory, and Los Alamos National Laboratory. We also obtained information from operators of large data centers, including companies such as Google, about their experiences.

We made two key observations: First, although most personal-computer users blame system failures on software problems (quirks of the operating system, browser, and so forth) or maybe on malware infections, hardware was the main culprit. At Los Alamos, for instance, more than 60 percent of machine outages came from hardware issues. Digging further, we found that the most common hardware problem was faulty DRAM. This meshes with the experience of people operating big data centers, DRAM modules being among the most frequently replaced components.

An individual memory cell in a modern DRAM chip is made up of one capacitor and one transistor. Charging or discharging that capacitor stores one bit of information. Unlike static RAM, which keeps its data intact as long as power is applied to it, DRAM loses information, because the electric charge used to record the bits in its memory cells slowly leaks away. So circuitry must refresh the charge state of each of these memory cells many times each second–hence the appellation "dynamic." Although having to include refresh circuitry complicates the construction of memory modules, the advantage of DRAM is that the capacitors for the memory cells can be made exceedingly small, which allows for billions of bits on a single chip.

A DRAM error arises when one or more bits that are written in one way end up being read in another way. Most consumer-grade computers offer no protection against such problems, but servers typically use what is called an error-correcting code (ECC) in their DRAM. The basic strategy is that by storing more bits than are needed to hold the data, the chip can detect and possibly even correct memory errors, as long as not too many bits are flipped simultaneously. But errors that are too severe can still cause machines to crash.

Although the general problem is well known, what stunned us when analyzing data from Lawrence Livermore, Argonne, and Google (and a few other sites later on) was how common these errors are. Between 12 percent and 45 percent of machines at Google experience at least one DRAM error per year. This is orders of magnitude more frequent than earlier estimates had suggested. And even though the machines at Google all employ various forms of ECC, between 0.2 percent and 4 percent of them succumb to uncorrectable DRAM errors each year, causing them to shut down unexpectedly.

How can a problem whose existence has been known for more than three decades still be so poorly understood? Several reasons. First off, although these failures occur more

SOURCE, BOTH PAGES: HWANG, STEFANOVICI, AND SCHROEDER, *PROCEEDINGS OF ASPLOS XVII*, 2012.

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

often than anybody would like, they're still rare enough to require very large data sets to obtain statistically significant frequency estimates: You have to study many thousands of machines for years. As a result, most reliability estimates for DRAM are produced in the lab—for example, by shooting particle beams at a device to simulate the effects of cosmic rays, which were long thought to be the main cause of these errors. The idea was that a high-energy cosmic ray would hit a gas molecule in the atmosphere, giving rise to a zoo of other particles, some of which could in turn cause bit errors when they hit a DRAM chip.

Another reason it's so hard to get data on memory failures "in the wild" is that manufacturers and companies running large systems are very reluctant to share information with the public about hardware failures. It's just too sensitive.

DRAM errors can usually be divided into two broad categories: soft errors and hard errors. Soft errors occur when the physical device is perfectly functional but some transient form of interference—say, a particle spawned by a cosmic ray—corrupts the stored data. Hard errors reflect problems with the physical device itself, where, for example, a specific bit in DRAM is permanently stuck at 0 or 1.

The prevailing wisdom is that soft errors are much more common than hard errors. So nearly all previous research on this topic focused on soft errors. Curiously, before we began our investigations, there were no field studies available even to check this assumption.

To test this widely held belief—indeed, to answer many open questions about memory errors—we examined large-scale computing systems with a wide variety of workloads, DRAM technologies, and protection mechanisms. Specifically, we obtained access to a subset of Google's data centers, two generations of IBM's Blue Gene supercomputer (one at Lawrence Livermore and one at Argonne), as well as the largest high-performance cluster in Canada, housed at the SciNet supercomputer center at the University of Toronto. Most of these facilities were already logging the relevant data, although we needed to work with SciNet's operators on that. In total, we analyzed more than 300 terabyte-years of DRAM usage.

There was some unquestionably good news. For one, high temperatures don't degrade memory as much as people had thought. This is valuable to know: By letting machines run somewhat hotter than usual, big data centers can save on cooling costs and also cut down on associated carbon emissions.

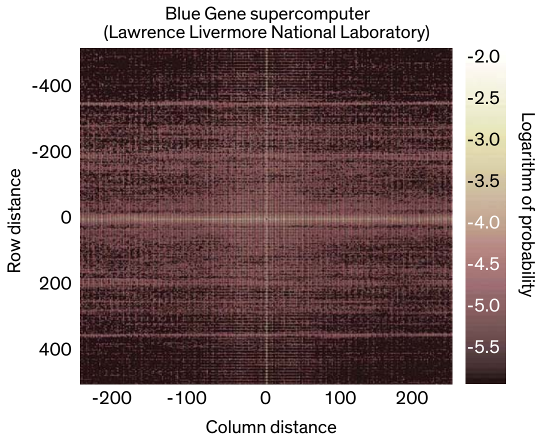

IN THE CROSSHAIRS: Once a DRAM error occurs, the likelihood of other errors in the same row or column of memory cells increases. The likelihood of errors even grows for nearby rows or columns, as seen in this "heat map" of error correlations.

One of the most important things we discovered was that a small minority of the machines caused a large majority of the errors. That is, the errors tended to hit the same memory modules time and again. Indeed, we calculated that the probability of having a DRAM error nearly doubles after a machine has had just two such errors. This was startling, given the prevailing assumption that soft errors are the dominant failure mode. After all, if most errors come from random events such as cosmic rays, each DRAM memory chip should have an equal chance of being struck, leading to a roughly uniform distribution of errors throughout the monitored systems. But that wasn't happening.

To understand more about our results, you need to know that a DRAM device is organized into several memory banks, each of which can be thought of as a 2-D array of memory cells laid out in neat rows and columns. A particular cell inside an array can be identified by two numbers: a row index and a column index.

We found that more than half the memory banks suffering errors were failing repeatedly at the same row and column address—that is, at the location of one iffy cell. Another significant fraction of the errors cropped up in the same row or same column each time, although the exact address varied.

It's unlikely, of course, that the same location on a device would be hit twice by an errant nuclear particle. Extremely unlikely. This is clear evidence that, contrary to prevailing notions, hard DRAM errors are more common than soft ones.

What are we to make of the unexpected prevalence of hard errors? The bad news is that hard errors are permanent. The good news is that they are easy to work around. If errors take place repeatedly in the same memory address, you can just blacklist that address. And you can do that well before the computer crashes. Remember, the only errors that really matter are the ones that flip too many bits for ECC to correct. So errors that corrupt fewer bits could be used as an early warning that drastic measures should be taken before a crash occurs.

Our investigation showed that this strategy could indeed work well: More than half of the catastrophic multibit errors followed earlier errors that were | CONTINUED ON PAGE 50

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

SAFE-GUARD-ING THE SUPER-GRID

Inside the lab that pushes megavolt circuit breakers to the limit

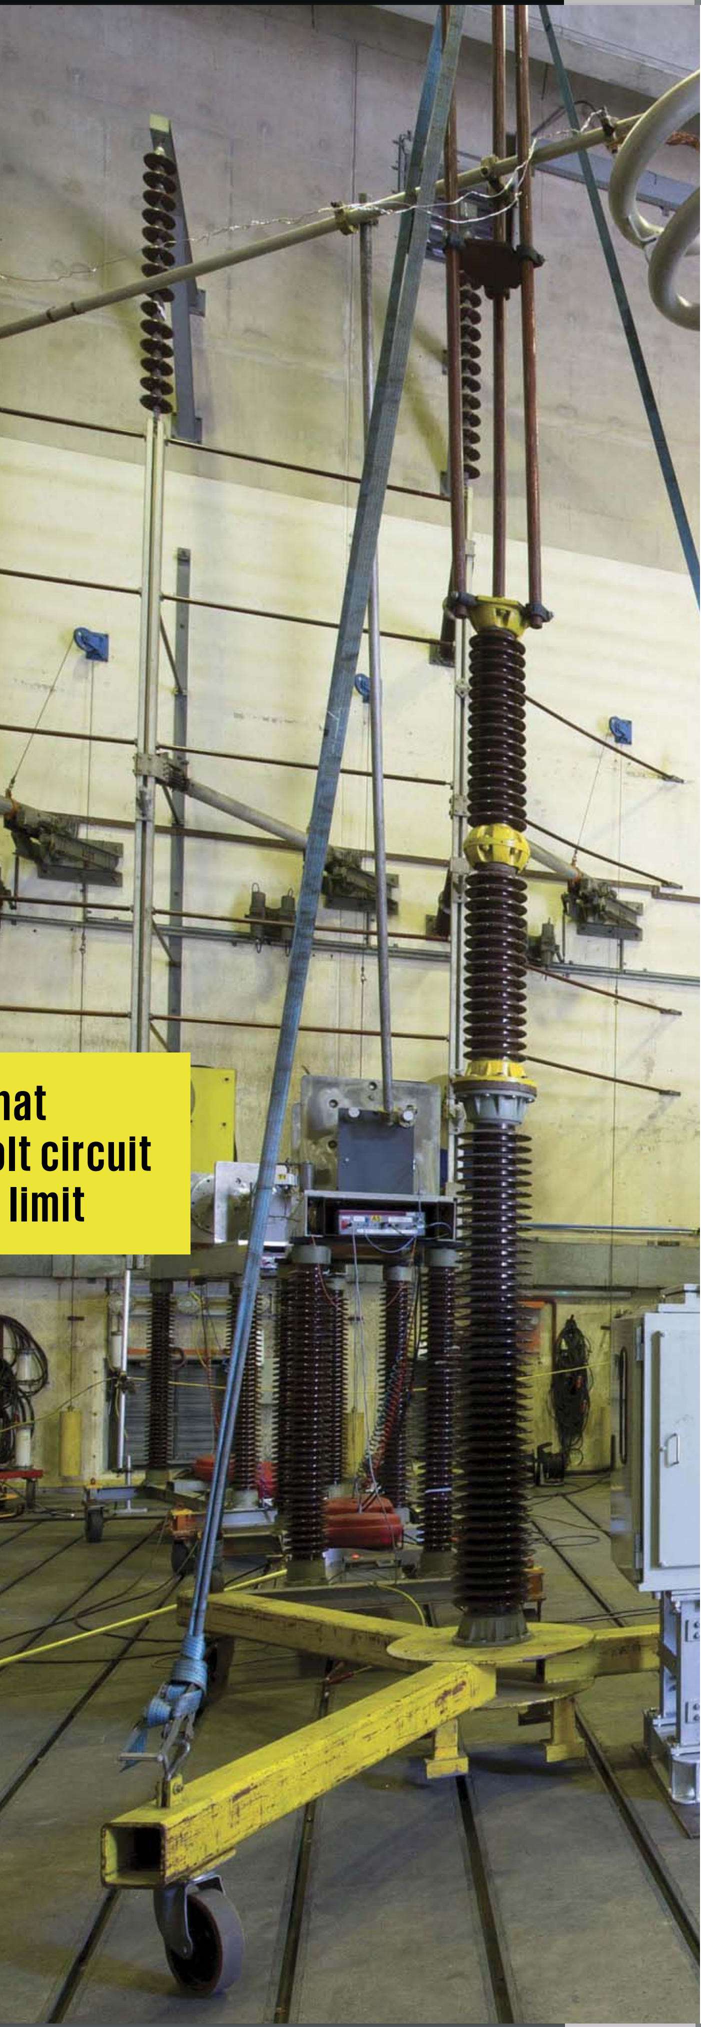

KEMA LABORATORIES

IMAGINE A BLAZINGLY HOT DAY IN CENTRAL CHINA, when all the air conditioners in every megacity are running at full blast. Through the remote mountains of Shanxi province, the major transmission lines that carry ultrahigh-voltage electricity to the cities are operating at close to maximum capacity. Heated by the sunshine and the flowing current, the transmission lines sag dangerously close to the treetops. Suddenly the current jumps from line to tree branch, finding the path of least resistance and pouring through the tree into the ground. There's a bright flash as the current ionizes the air. • During this short circuit, the abruptly unleashed current reaches 10 to 20 times its normal level within a blink of an eye. Now the power grid's protection system must act fast. Within milliseconds, protection relays must recognize the fault and command the circuit breakers »

BY RENÉ SMEETS

THE GUARDIAN: KEMA Laboratories tests a circuit breaker under extreme conditions to ensure it won't fail when it really matters.

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

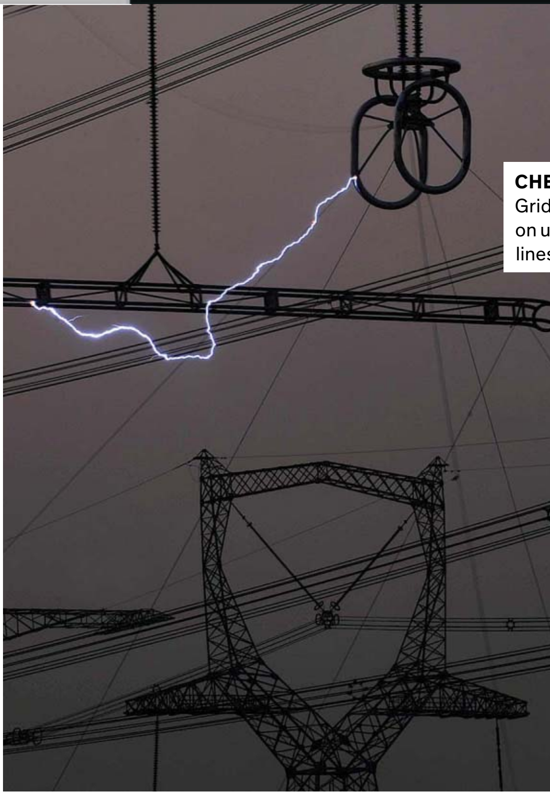

CHECKING THE LINES: The State Grid Corp. of China runs an experiment on ultrahigh-voltage transmission lines in Hubei province.

LI JINYONG/XINHUA

at both ends of the line to switch off the current, isolating the faulted line. The stakes are high: A sustained short-circuit current can trigger a chain reaction of failures throughout the grid and cause widespread blackouts, severely damaging expensive equipment in the process. The 2003 blackout in northeast North America was set off by a tree's contact with transmission lines in Ohio, which caused a cascade of failures that shut down more than 260 power plants, stopped the flow of 60,000 megawatts throughout the northeast grid, and darkened New York City.

