IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 11.16

The Last, Best Hope for the Stand-alone Pocket Camera

How smart algorithms let this little gizmo compete with a high-end DSLR P. 34

IEEE

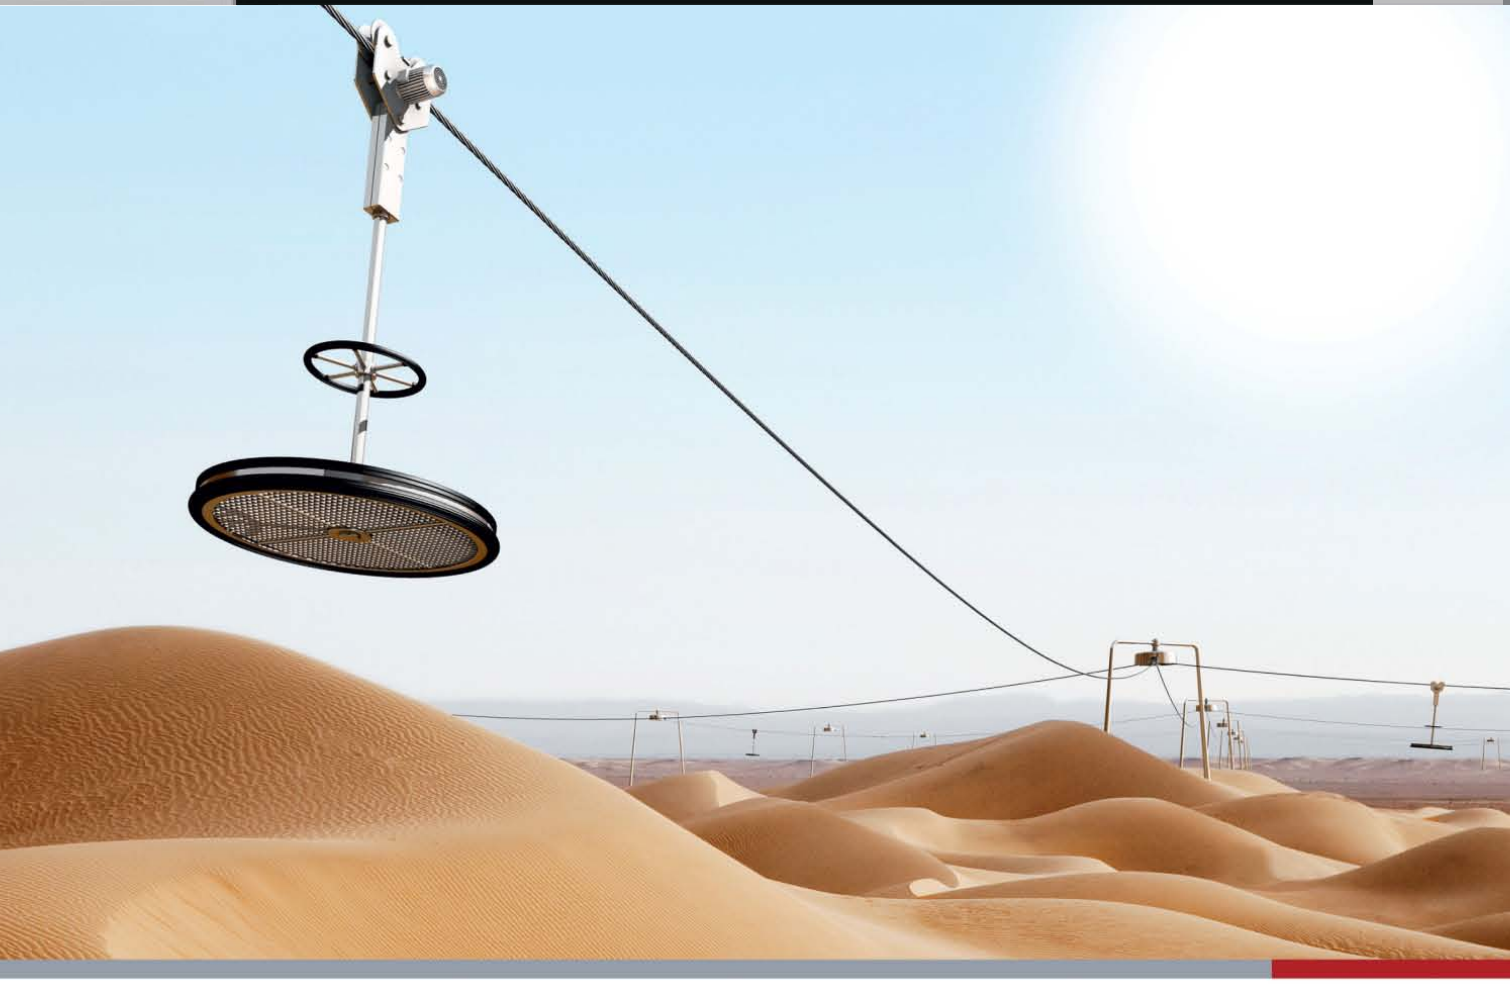

FEATURES_11.16

IEEE SPECTRUM

On the cover **Photograph for *IEEE Spectrum* by Sue Tallon**

PHOTOGRAPH BY **Corey Davis/Randels Media Group**

DEPARTMENTS_11.16

IEEE SPECTRUM

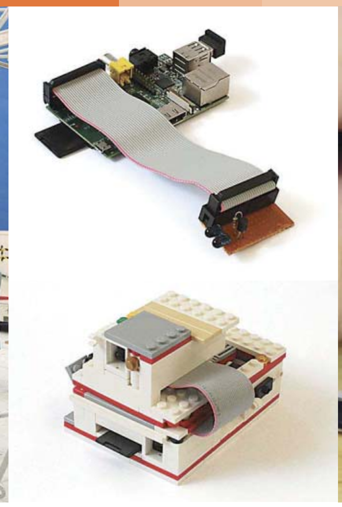

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Analyzing HF & LF RadHaz Scenarios With 3D EM Simulation—3 November
- Creative Statistics—4 November
- Realizing New Potentials in Miniature Subsea Robotics—8 November
- EMC Simulation for Space Applications With CST STUDIO SUITE—10 November
- Transformer Simulation—17 November
- Antenna Analysis and Design With Characteristic Mode Analysis—22 November

- **MASTER BOND RESOURCE LIBRARY**
  http://spectrum.ieee.org/static/master-bond-resource-library
- **NEW PRODUCT RELEASE LIBRARY**
  http://spectrum.ieee.org/static/new-product-release-library

The Institute

Available at theinstitute.ieee.org

- **BRAIN INITIATIVE** There's been a surge of activity around the world to develop tools and applications to help the brain function better. IEEE has formed its own initiative to advance brain research and neurotechnology development through workshops, standards, and partnerships with academia, government, and industry.

- **WESTON METERS** Weston Electrical Instrument Co. introduced the first portable and direct-reading current and voltage meters in the late 1880s. This breakthrough was recently honored with an IEEE Milestone.

- **IEEE DAY** Hundreds of groups around the world held events on 4 October to mark the annual celebration, which commemorates the founding of one of IEEE's predecessor societies.

FROM LEFT: XINHUA/ALAMY; DIOMIDIS SPINELLIS; WAZER

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2016 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 53, issue No. 11, 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.

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

BACK STORY_

The Startup Around the Corner

SENIOR EDITOR TEKLA S. PERRY HAS BEEN DRAWN to the Silicon Valley startup Light like a moth to...well, a light. Her first glimmer was a mention of Light in a 2014 report of venture investing; the company had closed a US $9.7 million funding round. The official word about what the company was doing was pretty vague: A "team of creative technologists [were] on a mission to reimagine the art and science of photography."

She started combing the Internet, and discovered that Light cofounder Rajiv Laroia had in 2013 been identifying himself at conferences as the founder of Tinz Optics. On a hunch, Perry walked the 10 minutes from her home in Palo Alto, Calif., to the building listed as Tinz headquarters—next to her favorite Caribbean restaurant—and indeed, there was a paper sign on the door that read "Light." (Later, next door, she bought a Ting soda and wondered if the name had been the inspiration for "Tinz.")

Peeking in the windows told her nothing. So she combed the online want ads, to see just what kind of engineers Tinz/Light was hiring. The company had been seeking an "imaging scientist with a deep understanding of computational optics" who could develop "innovative computation algorithms to combine multiple recorded images into a single image....and increase image dynamic range." The job had been filled; some LinkedIn sleuthing suggested that the hire was a research scientist from camera-maker Ricoh who had written papers about multi-aperture imaging.

Perry concluded that Light planned to use several small cameras to get better images—and she was right. This January, she held the prototype herself [above] at a private showing at the CES technology show, in Las Vegas. A few weeks earlier, Perry had finally met Laroia; in this issue, he reveals Light's details in "A Pocket Camera With Many Eyes."

That guess about Ting being the inspiration for Tinz? Way off base. The name comes from the first letters of the names of Laroia's four children. ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, Past Forward is in *IEEE Spectrum*, Vol. 53, no. 11 (INT), November 2016, p. 64, or in *IEEE Spectrum*, Vol. 53, no. 11 (NA), November 2016, p. 76.

11.16

IEEE
SPECTRUM

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

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

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

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

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

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

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

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

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

INTERN Andrew Silver

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

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

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

EDITORIAL ADVISORY BOARD
Susan Hassler, *Chair*; David Brock, Jason Cong*, Sudhir Dixit, Limor Fried, Robert Hebner, Joseph J. Helble, Grant Jacoby, Leah Jamieson, Deepa Kundur, Norberto Lerendegui, Steve Mann, Allison March, Jacob Østergaard, Umit Ozguner, Thrasos Pappas, John Rogers, Jonathan Rothberg, Umar Saif, Takao Someya, Yu Zheng, Kun Zhou, Edward Zyszkowski
* Chinese-language edition

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

LIGHT

CONTRIBUTORS_

Kevin Driscoll

Driscoll teaches media studies at the University of Virginia, in Charlottesville. His article about the history of computer bulletin-board systems [p. 54] is an outgrowth of his 2014 doctoral thesis at the University of Southern California. But his interest in computer BBSs goes all the way back to boyhood, when he became a computer bulletin-board user. Driscoll was born in 1980, shortly before this ad hoc "people's Internet" began gearing up. "I was on the tail end," he says, with some regret.

Benjamin Kroposki

Kroposki is director of the Power Systems Engineering Center at the National Renewable Energy Laboratory, in Golden, Colo. An IEEE Fellow, he writes in this issue about smarter solar inverters that can handle faults and help keep the power grid up and running [p. 42]. "As more wind, solar, and other inverter-based power generation is deployed, we need to make sure they work with existing generation so that the grid remains stable and reliable," Kroposki says.

Rajiv Laroia

Laroia is chief technology officer and cofounder of Light, based in Palo Alto, Calif. He previously founded Flarion Technologies, which developed the base technology for LTE, and before that he held R&D roles at Lucent Technologies' Bell Labs. Over the years, Laroia says, he's purchased lots of expensive camera equipment, but he found that his go-to camera for family photos was his cellphone. That realization inspired him to reengineer pocket-camera technology, as he describes in this issue [p. 34].

Ryan Randels

Photographing the Pikes Peak International Hill Climb, in Colorado, is one of the toughest assignments in motorsports. "You never want to be on the apex of a corner, where you're just asking to be hit," says Randels, of Randels Media Group, the event's official photographer. "Understanding the mountain itself and the paths of the competitors is critical." He and his colleagues shot the student team behind an all-electric motorcycle that raced against commercial rivals [see "Run Silent, Run Steep," p. 48].

Diomidis Spinellis

Spinellis is a professor of management science and technology at the Athens University of Economics and Business, in Greece. He is the editor in chief of the journal *IEEE Software*. In this issue he shows how to replace the pricey Intelligent Brick used in the Lego Mindstorms robotics kits with a much cheaper Raspberry Pi–powered remote control running child-friendly software [p. 21]. "I was impressed by how easy it was to seamlessly combine these diverse open-source components," says Spinellis.

11.16

IEEE

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

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

MANAGING DIRECTOR, PUBLICATIONS Michael B. Forster

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

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

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

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

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

DIRECTORS EMERITUS Eric Herz, Theodore W. Hissey

IEEE STAFF
EXECUTIVE DIRECTOR & COO James Prendergast
+1 732 502 5400, james.prendergast@ieee.org
CHIEF INFORMATION OFFICER Cherif Amirat
+1 732 562 6399, c.amirat@ieee.org
PUBLICATIONS Michael B. Forster
+1 732 562 3998, m.b.forster@ieee.org
CHIEF MARKETING OFFICER Karen L. Hawkins
+1 732 562 3964, k.hawkins@ieee.org
CORPORATE ACTIVITIES Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
HUMAN RESOURCES Shannon Johnston, SPHR
+1 732 562 6343, s.johnston@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Eileen M. Lach, +1 212 705 8990, e.m.lach@ieee.org
EDUCATIONAL ACTIVITIES Jamie Moesch
+1 732 562 5514, j.moesch@ieee.org
CHIEF FINANCIAL OFFICER Thomas R. Siegert
+1 732 562 6843, t.siegert@ieee.org
TECHNICAL ACTIVITIES Mary Ward-Callan
+1 732 562 3850, m.ward-callan@ieee.org
MANAGING DIRECTOR, IEEE-USA Chris Brantley
+1 202 530 8349, c.brantley@ieee.org
IEEE PUBLICATION SERVICES & PRODUCTS BOARD
Sheila Hemami, *Chair*; John Baillieul, Babak D. Beheshti, Herbert S. Bennett, Jennifer T. Bernhard, Eddie Custovic, Jean-Luc Gaudiot, Ron B. Goldfarb, Glenn Gulak, Lawrence Hall, Christofer Hierold, James M. Irvine, Hulya Kirkici, Carmen S. Menoni, Paolo Montuschi, Thrasos Pappas, Michael Pecht, Michael Polis, Sorel Reisman, Tariq Samad, Fred Schindler, Gianluca Setti, Gaurav Sharma, Curtis A. Siller, David Tian, Ravi Todi, H. Joel Trussell, Leung Tsang, Stephanie M. White, Steve Yurkovich

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

PHOTO-ILLUSTRATIONS BY **Gluekit**


IF WE™ touch the screen,
can we feel the cat?
Technology that simulates touch – a dream our software could bring to life.
Innovative thinkers everywhere use INDUSTRY SOLUTION EXPERIENCES from Dassault Systèmes to explore the true impact of their ideas. Insights from the 3D virtual world help companies develop technologies that use multiple human senses. Can we one day make a lion purr without entering the den?
3D V+R
3DEXPERIENCE
It takes a special kind of compass to understand the present and navigate the future.
3DS.COM/HIGHTECH
DASSAULT SYSTEMES | The 3DEXPERIENCE Company

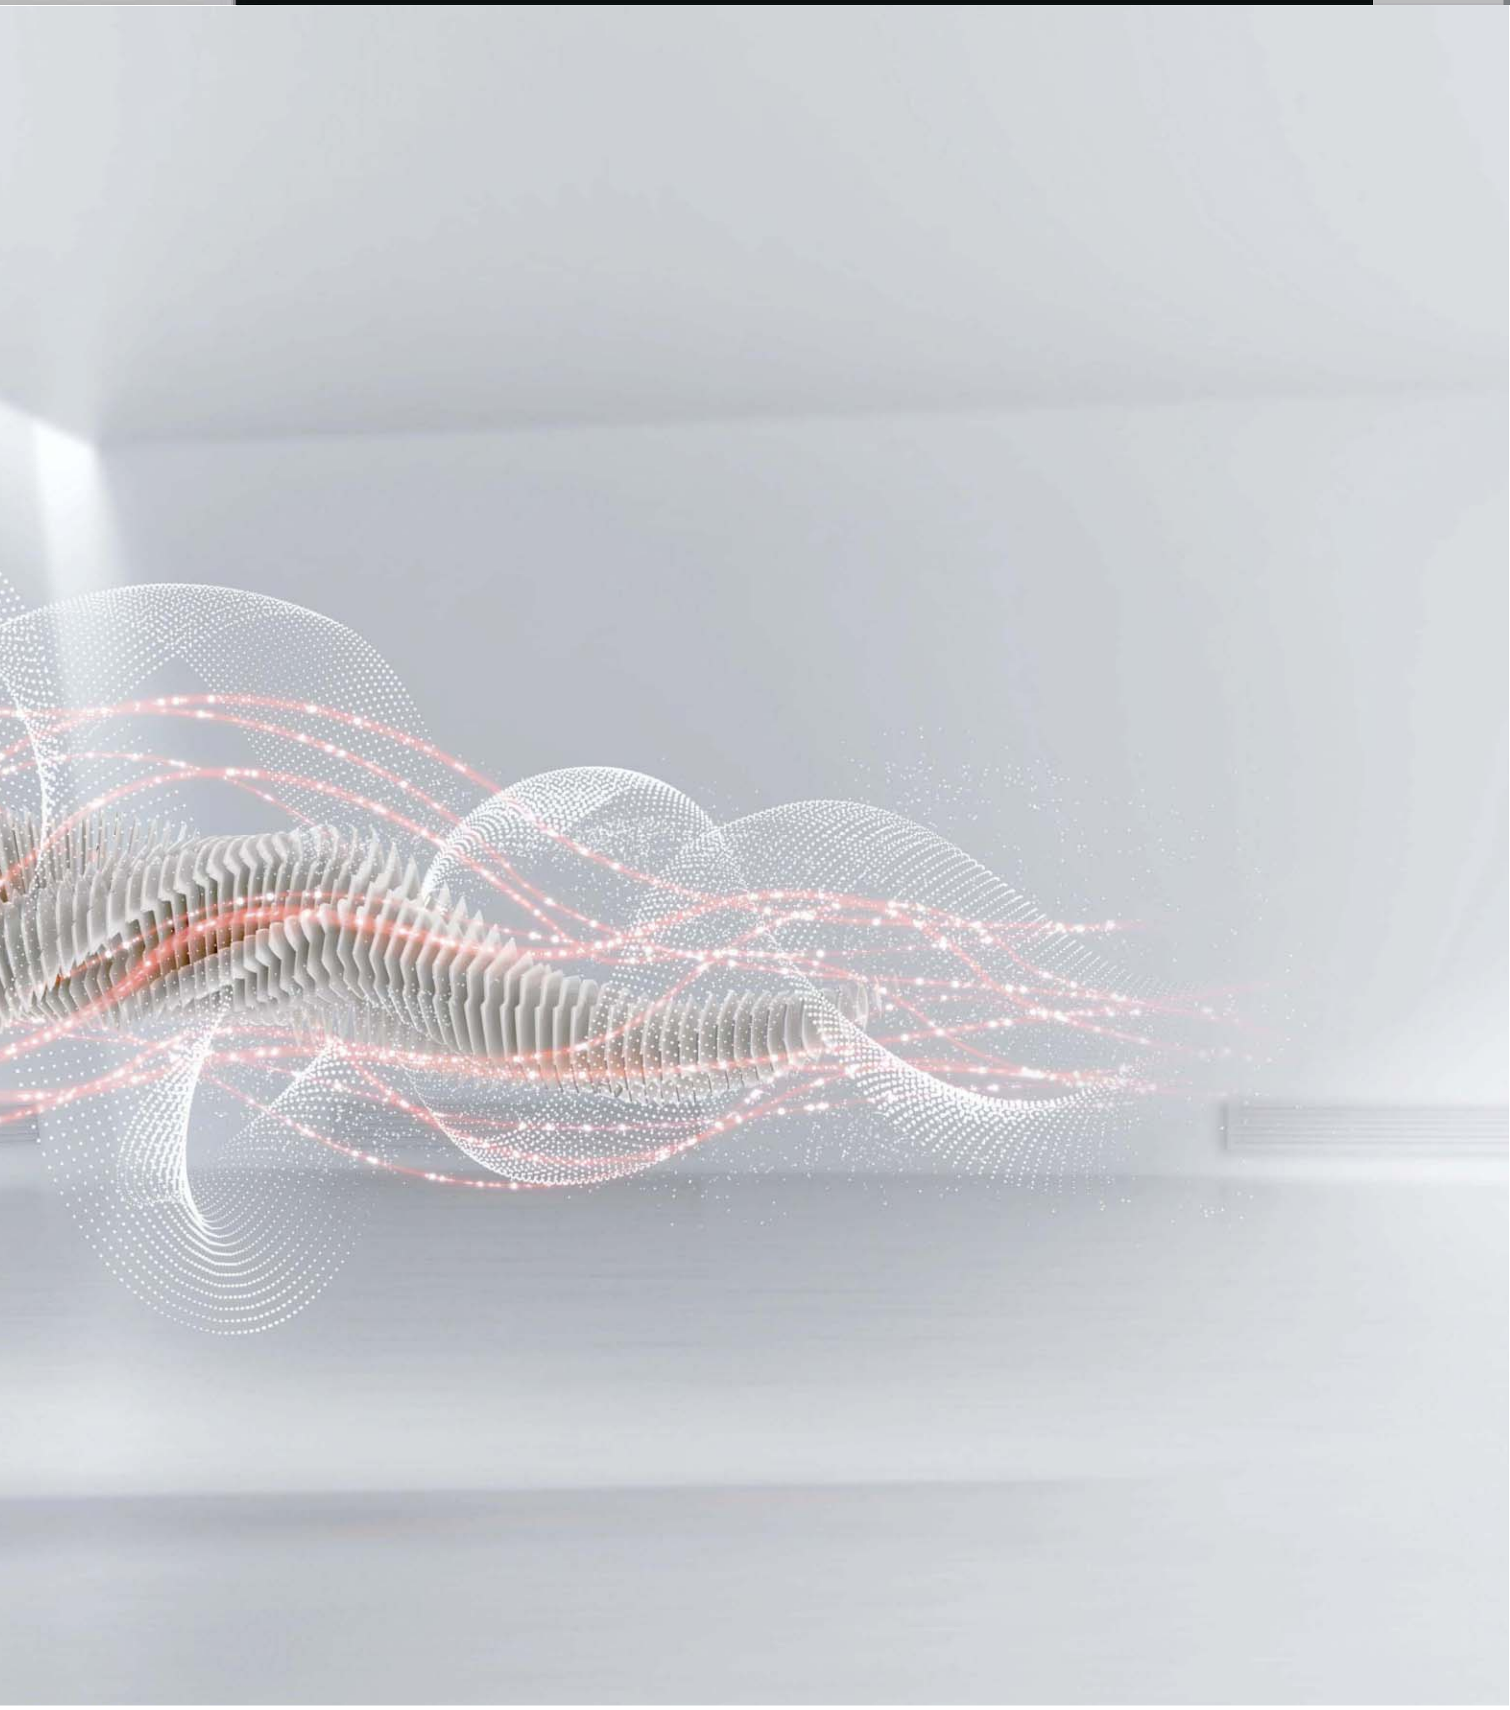

The Invisible Technology

We used to go to stores to buy it. Now software is so ubiquitous, we don't even notice it

Five years ago Marc Andreessen, the Web pioneer and celebrated tech investor, predicted software would eat the world.