In our hypothetical Chinese short circuit, everything is riding on the action of huge circuit breakers. Just like a household circuit breaker, these industrial breakers open their contacts within a fraction of a second, but because of the enormous amount of energy in the system, just separating the contacts doesn't stop the current. Instead, the current creates an electrical arc inside the breaker. That small space, which has a volume of just a few liters, now contains a roiling plasma that may reach temperatures of many thousands of degrees Celsius. The breaker can't contain that plasma for long; if it's not cleared away quickly, there will be a terrible explosion.

Now the alternating nature of the AC current comes into play: Each time it changes direction (every 10 milliseconds in China's 50-hertz system), the current temporarily becomes zero, and the energy supply to the arc plasma momentarily halts. It's at one of these "current zero" moments that the fault current must be interrupted. At that crucial moment, a cooling system inside the circuit breaker injects a high-pressure jet of gas into the gap, blasting away any residue of the hot arc plasma.

Immediately after the arc disappears and the fault is cleared, the power system ramps up again. In this recovery process, the voltage across the gap steeply rises to over 1 million volts before settling to its normal operational level. So in the microseconds before and after current zero, the contacts need to change over from channeling approximately 50 kiloamperes of current through the arc plasma to withstanding 1 megavolt of voltage. This rapid change puts enormous strain on the breakers' components.

Yet the circuit breakers must perform flawlessly, because the transmission line needs to go back into operation. They must work even though they may have been inactive for long stretches of time and through all kinds of weather. So how can the grid operator in our Chinese example trust that these breakers will do their jobs, and ensure that a megacity won't end up in the dark? Only rigorous testing can provide that peace of mind. I serve as innovation director for KEMA Laboratories, a Dutch division of the Norwegian consulting and certification company DNV GL. Our task: mimicking the stressful operating conditions of an ultrahigh-voltage AC system in extremis. Replicating this environment is a tremendously difficult engineering challenge, but it's one that must be met if we're to satisfy the energy demands of the coming decades.

Tomorrow's power grid will likely rely on large-scale renewable energy facilities such as hydropower plants, solar parks, and offshore wind farms, located far from power-hungry cities. To transport that energy across long distances, system operators are planning and constructing massive transmission lines. These lines must be high voltage, so they'll lose only a small fraction of energy through resistance in the lines. Building these cutting-edge high-voltage systems is quite expensive. But many power companies are deciding that the ability to move huge amounts of energy across vast distances justifies the costs.

Choosing to construct a high-voltage transmission system is the first step. The next step is to decide: DC or AC? High-voltage DC transmission systems are an increasingly attractive option, as DC overhead transmission lines require less space and lose less power than do AC

CIRCUIT BREAKERS 101

Most high-voltage transmission systems are protected from short circuits by gas circuit breakers. This simplified diagram shows the step-by-step process by which a breaker interrupts the dangerously high current flow created by a short circuit.

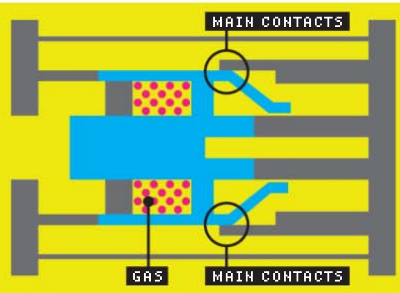

CLOSED: During normal operation, current flows through the circuit breaker's main contacts. When the power grid's protective system detects a short circuit, the breaker must open within milliseconds.

lines. But AC technology is more mature, and the world's most powerful transmission systems are still designed for AC. The latest AC supergrids use ultrahigh voltage (UHV) of at least 1,000 kilovolts, a staggering level not yet realized in DC. In this article, I'll focus on the gear required for AC networks.

The first commercial UHV-AC grid segment went into operation in China in January 2009. The State Grid Corp. of China spent 5.7 billion Chinese yuan (about US $900 million) on this 1,100-kV project, a 640-kilometer overhead line that connects China's north and central power grids. A total of 1,284 towers, each about 10 times as tall as the Great Wall, rise above the landscape of the Chinese interior and send electricity across both the Yellow and Han rivers. The towers support 25,000 metric tons of steel-reinforced aluminum conductors that can transport 5,000 MW. The system's three substations contain circuit breakers capable of interrupting a short-circuit current of up to 63 kA. In 2013, State Grid commissioned a similarly impressive east-west UHV-AC line stretching 650 km between Huainan and Shanghai, which will bring power from inland coal plants to the coastal cities.

Meanwhile, India is seeking to set a new voltage record by constructing a 1,200-kV AC supergrid. In 2012, India's Power Grid Corp. commissioned a test station for its UHV equipment, and a 350-km section of the 400-kV Wardha-Aurangabad line is now being upgraded to 1,200 kV. The country's power supply imbalances make such UHV projects particularly beneficial. The Wardha-Aurangabad line will bring electricity from coal-fired power plants in the center of the country to the city of Aurangabad, an emerging IT and manufacturing hub that struggles with power shortages.

We will also see pioneering projects in the decades to come that combine high-voltage AC and DC transmission networks into a "hybrid grid." (*IEEE Spectrum* explored this possibility in its August 2015 article "A Globe-Spanning Supergrid.") In Europe and China, and to a lesser extent in the United States, there's talk of hybrid systems that would use the advantages of both approaches. For example, a high-voltage undersea DC cable could efficiently take energy from a remote North Sea wind farm and bring it to an AC supergrid, where that energy could easily be transformed and shunted around Europe as needed.

In UHV transmission systems, the most crucial piece of technology is the circuit breaker. The breaker is the system's guardian: It must be eternally vigilant and prepared to act instantly. And it must function in all environmental conditions and despite great systemwide stress. At the KEMA test facility, in Arnhem, Netherlands, we put these breakers under extraordinary strain to provide an independent assessment of their performance. There's a clear need for this service: About a quarter of the circuit breakers brought to our labs fail to pass their tests.

Why not rely on simulations to study the stresses at work? Unfortunately, computer models aren't yet up to the task of simulating microsecond-scale interactions between electrical circuits and extremely hot and chemically complex plasmas. A study carried out by CIGRÉ, the International Council on Large Electric Systems, evaluated the simulation tools used by seven major manufacturers. First, the good news: These different tools did model the electrical fields at critical locations inside a circuit breaker with great accuracy and agreement. But when the tools modeled a breaker's failure—the point at which it succumbed to electric stress—they produced values quite different from each other and from the true tested value. It's like modeling the bending of a toothpick: It's easy to calculate the internal stresses, but the moment and the location of the wood's fracture can't be precisely predicted.

Therefore, at our lab we create real-world conditions to determine how the breakers will behave in the field. Of course, the electricity coming into our lab isn't quite what our tests demand, so we had to develop a few clever tricks that allow us to unleash powerful surges of both current and voltage. In a two-step test process, we mimic the two potentially catastrophic electrical stresses on a UHV circuit breaker in a precisely timed sequence.

First, short-circuit current must flow through the breaker while its contacts separate, thereby establishing an electrical arc inside. At present, we generate this current with four generators, each of which has a 54-metric-ton rotor moving at a speed matching the desired AC frequency, from 16.7 to 60 Hz; very soon

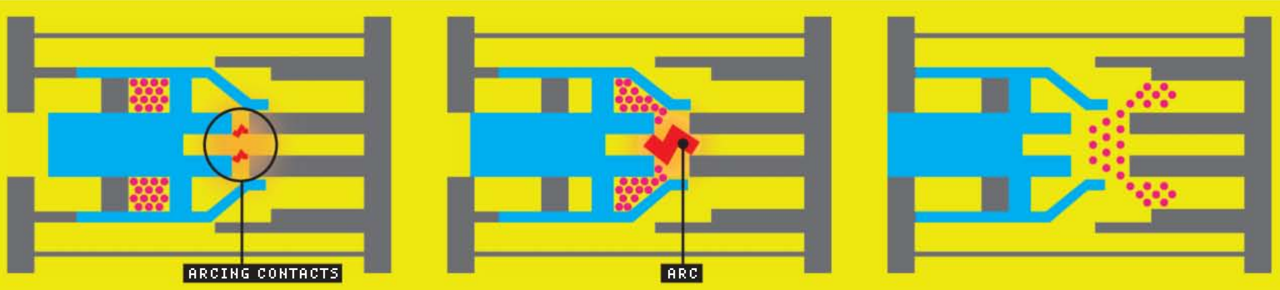

OPENING: First the breaker opens to interrupt the flow of current between the main contacts. With nowhere else to go, the current forms an electric arc in the gap between the arcing contacts.

OPEN: As the breaker's movable component slides over to create the gap, it compresses gas inside a chamber. That compressed gas then blasts through the gap.

RESET: The jet of gas clears away the plasma that constitutes an electric arc, ending the short circuit. The breaker will then close and resume normal operation.

ILLUSTRATION BY Erik Vrielink

1

Inside KEMA Labs [1], a circuit breaker (blue cylinder) is prepared for a test. Two capacitor banks (tall foreground structures) provide the ultrahigh voltage required. **A specialized transformer [2]** intended for KEMA's new lab expansion is checked on a barge before installation. **A short circuit [3]** at high current stresses a transmission line insulator in a lab demonstration. **The lab facility [4]**, on the banks of the Rhine River, offers a safe place to test equipment before it's tested in real-world conditions. **Powerful generators [5]** supply the current necessary to replicate the type of short circuits that can ruin today's most expansive and expensive transmission networks.

another two will be spinning. To initiate the short circuit, 12 synchronized switches energize the circuit, converting mechanical energy stored in the generators' rotors into electrical energy. We can draw a maximum current of 100 kA from each generator, enough to match the 80- to 90-kA short-circuit currents seen in the world's most powerful transmission networks.

So now we've reproduced the ruinously high currents seen in the first instant of a short circuit. But the breaker's work is not yet done. It must use its jet of gas to clear the arc from the gap at the critical current-zero moment and then, within microseconds, begin the circuit's recovery. During the recovery step, the components must withstand a transient voltage surge that far exceeds the circuit's typical voltage.

The 100-kA short-circuit current that we produce is available at only 17 kV. Half a dozen specialized transformers raise this voltage to 250 kV, but this is still far too low to properly test a UHV-rated circuit breaker. Using additional transformers to step up the voltage further doesn't make sense because the current would decrease accordingly. So we had to find another way.

For our second trick, we employ capacitor banks as large as four-story buildings, which are precharged to some 700 to 800 kV. At the critical moment, we trigger a spark gap, which discharges the capacitors in sequence. The first bank of capacitors supplies an initial megavolt, and then a few hundred microseconds later, a second capacitor bank adds another megavolt. That's how we strike the UHV breaker with a voltage that replicates what it would encounter in the field.

In 2008, we tested circuit breakers matching the Chinese supergrid's specifications using a pilot UHV installation, which supplied 2 MV within the millisecond following a short-circuit interruption. Now we are building a permanent installation, at a cost of $80 million, that will allow us to test both UHV circuit breakers and another key component of the supergrid: high-voltage transformers. Our tests have shown that around 25 percent of these transformers are damaged internally by the tremendous electrodynamic forces associated with short circuits. These transformers must survive the short-circuit current that flows through them for the brief span of time before the breakers do their job, which is no easy task.

It may surprise some electrical engineers to learn that such high-tech test equipment must be brought to bear on something as common as a circuit breaker. Didn't we figure out everything there is to know about breaker technology decades ago? In fact, the technology is still evolving, and so must our ability to evaluate its performance.

Oil-based circuit breakers dominated in the early 20th century. In these devices, the contacts sit inside a tank filled with oil; when the arc forms, it turns some of the oil into a high-pressure gas bubble, which surrounds and extinguishes the arc. But these oil tanks are unwieldy and dangerous contraptions. The 1970s saw the rise of circuit breakers

KEMA LABORATORIES

2

3

4

5

that use sulfur hexafluoride, an inert gas with good insulating properties, which is blasted through the gap to quench the arc. However, SF_6 is an extremely potent greenhouse gas, so the electricity industry is now developing alternative technologies.

Many researchers are investigating a type of circuit breaker that interrupts the current in a vacuum environment. The main difficulty here is managing the electric field in the vacuum. Because no gas or liquid is present, the arc plasma creates its own medium by releasing and ionizing metal vapor from the contacts themselves. As the hot plasma burns the contacts, it deforms their surfaces, creating microscopic dips and spikes. The protrusions jutting from a contact's surface are analogous to tall trees that tower above Earth's surface and are thus more likely to be struck by lightning. The roughened contacts may continue to channel current when they shouldn't, namely when the breaker is attempting to clear the arc.

To further develop these vacuum-based breakers, we need to test them under full-power conditions and study the current of mere tens of microamperes that is drawn out of the contacts' protrusions by the quantum mechanical tunneling effect. Do these tiny currents signal an imminent resumption of the electrical arc and thus the failure of the breaker? That's a heavily debated question in the scientific community. At KEMA Labs, we're looking for answers by evaluating the effect of these tiny currents in real equipment under full stress.

Indeed, many of our research projects involve working on the smallest scales imaginable. Surprisingly, very fast processes occurring on the microscale often determine whether the massive components of a transmission grid will malfunction–and possibly whether an entire city will go dark.

We must be able to study, for example, the events within the circuit breaker during the few microseconds surrounding the current-zero moment. In that minuscule time span, the breaker must change from a very good conductor to a nearly perfect insulator. With the newest tools, we are now able to monitor this transition. One high-resolution method we've developed can detect currents smaller than a single ampere and lasting just a few microseconds during full-scale short-circuit tests, which use currents that are measured in hundreds of kiloamperes. We look for these minute currents after the current-zero moment during which the arc should have been totally extinguished. If we find them, we have an indication that something is wrong with the breaker's recovery, and that the full electric arc may blaze anew inside its gap.

The electricity supergrids of the 21st century will rely on these circuit-breaking guardians, stationed like sentinels along continent-spanning transmission lines. And the breakers in turn rely on engineers in the test labs who unleash powerful floods of electricity to stress these components to the max, while also probing the intimate processes within the breakers' hearts. Only such top-to-bottom testing will ensure that the equipment can keep a Chinese megacity cool on even the hottest day. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/supergrid1215

KEMA LABORATORIES (4)

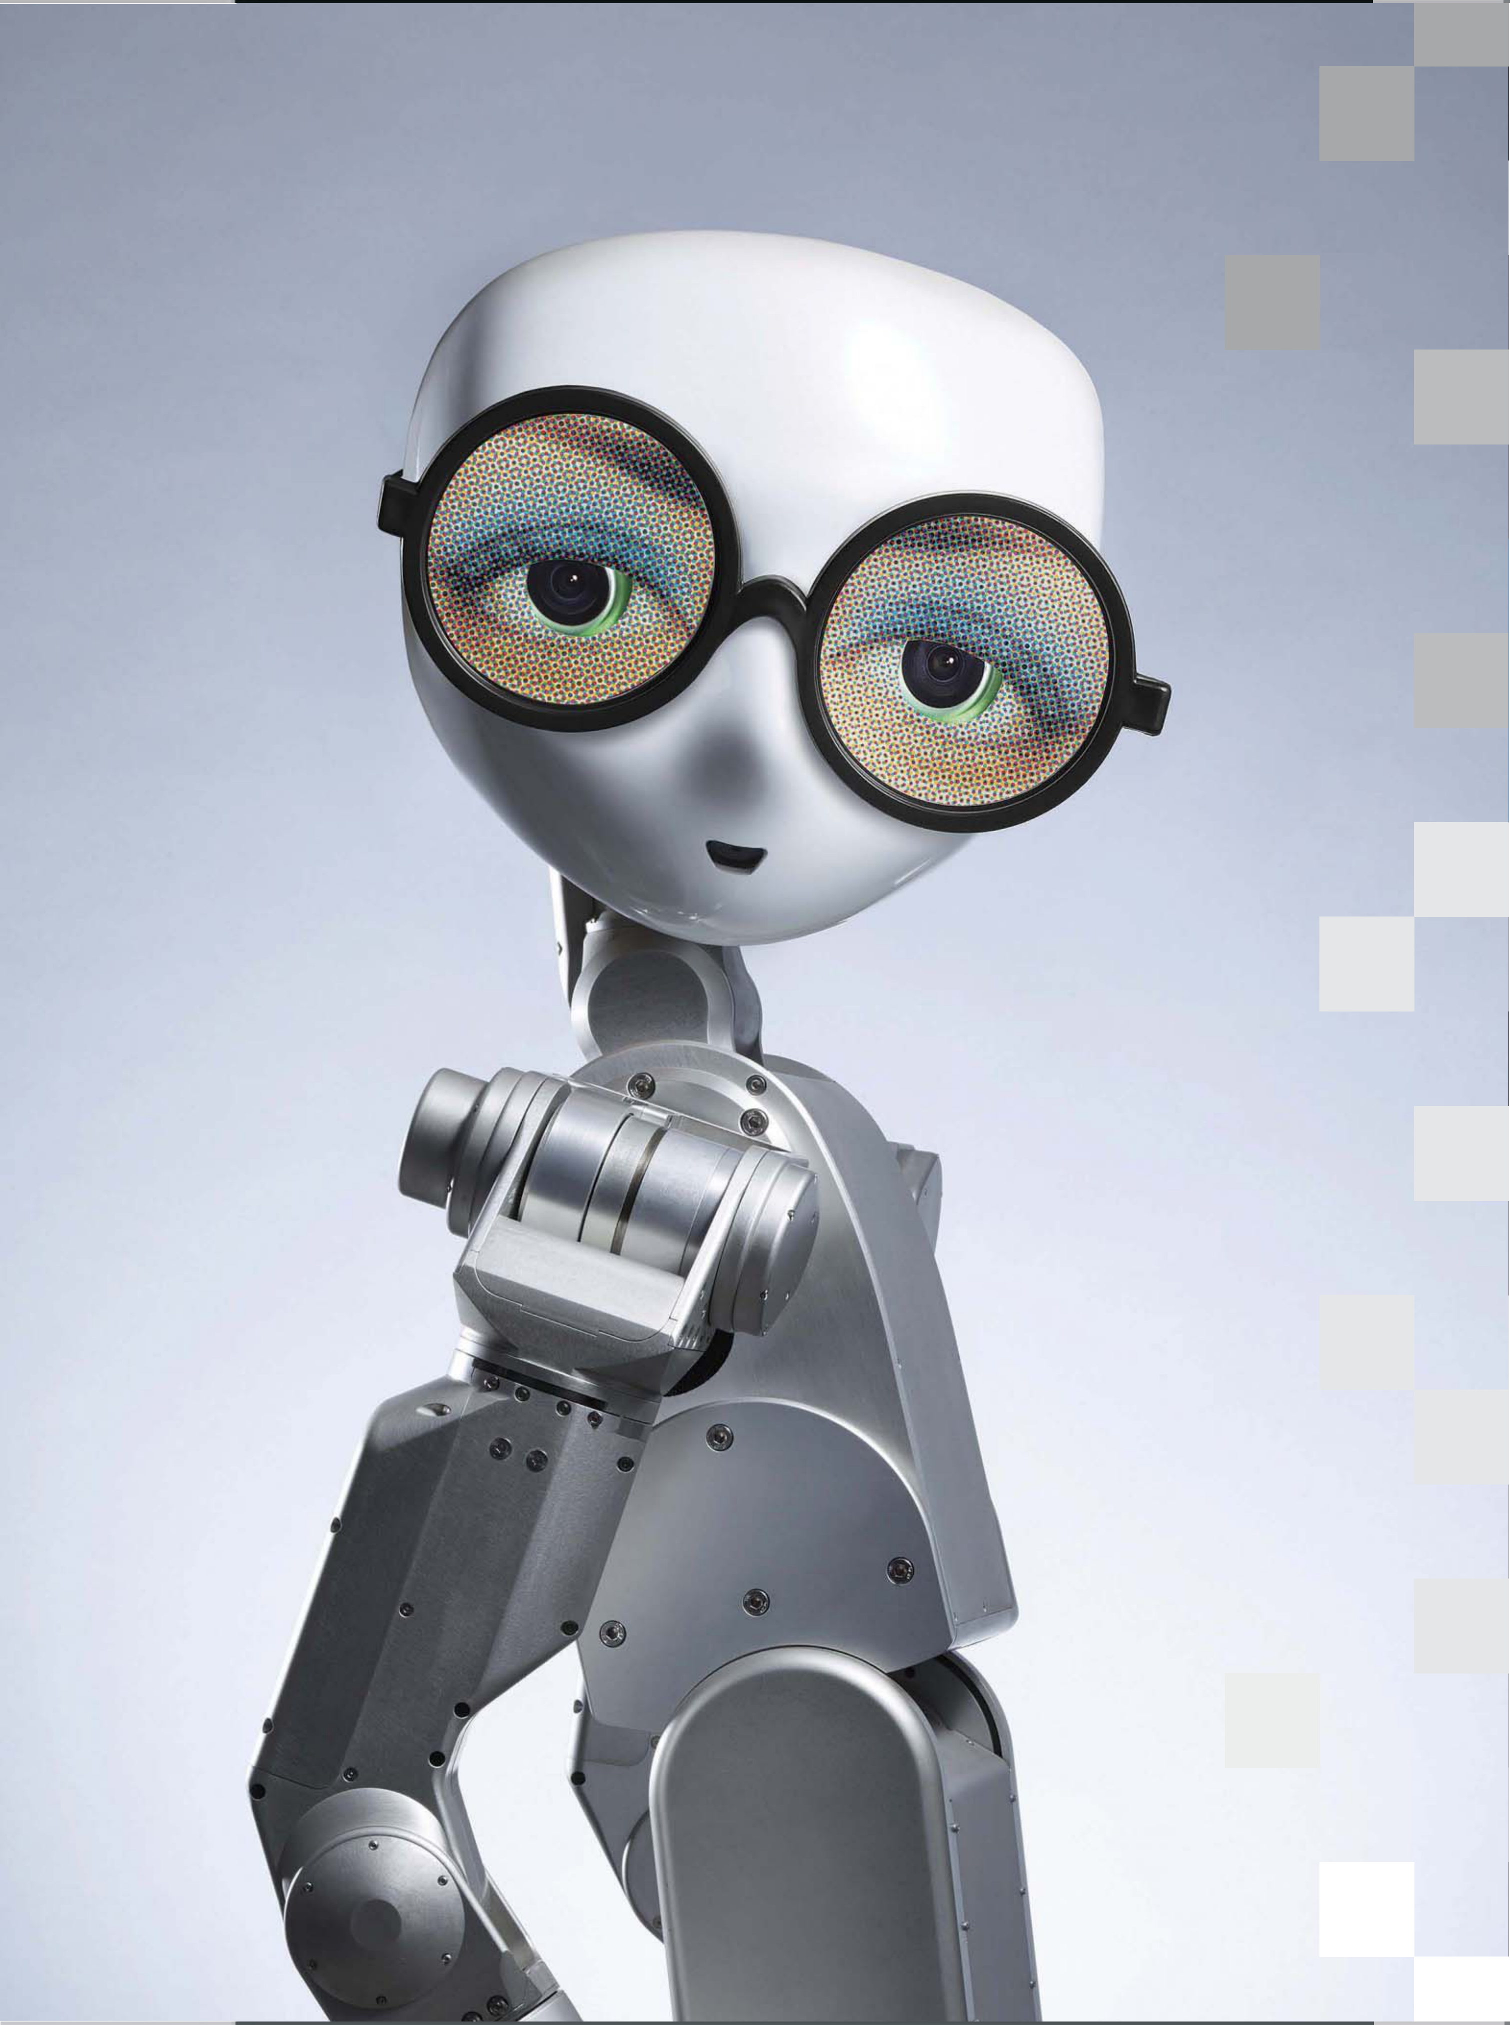

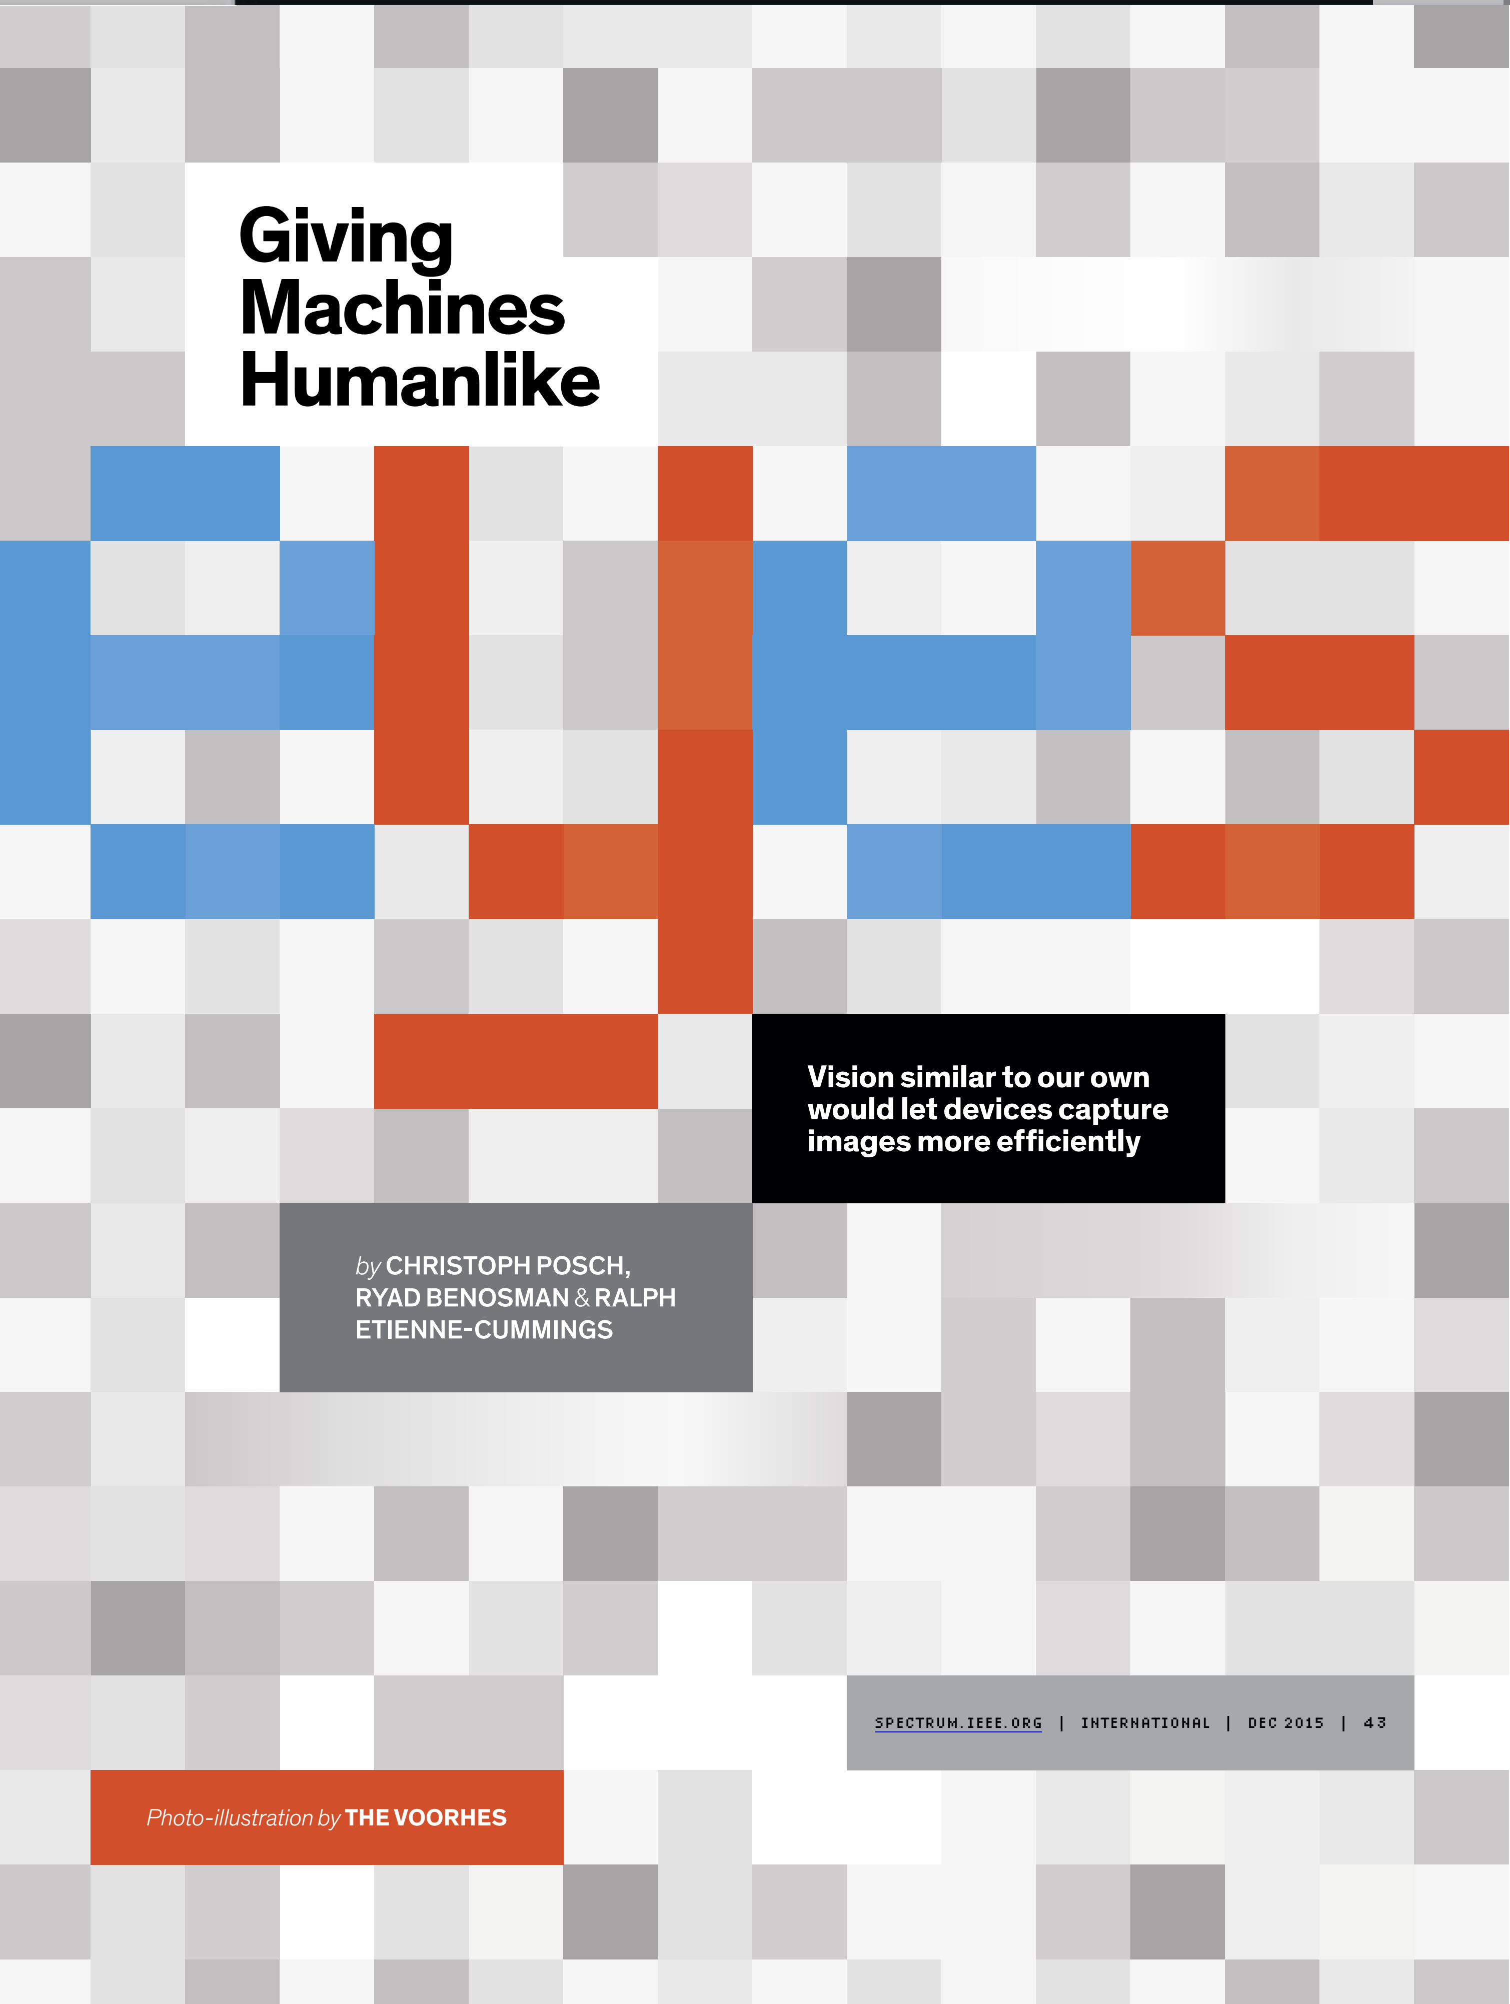

Giving Machines Humanlike Eyes

Vision similar to our own would let devices capture images more efficiently

by CHRISTOPH POSCH, RYAD BENOSMAN & RALPH ETIENNE-CUMMINGS

Photo-illustration by **THE VOORHES**

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

When Eadweard Muybridge set up his cameras at Leland Stanford's Palo Alto horse farm in 1878, he could scarcely have imagined the revolution he was about to spark. Muybridge rigged a dozen or more separate cameras using trip wires so that they triggered in a rapid-fire sequence that would record one of Stanford's thoroughbreds at speed. The photographic results ended a debate among racing enthusiasts, establishing that a galloping horse briefly has all four legs off the ground—although it happens so fast it's impossible for anyone to see. More important, Muybridge soon figured out how to replay copies of the images he took of animal gaits in a way that made his subjects appear to move.

Generations of film and video cameras, including today's best imaging systems, can trace their lineage back to Muybridge's boxy cameras. Of course, modern equipment uses solid-state detectors instead of glass plates, and the number of frames that can be taken each second is vastly greater. But the basic strategy is identical: You capture a sequence of still images, which when played back rapidly gives the viewer the illusion of motion. If the images are to be analyzed by a computer rather than viewed, there's no need to worry about whether the illusion is a good one, but you might still need to record lots of frames each second to track the action properly.

Actually, even with a high frame rate, your equipment may not be up to the task: Whatever you are trying to analyze could be changing too quickly. What then do you do? Many engineers would answer that question by looking for ways to boost the video frame rate using electronics with higher throughput. We argue that you'd be better off reconsidering the whole problem and designing your video equipment so it works less like Muybridge's cameras and instead functions more like his eyes.

The general strategy of creating electronic signal-processing systems inspired by biological ones is called neuromorphic engineering. For decades, this endeavor has been an exercise in pure research, but over the past 10 years or so, we and other investigators have been pursuing this approach to build practical vision systems. To understand how an artificial eye of the kind we've been investigating can outperform even a high-speed video camera, let us first disabuse you of the idea that the way modern video gear operates is sensible.

Imagine for a moment that you're trying to analyze something that happens really fast, say, a pitcher throwing a baseball. If you try to use a conventional video camera, which records at something like 30 or perhaps even 60 frames per second, you'll miss most of the movement of the pitcher's arm as he whips the ball toward the plate. Perhaps some frames will catch his arm in different positions. But you'll capture relatively little information of interest, along with much redundant imagery of the pitcher's mound, the infield turf, and other unchanging parts of the background. That is, the scene you record will be under- and oversampled at the same time!

There's no way to avoid that problem given that all parts of the image sensor in your camera share a common timing source. While this weakness won't be a problem for a casual viewer, if you wanted a computer to analyze nuances of the pitcher's arm motion, your data will be woefully inadequate. In some cases, sophisticated postprocessing might let you derive the results you wanted. But this brute-force approach would fail you in environments with limited power, bandwidth, and computing resources such as on mobile devices, multicopter drones, or other kinds of small robots.

The machine-vision community has been stuck with this basic problem for decades. But the situation may soon be changing for the better as we and other researchers develop equipment that samples different parts of the scene at different rates, mimicking how the eye works. With such gear, those parts of the scene that contain fast motions are sampled rapidly, while slow-changing portions are sampled at lower rates, going all the way down to zero if nothing changes.

OUT THE GAIT: A sequence of frames, like those of Eadweard Muybridge's animal-locomotion studies (such as the one shown above), undersamples parts of the scene—like the horse's swiftly moving legs here—while oversampling the parts that remain static.

CORBIS

Getting video cameras to work this way is tricky, because you don't know beforehand which parts of the scene will change and how rapidly they will do so. But as we describe below, the human eye and brain deal with this problem all the time. And the rewards of copying how they work would be enormous. Not only would it make fast-changing subjects–explosions, insects in flight, shattering glass–more amenable to analysis, it would also allow the video cameras on smartphones and other battery-operated devices to record ordinary motions using much less power.