He turned out to be right. Too right. Software is eating the world, and also eating itself.

The cannibalization of software defies easy explanation. Software is the motor of the world's digital economy. Code is the ground of our computationally rich existence. Software applications and platforms are the source of vast wealth for Apple, Google, Facebook, Amazon, and many other tech titans. Yet even as software grows in importance, code becomes less visible, less tangible, less understood, and–perhaps most paradoxically–less valuable in monetary terms.

How has this great shift happened? Software originally coevolved with computers themselves. The IBM System/360, the business computer of choice in the 1960s, came bundled with code, and if customers needed more or different programs, they asked IBM. In 1980, IBM chose to rely on outsiders, notably Microsoft, for PC code, igniting a ferocious race to sell programs as distinct products.

Over the past 20 years, the program as artifact has vanished. Consumers download new versions, patches, and feature improvements as easily as switching on their devices. They rarely pay for this. Software battles now occur over platforms, which define an experience such as shopping (Amazon), searching (Google), or networking (LinkedIn). Competitive advantage is now achieved through superior software, but software supremacy is neither the aim nor the result of the new game.

"Software on wheels," for instance, now defines the car of the future, and is the reason why Google and Apple, the reigning kings of code, each are pushing for a big role in next-gen autos. No surprise that Toyota, Ford, and Daimler want their own software expertise, not to peddle programs to drivers but to enhance their new models. Without cool software, these venerable automakers might not even survive (see Tesla's and Google's self-driving cars).

View Uber in this same light. The company is best understood not as a taxi service on steroids but as a software-management system for personal transportation.

The bottom line: No one gets rich making software anymore. The days of Bill Gates building a fortune on the strength of shrink-wrapped programs sold like disposable diapers is gone. Software today can make you rich only by enabling you to do something else that people pay for.

Netflix, Facebook, and Google don't get a dime from selling software, yet their revenue-producing services depend on continuously and seamlessly improving their code. Similar examples are legion.

The new logic of software has different implications for those who make it, sell it, and use it. For makers, code no longer spawns tycoons and celebrities. The last person to get famous from writing software was Linus Torvalds; and this year he celebrates the 25th anniversary of his seminal achievement. Today's top coders are largely if not wholly unknown by the wider public; at best, they are cult figures, revered in underground communities.

Another paradox: As software becomes ever more essential to creating the digital experience, the invisibility of software is a victory for the apostles of "the free." Hippies, misfits, and dropouts improbably created fabulous wealth through pricey apps in the last two decades of the 20th century. Some among these characters, notably Richard Stallman, promoted a counter-ethos that conceived of software as a public good, available without charge to all.

That free software sits at the heart of an explosion of profitable digital services is the latest, greatest riddle of global capitalism–and a sobering message for the poets of programming. Code writers are essential and well paid but increasingly interchangeable and anonymous.

For this, we should not pity the programmer, and we should remember: Hope springs from the unseen. –G. PASCAL ZACHARY

G. Pascal Zachary is a professor of practice at Arizona State University's School for the Future of Innovation in Society.

ILLUSTRATION BY **Greg Mably**

NEWS

CHINA'S 2,000-KM QUANTUM LINK IS ALMOST COMPLETE

The Beijing–Shanghai project pushes the limits for "unhackable" links

XINHUA/ALAMY

QUANTUM SPACE LINK: Staff work on China's quantum research satellite, which launched in August. It is part of a larger effort in the country to push the limits of quantum key distribution.

By the end of this year, a team led by researchers from the University of Science and Technology of China, in Hefei, aims to put the finishing touches on a 2,000-kilometer-long fiber-optic link that will wind its way from Beijing in the north to the coastal city of Shanghai.

What will distinguish this particular stretch of fiber from myriad other long-distance links is its intended application: the exchange of quantum keys for secure communication—a sophisticated gambit to protect data from present and future hackers. If all goes according to plan, this Beijing-Shanghai line will connect quantum networks in four cities. And this large-scale terrestrial effort now has a partner in space: A quantum science satellite was launched in August with a research mission that includes testing the distribution of keys well beyond the country's borders.

With these developments, China is poised to vastly extend the reach of quantum key distribution (QKD), an approach for creating shared cryptographic keys—sequences of random bits—that can be used to encrypt and decrypt data. Thanks to the fundamental nature of quantum mechanics, QKD has the distinction of being, in principle, unhackable. A malicious party that attempts to eavesdrop on a quantum transmission won't be able to do so without creating detectable errors.

QKD has already made its way into the real world. In 2007, the scheme was used to secure the transmission of votes in a Swiss election. Several years ago, the U.S.-based firm Battelle began to use the approach to exchange information securely over kilometers of fiber between its corporate headquarters in Columbus, Ohio, and a production facility in Dublin, Ohio.

But despite great progress, there has been a stumbling block to wide distribution. "The problem we've got is »

RIVER OF SECRETS: A 2,000-kilometer-long backbone will connect quantum networks in four cities. The line uses a number of secured nodes as relays along the route.

distance," says Tim Spiller, director of the United Kingdom's Quantum Communications Hub, a nationally funded project that is building and connecting quantum networks in Bristol and Cambridge, in England.

The challenge is that QKD encodes information in the states of individual photons. And those photons can't travel indefinitely in fiber or through the air; the longer the distance, the greater the chance they will be absorbed or scattered.

This characteristic has a direct impact on how quickly a quantum key can be generated, explains physicist Jian-Wei Pan, who leads the Chinese projects. If researchers attempted to send signals directly down 1,000 kilometers of fiber, Pan says, "even using all the best technology, we would only manage to send 1 bit of secure key over 300 years."

Instead, QKD fiber links must have a way to refresh the signal every 100 km or so to maintain a reasonable bit rate. But this can't be done with conventional telecommunications equipment. The same rules that protect quantum transmission against eavesdropping also prohibit a quantum key from being copied without corrupting it. The solution has been to concatenate, creating a daisy chain of individual quantum links connected by physically secured spots, or "trusted nodes." Each intermediate node measures the key and then transmits it with fresh photons to the next node in the chain.

The Beijing-Shanghai line will use 32 trusted nodes to create the 2,000-km line. This approach isn't ideal for security. Because each trusted node has to convert the quantum key back into classical (nonquantum) information before passing it on, an eavesdropper at the node could potentially hack the data stream there undetected. "That's the drawback," Pan says. But the approach is "still much better than traditional communications… [where] there is the possibility of performing eavesdropping" at every point along the route, he says. Here, the problem is limited to 32 spots under lock and key.

"A long-distance chain link like this, [it's] really the first time it's been done," says Grégoire Ribordy of ID Quantique, based in Geneva, which makes hardware for QKD networks. "It's inspiring other people to try to do similar things around the world."

If you want to avoid even the small vulnerability of trusted nodes, Spiller says, long-distance QKD must use quantum entanglement, a property that can link the states of photons separated by great physical distance and that can be exchanged between photons. "Quantum repeaters," used in place of trusted nodes, could take advantage of this phenomenon to relay a quantum key without having to measure it. But this technology is still in an early stage of development, says Spiller; among other things, a quantum repeater will likely require a form of quantum memory to help coordinate communication.

"[If you] don't have to trust any of the nodes along the network, that will broaden the applicability of QKD," says Michele Mosca, cofounder of the Institute for Quantum Computing at the University of Waterloo, in Ontario, Canada. One reason to improve QKD's reach is to protect communications from tomorrow's quantum computers, which could make short work of the public-key cryptography that underpins Internet security and many other applications.

But Mosca notes that QKD is not the only possible way to address this threat; many cryptographers are exploring new "postquantum" algorithms to replace our existing public-key systems. QKD offers an "extra degree of assurance," he says, but improved conventional cryptography will be a cheaper and more practical solution for many applications. Both will likely have a role to play in the coming years. –RACHEL COURTLAND

ILLUSTRATION BY Erik Vrielink

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

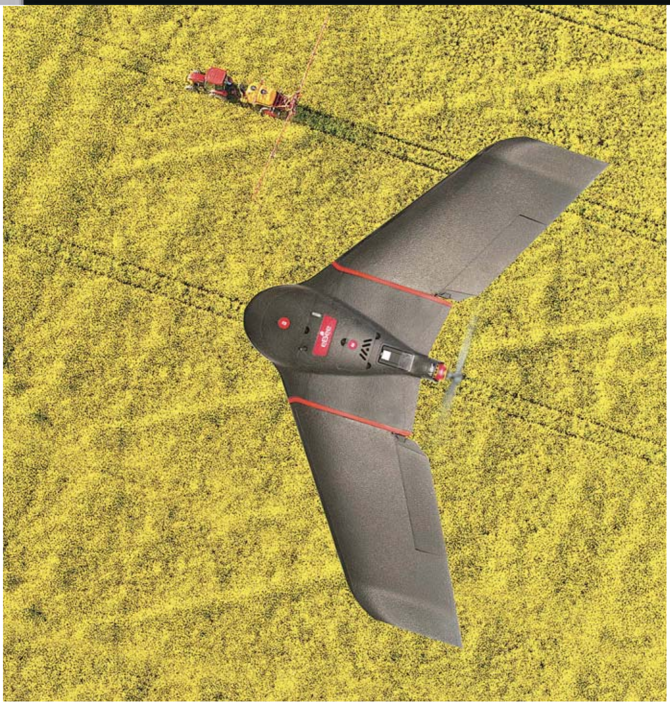

ABOVE IT ALL: SenseFly's eBee SQ can fly over hundreds of hectares of farmland on a single charge. The drone's multispectral camera scans in color and in four nonvisible wavelengths that are important to crop health.

SENSEFLY

AGRICULTURE DRONES ARE FINALLY CLEARED FOR TAKEOFF

New U.S. rules for commercial drones will benefit farmers and the drone industry

Tech-savvy farmers have been some of the earliest commercial adopters of drone technology, purchasing 45,000 drones last year alone. But if they were using the drones to check on the condition of their fields, spraying their crops, or keeping tabs on livestock, most of them were technically breaking the law. New U.S. federal rules that went into effect this summer, however, should make it easier for farmers to get a drone's-eye view of their fields.

The new rules allow commercial drone operators to get certified via a written test, so long as they fly drones that meet certain weight and altitude guidelines. Before this, operators had to pay for a pilot's license and get a special exemption to use a drone, a slow and cumbersome process.

The rules are a win for farmers, drone companies, and everyday consumers. Worldwide, agriculture is the largest commercial market for unmanned aerial vehicles, or UAVs. And it is expected to bag 80 percent of all commercial drone use in the United States. About half of the 5,500 exemptions that the U.S. Federal Aviation Administration (FAA) has approved up to now were for agricultural purposes, says Brian Wynne, president and CEO of the Association for Unmanned Vehicle Systems International. With the new rules, the FAA expects more than 600,000 commercial UAVs to be flying during the next year. "These numbers foreshadow the exciting potential that UASs [unmanned aircraft systems] have for the agriculture industry now that the small UAS rule is in place," Wynne says.

By using drones to scout for weeds and pests, spot diseased plants or dry areas, and spray the right amount of fertilizer and pesticide, farmers can increase yield with less resources and environmental harm. Returns are especially steep for high-value crops like wine grapes. Drones are cheaper than hiring a small plane and, in contrast to satellites, work in cloudy conditions and give higher-resolution images. Prices for ready-to-fly agricultural UAV systems–hardware, sensors, and software included–range from US $1,500 to over $25,000.

"This is going to have a big impact on farmers," says Kyle Miller, an applications engineer at agricultural drone maker AgEagle Aerial Systems, in Neodesha, Kan. Miller, who has been flying robotic aircraft over his family's soybean farms as a hobbyist for years, says interest in the technology among farmers he knows is soaring. "Now that the new rules are out, farmers know drones are coming to their area," he says.

Until now, many drone technology providers sat on the sidelines, not wanting to go ahead with drone testing or use given the lack of a clear regulatory path, says Nathan Stein, agricultural solutions manager at Swiss drone maker

NEWS

senseFly. "Whether it was a seed company or chemical company, their legal departments wouldn't let them move forward," he says. "Anything with the word *drone* in it was ominous."

Several U.S. businesses have focused on operations in other countries. Drones have been used legally for agriculture for more than 20 years in countries such as Australia, Japan, and South Korea, according to Wynne. "Now companies can bring business and technology back into the United States," Stein says.

Plus, American farmers might reap more benefits from drones than farmers in other countries. "The United States has incredibly diverse agriculture, from peaches to corn to grapes to cotton," Stein says. What's more, he says, U.S. farmers have "an adventurous spirit, so I expect a large market. People might come up with more creative ways to use drones on the farm."

With pressure mounting on the FAA from stakeholders, UAVs are poised to have an even bigger impact on agriculture. Current rules limit commercial drones to less than 55 pounds (25 kilograms), a maximum speed of 100 miles per hour (160 kilometers per hour), and an altitude of 400 feet (122 meters). They also must not fly above people and remain within the operator's visual line of sight.

That last rule is a killer for agriculture and other economic uses. Farms in the United States—think midwestern corn and soy expanses—are often much bigger than line-of-sight operation will allow for. There are around 2.1 million farms in the country. Of those, the 1.85 million or so smaller family farms (averaging 93 hectares) that make up about half the total U.S. farm acreage can benefit from drones right now. The same isn't true for larger farms, which account for 36 percent of total U.S. farmland.

The FAA is moving in the right direction, but it needs to do more, says Darryl Jenkins, chairman of the American Aviation Institute, an independent think tank for commercial aviation based in Washington, D.C. "Things move at a glacial speed at the FAA," Jenkins says. "To really get economic benefits from drones, we need beyond-line-of-sight operations, flights over people, and access to higher altitudes."

Perhaps then, like tractors and combines, drones might become standard farm equipment.

—PRACHI PATEL

1.85 million farms, or about half the total U.S. farm acreage, can benefit from drones right now

4 STEPS TO DUST-SIZE NEURAL IMPLANTS

Tiny ultrasound-powered motes could record and adjust nerve activity

Mainstream medicine is making increasing use of electronics inside the body, deploying implanted gadgets both to measure internal conditions and to provide stimulating jolts of electricity to nerves and muscles. But turning a human into a proper cyborg will require many minuscule devices that can be scattered throughout the body. As a step toward that goal, a team of bioengineers has built speck-size wireless electrodes that can be affixed directly to nerves—and that may one day be nestled inside the brain.

The engineers from the University of California, Berkeley, implanted one mote of what they call "neural dust" inside an anesthetized rat, and demonstrated that the electrode could record signals from the rat's sciatic nerve and wirelessly transmit the information. This experiment was a proof of concept, says Jose Carmena, who co-led the research at UC Berkeley's Center for Neural Engineering and Prostheses, where he is codirector. If the neural dust can be adapted for the human body and brain, doctors could have an intimate new interface with the human nervous system.

But first, it had to work in a conked-out rat.

To power the neural dust, a transducer outside the animal's body emits ultrasound vibrations that pass through skin and tissue. When the sound waves reach the implanted mote, its piezoelectric crystal converts the vibrations' mechanical energy into electricity, providing power to a tiny transistor pressing against the nerve. As natural electrical activity in the nerve varies, it changes the current passing through the transistor, thus providing a read-out mechanism for the nerve's signal.

To send the information back out of the body, the neural dust system also uses ultrasound. The external transducer alternates between sending ultrasound vibrations to power the mote and listening

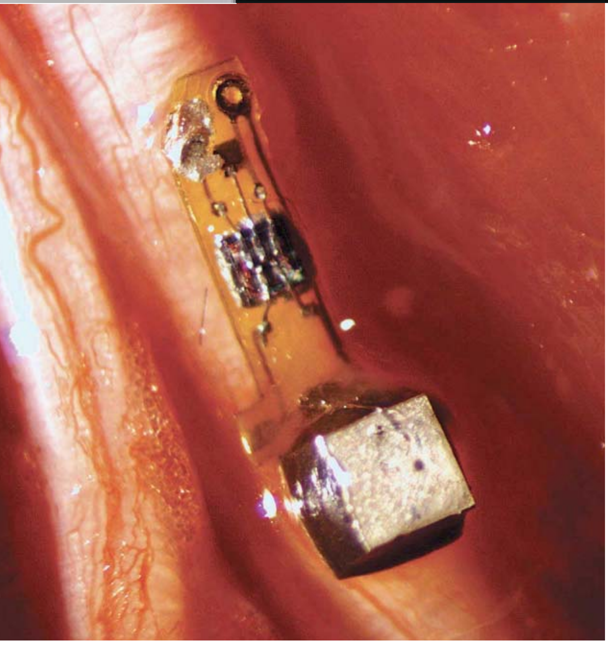

YOU SOUND NERVOUS: This tiny sensor, attached to a rat's nerve, is powered by ultrasound. Nerve signals change the way the device reflects ultrasound, and an external sensor can hear the change.

RYAN NEELY

for the returning echo as some of those vibrations bounce back. The changing current through the transistor alters the piezo crystal's mechanical impedance, thereby modulating how much bounce back the transducer receives.

To make the system viable for humans, Carmena and his colleagues are tackling four distinct technical challenges:

Challenge 1: Make the neural dust system work in rats that are awake and scurrying around.

The use of ultrasound to power the mote has the big advantage of making the system wireless, while other neural recording systems often use bulky batteries or ungainly wires that protrude through the skin. But going wireless has a downside. The external transducer must be aligned precisely with the implanted dust mote so the ultrasound vibrations hit its piezoelectric crystal dead-on. When an animal is unconscious, it's easy to keep that alignment, but it becomes a trickier problem when the rat is awake and on the move.

So study coleader Michel Maharbiz, in Berkeley's electrical engineering department, is now working on a rodent-wearable transducer. A tiny backpack holding the gear could be positioned over the implanted mote to keep the proper alignment, he says.

Still, some experts say the alignment problem may make the neural dust system impractical in human patients. Complicated external components could cancel out benefits gained by the simple wireless implants, says Dustin Tyler, a biomedical engineering professor at Case Western Reserve University, in Cleveland, who works on cuff electrodes that wrap around nerves. "Daily donning and doffing of arrays of [ultrasonic] transducers that need to be properly aligned can be quite daunting," he says, "and often prevent patients from using the system or from using it properly."

Challenge 2: Use multiple motes inside a single animal's body.

In the proof of concept, the researchers not only recorded from a rat's nerve, they also recorded electric signals from a muscle. By attaching motes to many muscles and nerves, the neural dust could be used as an internal sensor network, says Maharbiz. "Deep-tissue temperature sensors would enable monitoring organ function in a way not possible today," he says.

Powering and recording signals from multiple motes will require new techniques and better signal processing. The researchers may use an array of ultrasound transducers that sweep their beams across an animal's body, hitting all the motes in turn. The engineers are also writing new signal processing algorithms to make sense of returning echoes from multiple sources. Carmena is optimistic, noting that bats use the similar process of echolocation to navigate in caves filled with thousands of animals. "If the bats can do it," he says, "we should be able to do it too."

Challenge 3: Use the neural dust not only to record but also to stimulate nerves.

While recording signals from the nerves and muscles provides information about the body, researchers in the hot new field of "electroceuticals" are more interested in stimulating nerves with pulses of electricity to change the body's operations. Electroceutical companies are developing medical treatments that involve stimulating the vagus nerve in the neck, for example, showing progress on reducing inflammation in rheumatoid arthritis and blocking the pain of migraine headaches.

But the current neural dust mote is a very low-power device, drawing only 0.12 milliwatt. Chad Bouton, who heads the Center for Bioelectronic Medicine at the Feinstein Institute, in Manhasset, N.Y., says a stimulating mote will require more power. "They have the right idea, but it's going to take some work to develop that stimulating ability," he says. Bouton's lab has developed another nerve cuff with densely packed electrodes to provide powerful stimulating bursts, but that device currently requires penetrating wires.

Challenge 4: Make the motes small enough to be embedded in the brain.

The name "neural dust" is a bit aspirational: The researchers' latest mote is about 2.4 cubic millimeters, which is far bigger than your average dust particle. If they want to scatter their dust in the human brain, they'll need to make more-minuscule motes.

The researchers say their prototype mote used a commercially available circuit board. They think they can get the size down to 1 mm^3, about the size of the piezoelectric crystal in the prototype, by developing a customized board. And they're already working toward a more ambitious goal: a mote that measures 50 cubic micrometers. At that size, motes could be tucked into the deep folds and wrinkles of the brain, letting us peer into the mysterious, dusty corners of the mind.

—ELIZA STRICKLAND

NEWS

BREAKING THE MULTICORE BOTTLENECK

Simple hardware speeds core-to-core communication

Engineers at North Carolina State University and at Intel have come up with a solution to one of the modern microprocessor's most persistent problems: communication among the processor's many cores. Their answer is a dedicated set of logic circuits they call the Queue Management Device, or QMD. In simulations, integrating the QMD with the processor's on-chip network at a minimum doubled core-to-core communication speed and, in some cases, boosted it much further. Even better, as the number of cores was increased, the speedup became more pronounced.

In the last decade, microprocessor designers started putting multiple copies of processor cores on a single die as a way to continue the rate of performance improvement computer makers had enjoyed without causing chip-killing hot spots to form on the CPU. But that solution comes with complications. For one, it means that software programs have to be written so that work is divided among processor cores. The result: Sometimes different cores need to work on the same data or must coordinate the passing of data from one core to another.

To prevent the cores from wantonly overwriting one another's information, processing data out of order, or committing other errors, multicore processors use lock-protected software queues. These are data structures that coordinate the movement of and access to information according to software-defined rules. But all that extra software comes with significant overhead, which only gets worse as the number of cores increases. "Communications between cores is becoming a bottleneck," says Yan Solihin, a professor of electrical and computer engineering who led the work at NC State, in Raleigh.

The solution—born of a discussion with Intel engineers and executed by Solihin's student, Yipeng Wang, at NC State and at Intel—was to turn the software queue into hardware. This effectively turned three multistep software-queue operations into three simple instructions: Add data to the queue, take data from the queue, and put data close to where it's going to be needed next. Compared with just using the software solution, the QMD sped up a sample task such as packet processing—like network nodes do on the Internet—by a greater and greater amount the more cores were involved. For 16 cores, QMD worked 20 times as fast as the software could.