Engineers often liken the eye to a video camera. There are some similarities to be sure, but in truth the eye is a much more complicated creation. In particular, people's retinas don't just turn light into electrical signals: They process the output of the eye's photoreceptor cells in sophisticated ways, capturing the stuff of interest–spatial and temporal changes–and sending that information to the brain in an amazingly efficient manner.

Knowing how well this approach works for eyes, we and others are studying machine-vision systems in which each pixel adjusts its own sampling in response to changes in the amount of incident light it receives. What's needed to implement this scheme is electronic circuitry that can track the amplitudes of each pixel continuously and record changes of only those pixels that shift in light level by some very small prescribed amount.

This approach is called level-crossing sampling. In the past, some people have explored using it for audio signals–for example, to cut down on the amount of data you'd have to record with the usual constant-rate sampling. And academic researchers have been building electronic analogues of the retina in silicon for research purposes since the late 1980s. But only in the past decade have engineers attempted to apply level-crossing sampling to the practical real-time acquisition of images.

Inspired by the biology of the eye and brain, we began developing imagers containing arrays of independently operating pixel sensors in the early 2000s. In our more recent cameras, each pixel is attached to a level-crossing detector and a separate exposure-measurement circuit. For each individual pixel, the electronics detect when the amplitude of that pixel's signal reaches a previously established threshold above or below the last-recorded signal level, at which point the new level is then recorded. In this way every pixel optimizes its own sampling depending on the changes in the light it takes in.

With this arrangement, if the amount of light reaching a given pixel changes quickly, that pixel is sampled frequently. If nothing changes, the pixel stops acquiring what would just prove to be redundant information and goes idle until things start to happen again in its tiny field of view. The electronic circuitry associated with that pixel outputs a new measurement just as soon as a change is detected, and it also keeps track of the position in the sensor array of the pixel experiencing that change. These outputs, or "events," are encoded according to a protocol called Address Event Representation, which came out of Carver Mead's lab at Caltech in the early 1990s. The train of events such a vision sensor outputs thus resembles the train of spikes you see when you measure signals traveling along a nerve.

The key is that the visual information is not acquired or recorded as the usual series of complete frames separated by milliseconds. Rather, it's generated at a much higher rate–but only from parts of the image where there are new readings. As a result, just the information that is relevant is acquired, transmitted, stored, and eventually processed by machine-vision algorithms.

We designed the level-crossing and recording circuits in our camera to react with blazing speed. With our equipment, data acquisition and readout times of a few tens of nanoseconds are possible in brightly lit scenes. For standard room-light levels, acquisition and readout require a few tens of microseconds. These rates are beyond all but the most sophisticated high-speed video cameras available today, cameras costing hundreds of thousands of dollars. And even if you could afford such a camera, it would deluge you with mostly worthless information. Sampling different pixels at different rates, on the other hand, reduces not just equipment cost but also power consumption, transmission bandwidth, and memory requirements–advantages that extend well beyond the acquisition stage. But you'll squander those benefits if all you do is reconstruct a series of ordinary video frames from the data so that you can apply conventional image-processing algorithms.

To fully unlock the potential of eyelike vision sensors, you need to abandon the whole notion of a video frame. That can be a little hard to get your head around, but as soon as you do that, you become liberated, and the subsequent processing you do to the data can resolve things that you could otherwise easily miss–including the detailed arm motions of our hypothetical baseball pitcher.

To do this, though, you'll have to rethink how you process the data, and you'll probably have to write new code instead of using a standard video-analysis library. But the mathematical formulations appropriate for this new kind of video camera are simple and elegant, and they yield some very efficient algorithms. Indeed, in applying such algorithms to the output of our autosampling vision sensors, we were able to show that certain real-time vision tasks could be run at a rate of tens to even hundreds of kilohertz, whereas conventional frame-based video-analysis techniques applied to the same situation topped out at a painfully slow 60 hertz.

Another advantage of analyzing the nearly continuous data streams from our eyelike sensors instead of a series of conventional video frames is that we can make good use of signal timing, just as biological neurons do. This is perhaps best explained with a specific example.

Event-Based Video at a Crossroads

AN EVENT-BASED vision sensor takes in a traffic intersection [top]. The sensor outputs a measurement only when the light striking a pixel shifts by a preset amount. When the light level crosses that threshold, circuitry attached to the pixel produces a "change-detection event" and triggers an illuminance measurement, encoded by the interval between two more events. The results for the entire pixel array can be viewed for a given time slice, but only changing pixels contribute to these images: the positive and negative level-crossing events [white and black, respectively, middle right] and associated illuminance measurements [nonblack areas, bottom right].