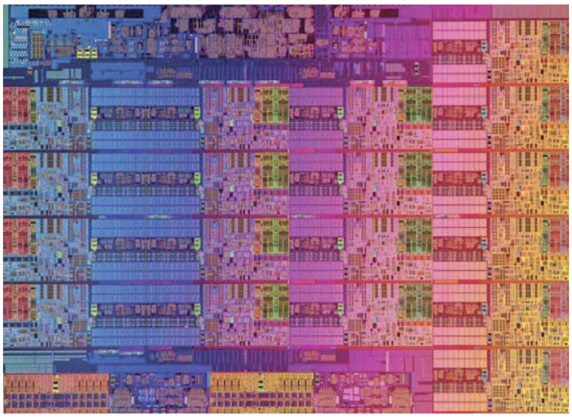

IT'S GETTING CROWDED: This Intel Haswell-EX Xeon E7 V3 processor has 18 cores trying to work together without messing up one another's calculations. A bit of additional hardware could speed up communication among the cores.

INTEL

Once they achieved this result, the engineers reasoned that the QMD might be able to do a few other tricks—such as turning more software into hardware. They added more logic to the QMD and found it could speed up several other core-communications-dependent functions, including MapReduce, a technology Google pioneered for distributing work to different cores and collecting the results.

Srini Devadas, an expert in cache control systems at MIT, says the QMD addresses "a very important problem." Devadas's own solution for the use of caches by multiple cores—or even multiple processors—is more radical than the QMD. Called Tardis, it's a complete rewrite of the cache management rules, and so it is a solution aimed at processors and systems of processors further in the future. But QMD, Devadas says, has nearer-term potential. "It's the kind of work that would motivate Intel—putting in a small piece of hardware for a significant improvement."

The Intel engineers involved couldn't comment on whether QMD would find its way into future processors. However, they are actively researching its potential. (Wang is now a research scientist at Intel.) The engineers hope that QMD, among other extensions of the concept, can simplify communication among the cores and the CPU's input/output system.

Solihin, meanwhile, is inventing other types of hardware accelerators. "We have to improve performance by improving energy efficiency. The only way to do that is to move some software to hardware. The challenge is to figure out which software is used frequently enough that we could justify implementing it in hardware," he says. "There is a sweet spot." —SAMUEL K. MOORE

PHOTOGRAPH BY Jason Henry/The New York Times/Redux

SAILORS NEED NOT APPLY

SELF-DRIVING CARS might not be here quite yet, but self-driving–well, self-sailing–is already happening out on the water. Vessels like the Saildrone, pictured here traversing San Francisco Bay in August, are already proving they can do things (like monitoring the melting of the polar ice caps) that are difficult and costly to do on a traditional boat with a human crew. The wind-powered Saildrone is controlled via satellite by operators at the eponymous company's Alameda, Calif., headquarters. The craft looks like a miniature America's Cup racing yacht. But the Saildrone's ability to haul up to 160 kilograms of equipment–which lets it capture a raft of information for better understanding the oceans and improving weather and climate prediction–limits its top speed to roughly 5 knots.

THE BIG PICTURE NEWS

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

RESOURCES

A DIY LEGO CONTROLLER

A LOW-COST WAY TO PROGRAM LEGO MACHINES

DIOMIDIS SPINELLIS

RESOURCES_HANDS ON

If you want to explore coding with Lego bricks, there's one major option: to use a kit from the company's well-known Mindstorms robotics line. Mindstorms-based machines are built around the Intelligent Brick, which can be programmed using Lego's graphical programming environment or one of a number of third-party alternative languages. But Lego also makes a collection of motors, connectors, lights, and infrared receivers collectively sold under the label of Power Functions. In place of a programmable brick, the Power Function line includes a handheld controller for transmitting command signals. • I wondered if it was possible to use a Raspberry Pi to replace the handheld controller, taking on the role of an Intelligent Brick. This would have some advantages. With programs being created on the same device used to control Lego constructions, it would eliminate the need to download the programs to the brick, speeding up development. The US $40 Pi is also a lot cheaper than the $190 Intelligent Brick. I also wondered if such a setup could be used with MIT's Scratch, a free visual programming environment aimed at children. Scratch extensions are available for use with the Mindstorms brick, but they require altering the brick's firmware, and I wanted to try something simpler. • As I discovered, most of the code required for controlling Lego toys using Scratch is already available as open source software. What was ▸

needed was integration, configuration, and some glue software.

First, I needed to build an infrared control link, which is basically two infrared LEDs operated via the Raspberry Pi's general purpose input/output (GPIO) connector and Lego's receiver. I used schematics and instructions by Alex Bain to build the hardware. For the software, I downloaded and installed LIRC, a package that has support for decoding and transmitting signals used by over 2,500 different infrared remote controls.

Getting the LIRC package to work with my home-brew infrared link was a simple matter of editing some configuration files and specifying which GPIO pins I had wired up for input and output.

Now I needed to get LIRC to send valid Lego command signals. This means specifying the waveform—a pattern of infrared pulses—that must be sent for each Lego command. Fortunately, Lego has released a document specifying the protocol and format of all commands (for example, a binary value of 1 is transmitted by six pulses of IR light at a frequency of 38 kilohertz, followed by a pause of 553 microseconds). The Lego Power Functions system supports up to four receivers working on different channels, and each receiver has a red side and a blue side, each of which can independently control a motor.

Building on this information, Conor Cary created lego-lirc, a Java program that generates command waveforms, complete with the correct checksums, in a format that LIRC understands. I downloaded lego-lirc and, with the Lego documentation in hand, created additional waveforms that allow the transmission of PWM (pulse-width modulation) commands. These commands allow precise speed adjustment of Power Function motors without requiring timing loops in the application software. (To avoid the hassle of running lego-lirc, you can just download my file of generated LIRC waveforms directly from my GitHub repository under the username of dspinellis.) To configure LIRC to use the Lego commands, I copied the waveform to the LIRC configuration directory. I could then send Lego commands from the Pi's command line through LIRC's irsend program.

BLOCK BY BLOCK

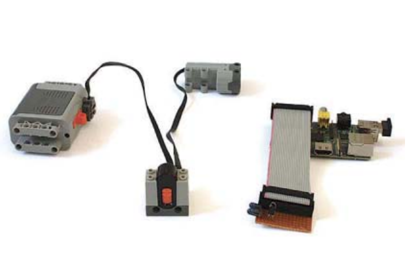

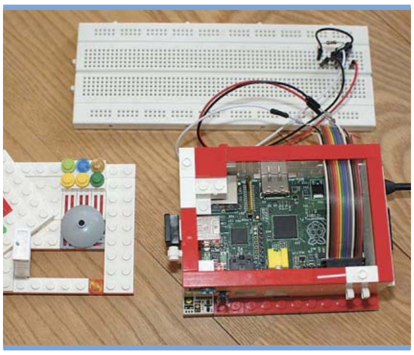

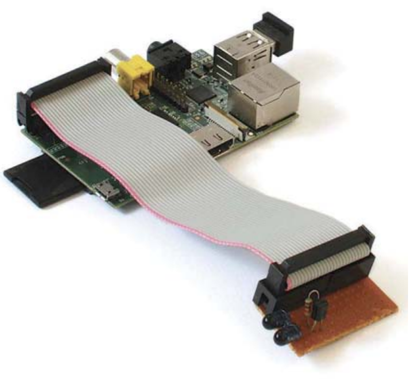

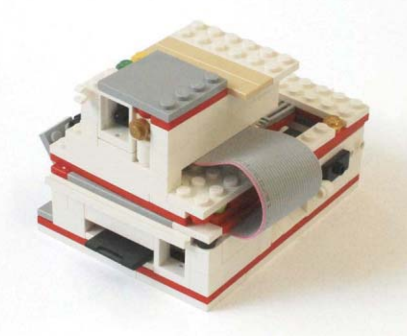

Lego Power Functions allow motors to be controlled with infrared signals [top]. Signals can be generated by connecting infrared LEDs to a Pi [middle images]. A Lego enclosure holds the components [bottom].

DIOMIDIS SPINELLIS (4)

The final step was to issue the LIRC commands from the Scratch environment. I enabled "remote sensor connections" in Scratch. This makes Scratch behave like a local server running on the TCP port 42001. Client software can connect to Scratch using this port and listen for messages from Scratch programs. (It's also possible to have the client software and Scratch environment run on separate machines, so you could have the Raspberry Pi–based infrared interface controlled by a Scratch program running on a desktop computer, for example.) I then installed Phillip Quiza's excellent scratchpy library, which allows you to write Scratch clients in the Python programming language.

Finally, I wrote a Python script that receives Scratch broadcast messages specifying Lego remote commands, and runs the LIRC command-line client to send them (this is also available from my lego-power-scratch GitHub repository). To run the script, run the control.py program in a separate terminal window and launch the Scratch environment. While control.py is running, it will display on its standard output the remote control messages it sends or the errors it detects on the incoming Scratch messages.

In Scratch, programs are constructed by chaining together graphical blocks on screen. Blocks perform functions such as program-flow control and graphics manipulation. To send a message to a Lego Power Functions receiver, a "broadcast" block is used, with a simple text string of the form "Lego <channel> <Blue|Red> <power level>." So, for example, the message "Lego 2 blue -7" will send a signal by way of the Python client and my transmitter to turn the motor connected to the blue side of the receiver on channel 2 at full speed, backward.

How does the system work in practice with its intended audience? I tried it out with a young budding engineer—who quickly wrote a Scratch program to control Lego's Volvo Wheel Loader kit with a computer's arrow keys. **—DIOMIDIS SPINELLIS**

PROFILE: CYPHER
THIS STARTUP PIVOTED FROM CLEANING UP CALLS TO HELPING MACHINES HEAR

JAN STROMME/GETTY IMAGES

When John Walker founded Cypher four years ago, he had a simple premise: No one wants to waste time on a noisy phone call. So he built software for smartphones, incorporating a deep neural network, which could apply machine learning to deliver crystal-clear conversations devoid of background noise.

To show it off, John Yoon, Cypher's head of strategy, recently stood at the busy intersection outside *IEEE Spectrum*'s office building in New York City. Yoon called CEO Walker, who was waiting in our office. They spoke on speakerphone for a few minutes as car horns honked and sirens blared.

Once Yoon switched Cypher's demo program on, the call became as quiet and clear as if he had dialed from a conference room. The company says it can cut out 99 percent of background noise at the cost of introducing a delay of just 24 milliseconds (far below the 200 ms that would be noticeable to a human listener).

To achieve clarity in conversations, Cypher's program is primed to recognize a voice based on characteristics of human speech that algorithms can easily trace. For example, "all human speech incorporates vowels," Yoon explains. "If you look at them on a spectral graph, because of the construction of your nasal cavity, tongue, and teeth, you get these really nice harmonic analyses. A jackhammer does not have the same type of harmonics."

Once the program has picked out a voice, its algorithms strip away everything else. This technique departs from many other software-based approaches, which have often focused on first identifying noise and then subtracting it from speech. However, because noise is much more varied than speech, it's difficult to correctly classify all unwanted sounds.

Deep neural networks have become very fashionable due to their ability to extract useful information from messy data. Sy Choudhury, senior director of product management at Qualcomm, says a wave of startups and manufacturers are working to integrate these networks into smartphones for a variety of purposes. In Cypher's case, the greatest challenge, however, has not been in developing a successful program but in figuring out how to sell it.

Cypher had hoped that the promise of exceptional voice quality would be enough to persuade smartphone manufacturers to install the program. It argued that the technology could set devices apart for consumers in a crowded market. "As phones become commoditized, who's to say Huawei couldn't use this as a differentiator?" says Yoon.

But customers care much more about camera quality and battery life than the clarity of calls. So Cypher tried another approach: convincing manufacturers to replace the dedicated hardware currently used for active noise cancellation with its algorithms, which can be incorporated into the phone's operating system. Cypher estimates that doing so could take 50 U.S. cents to $1 off the cost of making each smartphone.

However, that argument also failed to gain much traction. Though Cypher has completed trials with Samsung, LG, and Huawei, it hasn't inked a single licensing deal since its program launched last fall. Ronan de Renesse, a consumer technology analyst at Ovum, says top-of-the-line smartphone models cost $200 to $400 to manufacture, so saving 50 cents per phone isn't enough to pique manufacturers' interest.

Now, Cypher is again rethinking its strategy, to focus on helping the growing number of voice-activated digital assistants (such as Ok Google, Siri, and Cortana) to hear commands in noisy households. Recently, the team applied their software to Alexa, the digital assistant inside Amazon Echo, the company's at-home taskmaster. In a test run by Cypher, the program improved Alexa's word recognition (as measured by its error rate) by 116 percent as it analyzed hundreds of queries such as, "Alexa, what's the weather like in Reno?"

The company says this improvement makes the difference in whether Alexa, if placed in a noisy kitchen, understands a command or not. But it remains to be seen whether customers think that background noise is a true hindrance for digital assistants and if manufacturers will pay Cypher to reduce it. —**AMY NORDRUM**

Location: Salt Lake City **Founded:** 2012
Employees: 11 **Funding:** US $10 million

TALES OF ATARI
NEW STORIES OF A LEGENDARY COMPANY COME TO LIGHT

SILICON VALLEY LEGENDS: [From left] Former Atari engineer Owen Rubin, head of engineering Al Alcorn, and cofounder Nolan Bushnell.

TEKLA S. PERRY

I thought I'd heard them all. Atari stories, that is. I started covering the company in 1981, followed cofounder Nolan Bushnell and first engineer Al Alcorn through their later adventures, and became personal friends with more than a few Atari alumni.

But in September the IEEE Silicon Valley History Committee hosted a sold-out 100-person event, in Santa Clara, Calif. With key Atari players on hand, behind-the-scenes stories came out that were new to me—and even to some of the players. It's hard to get startups off the ground, particularly those trying to do something as revolutionary as start a new industry. So let's just say the truth, at times, was stretched back then—or simply ignored—in order to make things happen. A few examples:

Bushnell began moonlighting on coin-operated video games while working at Ampex, a Silicon Valley electronics company. He and some colleagues developed *Computer Space* in 1971, an arcade version of *Spacewar*, a game written for the DEC PDP-1 minicomputer. Recalled Atari's head of engineering, Alcorn: "You were doing this [on the side] in the evening, borrowing parts."

"Liberating parts," corrected Bushnell.

"A fine tradition," said Alcorn, "that Atari adopted. That's how the Apple II was made, with Atari's parts."

Steven Mayer, who was Atari's chief architect for home games and computer systems, chimed in: "There was a long tradition at Ampex of supporting people [to go out and do other things]. A guy working on a database at Ampex was Larry Ellison. A guy working in audio was Ray Dolby. Ampex was incredibly generous about letting these people start their own companies."

Computer Space was a failure, but Atari's second game, *Pong*, turned the company into a huge success in 1972. "I wanted world domination," said Bushnell. "And it turns out that there are two coin-op [game] distributors in every city. One would have Gottlieb pinballs, one Williams. We had chosen the best distributors, but the [distributors] who didn't have the Atari brand were doing everything they could to spawn a competitor. So I thought, let's make that happen."

So in 1973 Atari secretly started a second company, Kee Games, with Bushnell's next-door neighbor, Joe Keenan, at the helm. "We took our No. 2 engineer, our No. 1 manufacturing guy, and every other game in our lineup, and gave it to Kee.... We told the distributors, 'Those bastards are stealing our games!' and then the distributors would run off and grab the games."

Moving on to how Atari operated internally, "Nolan [Bushnell] and I had a vibrant, contentious relationship," said Alcorn. "Nolan is this dreamer doing all this crazy stuff, and I'm the one who gets the short end of the stick and has to make it happen."

"Nolan would come into engineering," Alcorn continued, "where we had three teams who had been working on things for months. And he would say, 'This game is s—t; let's do this instead.' That got to be a problem, so I gave our secretary a pager only to be used to page me in the event Nolan got into engineering. I would go in and stand behind him, and as soon as he left I would remind engineers who they worked for."

The first major customer for the home version of Atari's seminal *Pong* was Sears. Says Bushnell, "We were in a conference room with Tom Quinn [the Sears executive], and he says, 'How many can you build?' We had no clue. A big run for us at the time was 10,000 units. So I went out and asked our manufacturing guy. He said 25,000. Then I went back to Quinn and told him 75,000. He gave us an order for 150,000."

The Atari alumni also mentioned the importance of location for their entrepreneurial success: "One of the nice things about Silicon Valley is you all have worked next to an idiot who has gone off and been successful," said Bushnell. "So you think if he did it, I can do it too.

"And, of course, Steve Jobs was working for me, so it goes along." **—TEKLA S. PERRY**

An extended version of this story is available online.

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

SOLAR CHARGING KIT SHOWDOWN, HIMALAYAS EDITION

VOLTAIC SYSTEMS' ARC 20W OUTSHINES GOAL ZERO'S SHERPA 100

FOLDABLE SOLAR POWER: Goal Zero [left] and Voltaic Systems [right] offer portable solar panels that can fit into a backpack and charge battery packs with enough capacity for smartphones, digital cameras, and laptops.

LEFT: GOAL ZERO; RIGHT: VOLTAIC SYSTEMS

I'm not an outdoorsy person. But when a reporting assignment came up recently that required trekking in the Himalayas, there was really only one response: Yes!

The assignment would take me to a village that had no grid power and was a day's hike from the nearest road. My photographer and I would be carrying two laptops, three digital cameras, and a smartphone between us. So finding a reliable and portable source of power was essential.

We were heading to the Ladakh region, a high-altitude desert blessed with a spectacular amount of solar radiation. Looking around for lightweight solar panels and batteries, I found plenty of options. I ended up trying out two units, loaned to us by the manufacturers: Goal Zero's Sherpa 100 Solar Kit and Voltaic Systems' Arc 20W Solar Charger Kit.

Goal Zero is a well-established company, and its US $550 Sherpa 100 kit comes with an 8,800-milliampere-hour lithium-ion battery, a detachable 110-volt AC inverter, and a 20-watt solar panel that folds down to 33 by 21 by 2.5 centimeters. The battery has two USB ports, a laptop port, and a 12-volt port. I never found a use for the latter, but it's designated for lighting. Also included are an AC wall-outlet adapter and car-lighter adapter. The entire kit weighs 2 kilograms.

Voltaic Systems is an up-and-coming player. Its $278 kit comes with a 2,200-mAh lithium-polymer battery, a 20-W solar panel that folds to 19 by 25.5 by 2 cm, plus an AC adapter and lighter adapter. The battery has one 5-V USB port and a DC port for laptop charging. The kit weighs 1.5 kg.

One key difference between the kits is the inverter—Goal Zero's has one, Voltaic's does not. The inverter screws into the battery and lets you pretend you have a wall outlet (within reason, of course; you wouldn't use it to run a refrigerator). But it adds bulk—159 grams—and you lose some energy in converting DC to AC.

Rather than an inverter, Voltaic's system relies on an impressive array of DC adapters and USB chargers. I got USB battery chargers for our Sony and Canon cameras ($7.50 to $10 each) and a Magsafe 2 adapter ($20) for our Apple MacBook Pros.

The solar panels are light and flexible enough to be draped from backpacks, so in theory I could have charged as I hiked. I didn't actually do that. Instead, once we reached our destination, I located a flat, unshaded spot. When the morning sun emerged, I laid out the panels, angled them toward the light, and connected the batteries. Goal Zero's instructions caution against letting the battery get too hot, but I didn't find that to be a problem. What *was* an issue was the ever-present dust, which had me wiping off the panels several times a day.

My photographer and I never wanted for power. Every night she downloaded photos and video from her cameras and her iPhone to her laptop, backed them up to an external drive, and recharged the occasional camera battery. I used my computer and camera more sparingly, but I obliged several requests from others to recharge phones, cameras, and a GoPro.

The panels got about 10 hours of sun daily. The Goal Zero battery is supposed to take 10 to 20 hours to fully charge, but even when the battery started from a partial charge, it never reached 100 percent. Several times scattered clouds caused the unit to stop charging. Those same clouds did not pose a problem for the Voltaic system. Voltaic's battery, which the company says takes 7 to 10 hours to charge in full sun, did indeed fully charge every day.

So which system did I prefer? Both seem well made and engineered, but the Voltaic system had the edge in solar charging. And factoring in the significant difference in price and weight, I'll opt for Voltaic for my next off-grid adventure. **—JEAN KUMAGAI**

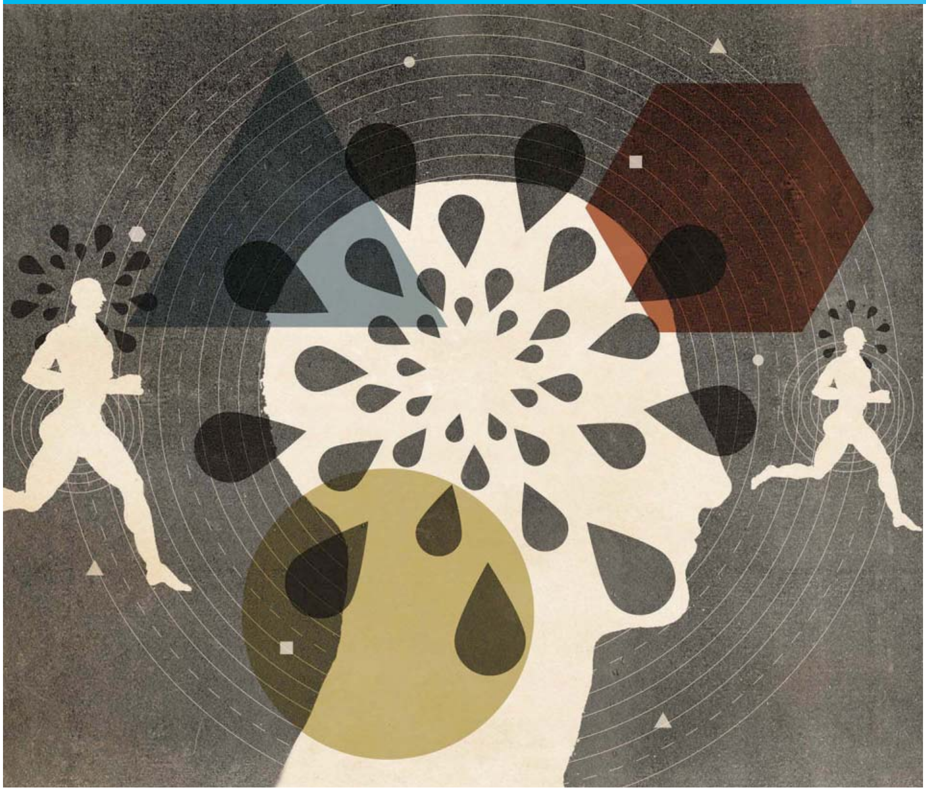

RUNNING, SWEATING, AND PERSISTENCE HUNTING

DURING THE TWO YEARS OF ITS MONTHLY APPEARANCE, this column has looked at many objects–cars, turbines, airplanes, windows, mobile phones, and nuclear reactors–made by humans. Today's focus is on the human body, specifically the way it keeps itself cool. • Before the development of long-range projectile weaponry some tens of thousands of years ago, in Africa, our ancestors had only two ways to secure meat: by scavenging the leftovers of mightier beasts or by running down their own prey. Humans were able to occupy the second of those ecological niches thanks, in part, to two great advantages of bipedalism. • The first advantage is in how we breathe. A quadruped can take only a single breath per locomotive cycle because its thorax must absorb the impact on the front limbs. We, however, can choose other ratios, and that lets us use energy more flexibly. The second, and greater, advantage is in our extraordinary ability to regulate our body temperature, which allows us to do what lions cannot: to run long and hard in the noonday sun. • It all comes down to sweating. The two large animals we have mainly used for transport perspire profusely, compared to other quadrupeds: In one hour a horse can lose about 100 grams of water per square meter of skin, and a camel can lose up to 250 g/m^2. However, a human being can easily shed 500 g/m^2, enough to remove 550 to 600 watts' worth of heat. Peak hourly sweating rates can surpass 2 kilograms per square meter, and the highest reported short-term sweating rate is twice that high.

We are the superstars of sweating, and we need to be. An amateur running the marathon at a slow pace will burn 700 to 800 W, and an experienced marathoner who covers the 42.2 kilometers in 2.5 hours will burn about 1,300 W.

And we have another advantage when we lose water: We don't have to make up the deficit instantly. Humans can tolerate considerable temporary dehydration providing that they make up the deficit within a day or so. In fact, the best marathon runners drink only about 200 milliliters per hour during the race.

Together these advantages allowed our ancestors to become the unrivaled diurnal, high-temperature predator. They could not outsprint an antelope, of course, but during a hot day they could dog its heels until it finally collapsed, exhausted.

Documented cases of such long-distance chases come from three continents and include some of the fleetest quadrupeds. In North America, the Tarahumara of northwestern Mexico could outrun deer. Further north, Paiutes and Navajos could exhaust pronghorn antelopes. In South Africa, Kalahari Basarwa ran down a variety of antelopes (mostly duikers, gemsbok, and kudus but also larger eland) and during the dry season even wildebeests and zebras. In Australia, some Aborigines would outrun kangaroos.

These runners even had an advantage over modern runners using expensive athletic shoes: Their barefoot running not only reduced their energy costs by about 4 percent (a nontrivial advantage on long runs) but it also exposed them to fewer acute ankle and lower-leg injuries.

In the race of life, we humans are neither the fastest nor the most efficient. But we are certainly the most persistent. ■

ILLUSTRATION BY Chad Hagen

THE EVER-EVOLVING FIELD OF ELECTRICAL ENGINEERING

I WAS AT A LARGE HARDWARE STORE, holding a power cord that I thought would fit my new home generator, when another shopper pointed to the cord in my hand and said that I had chosen the wrong connector. When I looked hesitant, he added, "I'm an electrical engineer, so I know things like that." ● Needless to say, I quietly changed my choice. But I started thinking: What are those things that *all* EEs know? What is the commonality of training and experience that holds us together as a profession, and how is it changing? What will it mean to be an EE in the future? ● I looked at the EE curricula in a number of universities. In the first two years of study, there is a lot of uniformity, with courses in basic math, physics, and computing. The EE-specific part begins with circuit analysis and design. After those first two years of school, however, EE students branch out in perhaps a half dozen different directions. A student who studies the physics of electron devices might have little in common with one who studies information theory. But each would be an EE, and almost the only specialized knowledge that they would hold in common other than basic science and math would be the principles of circuit design. So I decided: You are an EE if you know–or once knew–circuit design. ● After college, EEs enter so many different specialties and occupations that they are almost impossible to categorize. Nonetheless, I looked at the IEEE organization as a framework for professional practice. There are currently 39 societies within the IEEE that serve to guide publications and conferences. As engineering practice changes, so must the societies, yet over the last 20 years only a few societies have been added and one has disappeared, which indicates only moderate evolution. However, even the oldest societies mutate their domains while maintaining their descriptive legacy names. In addition, many societies have generic names denoting a function, rather than an underlying technology, such as the IEEE Communications Society.

Wikipedia defines electrical engineering as "a field of engineering that generally deals with the study and application of electricity, electronics, and electromagnetism." But it seems to me that we have moved well beyond that classical role. Many of the problems today are at the layers above the physical level, involving software, systems, algorithms, design, and cognition, and are usually cross-disciplinary. For example, the most hyped technologies today include machine learning, big data, security, autonomous vehicles, robots, and the blockchain. While electronics is the enabler in all of these, it is not where the important problems are. Nonetheless, we electrical engineers have already occupied this higher ground, as our training, aptitudes, and inclinations seem well suited for this work.

Practicing engineers will go wherever there is interesting work. My concern is more with the evolution of the educational and professional institutions that support engineers, and with the perception of electrical engineering held by potential students. The IEEE has already had at least two episodes of soul-searching over the scope of its domain–in 1963 when it incorporated "electronics" into its name, and some years later when it wrestled with the role that computer science would play.

As circuit design and other classic electrical engineering tasks are increasingly performed by computers and electrical engineers move to ever higher and more functional design, what will be the commonality that holds us together in the future? And how will electrical engineering be distinguished from other branches of engineering?

When I first entered electrical engineering in college, I had very little idea of what EEs actually do. After all these years, I'm still not sure. ■

ILLUSTRATION BY Dan Page

The Molten Tin Solution

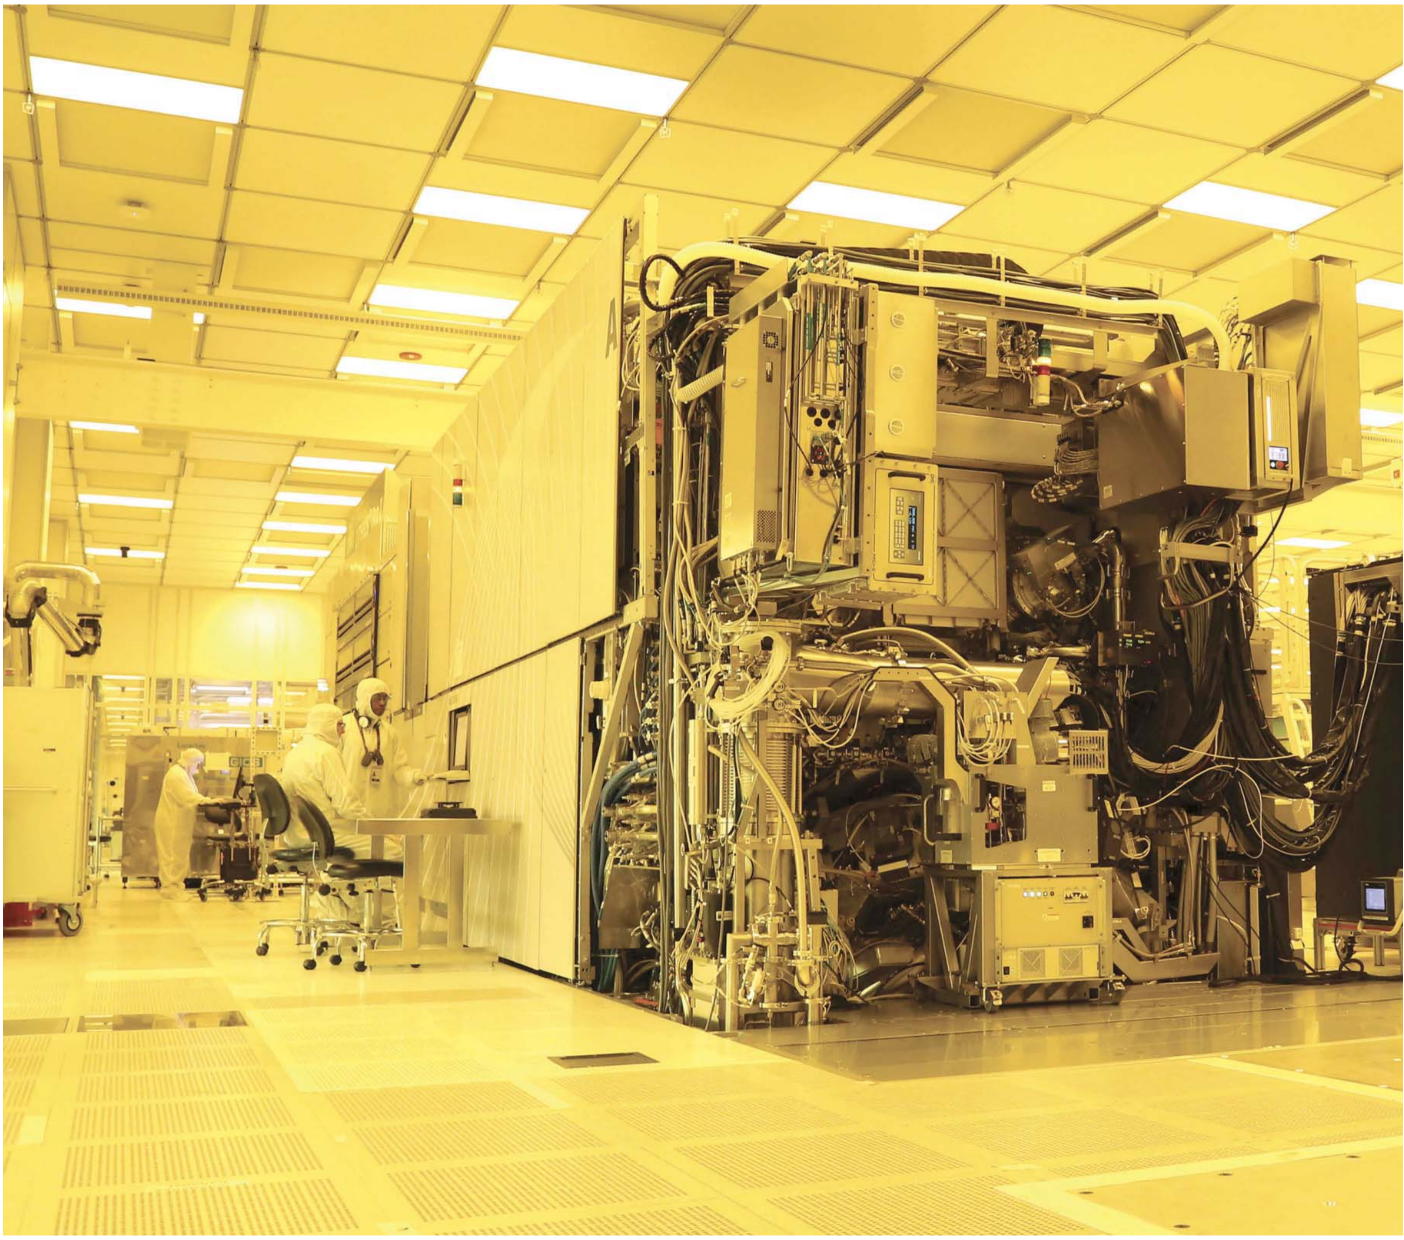

PUTTING EUV TO THE TEST: This EUV scanner (an ASML NXE:3300B) is used to print chip features at a SUNY Polytechnic Institute facility in Albany, N.Y. The EUV light needed to expose wafers is created near the bottom of the scanner, on the side visible in the foreground of this photograph. The far end of the machine is attached to a "track" that coats the wafers before exposure and processes them once they are done.

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

After a 30-year journey, EUV chip-printing technology is now poised to help keep Moore's Law on track

BY RACHEL COURTLAND

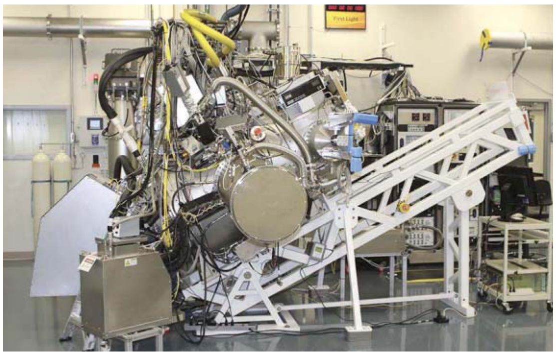

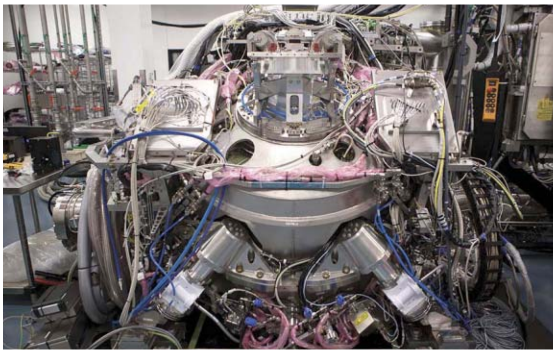

INSIDE THE MACHINE: To generate EUV, pulses of CO_2 laser light are sent into a vessel [top and middle] where they collide with tiny tin droplets to create plasma. This partially assembled EUV scanner [bottom] at ASML's headquarters in Veldhoven, Netherlands, is one of the company's more recent models.

LEFT: IBM RESEARCH; RIGHT: ASML (3)

Even after you don a bunny suit and get deep inside Fab 8, it's hard to get a sense of scale. Rows upon rows of tall machines, known as tools, dominate this US $12 billion GlobalFoundries facility, built amid forest north of Albany, N.Y. Carriers containing silicon wafers zip overhead along ceiling-mounted tracks, like tiny inverted roller coasters. If your timing is good, you'll be standing by a tool when one of those carriers descends to join it, moving a wafer along to the next step in the three-month-long process it takes to turn a dinner-plate-size disk of raw silicon into chips that could be used inside smartphones, personal computers, and servers. That's right: Begin making a microprocessor here on New Year's Day and it may just be finished by the start of spring.

More than 60 times in the course of this advanced manufacturing process, a wafer will be coated with a light-sensitive substance and enter a light-tight box called a scanner. There, in a process called photolithography, laser light shines through a patterned surface and casts shrunken versions of that pattern onto the wafer, creating the ultrafine features that are needed to build the minute transistors and wiring inside cutting-edge processors.

There is little to distinguish these lithography machines from the myriad other tools in this vast ocean of automation. There is no big red sign that flashes "Critically important step here!" But lithography, explains Fab 8 general manager Tom Caulfield, "is the heartbeat of the fab."

Think of these scanners as the front line of Moore's Law, the repeated doubling of the density of integrated circuit components that has defined more than 50 years of astounding technological progress. For decades, a steady series of remarkable breakthroughs, many of them in photolithography, have enabled chipmakers to repeatedly shrink chip features, keep the length of R&D cycles under control, and economically pack more and more transistors on a chip. Those advances have taken us from chips with thousands of transistors in the early 1970s to billions today.

But to keep the good times rolling, GlobalFoundries and other leading-edge chipmakers won't be able to rely on the brilliant lithographic advances of the past. And so they're contemplating another radical shift, one that could prove to be the most challenging yet.

For the entirety of its existence, semiconductor lithography has been done with electromagnetic radiation that was more or less recognizable as light. But for the change chipmakers are now weighing, the radiation is something else altogether. It's called extreme ultraviolet (EUV) radiation, but don't let that name fool you. Unlike the ultraviolet light used in today's scanners, EUV can't travel in air, and it can't be focused by lenses or conventional mirrors.

And it's also difficult to produce; the process begins by firing laser light at a rapid-fire stream of tiny molten tin droplets. The hope is that scanners built to use the resulting 13.5-nanometer light–a wavelength that is less than a tenth of what is used in today's most state-of-the-art machines–will save chipmakers money by allowing them to print in a single step layers that would otherwise require multiple exposures.

But creating EUV systems that are bright and reliable enough to operate in the fab–nearly 24 hours a day, 365 days a year–has proved to be a monumental engineering challenge. For many years, EUV faced significant skepticism and repeatedly failed to live up to predictions that it was almost ready for prime time.

Now, though, the technology really does seem to be turning a corner. The brightness of the EUV light source made by Dutch lithography-tool manufacturer ASML Holding seems to be closing in on a figure long targeted for commercial production. ASML, which has emerged as the technology's standard-bearer, is now shipping EUV scanners that it says should be ready to mass manufacture leading-edge microprocessors and memory starting in 2018. The world's most advanced chipmakers are working hard to determine when and how these machines will be incorporated into their production lines.

The stakes are high. Moore's Law is facing significant challenges, and no one is sure how the semiconductor industry–which grossed more than $330 billion last year–will navigate the next five or 10 years or what a post-Moore's-Law semiconductor industry will look like. A decline in revenues might be inevitable. But if keeping the "law" in effect avoids, say, a 15 percent drop in the industry's income, that would keep an amount of money flowing that is twice as great as the total revenues of the U.S. video game industry.

The fineness of the details that can be made with a photolithographic system depends on several factors. But a powerful way to make dramatic improvements is to shorten the wavelength of the light it uses. For decades, lithographers have done just that, shifting their wafer-exposing tools from operation at the blue edge of what's visible to the human eye down to successively shorter wavelengths in the ultraviolet part of the spectrum.

In the late 1980s, the semiconductor industry was beginning the process of shifting from mercury lamps to lasers as the light source of choice, reducing the wavelength from 365 nm to 248 nm in the process. But some researchers were already contemplating a far bigger jump, into the X-ray range. Hiroo Kinoshita, then at the Japanese telecommunications firm NTT, reported the results of early work on this idea way back in 1986, using 11-nm radiation. Others, at AT&T Bell Laboratories and at Lawrence Livermore National Laboratory, also explored the technology independently. In 1989, some of these researchers met and traded notes at a lithography conference. In ensuing years, research

Inside the Scanner

To make patterns with EUV, engineers had to leave lenses behind. A series of mirrors brings EUV radiation from the scanner's light source [bottom right] to a mask, which carries the patterns to be printed, and then on to the wafer. An attached "track" [left side, not shown] brings wafers in and out of the scanner. The masks have a separate entrance.

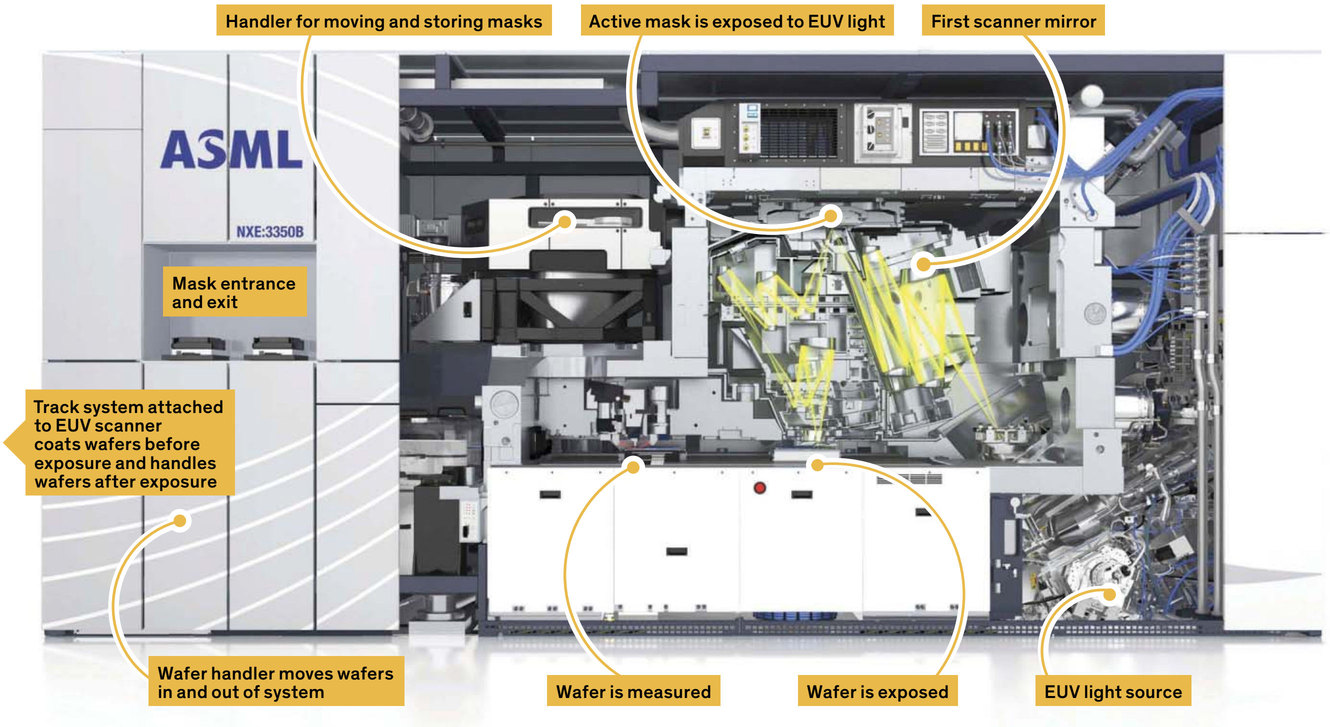

into the notion got infusions of investment from government and industry.

ASML and several partners began work on what was by then known as EUV lithography in the late 1990s. That was when Anton van Dijsseldonk, who grew up in Veldhoven, the Dutch town where ASML is headquartered, became the company's first full-time employee on the project. "The end of Moore's Law was predicted," he recalls, and the semiconductor industry was hunting for ways to keep resolution improvements going. Chipmakers were also struggling to improve overlay—the ability to put a wafer back in the scanner again after its peregrinations through the fab and have the next set of patterns print in exactly the right place. "People were looking in those days at the alternatives," van Dijsseldonk says, "and EUV was the exotic one."

But from the beginning, ASML's EUV researchers were convinced they could make the technology work—and that it would be the most economical option for chipmakers. Before the decade was out, the company had decided to build demonstration scanners that could be used by other researchers to test the approach.

Nothing about it was easy. Physics offers few favors for engineers hoping to cast patterns with what are essentially X-rays. At 13.5 nm, the wavelength the company ultimately chose, light is readily absorbed by many materials. Even the air we breathe "is absolutely black," absorbing every last bit of the radiation, van Dijsseldonk notes. So he and his colleagues realized early on that the only way an EUV scanner could work was in vacuum, with each wafer entering and leaving the scanner through an air lock.

And then there's the problem of bending the radiation. EUV is also absorbed by glass, so directing it through the machine would require a shift from lenses to mirrors. And not just any mirrors. A simple polished surface would not be nearly reflective enough, so they'd have to use Bragg reflectors—multilayer mirrors that can constructively reinforce many small reflections into a single, reasonably strong one.