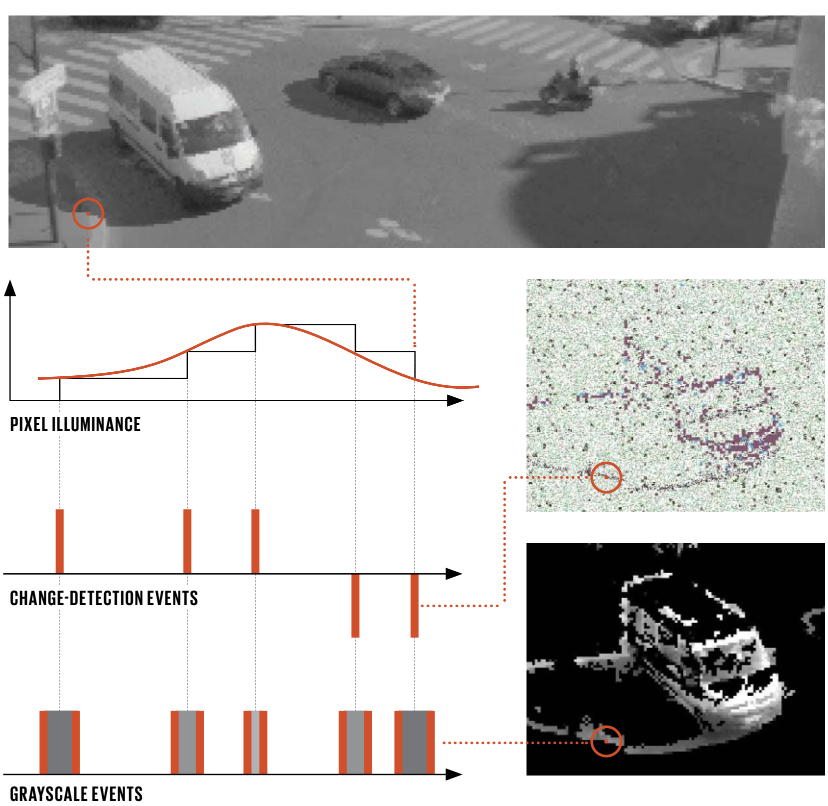

CHRONOCAM (3)

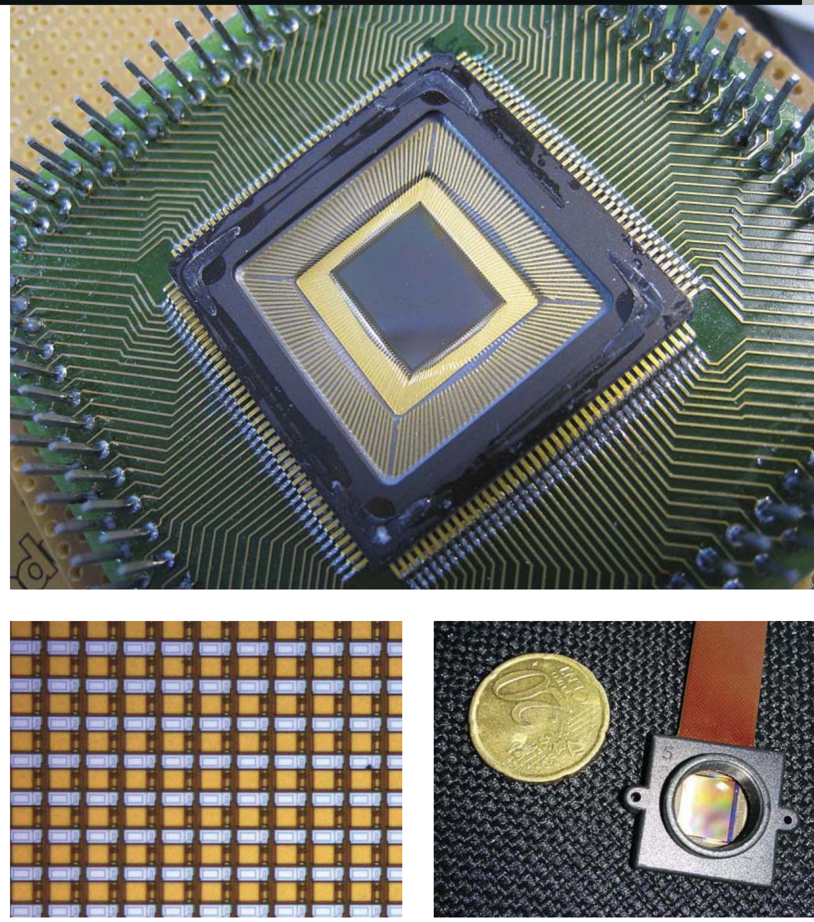

EYES OF SILICON: The author's asynchronous time-based image sensor is attached to a test fixture [top]. A close-up of a portion of the sensor's pixel array reveals its many photodiodes [blue areas, left]. A packaged vision sensor [right] serves as the input device for Pixium Vision's artificial retina.

CHRONOCAM (3)

Suppose you wanted to design a mobile robot that uses a machine-vision system to navigate its environment. Clearly, having a 3-D map of the things around it would be helpful. So you'd no doubt outfit the robot with two somewhat separated cameras so that it had stereo vision. That much is simple enough. But now you have to program its robotic brain to analyze the data it receives from its cameras and turn that into a representation of 3-D space.

If both cameras record something distinct–let's say it's a person stepping in front of the robot–it's easy enough to work out how far away the person is. But suppose two different people enter the robot's field of view at the same time. Or six people. Working out which one is which in the two camera views now gets more challenging. And without being able to ascertain identities for certain, the robot will not be able to determine the 3-D position of each one of these human obstacles.

With vision sensors of the type we've been studying, such matching operations become simpler: You just need to look for coincidences in the readings from the two cameras. If pixels from separate cameras register changes at the very same instant, they are almost certainly observing the same event. Applying some standard geometrical tests to the observed coincidences can further nail down the match.

Tobi Delbrück and others at ETH Zurich demonstrated the power of this approach in 2007 by building a small-scale robotic soccer goalie using an eyelike sensor that was broadly similar to ours. It had a reaction time under 3 milliseconds. (Peter Schmeichel, eat your heart out.) Were you to try to achieve that speed using a conventional video camera, you'd need to find one that could record some hundreds of frames per second, and the computational burden would be enormous. But with Delbrück's neuromorphic Dynamic Vision Sensor, the computer running his soccer goalie was loping along at a mere 4 percent CPU load.

Compared with standard video techniques, neuromorphic vision sensors offer increased speed, greater dynamic range, and savings in computational cost. As a result, demanding machine-vision tasks–such as mapping the environment in 3-D, tracking multiple objects, or responding quickly to perceived actions–can run at kilohertz rates on cheap battery-powered hardware. So this kind of equipment would allow for "always-on" visual input on smart mobile devices, which is currently impossible because of the amount of power such computationally intense tasks consume.

Another natural application of neuromorphic vision sensors is in electronic retinal implants for restoring sight to those whose vision has been lost to disease. Indeed, two of us (Posch and Benosman) helped to found Pixium Vision, a French company that has developed a neuromorphic retinal implant, which is now undergoing clinical trials. Unlike competing implants under development, which are frame based, Pixium's products use event-based sampling to provide patients with visual stimulation. Right now, these implants are able to give patients only a general ability to perceive light and shapes. But the technology should improve swiftly over the next few years and perhaps one day will be able to offer people who have lost their natural vision the ability to recognize faces–all thanks to artificial retinas inspired by real ones.

You can expect eyelike vision sensors to evolve from the pioneering designs available today into forms that eventually play a big role in medical technology, robotics, and more. Indeed, it wouldn't surprise us if they proved just as seminal as Muybridge's wooden cameras. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/machinevision1215

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

The Quest for the Ultimate Vacuum Tube

CONTINUED FROM PAGE 31

and Shoulders described nearly 50 years ago is still the most promising: field emission. All you need to make it work is a strong electric field at the surface between the emitter and the vacuum. Ordinarily, free electrons in a metal at room temperature encounter an energy barrier near the metal's surface that keeps them confined. But when you heat the metal–as in a thermionic emitter–you're giving some of those electrons enough energy to overcome the barrier.

Instead, let's apply an inward-directed electric field at the surface, which exerts an outward force on the electrons and thins the barrier; as the field increases, the barrier is thinned still more. To achieve high field emission at the level you would need in a traveling-wave tube, the barrier must be so thin that the electrons leak right through it. In quantum mechanical terms, the electrons have "tunneled" through the barrier. To get significant numbers of electrons to escape, you'll need to apply a very high electric field, of more than a billion volts per meter. For comparison, during a thunderstorm a local electric field of around 2,000 volts per meter is considered dangerous.

So how do you get a high field without having to apply excessively high voltage? The answer lies in nanofabrication. We know that the electric field between two conducting electrodes gets stronger as the electrodes are brought closer together. So if we bring the positively charged gate electrode (that is, the electrode that modulates the emission from the cathode) and the negatively charged emitter very close together–as in less than a micrometer apart–a strong inward electric field results at the surface of the emitter. We also shrink the emitter features to the nanometer scale and take advantage of the field enhancement that occurs naturally at sharp edges and points, sort of like a nano lightning rod. Now you have a field-emitter cathode that operates at less than 100 V.

To date, the best-performing cold-cathode device, as measured by total emission current and current density, was demonstrated by Spindt, Christopher Holland, and Paul Schwoebel at SRI. Their cathode consists of tens of thousands of micrometer-size molybdenum cones deposited on a circular silicon substrate with an area of about one square millimeter. Each cone acts as an electron emitter and sits in its own well, carved into a 2.5-micrometer-thick silicon-dioxide insulating layer. A thin metal coating over the insulating layer serves as the gate electrode for the array of field emitters.

For more than a decade, my colleague David Whaley and I, working with the SRI team, have been trying to turn this laboratory device into something more practical. Milestones to date include the first operation of a 100-W cold-cathode TWT at a frequency of 4 to 6 GHz and the first 100-milliampere electron beam generated by a cold cathode. Whaley and I also demonstrated a TWT with current densities from the cold cathode exceeding 15 amperes per square centimeter, which far exceeds the typical 2 A/cm^2 from a thermionic cathode. The current density from our cathode was so high, in fact, that to get the beam to travel through the helix, we had to expand the beam first rather than compress it.

The cathode worked well for more than 150 hours before breaking down electrically. Breakdown is a major challenge for these field emitters because of the high fields within the structure and the thin-film gate electrode. An electrical short between the gate and any one of the thousands of emitter cones will burn out the entire array of emitters and render it unusable.

Recently, our colleagues at SRI have developed a feature to mitigate the damaging effects of such shorts. They interrupt the electrical breakdown path between the base of the cones and the gate by adding a dielectric layer between them–a shield, if you will. My company, L-3 Electron Devices, tested one of these SRI cathodes in a traveling-wave tube that generated up to 10 W at 18 GHz–a record-high frequency for a cold-cathode device. Testing showed that the cathode was much more resistant to, and tolerant of, individual emitter failures.

One challenge we're still addressing is the traveling-wave tube's high operating voltage. While the field-emitter cathode can emit electrons at around 100 V, you need a much higher voltage to accelerate and focus the electron beam coming from the cathode to interact with the helix circuit. But you also need to avoid any high-voltage arcs between, for example, the anode and the field-emitter cathode. We're also working on ways to prevent stray ions from traveling back upstream and striking the cathode's gate and field emitter tips. In our TWT, we've created an electrostatic ion barrier to prevent such ions from making it to the cathode.

While much work remains on the cold-cathode TWT, I'm confident we will have a practical device within five years. Advances like this one underscore the ongoing relevance and vitality of vacuum electronics, even after more than a century of development. Maybe that's why the Defense Advanced Research Projects Agency announced in August a pair of new programs aimed at pushing the field even further. Vacuum technology continues to smash through performance barriers and open up new areas of application. It refuses to fade into obscurity. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/travelingwavetube1215

NOTE: The views are those of the author.

BATTLING BORKED BITS

CONTINUED FROM PAGE 35

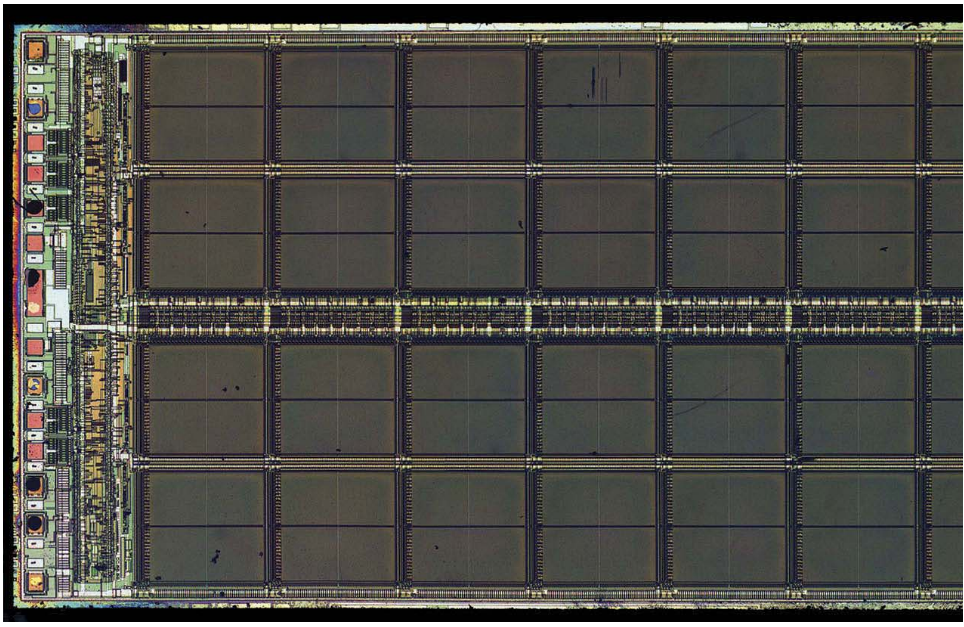

ZEPTOBARS

less severe and thus correctable. So with the proper mechanisms in place, most crash-inducing DRAM glitches could be prevented. Indeed, rather than just relying on ECC or other expensive kinds of correction hardware, computer operating systems themselves could help to protect against memory errors.

Operating systems typically divide the computer's main memory into areas known as pages (usually 4 kilobytes in size). And from the operating system's point of view, the majority of errors come from a very small fraction of the available pages. Indeed, more than 85 percent of the errors come from just 0.0001 percent of all the pages. By removing those problematic pages from use, the operating system could prevent most of the risk of errors without giving up much memory capacity.

Although others have previously proposed this very tactic, known as page retirement, it's seldom used. None of the organizations we worked with for our DRAM study employed it. One reason might be that people just didn't realize how often DRAM errors tend to crop up in the same page of memory. After our research brought that fact to light, Facebook adopted page retirement in its data centers. But it's still not widely used, perhaps because there's still some confusion about what sort of page-retirement schemes would work best.

To help clarify that issue, we investigated five simple policies. Some were as conservative as "Retire a page when you see the first error on it." Others involve trying to prevent row or column errors by retiring the entire row or column as soon as one memory cell starts to show problems. We used the data we collected to see how well such policies might protect real computer systems, and we found that almost 90 percent of the memory-access errors could have been prevented by sacrificing less than 1 megabyte of memory per computer–a tiny fraction of what is typically installed.

Sure, a technician could replace the entire memory module when it starts to have errors. But it would probably have only a few bad cells. Page retirement could isolate the regions of memory prone to errors without sacrificing the other parts of an otherwise functional module. Indeed, applying sensible page-retirement policies in large data centers and supercomputing facilities would not only prevent the majority of machine crashes, it would also save the owners money.

The same applies to consumer gadgets. With the growing ubiquity of smartphones, tablets, wearable electronics,

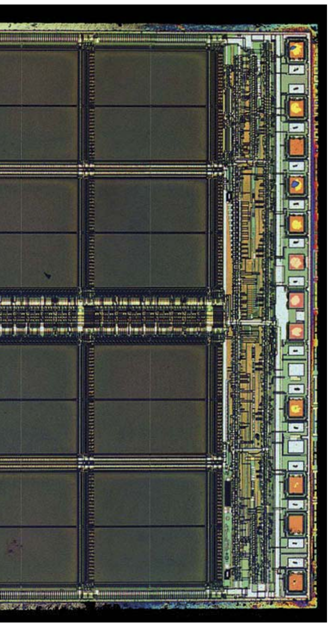

ROWS AND COLUMNS: This image shows the internal makeup of a typical DRAM chip, in this case Micron Technology's MT4C1024, which stores 2^{20} bits of information (1 mebibit).

and other high-tech gear, the number of devices that use DRAM memory is skyrocketing. And as DRAM technology advances, these devices will contain ever-larger quantities of memory, much of which is soldered straight onto the system board. A hard error in such DRAM would normally require replacing the entire gizmo. So having the software retire problematic parts of memory in such an environment would be especially valuable.

Had we accepted the received wisdom that cosmic rays cause most DRAM errors, we would never have started looking at how these chips perform under real-world conditions. We would have continued to believe that distant events far off in the galaxy were the ultimate cause of many blue screens and server crashes, never realizing that the memory errors behind them usually stem from underlying problems in the way these chips are constructed.

In that sense, DRAM chips are a little like people: Their faults are not so much in their stars as in themselves. And like so many people, they can function perfectly well once they compensate for a few small flaws. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/dram1215

SMU | LYLE
SCHOOL OF ENGINEERING

Southern Methodist University

The Southern Methodist University, Department of Electrical Engineering, invites applications for a full-time tenured-track faculty position at the rank of Assistant Professor. Candidates in all areas of interest within electrical engineering will be considered, including but not limited to the following: smart grids, power systems, signal processing, wireless communications and networking, electromagnetics, general analog/digital mixed-mode circuits, biomedical, photonics, MEMS, semiconductor devices, and control systems. Exceptional candidates in other closely related areas will also be considered. Candidates with exceptional prior records will also be considered at the Associate Professor or Professor levels.

Successful candidates must have a PhD degree in Electrical Engineering or a related field, and are expected to teach undergraduate and graduate courses, as well as to develop an innovative, externally funded, research program. Successful candidates will also be expected to supervise Master's and Doctoral students and participate in interdisciplinary efforts within the School.

SMU is a private university dedicated to academic excellence. Located in Dallas, SMU maintains a moderate size of about 11,000 students. The Electrical Engineering Department resides within the Lyle School of Engineering. The department offers B.S., M.S., and Ph. D. degrees in Electrical Engineering and a M.S. degree in Telecommunications. Additional information is provided at: **http://lyle.smu.edu/ee**. To learn more about the rich cultural environment of SMU, please see: **http://smu.edu**.

SMU is designated as a preferred employer in the Dallas/Forth Worth metroplex, one of the most prolific high-tech industrial centers in the country. The Dallas/Fort Worth metroplex is a multi-faceted business and engineering community, offering exceptional museums, diverse cultural attractions and a vibrant economy. Dallas' quality of life is exceptional with a relatively low cost of living, upscale apartments and homes within walking distance of campus, the opportunity to live in the city or out in the country with a relatively short commute, and the availability of both mass transit systems and plentiful on-campus parking.

Interested and qualified individuals should send a letter of application, curriculum vitae, a statement of educational interests, a research plan, and a list of five references to **eejobs@lyle.smu.edu** (preferred) or by mail to: Dr. Dinesh Rajan, Electrical Engineering Faculty Search, Electrical Engineering Department, P.O. Box 750338, Dallas, TX 75275-0338. The position will begin the Fall 2016 semester. To ensure full consideration, applications must be emailed or postmarked prior to January 15, 2016. The committee, however, will continue to accept applications until the position is filled. Please reference position number 53724.

SMU will not discriminate in any employment practice, education program, education activity, or admissions on the basis of race, color, religion, national origin, sex, age, disability, genetic information, or veteran status. SMU's commitment to equal opportunity includes nondiscrimination on the basis of sexual orientation and gender identity and expression. The Executive Director of Access and Equity/Title IX* Coordinator is designated to handle inquiries regarding the nondiscrimination policies, including the prohibition of sex discrimination under Title IX. The Executive Director/Title IX Coordinator may be reached at the Perkins Administration Building, Room 204, 6425 Boaz Lane, Dallas, TX 75205, 214-768-3601, **accessequity@smu.edu**. Inquiries regarding the application of Title IX may also be directed to the Assistant Secretary for Civil Rights of the U.S. Department of Education. Hiring is contingent upon the satisfactory completion of a background check. SMU is committed to achieving excellence through diversity. The university actively encourages applications and/or nominations of women, persons of color, veterans and persons with disabilities.

*Title IX of the Education Amendments of 1972, 20 U.S.C. §§ 1681-1688.

BAYLOR UNIVERSITY

Faculty Position

The Electrical and Computer Engineering Department of Baylor University seeks faculty applicants for a tenured/tenure-track Faculty Position at any level. Any area of expertise will be considered but applicants in computer engineering will be given special consideration. Applicants for assistant professor must demonstrate potential for sustained, funded scholarship and excellent teaching; applicants for associate or full professor must present evidence of achievement in research and teaching commensurate with the desired rank. The ECE department offers B.S., M.S., M.E. and Ph.D. degrees and is rapidly expanding its faculty size. Facilities include the Baylor Research and Innovation Collaborative (BRIC), a newly-established research park minutes from the main campus.

Chartered in 1845 by the Republic of Texas, Baylor University is the oldest university in Texas. Baylor has an enrollment of over 15,000 students and is a member of the Big XII Conference. Baylor's mission is to educate men and women for worldwide leadership and service by integrating academic excellence and Christian commitment within a caring community. The department seeks to hire faculty with an active Christian faith; applicants are encouraged to read about Baylor's vision for the integration of faith and learning at **www.baylor.edu/profuturis/**.

Applications will be considered on a rolling basis until the **January 1, 2016** deadline. Applications must include:

1) a letter of interest that identifies the applicant's anticipated rank,

2) a complete CV,

3) a concise statement of teaching and research interests,

4) the names and contact information for at least four professional references.

Additional information is available at **www.ecs.baylor.edu**. Send materials via email to Dr. Keith Schubert at keith_schubert@baylor.edu. Please combine all submitted material into a single pdf file.

Baylor University is affiliated with the Baptist General Convention of Texas. As an Affirmative Action/Equal Employment Opportunity employer, Baylor encourages candidates of the Christian faith who are minorities, women, veterans, and persons with disabilities to apply.

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

香港中文大學(深圳)
The Chinese University of Hong Kong, Shenzhen

Professor/Associate Professor/ Assistant Professor

The School of Science and Engineering

Located in Longgang District of Shenzhen, CUHK(SZ) is a research-intensive university established through a Mainland–Hong Kong collaboration, with campus and infrastructure provided by the Shenzhen Government. It will develop its academic programmes in phases and offer courses in Schools of Science and Engineering, Management and Economics, and Humanities and Social Science. The language of instruction will be in both English and Chinese, and the students will receive degrees of The Chinese University of Hong Kong.

The School of Science and Engineering has admitted its first batch of students in four undergraduate programmes (Computer Science and Engineering, Electronic Information Engineering, New Energy Science and Engineering, and Statistics) and one graduate programme (Financial Engineering) in 2015. Other new programmes will be gradually launched in the next few years. More information on the School can be found at **www.cuhk.edu.cn/en/Xueyuan/ligong.html**.

Faculty positions are available in all related fields of Computer Science and Engineering, Electrical and Information Engineering, New Energy Science and Engineering, Statistical Science, Data Science, Mathematics, Bioinformatics and Genomics, Financial Engineering and Quantitative Finance. Applications in areas of Design and Manufacturing Systems, Industrial Engineering and Operations Research, Physics, Chemistry, Biology, Material Science and Engineering, Environmental Science and Engineering, etc., will also be considered.

Junior applicants should have (i) a PhD degree (by the time of reporting for duty) in related fields; and (ii) high potential in teaching and research. Candidates for Associate and Full Professor posts are expected to have demonstrated academic leadership and strong commitment to the highest standards of excellence. Appointments will normally be made on contract basis for up to three years initially, leading to longer-term appointment or tenure later subject to review. Exceptionally, appointment with tenure can be offered forthwith to candidates of proven ability.

Applications (with CV, and three references) should be emailed to **hr-1@cuhk.edu.cn**.

SMU | LYLE
SCHOOL OF ENGINEERING

Southern Methodist University

The Southern Methodist University, Department of Electrical Engineering, invites applications for one full-time tenured Full Professor. The successful candidate will also hold the title of Mary and Rich Templeton Centennial Chair in Electrical Engineering. Candidates in all areas of interest within electrical engineering will be considered, including but not limited to the following: smart grids, power systems, signal processing, wireless communications and networking, electromagnetics, general analog/digital mixed-mode circuits, biomedical, photonics, MEMS, semiconductor devices, and control systems. Exceptional candidates in other closely related areas will also be considered.

Successful candidates must have a PhD degree in Electrical Engineering or a related field, and are expected to teach undergraduate and graduate courses, as well as to develop an innovative, externally funded, research program. Successful candidates must have demonstrated their ability to build and sustain a world class research program. Successful candidates will be expected to supervise Master's and Doctoral students and participate in interdisciplinary efforts within the School.

SMU is a private university dedicated to academic excellence. Located in Dallas, SMU maintains a moderate size of about 11,000 students. The Electrical Engineering Department resides within the Lyle School of Engineering and is located in the Jerry R. Junkins Building, completed in August 2002. The Jerry R. Junkins Building houses a research laboratory complex with a 2,800 square foot Class 10,000 cleanroom. The department offers B.S., M.S., and Ph. D. degrees in Electrical Engineering and a M.S. degree in Telecommunications. Additional information is provided at: **http://lyle.smu.edu/ee**. To learn more about the rich cultural environment of SMU, please see: **http://smu.edu**.

SMU is designated as a preferred employer in the Dallas/Forth Worth metroplex, one of the most prolific high-tech industrial centers in the country. The Dallas/Fort Worth metroplex is a multi-faceted business and engineering community, offering exceptional museums, diverse cultural attractions and a vibrant economy. Dallas' quality of life is exceptional with a relatively low cost of living, upscale apartments and homes within walking distance of campus, the opportunity to live in the city or out in the country with a relatively short commute, and the availability of both mass transit systems and plentiful on-campus parking.

Interested and qualified individuals should send a letter of application, curriculum vitae, a statement of educational interests, a research plan, and a list of five references to **eejobs@lyle.smu.edu** (preferred) or by mail to: Dr. Dinesh Rajan, Electrical Engineering Faculty Search, Electrical Engineering Department, P.O. Box 750338, Dallas, TX 75275-0338. The position will begin on or before the Fall 2016 semester. To ensure full consideration, applications must be emailed or postmarked prior to January 15, 2016. The committee, however, will continue to accept applications until the position is filled. Please reference position number 53659.

SMU will not discriminate in any employment practice, education program, education activity, or admissions on the basis of race, color, religion, national origin, sex, age, disability, genetic information, or veteran status. SMU's commitment to equal opportunity includes nondiscrimination on the basis of sexual orientation and gender identity and expression. The Executive Director of Access and Equity/Title IX* Coordinator is designated to handle inquiries regarding the nondiscrimination policies, including the prohibition of sex discrimination under Title IX. The Executive Director/Title IX Coordinator may be reached at the Perkins Administration Building, Room 204, 6425 Boaz Lane, Dallas, TX 75205, 214-768-3601, **accessequity@smu.edu**. Inquiries regarding the application of Title IX may also be directed to the Assistant Secretary for Civil Rights of the U.S. Department of Education. Hiring is contingent upon the satisfactory completion of a background check. SMU is committed to achieving excellence through diversity. The university actively encourages applications and/or nominations of women, persons of color, veterans and persons with disabilities.

**Title IX of the Education Amendments of 1972, 20 U.S.C. §§ 1681-1688.*

CityU
香港城市大學
City University of Hong Kong
專業 創新 胸懷全球
Professional · Creative
For The World

57th World University | 5th World's top 50 Universities under age 50
1st Engineering/Technology/Computer Sciences in Hong Kong | 2nd Business School in Asia

Worldwide Search for Talent

City University of Hong Kong is a dynamic, fast-growing university that is pursuing excellence in research and professional education. As a publicly-funded institution, the University is committed to nurturing and developing students' talents and creating applicable knowledge to support social and economic advancement. The University has seven Colleges/Schools. As part of its pursuit of excellence, the University aims to recruit **outstanding scholars** from all over the world in various disciplines, including **business**, **creative media**, **energy**, **engineering**, **environment**, **humanities**, **law**, **science**, **social sciences**, **veterinary sciences** and other strategic growth areas.

Applications and nominations are invited for :

Chair Professor/Professor/ Associate Professor/Assistant Professor Department of Electronic Engineering

[Ref. A/187/36]

Duties : Teach and supervise student projects at both undergraduate and postgraduate levels, conduct frontline research, and develop new research directions and courses, where applicable, in the areas of software systems, data structure and algorithm, object-oriented programming, SAS programming, digital or computer forensics, data centre, data mining and cloud computing.

Requirements : A PhD in Computer Engineering/Computer Science or related disciplines, with a good honours degree in Computer Engineering/Computer Science plus a strong research record.

Salary and Conditions of Service
Remuneration package will be driven by market competitiveness and individual performance. Excellent fringe benefits include gratuity, leave, medical and dental schemes, and relocation assistance (where applicable). Initial appointment will be made on a fixed-term contract.

Information and Application
Further information on the posts and the University is available at http://www.cityu.edu.hk, or from the Human Resources Office, City University of Hong Kong, Tat Chee Avenue, Kowloon Tong, Hong Kong [Email : hrojob@cityu.edu.hk/Fax : (852) 2788 1154 or (852) 3442 0311].