Today, the mirrors inside ASML's EUV machines consist of 40 pairs of alternating silicon and molybdenum layers—each just a few nanometers thick. Zeiss, the company that developed these mirrors, constructs their aspheric surfaces with great precision. But at the end of the day, van Dijsseldonk says, "if you do it fantastically [well], you get a mirror with a reflectivity of 70 percent." That level of reflectivity means that, for every pair of mirrors used in the system, the light is cut by half. And a scanner could easily require a dozen mirrors to take light from the source to the mask—itself a mirror—and then on to the wafer. After an EUV beam has traversed this gauntlet, less than 2 percent of the light may be left.

ASML

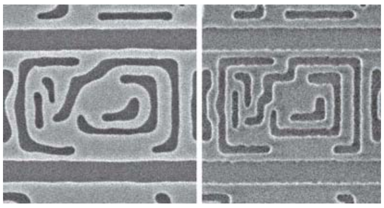

Curves and Corners

EUV promises to create sharper shapes [right] than those that can be created through multiple patterning with today's 193-nanometer light [left]. The lines in these micrographs have a minimum width of 24 nm.

HEADED EAST: A portion of an ASML NXE:3350B EUV scanner is loaded onto a 747-400 extended-range freighter at the Amsterdam Airport Schiphol in the Netherlands, for shipment to a customer. The scanner is split into nine parts (some larger than others) for shipping.

ASML (5)

The less light that reaches a wafer, the longer a wafer must remain in the scanner to be exposed. And in a fab, time means money. For EUV to make it into commercial use, it needs to be able to compete with the cost of existing lithographic methods. So the losses among the mirrors have to be compensated by a radiation source that is extremely bright. And that proved to be really, really hard to engineer.

In the early days, EUV researchers used just about everything they could think of to generate X-rays, including lasers and particle accelerators. But the method that won out, which seemed to offer a practical and economical way to get sufficient brightness, employs plasma. Zap the right material with a powerful enough laser or electric current and you can separate electrons from the atoms they're attached to. The resulting plasma will radiate EUV as the superhot blob cools back to its prezapped state.

The oft-repeated target for such plasma-based light sources has been 250 watts at the intermediate focus, the location where the EUV light exits the source and enters the scanner. That level of light would enable the machine to produce in the neighborhood of 125 wafers per hour, a mass production target that is about half the rate of today's advanced 193-nm machines.

But for a number of years, progress was slow, and brightness improvements consistently lagged behind predictions. By 2011, some five years after ASML delivered its first test scanners to two customers, one of the leading light-source developers, San Diego-based Cymer, had succeeded in developing a source that could consistently deliver 11 W. "We probably underestimated how difficult it would be," concedes ASML's Hans Meiling, who is responsible for EUV product marketing. In the end, in an effort to accelerate development, ASML put in a bid for Cymer, formally completing the reported €3.1 billion acquisition in 2013.

To create EUV light, Cymer uses an approach called laser-produced plasma, which fires 50,000 microscopic droplets a second of ultrapure molten tin across a vacuum chamber, hitting each with powerful CO_2-laser light generated by a series of amplifiers derived from a design originally used for metal cutting. When a laser pulse hits a molten tin droplet, it heats it up into an EUV-emitting plasma. A collector mirror reflects light created in this process and casts it into the scanner. Because the approach generates EUV light as well as tin debris, hydrogen gas constantly flows across the collector mirror to keep it from being rapidly covered with a layer of tin.

"The first time I heard about it, I thought it was insane," admits ASML's Alberto Pirati, who joined the company's EUV light-source program in early 2013. But little by little, the team achieved the seemingly impossible. One of the biggest breakthroughs came with the introduction of a technique the Cymer team began exploring before being acquired by ASML. They found that if they fired a "prepulse" before the main laser, they could flatten each tin droplet into a pancake, creating more surface area for the main laser to hit and increasing how much of the tin droplet was converted to plasma. The change has boosted the laser-to-EUV conversion efficiency from a meager 1 percent to some 5 percent. Earlier this year, thanks to the prepulse and other optimizations, ASML reported that it had reached 200 W in the lab. Another light-source developer, Gigaphoton, has also reported great progress. The long-awaited production target of 250 W no longer seems far off. But the true test of whether EUV is ready to go into production will happen in the labs and fabs–and spreadsheets–of ASML's chipmaking customers.

Nobody doubts that EUV machines can make fine features. Go to a semiconductor conference and you'll probably come across a presentation with sharp-looking micrographs of EUV-made pat-

The Evolution of a Droplet

To create EUV light, molten tin droplets are flattened by one laser pulse and then converted to light-emitting plasma by a second pulse.

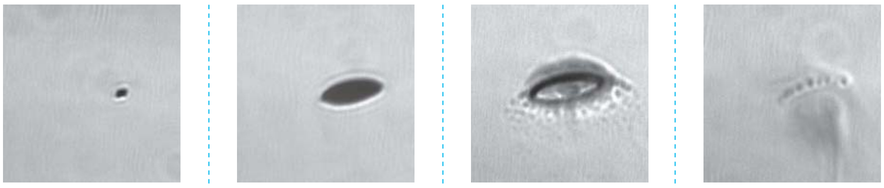

ASML (4)

terns, juxtaposed against blurrier ones made using today's conventional methods.

The question now is what role EUV will play in the mass production of major commercial chips–and when it will do it. Certainly, the cost of adopting EUV seems daunting. The list price of ASML's newest EUV machine exceeds €100 million, more than twice that of an average 193-nm scanner, says spokesperson Niclas Mika. It is about the height and width of a New York City bus and is shipped in multiple 747s. Customers estimate the electrical load could be some 1.5 megawatts, significantly more than that of a 193-nm machine.

But a simple comparison of specs won't capture all the costs of manufacturing. Today's state-of-the-art 193-nm lithographic systems can produce chip features with dimensions that are a small fraction of that wavelength. This feat was made possible by two major breakthroughs. The first was immersion lithography, which places water between a silicon wafer and the optics that project a pattern onto it. The second was multiple patterning, which splits the process of making the features in a layer into two or more steps. To create a set of closely spaced holes, for example, a wafer can be run once through the scanner to fashion half of them, and then again to fashion the other half, slightly offset. Because wafer position can be determined with great precision, multiple patterning lets engineers create features that are more closely spaced than would be possible in one step. In principle, the more steps used, the finer the features. But each addition makes the process of bringing a chip into the world costlier and more complex.

Today, GlobalFoundries uses triple patterning when it makes its 14-nm chips, the most advanced ones currently created in Fab 8. This means that, for certain critical layers, a chip takes two extra passes through a scanner–and every other tool that is used to make those layers. And the company anticipates going to quadruple patterning at 7 nm, its next chip generation, says George Gomba, who is leading the company's task of evaluating the technology at a SUNY Polytechnic Institute facility in Albany, along with colleagues from IBM.

For now, GlobalFoundries plans to roll out its 7-nm chips in 2018 without EUV, but it is reserving the option of pulling the technology in when it is ready. A key question for Gomba and his colleagues is when the cost of EUV will break even with multiple patterning. And it's a very tricky question to answer because it depends on a number of unknown factors, including how bright EUV light sources will become and the uptime of an entire EUV lithographic system–the percentage of time it's actually available to be used.

GlobalFoundries and IBM are not the only ones that have poured money into EUV. In 2012, Intel, Samsung, and Taiwan Semiconductor Manufacturing Co. (TSMC) committed a total of €1.38 billion in R&D funding to ASML for next-generation lithography research (the same deal garnered ASML €3.85 billion for nonvoting shares in the company). ASML's Meiling estimates about 4,000 people work on EUV for the company, a figure that does not include the researchers at leading chipmakers and research institutions with EUV programs of their own.

The reason for all this investment is not only that EUV is hard but that chipmakers are coming around to the idea that, soon, they may not be able to move forward without it. If you ask Anthony Yen, who leads EUV lithography development at TSMC, how critical EUV is to Moore's Law, he won't beat around the bush: "Totally critical. 100 percent critical. Very, very critical." TSMC expects to adopt EUV in 2020, when the company aims to begin producing chips on its 5-nm manufacturing line.

For now, though, there are still some engineering challenges. At the top of Yen's list is protecting the mask, the stencil-like surface that carries a pattern to be printed on the wafer. Just as the optical components of the EUV scanner are reflective, so too is the mask. And here, once again, the pesky matter of reflectance becomes an issue.

In a 193-nm immersion machine, the mask is protected by a thin film called a pellicle, which is suspended a short distance away and stretched across it like a piece of plastic wrap. At current feature sizes, a piece of dust too small to see with the naked eye is still big enough to blot out hundreds of transistors. Thanks to the magic of optics, if a piece of dust falls on the pellicle, it will be too out of focus to create a pattern on the wafer.

But these 193-nm pellicles aren't designed to be transparent to 13.5-nm light; EUV would quickly obliterate them. ASML planned to build the scanners without pellicles, but chipmakers weren't comfortable with the potential downside. "If one particle gets on a mask," Yen explains, "every single die on your wafer is damaged. Basically, you can end up with zero yield." Months of work and, depending on the number of wafers produced, tens or even hundreds of thousands of dollars' worth of chips would be lost.

| CONTINUED ON PAGE 41

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


light

A Pocket Camera With Many Eyes

By Rajiv Laroia

THIS CELLPHONE-SIZE CAMERA'S PICTURES ARE AS GOOD AS THOSE FROM TODAY'S DSLRS

Photograph by Sue Tallon

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

THE BEST DIGITAL CAMERAS TODAY ARE SLRs (single-lens reflex cameras), which use a movable mirror to guide the same light rays that fall on the sensor into the viewfinder. These cameras normally have precisely ground glass lenses and large, high-quality image sensors. In the right hands, they can shoot amazing pictures, with brilliant colors and pleasing lighting effects, often showing a crisply focused subject and an aesthetically blurred background. • But these cameras are big, heavy, and expensive: A good digital SLR (DSLR) with a decent set of lenses–including a standard 50 mm, a wide angle, and a telephoto, for example–can easily set you back thousands of dollars. • So most photos today aren't being shot with DSLRs but with the tiny camera modules built into mobile phones. Nobody pretends these pictures match the quality of a photograph taken by a good DSLR; they tend to be grainy, and the camera allows very little artistic control. But smartphone cameras certainly are easy to carry around. • Can't we have it both ways? Couldn't a high-quality yet still-tiny camera somehow be fit into a mobile device? • That's the question I asked myself five years ago. And the very positive answer, announced last October and shipping early in 2017, is coming from a company I started: Light. • The Light camera starts with a collection of inexpensive plastic-lens camera modules and mechanically driven mirrors. We put them in a device that runs the standard Android operating system along with some smart algorithms. The result is a camera that can do just about everything a DSLR can–and one thing it can't: fit in your pocket. More on how it works later. First, let me tell you a little bit about my background, because that helps explain how–and why–I came up with this approach.

In 2011, I was looking for my next challenge. I had just left Flarion Technologies, a company I had founded and later sold. My engineering career–as a researcher at Lucent Technologies' Bell Labs, before I started Flarion–had been in communications and information theory, so I'd expected to stay involved in those technologies. But I found myself instead thinking more and more about cameras.

I had taken photographs with film cameras as a child in India, but I never had any particularly good equipment. I was intrigued when digital photography started taking off, but I didn't buy my first digital camera until 1999, when my daughter was born. I took a lot of pictures with that Kodak camera, then moved on to various Sony digital cameras. Eventually, when DSLRs came out, I went all out, purchasing a couple of DSLR cameras, a bunch of Canon lenses, and all sorts of other equipment–three camera bags' worth. And I took lots of pictures with this gear.

Ten-plus years later, I got an iPhone and also started taking pictures with it, not because the quality was really there but because of the convenience: I had to plan ahead if I wanted to take pictures with one of my good cameras, which I did less and less. As I talked with other avid photographers, I discovered I wasn't the only one who had expensive camera gear gathering dust. It's not that any of us were happy with the quality of the pictures we were taking with our phones–indeed, we were all frustrated by it. But at the end of the day, convenience always won out.

So when I was looking around for a new professional challenge, I realized I could attack a problem I was dealing with myself. But I didn't start immediately; I didn't know optics well, and I assumed that if there was to be a technical solution it would come from somebody who did.

But I kept investigating, and I soon discovered that most experts in optics don't really understand digital-image processing, and electrical engineers and computer scientists who understand digital images don't generally

PHOTOGRAPH BY Gabriela Hasbun

know much about optics. The solution, I later realized, straddled both the digital and the optical worlds. And I had the luxury of being able to sit at home and teach myself optics for a year.

I had only a rough idea of how the problem might be solved at that point. But as I dug in, I realized not only that there might be a technical solution but also that this was the perfect time to try to build a different kind of commercially viable camera.

I knew that cellphone camera lenses were molded out of plastic–that's why they were so inexpensive. And thanks to cellphones, molded plastic lens technology had been nearly perfected over the previous five years to the point where these lenses were "diffraction limited"–that is, for their size, they were as good as the fundamental physics would ever allow them to be. Meanwhile, the cost had dropped dramatically: A five-element smartphone camera lens today costs only about US $1 when purchased in volume. (Elements are the thin layers that make up a plastic lens.) And sensor prices had plummeted as well: A high-resolution (13-megapixel) camera sensor now costs just about $3 in volume.

While smartphone lenses have become extremely good, the quality of the smartphone camera today is nowhere comparable to that of a high-end DSLR. There are four main reasons: First, the lenses of smartphone cameras are small and collect very little light. And you can't produce good pictures without capturing enough light energy. So smartphone photos will often be "noisy" or grainy, particularly in low light. Second, small sensors and a high pixel count mean that the individual pixels are tiny (approximately 1 micrometer across), and therefore hit a saturation point after receiving just a little bit of light. This results in pictures with very limited dynamic range (limited differences between the darkest darks and the lightest lights). Third, smartphone-camera lenses have a fixed focal length and so can't zoom. Finally, because of the small lens aperture, smartphone pictures have a very large depth of field–that is, they are sharp over a very large range of subject distances from the camera. That might sound like a good thing, but it's not, because controlling the depth of field is essential for artistic photography.

I thought that all these shortcomings of smartphone cameras might be overcome by using multiple smartphone camera modules to take multiple pictures simultaneously, which could then be combined digitally. By using many modules, the camera could capture more light energy. The effective size of each pixel would also increase because each object in the scene would be captured in multiple pictures, increasing the dynamic range and reducing graininess. By using camera modules with different focal lengths, the camera

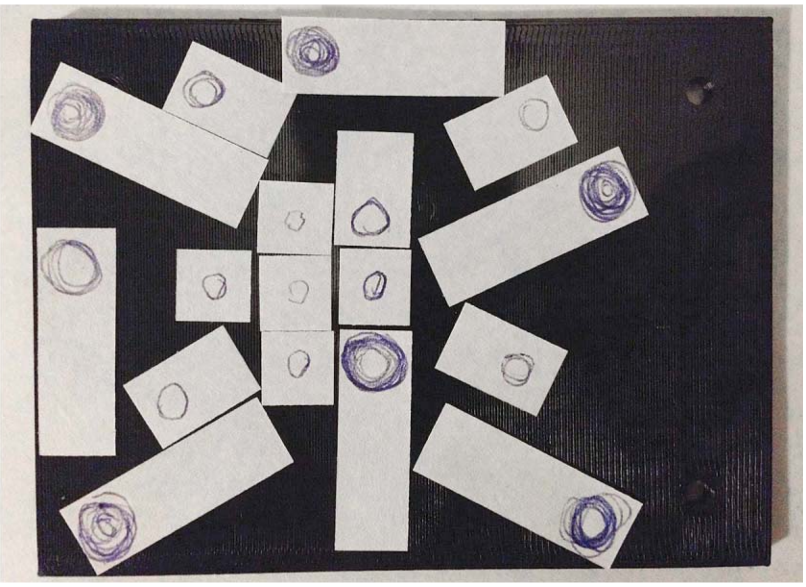

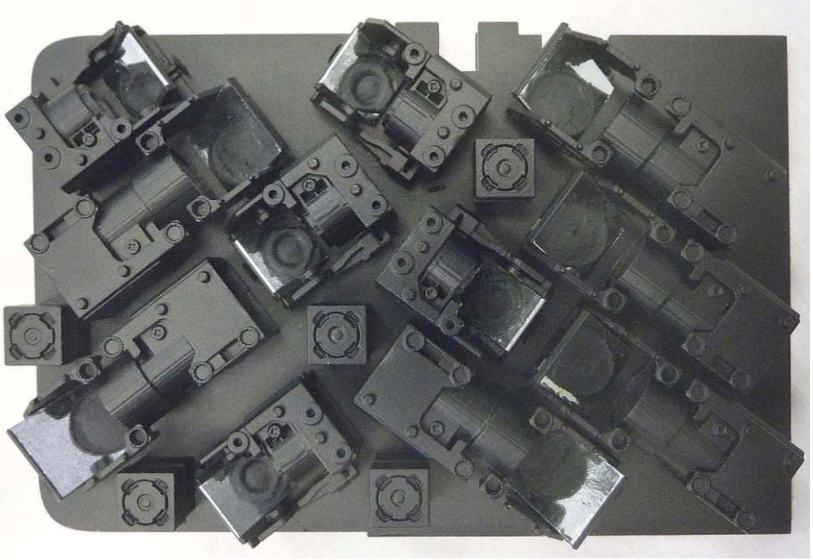

PUZZLING IT OUT: Fitting the 16 cameras [above] into a box of a reasonable size turned out to be a challenge. Light founder Rajiv Laroia [opposite] did paper sketches and then tried to figure it out on a computer-assisted-design system. Eventually, he cut out bits of paper [top] and worked with them as if they were part of a jigsaw puzzle.

TOP: RAJIV LAROIA; BOTTOM: TEKLA S. PERRY


Standard image

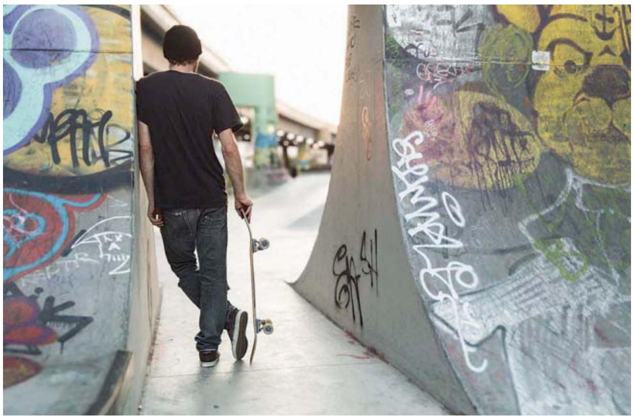
Controlled depth of field


28 mm


70 mm


150 mm

CHOICES, CHOICES: Though the individual camera modules have fixed focal lengths and aperture sizes, combining the images from the lenses in different ways allows photographers to change the depth of field and zoom using the focal-length equivalents of 28 mm through 150 mm.

would also gain the ability to zoom in and out. And if we arranged the multiple camera modules to create what was effectively a larger aperture, the photographer could control the depth of field of the final image.

While at that point I knew it was theoretically possible to use multiple modules to overcome the shortcomings of a small camera, I needed to confirm that such a strategy was truly practical. So I figured I'd better check with some optics experts. In the first half of 2013, I cold-called Julie Bentley, a lens-design expert at the University of Rochester, in New York, and drove 5 hours from where I was living in New Jersey to meet with her.

It started as a contentious conversation: She basically told me I was crazy. But she kept asking questions. I answered them, apparently convincingly, because after about 45 minutes, she became supportive of the approach I was proposing and sent me to Moondog Optics in Fairport, N.Y., a company that does optical design. I met with the CEO, Scott Cahall, who thought what I was suggesting was doable.

And so later in 2013, I joined with Dave Grannan, who had just left Nuance Communications, and officially started Light.

The first and current version of the Light camera—called the L16—has 16 individual camera modules with lenses of three different focal lengths—five are 28-mm equivalent, five are 70-mm equivalent, and six are 150-mm equivalent.

"Equivalent" means that the lens achieves the same field of view as a lens of the specified focal length in a conventional film camera. For a simple single-lens element, when light comes from far away and hits the lens, it converges at a point. The focal length represents the distance of that point from the lens. For example, the equivalent focal length of a human eye is around 50 mm. A lens with a larger focal length has a higher magnification and makes the subjects appear larger.

All the lenses in the L16 are molded plastic, and all the camera modules capture images using standard CMOS sensors, similar to those used in smartphone cameras. Each camera module has a lens, an image sensor, and an actuator for moving the lens to focus the image. Each lens has a fixed aperture of F2.4—that is, the focal length of the lens divided by its diameter is equal to 2.4. (A lens with a lower F number has a larger aperture and so lets in more light.)

Five of these camera modules capture images at what we think of as a 28-mm field of view; that's a wide-angle lens on a standard SLR. These camera modules point straight out. Five other modules provide the equivalent of 70-mm telephoto lenses, and six work as 150-mm

LIGHT

equivalents. These 11 modules point sideways, but each has a mirror in front of the lens, so they, too, take images of objects in front of the camera. A linear actuator attached to each mirror can adjust it slightly to move the center of its field of view.

Each image sensor has a 13-megapixel resolution. When the user takes a picture, depending on the zoom level, the camera normally selects 10 of the 16 modules and simultaneously captures 10 separate images. Proprietary algorithms are then used to combine the 10 views into one high-quality picture with a total resolution of up to 52 megapixels. The image fusion can be done either in the camera or on another computer.

When you press the shutter button to take a 28-mm image, all five 28-mm modules fire simultaneously, recording five images of the same thing but from slightly different perspectives. All five 70-mm modules also record images. Normally, a 70-mm lens can capture approximately a quarter of the scene that a 28-mm lens takes. But our camera adjusts the mirrors in front of four of those lenses so that different modules point at each of the four quadrants of the 28-mm frame we're trying to take, so these four 70-mm images effectively end up covering most of the 28-mm frame.

Because four 70-mm images are recorded, each at a 13-megapixel resolution, the camera captures 52 megapixels of information. The fifth 70-mm module points at the center of the 28-mm frame to ensure the best picture quality at the center of the final image. The software in the camera uses information from the 28-mm modules to precisely stitch together the 70-mm images and then combines all the data into one high-resolution, high-quality picture. It's easy to see how this process lets us capture much more light than if we'd just used one camera module.

To take pictures at 70 mm, we move the mirrors so that the five 70-mm modules now point straight out from the camera; all of them cover approximately the same field of view but from slightly different perspectives. We now enlist the 150-mm modules as well, adjusting their mirrors so that they capture four images that tile the 70-mm modules' field of view. And once again, we combine the many images digitally to provide a better picture than a single-module camera could possibly take, one that rivals a DSLR image.