To apply, please submit an online application at http://jobs.cityu.edu.hk, and include a current curriculum vitae. Nominations can be sent directly to the Human Resources Office. **Applications and nominations received before 31 January 2016 will receive full consideration.** Only shortlisted applicants will be contacted; and those shortlisted for the post of Assistant Professor will be requested to arrange for at least 3 reference reports sent directly by the referees to the Department, specifying the position applied for. The University's privacy policy is available on the homepage.

City University of Hong Kong is an equal opportunity employer and we are committed to the principle of diversity. Personal data provided by applicants will be used for recruitment and other employment-related purposes.

Worldwide recognition ranking 57th (QS survey 2015); 5th among top 50 universities under age 50 (QS survey 2014); 1st in Engineering/ Technology/Computer Sciences in Hong Kong (Shanghai Jiao Tong University survey 2015); and 2nd Business School in Asia-Pacific region (UT Dallas survey 2014).

The Ohio State University

The Ohio State University invites applications for multiple tenure track faculty positions in the Department of Electrical and Computer Engineering. All areas and ranks will be considered. We are especially interested in (i) control systems or robotics for intelligent vehicle and transportation systems, (ii) biomedical applications including cancer detection and imaging, tumor cell biology, or the genomics of cancer, (iii) electronic materials and devices, emphasizing low energy electronics or wide bandgap semiconductors, (iv) neural electronics, and (v) senior candidates in electromagnetics, microwave theory, RF systems, or remote sensing. All positions may involve joint appointments with other engineering departments. Applicants must have a Ph.D. and outstanding academic credentials. Successful candidates will be expected to develop a vigorous externally-funded research program, show excellence and leadership in academic and scholarly activities, and demonstrate outstanding teaching at the undergraduate and graduate levels. Applicants are requested to send a letter of application, curriculum vitae, statement of research plans, brief statement of teaching philosophy, and name, address, and email of four references to Professor John L. Volakis at search@ece.osu.edu. Several of these positions are partially funded by Ohio State's Discovery Themes Initiative, a significant faculty hiring investment in key thematic areas that build on the university's culture of academic collaboration to make a global impact. The Ohio State University is committed to establishing a culturally and intellectually diverse environment, encouraging all members of our learning community to reach their full potential. We are responsive to dual-career families and strongly promote work-life balance to support our community members through a suite of institutionalized policies. We are an NSF Advance Institution and a member of the Ohio/Western Pennsylvania/West Virginia Higher Education Recruitment Consortium.

FACULTY POSITIONS

Electrical Engineering

NYU SHANGHAI

NYU Shanghai invites outstanding applications from all ranks to apply for tenure-track or tenured faculty appointments in Electrical and Computer Engineering (ECE) with special emphasis in Computer Engineering, Multimedia Analytics, Neural Engineering, and Robotics. Exceptionally strong candidates in other areas and academic ranks in ECE will also be considered. The applicants must have demonstrated abilities in both research and teaching. In describing teaching experience, candidates should identify courses they could teach both within and outside their specialty. Candidates must have completed a Ph.D. or equivalent by the time of appointment.

The terms of employment in NYU Shanghai are comparable to U.S. institutions in terms of research start-up funds and compensation, and include housing subsidies and educational subsidies for children. Faculty may also spend time at NYU New York and other sites of the NYU global network, engaging in both research and teaching.

Applications are due no later than February 1, 2016 and will be reviewed until the position is filled. To be considered, candidates should submit a curriculum vitae, separate statements of research and teaching interests (no more than three pages each), and electronic copies of up to five recent, relevant publications. To complete the online process, applicants will be prompted to enter the names and email addresses of at least three referees. Each referee will be contacted to upload their reference letter. Senior candidates are not required to submit references at this time. Please visit our website at **http://shanghai.nyu.edu/about/work/faculty-positions** for instructions and other information on how to apply. If you have any questions, please e-mail **shanghai.engineering.recruitment@nyu.edu.**

About NYU Shanghai:
NYU Shanghai is the newest degree-granting campus within the NYU Global Network University. It is the first Sino-US higher education joint venture to grant a degree that is accredited in the U.S. as well as in China. All teaching is conducted in English. A research university with liberal arts and science at its core, it resides in one of the world's great cities with a vibrant intellectual community. NYU Shanghai recruits scholars who are committed to NYU's global vision of transformative teaching and innovative research.

New York University has established itself as a Global Network University, a multi-site, organically connected network encompassing key global cities and idea capitals. There are three degree-granting campuses in New York, Shanghai, and Abu Dhabi, and complemented by eleven additional academic centers across five continents. Faculty and students circulate within the network in pursuit of common research interests and cross-cultural, interdisciplinary endeavors, both local and global.

上海纽约大学
NYU SHANGHAI

EOE/Minorities/Females/Vet/Disabled

University of Minnesota
Driven to Discover℠

Electrical and Computer Engineering, University of Minnesota – Twin Cities, invites applications for faculty positions in (1) power and energy systems, (2) communications and signal processing, (3) RF/microwave/terahertz circuits and systems, and as part of the MnDRIVE Initiative https://mndrive.umn.edu/robotics (4) embedded systems. Women and other underrepresented groups are especially encouraged to apply. An earned doctorate in an appropriate discipline is required. Rank and salary will be commensurate with qualifications and experience. Positions are open until filled, but for full consideration, apply at **http://www.ece.umn.edu/** by December 15, 2015. The University of Minnesota is an equal opportunity employer and educator.

College of Engineering, Michigan State University

Faculty Positions in Coatings and Diamond Technologies

Join Michigan State University's Global Impact Initiative, designed to address the grand challenges through the creation of over 100 new faculty positions in some of the most promising and exciting fields of research. We welcome applicants from diverse backgrounds. MSU offers an inclusive and collaborative work environment. To learn more visit **http://research.msu.edu/global-impact**.

The Departments of Electrical and Computer Engineering (ECE) and Chemical Engineering and Materials Science (CHEMS) at Michigan State University invite applications for a junior level tenure-system faculty position and an open-rank, fixed-term faculty position in the areas of electronic materials, devices or functional coatings. The tenure-system position is targeted at the junior faculty level, but senior candidates with outstanding credentials may be considered. These faculty positions are part of an expansion of the MSU/Fraunhofer Center for Coatings and Diamond Technologies (CCD), which is located on the MSU campus and includes 30 individuals consisting of faculty, graduate students, Fraunhofer USA engineers and staff, and student interns (**http://www.egr.msu.edu/fraunhofer-ccd**). CCD is looking to expand its research and development activities in the areas of diamond and ultra-wide bandgap electronics, devices, and sensors and in the area of hard functional coatings. The current CCD activities include diamond synthesis, diamond synthesis machine design, diamond electronics, diamond device and sensor microfabrication, diamond-like carbon (DLC) coatings and their applications and other hard, functional coatings.

Exceptional candidates should have an established record of excellence in one or more of the following areas: wide bandgap electronics, electronic device design, modeling, simulation, fabrication and testing; electronic device packaging, especially high power and high frequency wide bandgap electronics; diamond photonics, optics and X-ray optics; diamond electro-chemistry and sensors; functional coatings; electronic materials characterization; wide bandgap electronics reliability and lifetime; or diamond and related wide bandgap semiconductor synthesis and processing. Candidates with either experimental or computational expertise in wide bandgap material, device and sensors are sought. Also sought is a candidate with experimental characterization expertise to lead the materials characterization effort in the coatings and diamond areas.

The openings are available starting August 16, 2016. The tenure-system position include research, teaching and service responsibilities, while the fixed-term position has primarily research responsibilities. These appointments can be either in the ECE Department or CHEMS Department, or joint in both departments. Multidisciplinary research across a broad range of disciplines is strongly encouraged.

The ECE Department presently has 49 faculty members including two National Academy of Engineering members, 18 IEEE Fellows, and 13 NSF CAREER awardees and it administers B.S., M.S. and PhD programs. The Department has strong research programs in all major areas of electrical and computer engineering, with annual research expenditure of over $17M. The Department has accredited B.S. degree programs in both Electrical Engineering and Computer Engineering. The current enrollment is approximately 260 full-time graduate students and 770 undergraduate students.

The CHEMS Department consists of 28 faculty members, including three University Distinguished Professors and eight professional society fellows. Research programs span all major areas in chemical engineering and materials science, with annual research expenditures of $9M. The Department offers degree programs at the B.S., M.S. and PhD levels in both chemical engineering and materials science and engineering. Bachelor's programs in both CHE and MSE are fully accredited, and current enrollment for the department stands at approximately 110 graduate students and 730 undergraduates.

MSU enjoys a park-like campus with outlying research facilities and natural areas. The campus is adjacent to the city of East Lansing and the capital city of Lansing. The Lansing metropolitan area has a diverse population of approximately 450,000. Local communities have excellent school systems and place a high value on education. Michigan State University is pro-active in exploring opportunities for employment for dual career couples, both inside and outside the University. Information about MSU's dual career support can be found at **http://miwin.msu.edu/**. Information about WorkLife at MSU and the College of Engineering can be found at **http://www.egr.msu.edu/WE**.

Candidates should submit an application for these positions through: **http://jobs.msu.edu/** and refer to position #2293 for the tenure-stream position and #2257 for the fixed-term position. Applicants should include a cover letter, curriculum vitae, the names of at least three references, and statements of research and teaching interests. For full consideration, applications should be received before January 25, 2016. Applications will be reviewed thereafter on a continuing basis until the positions are filled. Nominations or questions are welcome by contacting the search committee chair through email at **ece-diamond-facultysearch@egr.msu.edu**.

Michigan State University has been advancing the common good with uncommon will for more than 160 years. A member of the Association of American Universities, MSU is a research-intensive institution with 17 degree-granting colleges.

MSU is an affirmative-action, equal opportunity employer. MSU is committed to achieving excellence through a diverse workforce and inclusive culture that encourages all people to reach their full potential. The university actively encourages applications or nominations of women, persons of color, veterans, and persons with disabilities.

UNIVERSITY OF MICHIGAN-DEARBORN

ECE Faculty Advertisement for Intelligent Systems and Power Electronics

The University of Michigan-Dearborn, Dept. of Electrical & Computer Engineering (ECE) invites applications for two tenure-track faculty positions, one as Assistant Professorship in Intelligent Systems, and another as Assistant/Associate/Full Professorship in Power Electronics & Motor Drives. Qualified individuals are invited to apply.

The applicant for Intelligent Systems should have research interests that fall within the areas of autonomous systems, machine learning, computer vision, pattern recognition, artificial intelligence, expert systems, and big data analytics. The applicant for Power Electronics & Motor Drives should have research interests in power electronics including electric machine drives and electric and hybrid vehicles' energy conversion systems. Selected candidates are expected to establish an excellent externally funded research program in their chosen area.

Qualified candidates must have, or expect to have, a Ph.D. in Computer Engineering, Electrical Engineering, or a closely related discipline by the time of appointment. Candidates will be expected to perform scholarly and sponsored research, as well as teaching at both the undergraduate and graduate levels. The ECE Department offers several BS and MS degrees, and participates in two interdisciplinary Ph.D. programs, Ph.D. in Automotive Systems Engineering and Ph.D. in Information Systems Engineering. The current funded research areas in the department include intelligent systems, power electronics, hybrid vehicles, battery management, computer networks, wireless communications, and embedded systems.

Applicants should submit a cover letter, curriculum vitae including current contact information: phone number, mailing address and e-mail address, teaching statement, research statement, and a list of three to five referees to the Department of Electrical and Computer Engineering, University of Michigan-Dearborn, 4901 Evergreen Road, Dearborn, Michigan, 48128, Phone: 313 593-5420, Fax: 313 583-6336 or email to: **umdintelligentsys@umich.edu** (Intelligent Systems) or **powerelectronics@umich.edu** (Power Electronics)

The University of Michigan-Dearborn is an equal opportunity/affirmative action employer, complies with all applicable federal and state laws regarding nondiscrimination and affirmative action.

POSITION OPEN:
Toyota Technological Institute has an opening for a tenured-, or tenure track associate professor, or 5-year tenure assistant professor position in the Department of Advanced Science and Technology. Applications are encouraged from all relevant areas. For more information, please refer to the website;

http://www.toyota-ti.ac.jp/english/employment/associate.html
Position: Associate Professor or Assistant Professor
Research field: Basic or application technologies for electronic devices with the use of compound semiconductors especially wide band gap semiconductors.
Qualifications: A Ph.D. in a relevant field. The successful candidate is expected to demonstrate the potential to develop strong and outstanding programs in the above research fields. It is also necessary for him or her to supervise students, and to teach advanced and basic courses both at the undergraduate and graduate levels.
Starting date: The earliest possible date of 2016.
Documents: (1) A curriculum vitae
(2) A list of publications
(3) Copies of 5 selected papers
(4) A brief description of research activities and future plans for research and education (3 pages each)
(5) The names of two references (telephone numbers and e-mail addresses required)
(6) An application sheet designated to our website
Deadline: February 29, 2016
Inquiry: Search Committee Chair Professor Yasutake Ohishi
Phone: +81-52-809- 1860 E-mail: y-ohishi1@toyota-ti.ac.jp

The above should be sent to:
Mr. Takashi Hirato
Administration Division
Toyota Technological Institute
2-12-1, Hisakata, Tempaku-ku
Nagoya, Aichi 468-8511 Japan
Applicants should please write "Application for Electronic Devices" in red on the return envelope.

香 港 中 文 大 學
The Chinese University of Hong Kong

Applications are invited for:-

Department of Mechanical and Automation Engineering
Professors / Associate Professors / Assistant Professors
(Ref. 1516/083(370)/2)

The Department is seeking excellent candidates to fill the above faculty positions in the following areas:

- robotics and automation;
- design and manufacturing, in particular in areas of 3D printing and CAD;
- energy and environmental engineering, including smart/green building, building automation and control, pollution measurement and monitoring, smart grid, and energy management.

Applicants should have (i) a PhD degree; and (ii) a proven track record or demonstrating potential for teaching and research excellence.

The appointees will (a) teach undergraduate and postgraduate courses; (b) develop an externally funded research programme; and (c) supervise postgraduate students.

Appointments will normally be made on contract basis for up to three years initially commencing August 2016, which, subject to mutual agreement, may lead to longer-term appointment or substantiation later. Outstanding candidates with substantial experience for Professor rank may be considered for substantive appointment forthwith. Applicants in robotics and automation areas for Research Assistant Professorship may also be considered.

Applications will be accepted until the posts are filled. Further information about the Department is available at http://www.mae.cuhk.edu.hk.

Salary and Fringe Benefits
Salary will be highly competitive, commensurate with qualifications and experience. The University offers a comprehensive fringe benefit package, including medical care, plus a contract-end gratuity for appointments of two years or longer, and housing benefits for eligible appointees. Further information about the University and the general terms of service for appointments is available at https://www2.per.cuhk.edu.hk/. The terms mentioned herein are for reference only and are subject to revision by the University.

Application Procedure
Application forms are obtainable (a) at https://www2.per.cuhk.edu.hk/, or (b) in person/ by mail with a stamped, self-addressed envelope from the Personnel Office, The Chinese University of Hong Kong, Shatin, Hong Kong.

Please send full resume, copies of academic credentials, publication list with abstracts of selected published papers, details of courses taught and evaluation results (if available), a research plan, a teaching statement, together with names, addresses and fax numbers/e-mail addresses of three to five referees to whom the applicants' consent has been given for their providing references (unless otherwise specified), to the Dean, Faculty of Engineering by e-mail to recruit@erg.cuhk.edu.hk.

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

Electrical and Computer Engineering, Michigan State University

Faculty Positions in Applied Electromagnetics

Join Michigan State University's Global Impact Initiative, designed to address the grand challenges through the creation of over 100 new faculty positions in some of the most promising and exciting fields of research. We welcome applicants from diverse backgrounds. MSU offers an inclusive and collaborative work environment. To learn more visit research.msu.edu/global-impact.

The Department of Electrical and Computer Engineering (ECE) at Michigan State University invites applications for two tenure-system faculty positions in the broad area of applied electromagnetics. Applicants must have a Ph.D. in Electrical Engineering or a closely related field, with expertise in, but not limited to:

Position-1: RF hardware, including materials and devices, integrated circuits, micro-systems and packaging; EMI/EMC for autonomous systems.

Position-2: Application driven design and implementation of RF sensors and systems, and development of multiphysics and multiscale algorithms for describing their behavior.

The successful candidate will be expected to collaborate with colleagues within the ECE Department and the College of Engineering to complement an internationally recognized program in applied electromagnetics. All applications must demonstrate evidence of strong research potential, interest in both undergraduate and graduate teaching, and a serious commitment to teamwork. These positions are targeted at the junior faculty level, but senior candidates with outstanding credentials may be considered.

The Electrical and Computer Engineering department has a strong interdisciplinary research and educational programs on a foundation of core electrical and computer engineering disciplines and provides first-class education while engaging in research at the frontiers of knowledge. The Department has 49 faculty members, including two National Academy of Engineering members, 18 IEEE Fellows, and 13 NSF CAREER awardees. The Department has strong research programs in all major areas of electrical and computer engineering, with annual research expenditure of about $18M. Faculty in the Department are leading several federal and industry-supported centers, including the NSF Science and Technology Center BEACON, and the Fraunhofer Center for Coatings and Laser Applications. The Department has accredited B.S. degree programs in both Electrical Engineering and Computer Engineering. The current enrollment is approximately 260 full-time graduate students and 770 undergraduate students. For additional information about the ECE Department, the College of Engineering or MSU, see: **http://www.egr.msu.edu/ece/**.

MSU enjoys a park-like campus with outlying research facilities and natural areas. The campus is adjacent to the city of East Lansing and the capital city of Lansing. The Lansing metropolitan area has a diverse population of approximately 450,000. Local communities have excellent school systems and place a high value on education. Michigan State University is pro-active in exploring opportunities for employment for dual career couples, both inside and outside the University. Information about MSU's dual career support can be found at **http://miwin.msu.edu/**. Information about WorkLife at MSU and the College of Engineering can be found at **http://www.egr.msu.edu/WE**.

Interested individuals should submit only one application through: http://jobs.msu.edu/ and indicate their preference by applying to #2248 for position-1 or #2247 for position-2. Applicants must submit a detailed resume, a cover letter summarizing their qualifications, vision statements for teaching and research, and the names and contact information for five references. Applications will be reviewed on a continuing basis until the position is filled. For full consideration, applications should be received before February 1st, 2016. Nominations or questions are welcome by contacting the search committee chair through email at **ece-appliedEM-facultysearch@egr.msu.edu**.

Michigan State University has been advancing the common good with uncommon will for more than 160 years. A member of the Association of American Universities, MSU is a research-intensive institution with 17 degree-granting colleges.

MSU is an affirmative-action, equal opportunity employer. MSU is committed to achieving excellence through a diverse workforce and inclusive culture that encourages all people to reach their full potential. The university actively encourages applications or nominations of women, persons of color, veterans, and persons with disabilities.

AND NOW THIS...

WHEN CARUSO WAS KING

Hit Victrola records made owning music fashionable

At the start of the 20th century, the phonograph industry was the second-fastest growing in the United States, trailing only automobiles. Victrola gramophones played black discs formulated from slate powder and the resinous secretions of lac insects. A standard 25-centimeter disc held 3 minutes of sound. But the real profits lay in Victor's 30-cm Red Seal recordings of classical performers; tenor Enrico Caruso [far left in this December 1920 ad] was a favorite. As music formats have morphed from vinyl discs to cassettes to CDs and MP3s, the classical recording industry is nearly extinct. With sound streamed from the cloud, do you really need to own the music you listen to? **—ALEXANDER B. MAGOUN, IEEE History Center**

PHOTOGRAPH BY Randi Klett

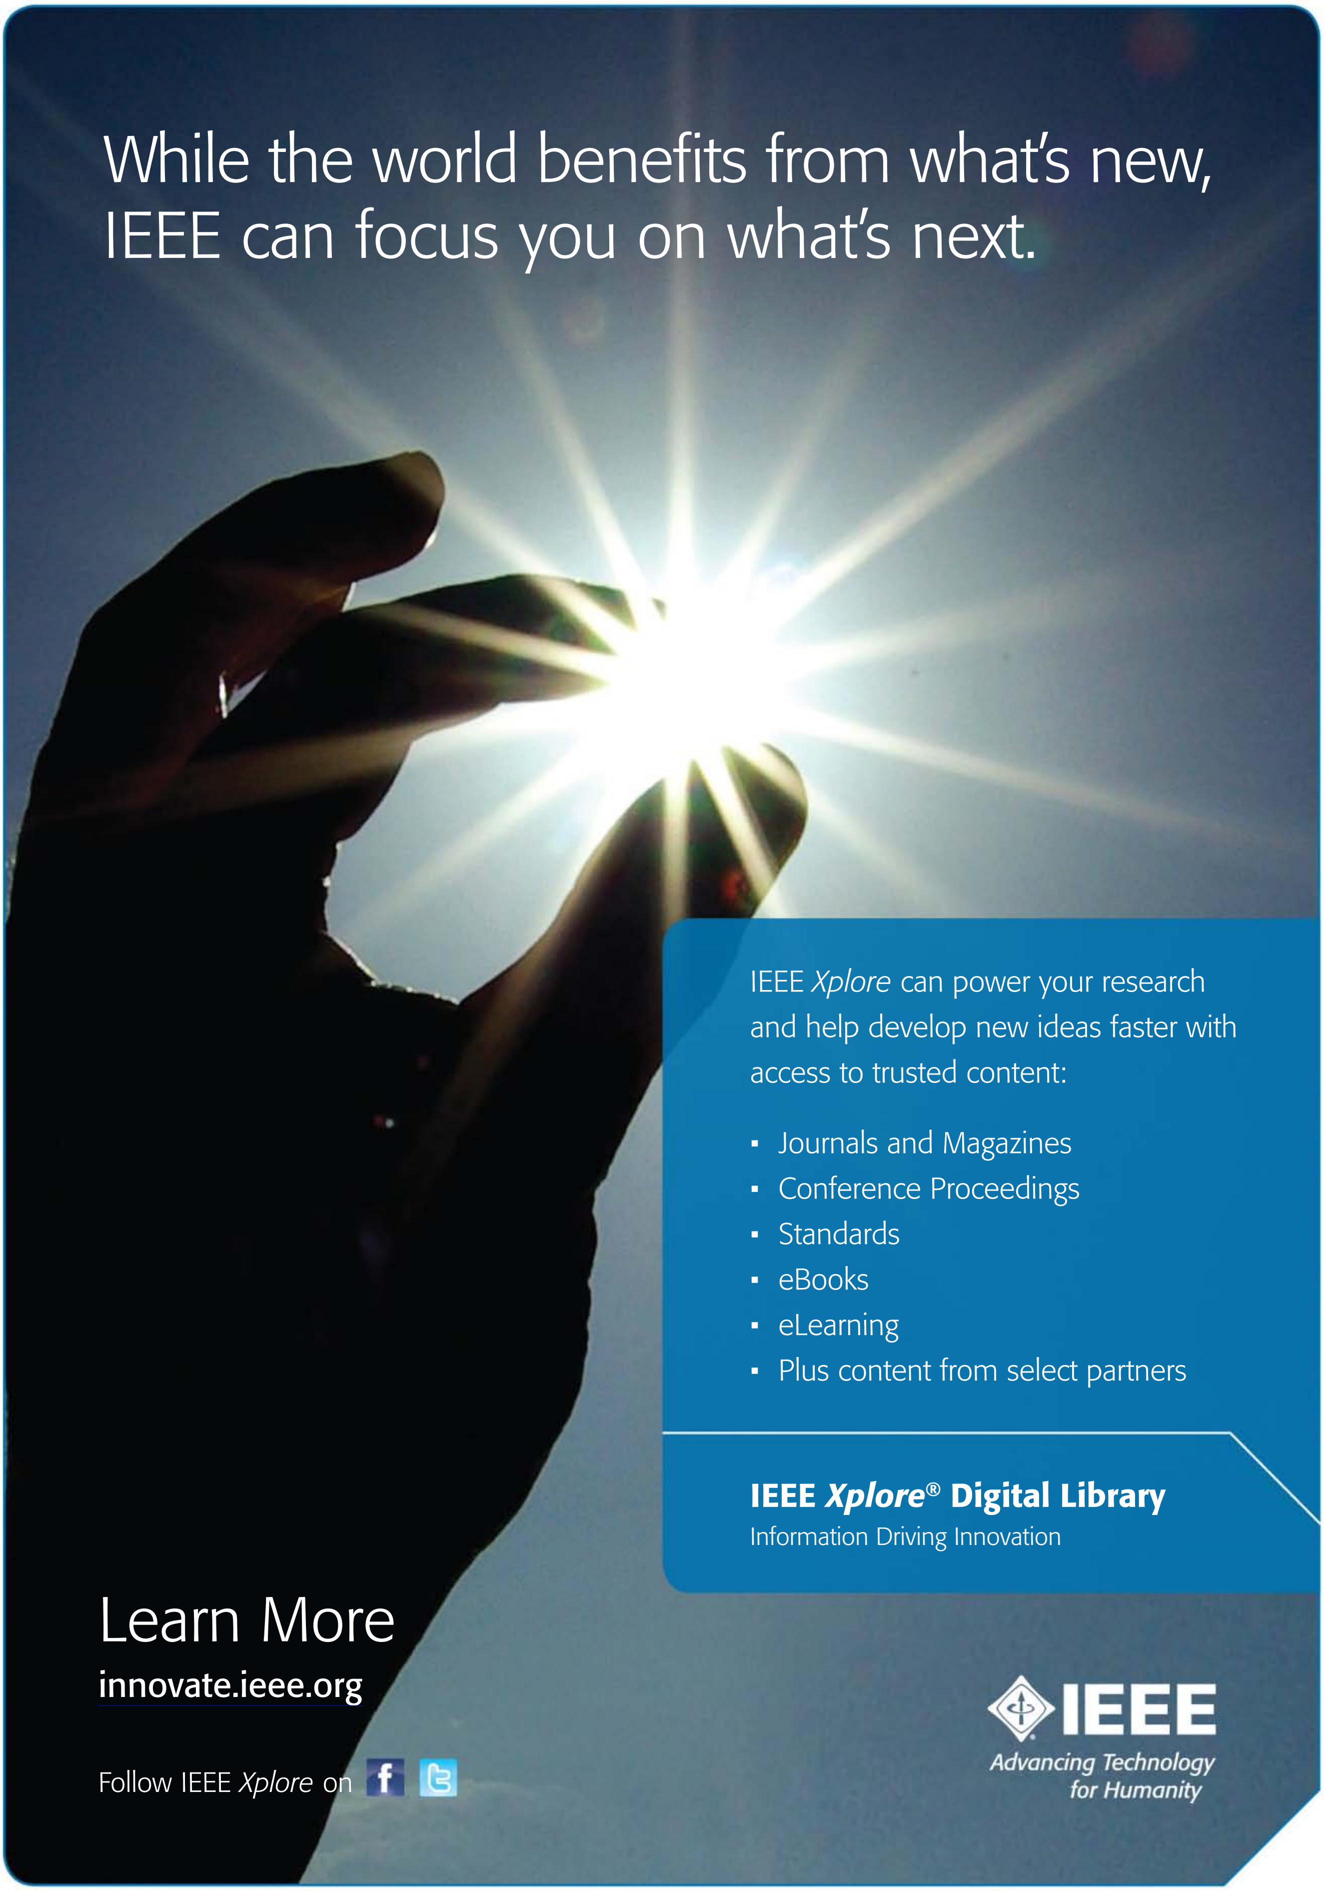
While the world benefits from what's new,
IEEE can focus you on what's next.
IEEE Xplore can power your research and help develop new ideas faster with access to trusted content:
▪ Journals and Magazines
▪ Conference Proceedings
▪ Standards
▪ eBooks
▪ eLearning
▪ Plus content from select partners
IEEE Xplore® Digital Library
Information Driving Innovation
Learn More
innovate.ieee.org
Follow IEEE Xplore on
IEEE
Advancing Technology for Humanity


MathWorks®
Accelerating the pace of engineering and science
Puhutko MATLABia?
Over one million people around the world speak MATLAB. Engineers and scientists in every field from aerospace and semiconductors to biotech, financial services, and earth and ocean sciences use it to express their ideas. Do you speak MATLAB?
Image of the Antenna Galaxies processed in MATLAB with the ADRIC galaxy matching program. Provided by Dr. Marianne Doyle, Univ. of Queensland.
Programs available at mathworks.com/ltc
SuperCOSMOS Image: WFAU, Univ. Edinburgh/STFC/AAO. Hubble Image: NASA ©2015 The MathWorks, Inc.
MATLAB®
The language of technical computing