We can even use our technology to zoom anywhere in the range of 28 to 150 mm. Traditional zoom lenses change focal length by physically moving the lens elements with respect to one another when you rotate the zoom-control ring. Our modules are too small to have either the space or the mechanical precision to accomplish this synchronously across multiple camera modules. So we took a systems approach to solving the problem, using fixed-focal-length lenses.

Suppose you wanted to capture an image with a 50-mm field of view, smaller than what a 28-mm lens captures but larger than that of a 70-mm lens. We activate all the 28-mm camera modules and crop each of the images to the 50-mm frame. (Cropping is not ideal because we lose some sensor area and light.) We also simultaneously use the 70-mm mod-

Zoom, Zoom

To shoot an image with the equivalent of a 28-mm focal length, the Light camera's 28-mm and 70-mm camera modules get into the act; the resulting 52-megapixel image is a combination of the data. To shoot at 70 mm, the 70-mm and 150-mm camera modules go to work and the 28-mm modules rest. The Light camera doesn't contain any 50-mm modules, so taking a 50-mm image involves the 28-mm and 70-mm modules (the former slightly cropped and the latter slightly overlapped), resulting in a 40-megapixel image.

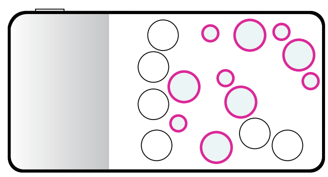


28 mm

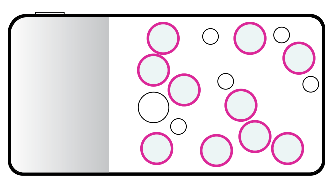


70 mm

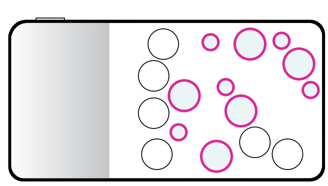
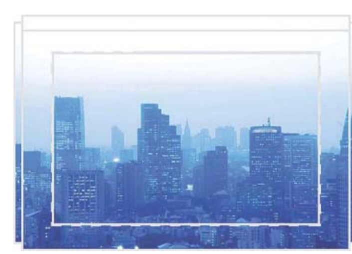

50 mm

PHOTOGRAPHS: LIGHT; ILLUSTRATIONS: ERIK VRIELINK ules. But before we do that, we move the mirrors on four of the 70-mm modules so the captured 70-mm images overlap enough to cover only the 50-mm frame. This way, we retain all of the light collected by the 70-mm modules.

There are other advantages to having multiple lenses. Because the different lenses are set slightly apart from one another, just as your eyes are, the Light camera can obtain images from multiple viewpoints and use them to generate a depth map of the scene. That's valuable because it allows the software to produce any desired depth of field by appropriately blurring those portions of the scene that are outside a selected range of depths.

The software can also change what is called the bokeh, which refers to the aesthetic quality of the blur that appears in the out-of-focus parts of an image. Traditional cameras adjust their lens apertures by opening and closing an iris of sorts made of plastic leaves, overlapped to try to mimic a circular opening. As a result, small, bright, out-of-focus objects appear as regular polygons or circular disks. This is the effect most people are used to. Many photographers consider the ideal bokeh as having a very gentle roll-off, with no sharp edges defining the circle—a Gaussian blur.

I want to use this technology to build a camera with a 600-mm lens equivalent

Photographers will pay a lot of money for a lens with the right bokeh. In our design, the camera uses software to add blur with the right bokeh to those parts of the scene that are outside the selected depth of field. This approach means that users can get whatever bokeh they want. They can choose the conventional disk-shaped bokeh or one with a Gaussian blur. Or they can get creative—for instance, picking a star-shaped bokeh for use in holiday photos, making small decorative lights appear as stars.

The Light camera also naturally allows for an increased dynamic range. That's because the modules don't all have to use the same exposure. We can deliberately overexpose pictures in some modules so that they can image the dark areas with less noise, while underexposing others to capture the highlights perfectly. Because the camera records redundant images, it still has all the information to reconstruct the final picture, but with a much larger dynamic range. Apple iPhone cameras can do this today in their HDR (high dynamic range) mode, but they achieve that by taking a sequence of pictures in time, which can cause motion artifacts.

Recording all these images and then properly combining them into a composite image takes a lot of processing power. We're using the Qualcomm Snapdragon 820 processor, which today is the state of the art in mobile processing. We also have a custom integrated circuit that enables us to interface 16 cameras with the Snapdragon processor. This allows near real-time processing of images in the camera, which results in a resolution of about three megapixels; that's good enough for sharing on social media.

We expect that most of our users will render full-resolution images on their computers, though. That will be faster and won't eat up the battery of the mobile unit. In the next generation of our camera, we will build in hardware acceleration of our processing algorithms, which will then enable the camera to process full-resolution images and manipulate the depth of field without unduly taxing the battery or the user's patience.

The rest of our camera hardware is a standard Android package, which can run Android apps; it's about the same length and width as a smartphone, but at 21 mm it's about two to three times as thick. This approach lets us easily write a user-friendly camera app to simplify the interface and make it less intimidating to use the camera's advanced features.

The Light camera can run in either auto or manual mode. The auto mode will allow you to tell the camera what you want to do—take a portrait, say, in which case the camera will activate the flash if it's dark, select a shallow depth of field to highlight the subject, and adjust the exposure. The manual mode will offer more controls, but we're not going to ask users to set 10 different exposures on 10 different camera modules. We will likely allow them to set everything that they could adjust on a DSLR—flash, aperture, shutter speed, ISO, and exposure—and let the camera's software figure out the details.

Our first-generation L16 camera will start reaching consumers early next year, for an initial retail price of $1,699. Meanwhile, we have started thinking about future versions. For example, we can improve the low-light performance. Because we are capturing so many redundant images, we don't need to have every one in color. With the standard sensors we are using, every pixel has a filter in front of it to select red, green, or blue light. But without such a filter we can collect three times as much light, because we don't filter two-thirds of the light out. So we'd like to mix in camera modules that don't have the filters, and we're now working with On Semiconductor, our sensor manufacturer, to produce such image sensors.

But beyond tweaking our current technology to improve performance, I have a more ambitious goal. I want to use this technology to build a camera with a 600-mm lens equivalent in something the size of a tablet computer, perhaps a little thicker. Today, a 600-mm lens is bigger than a rolling pin, weighs more than 4 kilograms, and costs upwards of $12,000. Only professional wildlife or sports photographers would ever buy one. But if consumers could afford a 600-mm camera, travel with it easily, and take high-quality pictures with it, that would be incredibly cool.

That's my next challenge. For now, however, I'm looking forward to seeing Light technology migrate into cellphones. Then the cameras we all carry with us everywhere will be as good as the ones we leave sitting at home. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/lightcamera1116

CONTINUED FROM PAGE 33 | As a result, ASML has been researching ways to make a pellicle that can withstand the harsh onslaught of EUV radiation. The pellicle has to be as transparent as possible, so that little additional light is lost before it gets to the wafer. And it's doubly difficult in this case: Because EUV masks are reflective rather than transmissive, light must go through a pellicle twice–once on the way in and once on the way out.

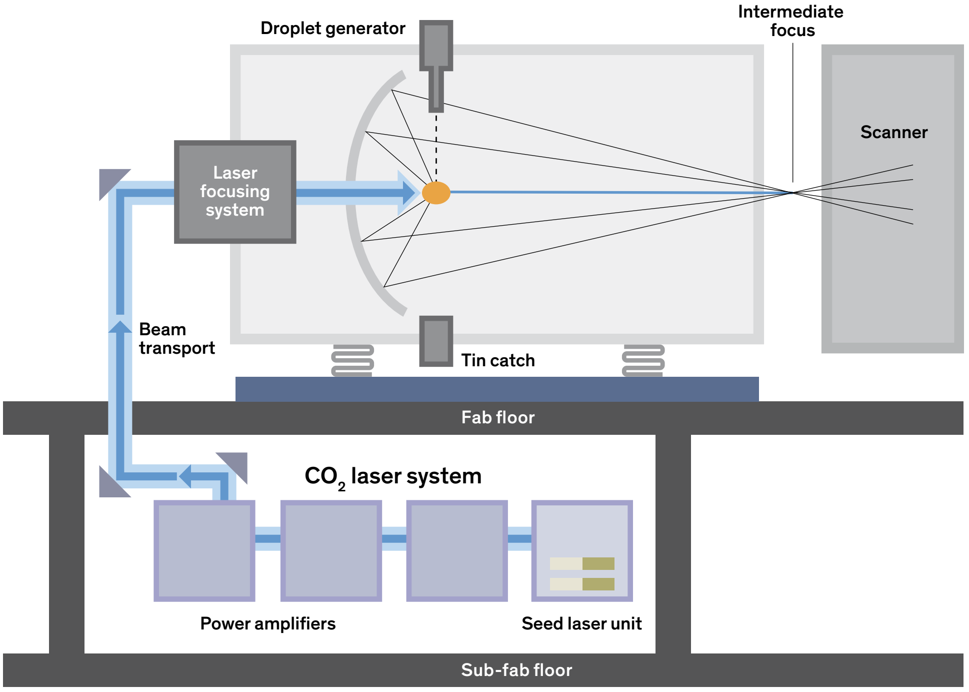

TIN POWER: To generate EUV, ASML's light source hits a rapid-fire stream of molten tin droplets with laser pulses. The process begins beneath the fab floor, in a system that generates two sets of laser pulses. Each tin droplet is flattened by a "prepulse" and then hit with a "main pulse," which heats the droplet to create a plasma. A collector mirror directs the resulting radiation into the scanner.

SOURCE: ASML; ILLUSTRATION: MARK MONTGOMERY

There are still other challenges with EUV that its potential customers want to see resolved before they commit to the technology. One is the ability to make EUV masks free of defects–and develop efficient ways to confirm that they are in fact defect free. Another piece of the puzzle is the photoresist–the light-sensitive layer that is coated on a wafer to take a mask's pattern.

Today's resists, known as chemically amplified resists, are made from polymeric chains of molecules that multiply the effect of incoming photons. But these materials are not especially efficient absorbers of EUV light, explains Andrew Grenville, CEO of the EUV-resist startup Inpria. What's more, the recorded image tends to blur slightly as the light-triggered amplification reaction diffuses through the material. To make finer features than we have today, "you really need much smaller and more robust building blocks," Grenville says. Inpria is working on a resist made of smaller tin-oxide components that are designed to absorb EUV five times as readily and create patterns without amplification.

Will it all be ready in time to keep Moore's Law from dissolving, either permanently or temporarily? Lithography expert Chris Mack doubts that all the pieces will come together for chipmakers by 2018. Planning for a new chip-manufacturing generation happens years in advance. Making a commitment to use EUV in just a few years' time, he says, "is too risky."

Mack, a notorious EUV skeptic who once wagered his Lotus Elise sports car against the technology, does, however, admit to a "glimmer of hope" for it. Chipmakers are struggling to keep miniaturization on track and costs under control; the timing between successive chip generations seems to be getting longer, and chip features are not shrinking as aggressively as they used to. These struggles could give EUV an opening, he says: "There's a real possibility that this slowdown in Moore's Law could enable EUV to have enough time to catch up."

Enough time, that is, before costs cause Moore's Law to grind to a halt. EUV may well reach a point where it is ready for adoption and will reduce manufacturing costs, Mack says. But by that point, he says, the overall cost of the next generation of advanced chips may prove too much–and the performance benefits not impressive enough–for chipmakers to pursue it. Already, Mack says, older chip-manufacturing generations are staying in the mix longer: "I think we're going to see the marketplace splitting up and lots of companies doing lots of different things."

As it has in the past, the fate of Moore's Law will hinge not only on how finely we can print features but on how well physicists and engineers can keep improving the resulting transistors and circuits. Even a rapid-fire sequence of tin plasma flashes won't be able to shed any light on exactly when the world's greatest technological winning streak will finally end. But it just might light the path ahead. ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/lithographyfuture1116

CAN SOLAR SAVE THE GRID?

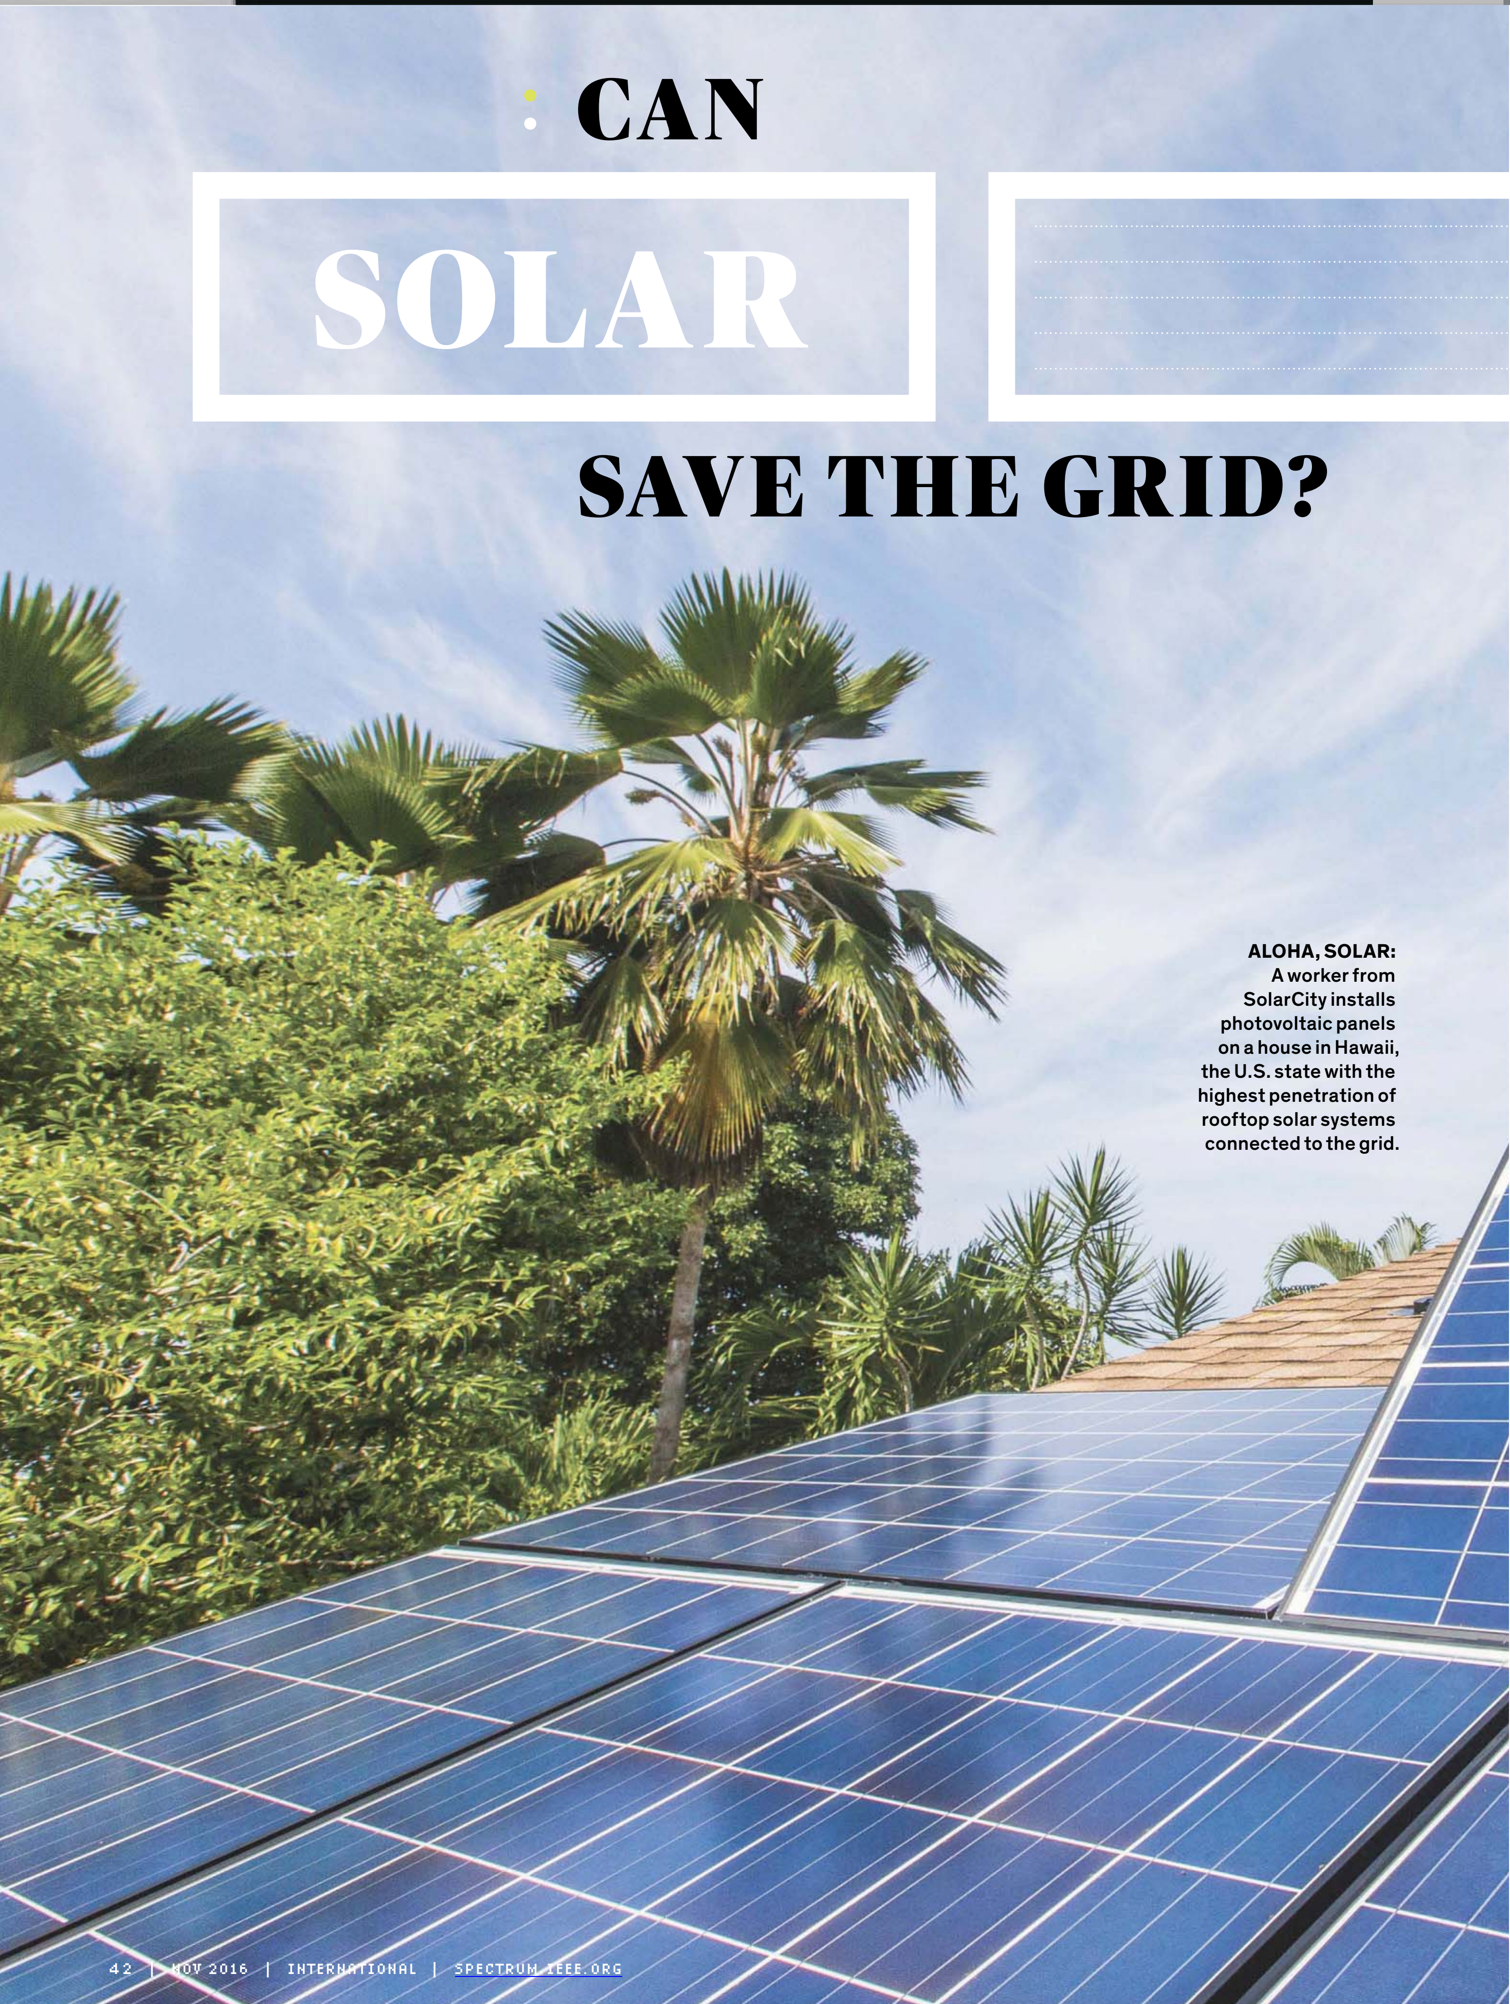

ALOHA, SOLAR: A worker from SolarCity installs photovoltaic panels on a house in Hawaii, the U.S. state with the highest penetration of rooftop solar systems connected to the grid.

USING A NEW TECHNIQUE CALLED VIRTUAL OSCILLATOR CONTROL, SOLAR PV SYSTEMS CAN HELP STABILIZE THE POWER GRID

By Benjamin Kroposki

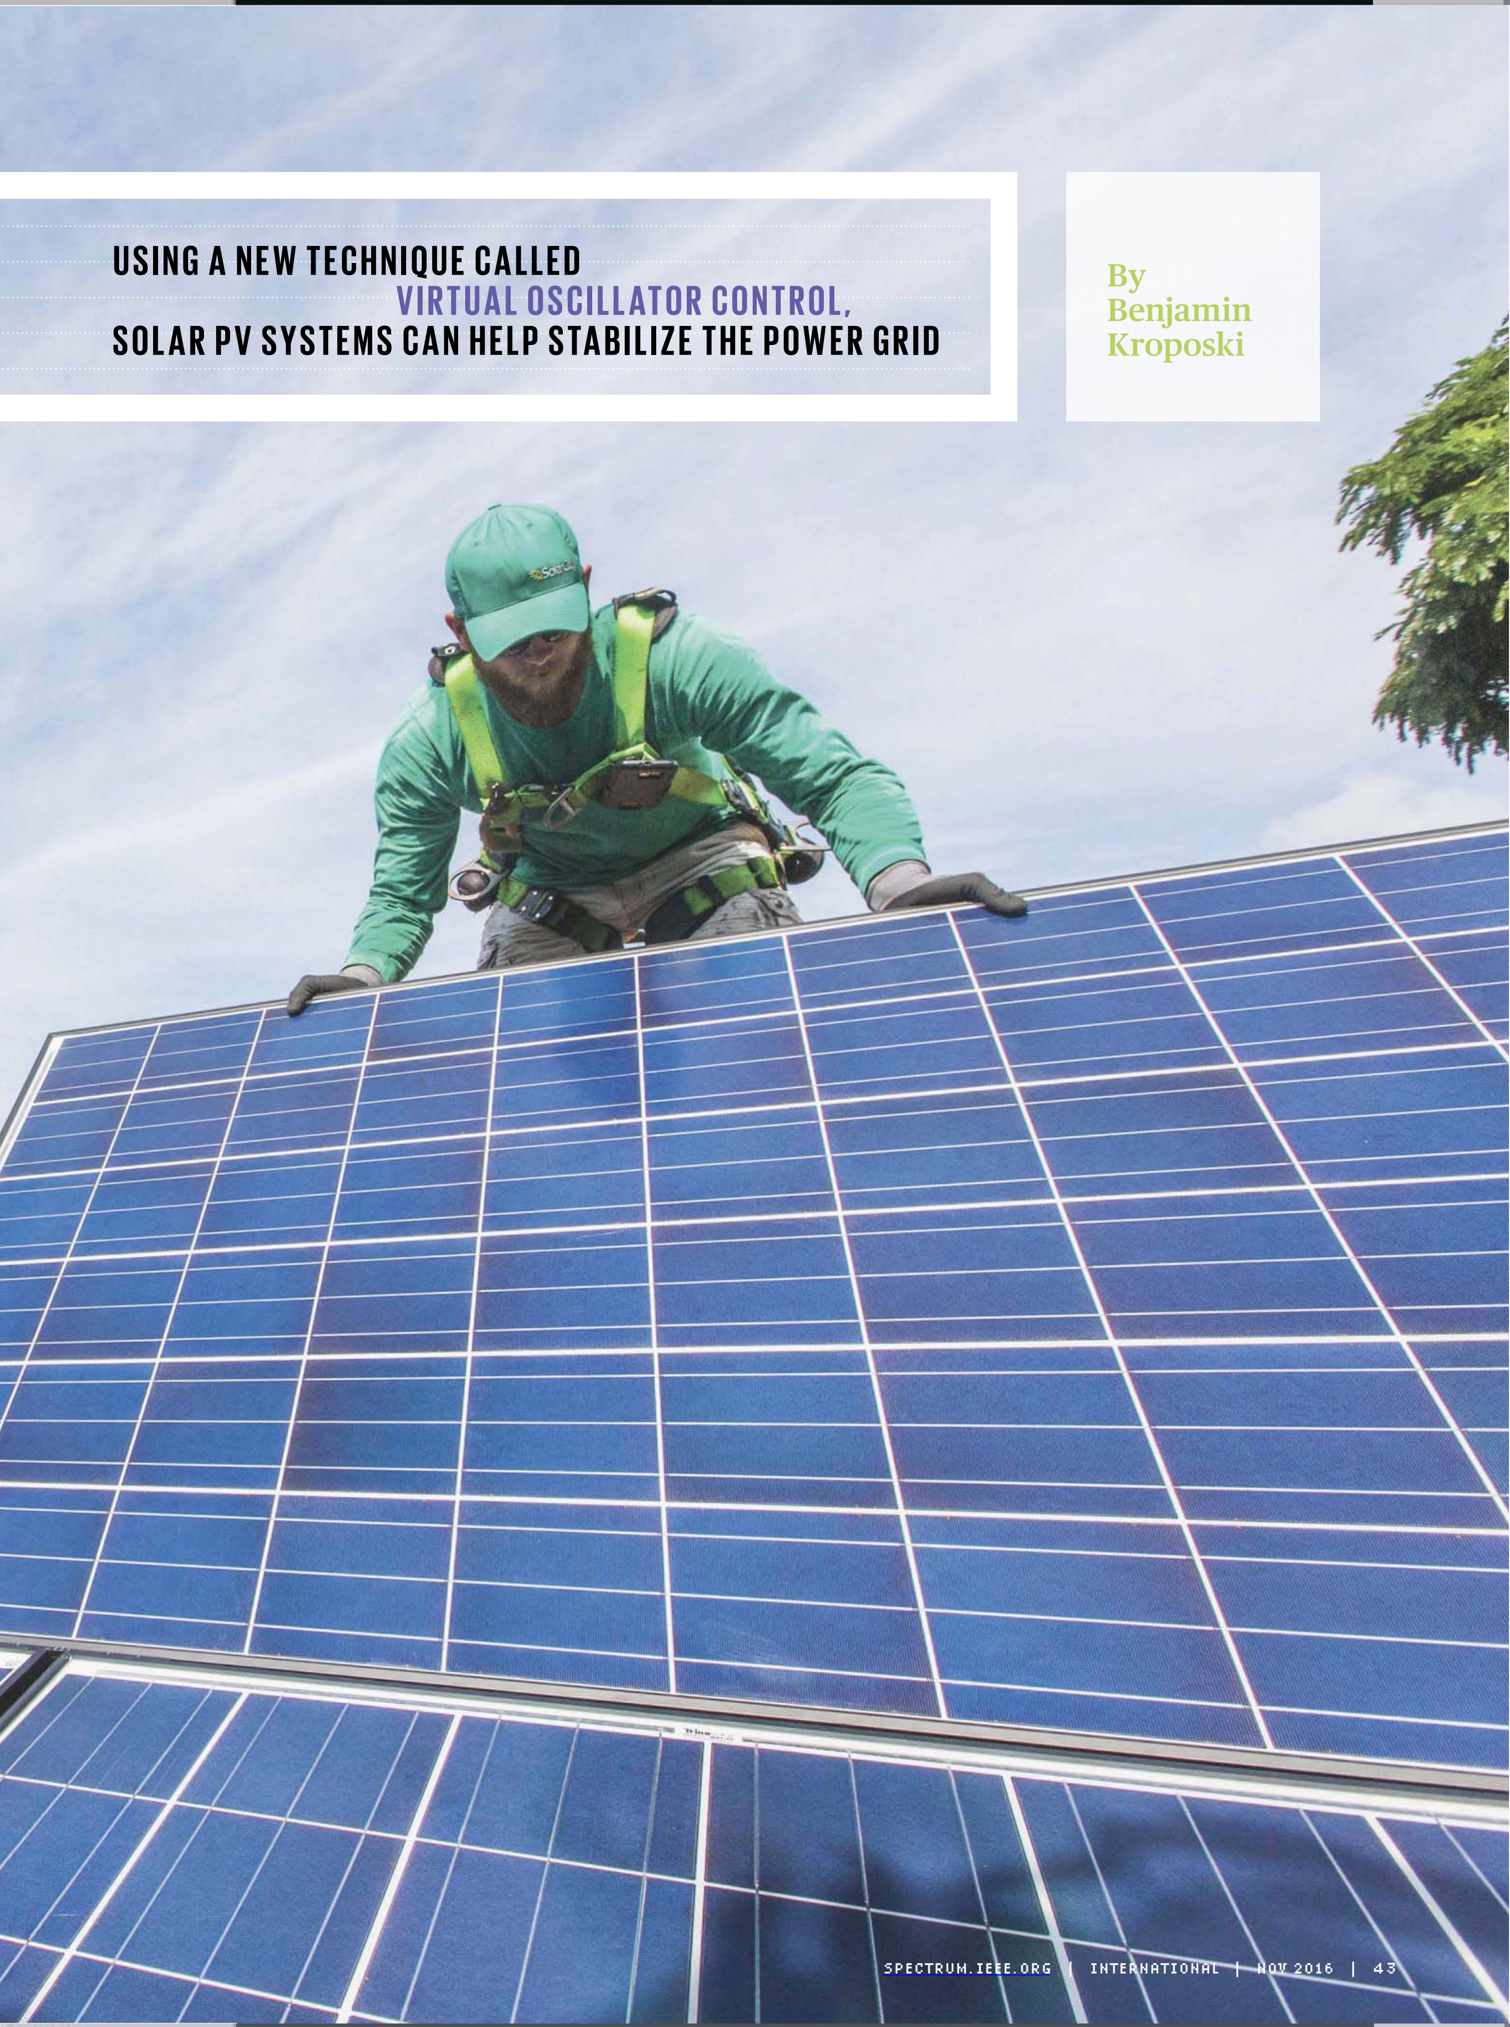

When Steve Johnson had solar panels installed on the roof of his Boulder, Colo., home several years ago, he considered it his personal contribution to making the world a little cleaner. And if his grid-connected photovoltaic (PV) system would occasionally cause his electricity meter to spin backward during the day and trim his utility bills, so much the better.

It was only later, during a tour of the National Renewable Energy Laboratory (NREL), in nearby Golden, that he learned about solar's potential downside: Most PV systems are set up to disconnect from the grid whenever they detect a significant fault. If a single home's PV system trips off-line, it's only a headache for the owner. But if hundreds or thousands of them do so simultaneously, it could upset the network's delicate balance, turning an otherwise small disturbance into an outage blacking out an entire city or county.

Throughout much of the developed world, electric utilities are facing an unprecedented challenge. Growing numbers of customers are installing solar PV systems on their homes or businesses. In the United States alone, the installation of PV systems has seen a compound annual growth rate of nearly 60 percent since 2010, resulting in an installed capacity of 32 gigawatts. The power they're injecting into distribution lines is causing voltage- and frequency-control problems that threaten to destabilize the grid. While this is not yet a major problem, it could become one as distributed solar systems proliferate.

The cause of the problem is the inverter, an electronic system that converts the direct current (DC) supplied by the PV panels into the alternating current (AC) that flows on the power grid. The vast majority of inverters sold to homes and businesses today, including the one on Johnson's house in Boulder, are "dumb" inverters. Although they supply AC at the right voltage and frequency to sync with the distribution grid, they are otherwise passive. They can't sense what is happening on the grid and adjust themselves accordingly.

But newer "smart" inverters can prevent a PV system from going off-line when it doesn't have to. By doing so, they can actually make the grid *more* stable, by preventing the sudden deterioration of voltage and frequency that would otherwise occur when hundreds or thousands of PV panels are suddenly taken off-line, according to recent research by NREL and its partners. My colleagues at NREL have also been developing an innovative method for controlling inverters that will keep the grid stable even when *all* of the power is coming from solar, wind, and other forms of generation that connect to the grid via an inverter.

Smart inverters are poised to fill a big need in the fast-evolving electric-utility industry. As more and more homeowners put PV panels on their roofs, the power they are supplying is reducing the need for big, centralized generating plants. The upshot is that increasing numbers of these traditional power plants are getting retired, and grid operators are scrambling for ways to keep their networks running with the same high level of reliability that their customers have long taken for granted. The combination of smart inverters and new control methods will be essential to helping utilities transition to the grid of the future, in which vast amounts of wind- and solar-generated electricity will be the norm.

STEADY STATE: A new technique known as virtual oscillator control allows smart solar inverters to sense and adjust to grid disturbances, such as a sudden change in frequency or voltage. As distributed solar grows and big power plants are retired, VOC will help keep the grid stable.

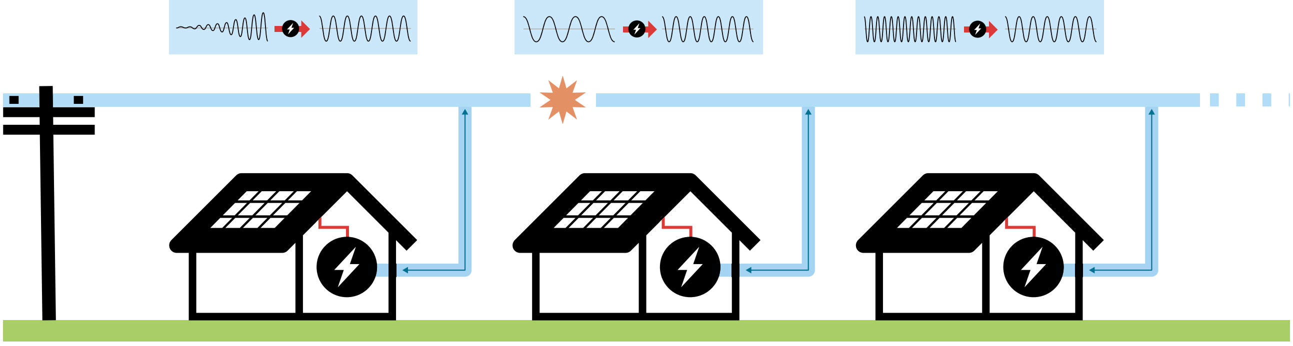

ILLUSTRATION: Mark Montgomery

PREVIOUS PAGES: SOLARCITY

Consider what can happen on a sunny day at noon when a coal-fired plant unexpectedly goes off-line or a transmission line suddenly goes down. The loss might cause the grid voltage to dip by up to 10 percent. To restore the voltage to normal, technicians must bring additional generation on line quickly—through, for instance, backup "spinning reserve" capacity that a utility maintains for just such a situation. Or they could curtail demand, perhaps through an automated mechanism known as demand response, which enlists customers to reduce their electricity consumption when power prices are high or when grid reliability is threatened.

What you *don't* want to happen in a situation like this is to lose even more generators—and yet that is exactly what *does* happen when a host of dumb inverters switches off en masse. Left unchecked, what starts as a minor voltage drop can cascade into a widespread fiasco.

To be sure, there's a good reason for inverters to disconnect when the grid is down: You wouldn't want PV systems to inadvertently energize a downed power line because that line could shock anyone who touched it. But in the first few seconds of a fault, it's critical to maintain PV generation in sections of the grid that have not been damaged.

In contrast to dumb inverters, a smart inverter can "ride through" voltage or frequency dips and other short-term grid disturbances. And if these inverters have communications capabilities, they can let grid operators monitor and control them in response to changing conditions.

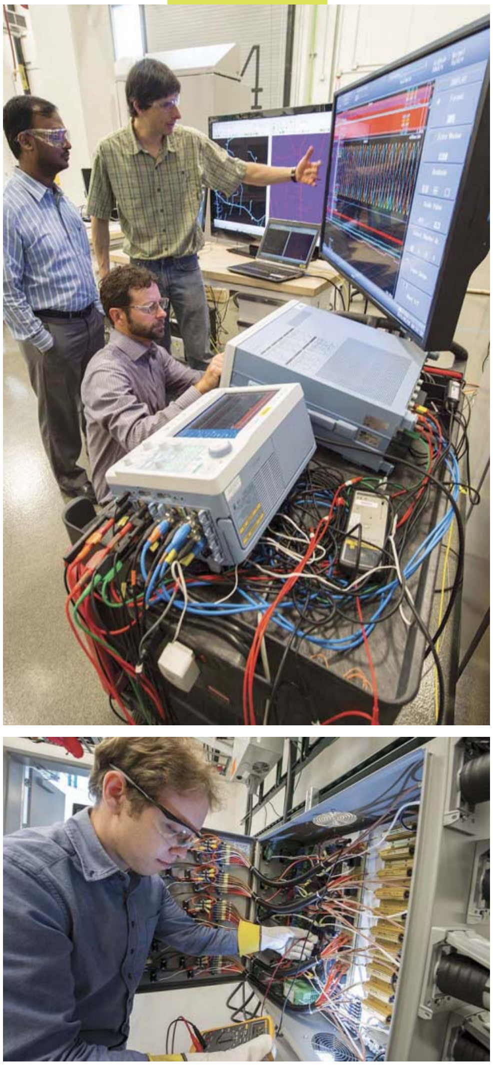

TOMORROW'S GRID: The National Renewable Energy Laboratory's Murali Baggu, Greg Martin, and Bryan Palmintier [top] test a smart PV inverter linked to a power grid simulator. Brian Johnson [bottom] studies a custom-built virtual oscillator control system.

NREL

To understand the effects of smart solar inverters on utility systems, there's no better place to study them than Hawaii. The state has the highest electricity rates in the United States, which, along with incentives from the utility, has given residents and businesses powerful motivation to install their own PV systems. The state now has about 20 times as many solar installations connected to its grid as the average mainland state does, and that trend is likely to continue. Last year the Hawaiian government set a goal to obtain 100 percent of its electricity generation from renewable energy sources by 2045. Some of that will be wind, but a lot of it will be distributed solar.

Even now, the solar penetration on the island of Oahu is high enough that the state's electricity utility, Hawaiian Electric Company (HECO), regularly contends with distributed solar systems feeding power back through distribution lines into substations. Such backfeeding in itself is not a problem. However, if a distribution line is suddenly knocked out of commission, that solar generation gets redirected toward the customers' facilities. Let's say a tree branch grounds a line and causes the line's circuit breaker to open; the backfed power can then flow back toward any homes or businesses up to the point where the line's breaker opened.

This surge of power causes the voltage to spike. If the spike is high enough and lasts long enough, it can damage motors, generators, and distribution equipment. To prevent such overvoltage situations, HECO and the Hawaiian Utility Commission decreed in 2014 that the amount of solar-generated power on any distribution line on Oahu could not exceed 120 percent of that line's minimum daytime load. The unsurprising result was that several thousand HECO customers were prohibited from connecting their PV systems to the grid. Many others who were contemplating installing new rooftop systems decided not to, and quite a few solar installers saw their business decline.

Clearly, a better solution was needed. Through the U.S. Department of

Energy's SunShot Initiative, NREL researchers Sudipta Chakraborty and Anderson Hoke worked with engineers at HECO and SolarCity, the country's largest solar-power provider, to study ways to allow more solar on the Oahu grid. They started by running inverter tests at NREL's Energy Systems Integration Facility to check different scenarios in which an overvoltage occurred. The goal was to see if smart inverters could respond to such overvoltages by quickly reducing the amount of power they were releasing onto the grid.

The researchers ran the tests using five commercially available advanced inverters ranging in capacity from 3 to 12 kilowatts. Among the types of inverters tested were string inverters, which draw DC power from one or more strings of solar panels, and microinverters, which draw from a single panel. They connected each inverter to a grid simulator and a load that simulated the home to which the panels would be supplying power. They then monitored the inverter's response when the grid simulator's voltage spiked. They repeated the tests for a range of ratios of solar generation to load, from an even match on up to 10 times as much power generation as load.

The results were reassuring: The inverters all responded immediately by cutting their output, and so the maximum overvoltage never exceeded twice the normal voltage in any test, with the typical overvoltage significantly lower. In addition, the voltage spikes lasted only microseconds to milliseconds, an acceptably short duration. The tests demonstrated that when smart inverters were installed, voltage spikes posed less of a problem than had been feared.

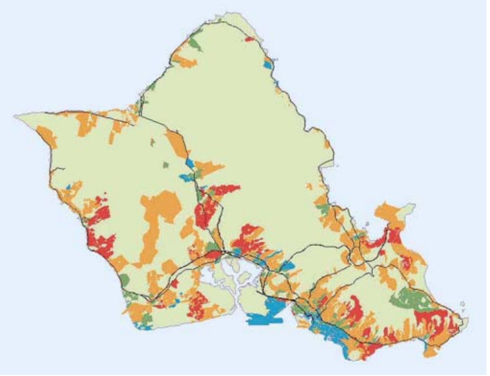

% OF DAYTIME MINIMUM LOAD ON OAHU

- Less than 50%
- 50 to 100%
- 100 to 250%
- 250% and greater

This map shows the relative density of grid-connected PV systems installed on the Hawaiian island of Oahu.

HAWAIIAN ELECTRIC COMPANY

Persuaded by these tests, HECO lifted its moratorium on solar systems on Oahu in early 2015, provided that new customers installed one of the advanced inverters vetted by NREL. The utility is now working with inverter makers to certify other models of inverters, using the NREL test methods. Separately, HECO worked with Enphase Energy, of Petaluma, Calif., to remotely reprogram 800,000 microinverters already installed on the Oahu grid, to allow them to ride through frequency and voltage anomalies and help improve grid stability during such disturbances. And HECO has more than doubled its limit on PV capacity on each distribution line, to 250 percent of the minimum daytime load. The utility's moves may eventually have a national or even global impact, as the IEEE and UL, which set international technical standards for inverters, are now preparing to include overvoltage protections in their smart-inverter standards.

While the current generation of smart inverters can solve some of the problems confronting utilities like HECO, an even larger issue looms. Today's power grids still rely on large, centralized power plants with a massive generator rotating at just the right speed to produce power at a certain frequency—50 cycles per second in most of the world, 60 cycles per second in the United States and parts of Asia. These generators collectively give the entire power grid mechanical inertia: They are very large and rigidly synchronized with each other, so only major disturbances can affect the grid's frequency.

But distributed energy sources, including solar, wind, fuel cells, and batteries, don't have synchronous generators; they supply variable and uncertain amounts of electricity, and the inverters they use to connect to the grid are designed to simply lock onto the grid's frequency and follow it. And as renewable-energy generation spreads, many large central power plants—especially coal—are being retired. Eventually, the grid will lack the inertia it has today to maintain a stable voltage and frequency in the event of a large disturbance.

For this reason, the power grids of the future may need advanced inverters that don't simply follow what the grid is doing but actually help *form* the grid, by responding instantaneously to disturbances and working in concert to help keep the grid stable. This approach, known as virtual oscillator control (VOC), was developed by an NREL team led by Brian Johnson, working with Sairaj Dhople of the University of Minnesota; Francesco Bullo at the University of California, Santa Barbara; and Florian Dörfler at ETH Zurich.

The basic idea behind VOC is to leverage the properties of oscillators—that is, anything that moves in a periodic fashion, such as a mechanical spring, a metronome, or a motor. When two oscillators are paired, or coupled, their motions will naturally align. Picture several weights hanging from springs that are connected to a rigid and fixed board. As the weights are given a pull to set the springs bouncing, the springs will tend to bounce in an uncoordinated fashion. However, if the board itself is suspended from the ceiling on springs, it becomes a feedback mechanism that responds inversely to the pushes and pulls from the bouncing weights. Within a short period of time, the weights will all start bouncing at the same frequency, all in lockstep.

For virtual oscillator control, the trick is to make each inverter respond to the grid much like a spring: If the grid voltage drops, the inverter adjusts its output such that it "pushes" against the voltage change. Likewise, a surge in grid voltage will induce an inverter to "pull" the voltage back to the nominal range. An increase or decrease in grid frequency will similarly cause the inverter to adjust its power output to compensate.

VOC is implemented in the software that controls each inverter. When VOC-enabled inverters are connected to the grid, the grid itself acts like that board hung from springs, coupling the springlike inverters together. Setting up hundreds or thousands of such inverters on the same circuit will cause them to all work together to maintain the system voltage and frequency within an acceptable range.

NREL researchers originally designed VOC to help stabilize relatively small, isolated electricity networks known as microgrids. A microgrid can be connected to the larger utility power grid, or it can disconnect and operate independently. When isolated from the larger grid, the microgrid must generate all of its own electricity. In this scenario it can be challenging to maintain stability because a microgrid, by its very nature, has a much smaller total load. So the sudden addition of a new load, such as when a large central air conditioning unit kicks on, can mean a relatively large jump in the total demand.

To test the VOC approach with a microgrid, NREL researchers used a setup with multiple load types and several types of small test inverters that were about the size that might be installed on a PV panel. The team added the VOC software to each inverter to allow it to constantly evaluate and respond to local conditions on the grid. The VOC inverters were able to act in concert to keep the microgrid stable without needing any deliberate communication among them or a dedicated supervisory control system. The team essentially created a self-regulating microgrid.

This approach provided a much faster response than the existing mechanism, called droop control, for controlling inverters and synchronous machines such as generators. Droop control is based on the idea that the frequency and voltage of a single generator varies, or droops, in proportion to its power output. Droop control is what allows multiple generators to be integrated into a system.

But when droop control is applied to inverter-based systems, it requires a control system that first measures the frequency and voltage and then calculates the appropriate response. These necessary computations result in a sluggish response. VOC, on the other hand, responds immediately to changes in grid conditions, just as a spring responds to any stimuli, with minimal need for measurements or calculations. The inverters don't even need to explicitly communicate with one another.

NREL is now testing whether the VOC approach could work on utility-scale grids, which connect thousands of homes and businesses. For obvious reasons, we can't experiment on a real-world grid, so our test will use a large collection of residential-scale inverters and loads connected to a simulated grid. The inverters and loads will interact realistically with this simulated grid, allowing a range of scenarios to determine if the grid voltage and frequency stays within the desired range under normal operating conditions as well as disturbances.

The VOC inverters being used in the NREL tests have been custom-built and programmed, but it should be straightforward to standardize the VOC software so that it can be used in off-the-shelf inverters. In our tests, we've also shown that VOC inverters are compatible with the droop control methods currently used in large power grids.

Further into the future, when the world's power grids move away from synchronous, rotating machinery and toward distributed generation controlled by power electronics, VOC may prove one of the best ways to maintain a stable grid. That transition is beginning to look inevitable, and now is the time to start preparing for it. ■

SOLARCITY

SMARTER SOLAR: Inverters convert the direct current from PV panels to the alternating current used by the power grid. Traditional inverters are designed to shut down when they detect a fault, but newer inverters, like these, can be remotely updated to ride through such disturbances.

POST YOUR COMMENTS at http://spectrum.ieee.org/smartsolarinverter1116

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

RUN SILENT,

Ohio State engineers optimized their bike for one thing: to climb Pikes Peak faster than any other bike, electric or gasoline

By **PHILIP E. ROSS**

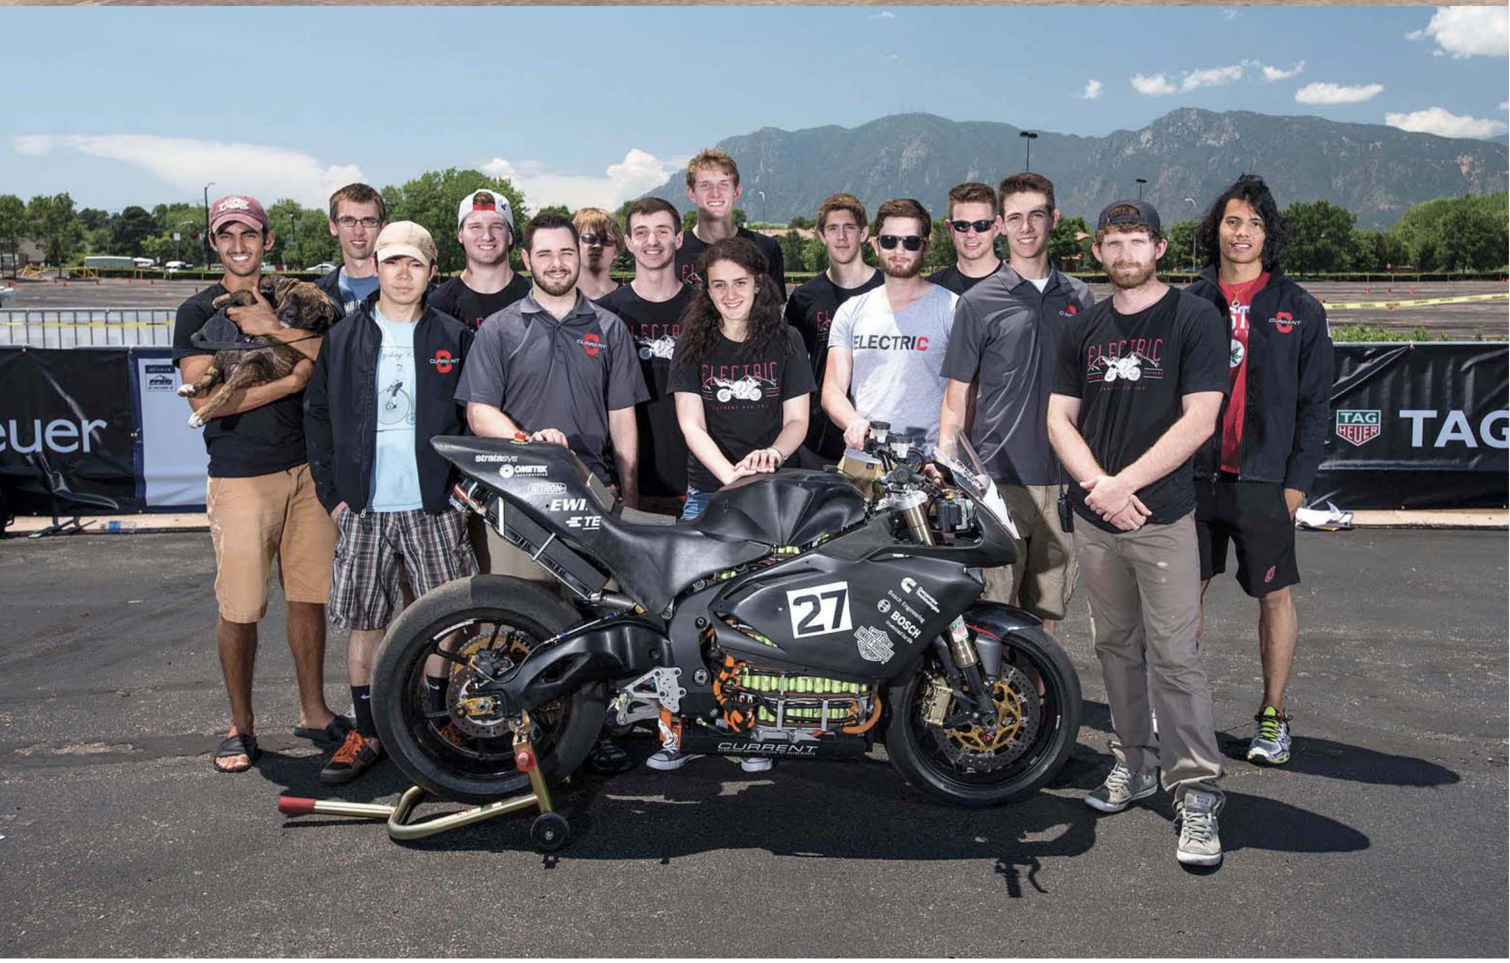

RUN STEEP

NOWHERE TO GO BUT UP: Pro rider Rob "The Bullet" Barber [here and opposite, top] makes runs on the grueling Pikes Peak racecourse. The 2016 Ohio State Buckeye Current motorcycle team was led by Aaron Bonnell-Kangas [opposite, in light gray T-shirt and sunglasses]. Next year's team will be led by Polina Brodsky [opposite, center], who joined years ago as a high school student. The Buckeye Current RW-3 [above] has a 7.7-kilowatt-hour battery pack, which includes 938 cells.

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

AN HOUR BEFORE THE QUALIFYING ROUND IS NO TIME TO BLOW OUT THE POWER INVERTER IN YOUR ELECTRIC RACING MOTORCYCLE.

It's a Friday morning in June and easily the worst day of the year for a score of engineering students from Ohio State University, in Columbus. Months of painstaking labor have gone into their all-electric cycle, the Buckeye Current RW-3, to handle one race and one race only: the Pikes Peak International Hill Climb, a motorsports race up Colorado's iconic mountain. It towers in the near distance; the students are at rock bottom.

In the past week, they've overcome adversities that include a burned-out noise maker (to alert pedestrians to the cycle's silent approach), a wobbly seat, a cooling-system snafu, and a run-in or three with the powers that be. And now, a dead power inverter.

And yet, if panic is breaking out, there's no sign of it. "This, too, is fun," insists team leader Aaron Bonnell-Kangas, speaking somewhat less animatedly than he usually does.

Punching numbers on his smartphone, he tries to raise somebody–anybody–at Tritium, maker of the inverter, which turns a battery's direct current into the alternating kind that the motor requires. Alas, it's the wee hours of a Saturday morning at Tritium's headquarters, outside Brisbane, Australia.

So everyone jumps. Off comes the bike's seat, out come its electronic guts, on go the probes of the circuit testers. But the inverter's troubles remain obscure. This isn't a mechanical quirk in your father's Ducati; this is a bug hiding somewhere among scores of chips and thousands of lines of code.

"It's talking to us, but it's not sending a signal to the motor," says Polina Brodsky, a mechanical engineering student. "The computer is booting up, but it's not doing anything."

With just minutes left before the qualifying round is set to begin, Rob "The Bullet" Barber, the team's pro driver, improvises a plan B. He hops onto a motorcycle belonging to one of the students–a gasoline-powered Kawasaki–and makes for the starting line. "Hope they [the race officials] count it, but it isn't the team's bike," he frets.

The officials do count the bike. So now, all the team has to do is repair the inverter on the RW-3, run tests on a local track, charge each of the bike's 938 A123 lithium-ion cells to capacity, roll the bike to the base of the mountain, and run the race. They've got 40 hours.

STUDENT PROJECT THOUGH IT IS, the Buckeye Current is regarded as a serious contender. The team has gone far to optimize the bike to put out crazy levels of power, something it has to do for only the 10 minutes it should take to round the course's 156 hairpin turns while skirting a skid into an abyss. Power is everything; energy storage, nothing. If the bike's batteries die at the summit, the students will happily walk it down.

Sheer performance is no mere ornamental attribute but rather the heart of the surging appeal of the electric vehicle–the feature that shattered the golf-cart image. From zero to fabulous in 3 seconds, promises the dual-motor Ludicrous Speed mode in the Tesla Model S P90D. And, like the Buckeye Current, the Tesla does it all without a gearbox.

"This motorcycle is our idea of the perfect bike: powerful and nimble," says Bonnell-Kangas, a grad student in electrical engineering. "We want to go up Pikes Peak ahead of the gasoline bike. We want to be the best."

The race plays to the strengths of an electric drivetrain. First, an electric power plant offers instant-on torque, and that comes in handy on the course's many switchbacks. Each one of them becomes a sort of mini drag race. Second, the EV will be indifferent to the thinning of the air as it makes its circuitous ascent–a rise of 1.4 kilometers (0.9 mile) over the 20-km course.

Why not use a compressor to load–or "supercharge"–the combustion chambers? "It's hard to tune a supercharger because the air is so different at the base and the summit," Brodsky explains.

This year's iteration of the bike, the RW-3, uses the same frame as last year's–a Honda sport model but minus the gearbox, engine, and fuel

PIECES OF A DREAM: Bottom, from left: With the bike's tail fairing removed, the silvery inverter housing (1) is visible, and with the inverter removed, the cooling-fluid lines (2) are exposed. The blue motor-cooling reservoir (3) perches above the black cooling pump, which is held by an orange bracket (4) above the bike's rear tire. A blown fuse was traced to a board (5) in the inverter. The Manzanita Micro battery charger (6) can provide up to 20 amperes.

PREVIOUS PAGES, TOP LEFT AND BOTTOM RIGHT: STEPHEN SULLIVAN/RANDELS MEDIA GROUP; TOP RIGHT AND BOTTOM LEFT: COREY DAVIS/RANDELS MEDIA GROUP


1


2


3

PHOTOGRAPHS BY Philip E. Ross

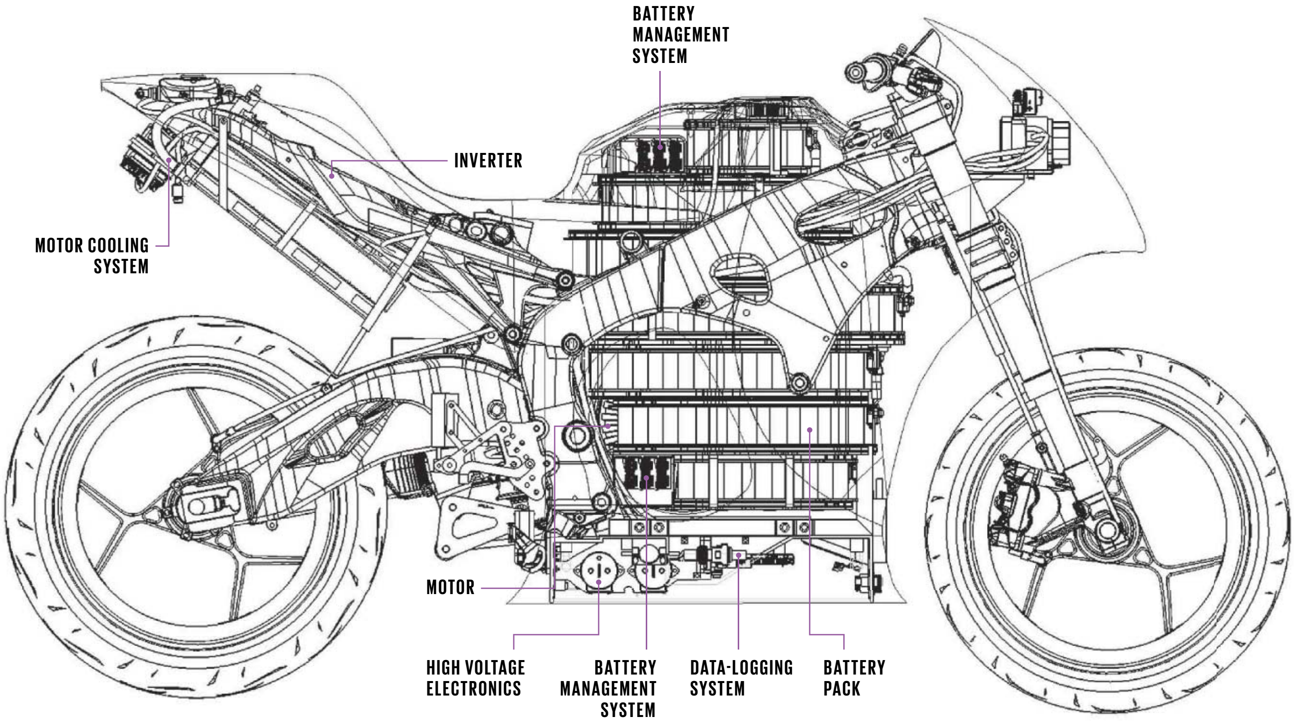

FAST FORWARD: The Buckeye Current RW-3 gets its oomph from an Enstroj Emrax motor, which is rated at 100 kilowatts (134 horsepower). The motor is powered by 938 A123 lithium-ion cells, each about the size of a D battery. They weigh 80 kilograms and can store 7.7 kilowatt-hours—a quarter the capacity of a Nissan Leaf.

tank. In their place are the green mass of thumb-size batteries that cluster under the seat, all the way to the bottom. They weigh 80 kilograms (176 pounds), hold 7.7 kilowatt-hours—a quarter the capacity of the Nissan Leaf and about a tenth that of the Tesla Model S—and feed current through the Tritium inverter to an Enstroj Emrax motor, a 40-kilogram whirlwind rated at 100 kilowatts (134 horsepower). And the current comes with a lot of waste heat.

"There are two cooling loops for the motor and the inverter, each with its own radiator and pump—each part optimizes at different temperatures," Bonnell-Kangas says. "A gas bike's more efficient to cool because when you're hotter, it's easier to purge heat out," typically through metal fins surrounding the cylinder chambers. "That's why we need water cooling."

The heat problems had actually started a couple of days earlier. After a time trial, Barber told the team that the engine had gotten a little hot. But the students couldn't get the temperature data because the sensor, a thermistor, had been out of commission for some time. "It was down on our list of bugs, and other things were more important," says Bonnell-Kangas, ruefully.

And putting sensors everywhere can increase the burden, adding more things that can go wrong. There are pressure gauges in the tires, flow gauges in the coolant pipes, voltmeters everywhere, all sending data to a central processor between test runs. That way, the team can fix glitches on the fly, a skill that comes in handy in a competition.

"We put a lot of effort into designing these electronic systems" to analyze data, says Sean Harrington, a team leader and an electrical engineering major. "We design with debugging in mind."

He who errs last, loses. Here, in design methodology, trial and error is as important as theory, if not more so—a lesson you might not find in the standard engineering curriculum.

TOP: BUCKEYE CURRENT RACING


4


5

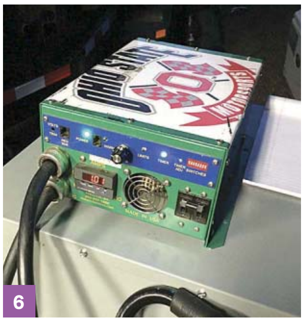
6

RACE TO THE CLOUDS: The course has many treacherous twists and hairpin turns [in purple] and had unpaved sections until 2012. Nevertheless, there have been only five fatalities since the race began in 1916. For the first time ever, an electric car won the race last year, when another electric vehicle took second place as well.

PPIHC

HOW MUCH EFFORT does this extracurricular activity consume? In some cases, more time than the coursework.

Aaronn Sergent, a mechanical engineering major, confesses to me that he puts in just 20 or 30 hours a week on the bike project. He's almost apologetic as he tells me this outside Uncle Sam's Pancake House, where the team has just breakfasted on eggs and pancakes. It's 8:00 a.m, but the team started this day at 3:00 in the morning, doing time trials and tune-ups, well before the onslaught of summer tourists. Now, bellies full, these kids are about to go right into a second day of work back at their rented house in Colorado Springs.

Or even a third day, maybe. Bonnell-Kangas gets very little shut-eye: He texts me on the coming day's plans well after I've fallen asleep, then does it again at 2:00 a.m., well before I wake up the next morning. Many hours later, as I stagger off to catch a ride to my hotel, I pass four of the students. It's early afternoon, and they're blowing off steam shooting hoops in the front yard.

Rest? Who needs rest?

The inverter problem defies an easy solution. The Buckeyes can't buy an inverter locally, and they haven't the time to airlift one in from the antipodes. So there the electrical subteam sits, hunched around the dining-room table of their rental house, tearing the broken inverter apart.

"The motor is synchronous, which means the inverter has to keep up," Bonnell-Kangas tells me.

The rotor of a synchronous motor moves in lockstep with the rotating magnetic field. That field, in turn, is generated by output from the inverter, which varies in a repeating pattern, or cycle. The faster the motor turns, the more cycles the inverter must provide. But this inverter can crank out no more than 500 cycles per second—enough to turn the motor at 3,000 rpm. That's the redline.

"So we designed a control system to keep the motor below that critical speed," Bonnell-Kangas says.

This is why the inverter's paralysis is such a big deal. As things now stand, the students can't even get the motor to turn at all. And if they can't fix it, then it's good-bye to a year's worth of dreams.

Just as I begin to head back to my hotel, the problem is apparently found: two very well-fried fuses in the power supply board. They got electrocuted earlier during a tune-up.

"We reversed the polarity [the direction of current]," Bonnell-Kangas tells me. "That burned the fuses, and that interrupted the power supply. We can fix this." And they did.

It's things like these that make you appreciate the sweat that goes into anything complicated: a motorcycle, a car, a jetliner. Fix one thing and you tee up another to fail.

Despite the frustrations, pressures, and compromises, they all can't imagine doing anything else with their free time. When I ask Marc Ahlborg, an aerospace engineering major, why he chose this team rather than something more aeronautical (surely there's a drone club?) he notes that he rides a bike himself, as do most team members. Though limited, the resources of this club are among the best at Ohio State, he adds. Most come as in-kind gifts from sponsors.

Take the 100-kW alternator that Cummins Generator Technologies provided to them gratis. The team had wanted a smaller one, but apparently Cummins doesn't do small.

But the food bill comes out of the students' pockets. So does the rent for the house, which they split evenly, and the transportation. Even bringing their own motorcycles out West is their own burden to bear: Ohio State has a rule against using school vehicles for private purposes. So while Buckeye Current and its appurtenances travel in a trailer worthy of a rock star, the personal gear mostly goes on the back of a pickup.

Even petty cash is scarce. When Brody Ringler, a mechanical engineer, argues that the team needs a $10 infrared thermometer gun to quickly check the tires and other parts, Bonnell-Kangas mulls it over for a good 3 seconds: "Okay, buy it."

On Friday afternoon, when the students figure out the inverter's problem and replace the blown fuses, I ask Bonnell-Kangas what he'd have done that morning if he'd gotten through to the Aussie company. Would he have maxed out his credit card to pay the AU $6,000 (US $4,592) for a new Tritium WaveSculptor200 inverter, plus the whatever-it-would-cost to fly it the 14,000 km from Australia to Colorado Springs?

"I'd have paid it," he says.

RACE DAY STARTS at the wretched hour of 1:30 a.m. Saying little, we drive in darkness, the mountain a huge black hulk in the distance.

Poking along the final stretch, we jostle along with vehicles from every category of racing–production cars, vintage cars, modified cars, dune buggies, trucks (even a semi), and of course, motorcycles of various weights, power sources, and designs. It's a surreal outdoor museum of automotive oddities, "Jay Leno's Garage" taking a run on the wild side. There are even low-riding buglike affairs with a sidecar, in which a passenger, called the monkey, kneels in order to hang off the side as a counterweight in a tight turn. (It's a living.)

The team pitches a tent in the pit area, turns on the generator, powers up the fluorescent lights, puts the bike up on stilts, applies the blue Chicken Hawk electric tire warmers. They matter because the tires are racing slicks, so-called for their lack of treads, and they provide traction only when warm. They cool down fast in this chill air and on the frigid surfaces; all around, the mountain walls are pocked with snow.

While warming the tires, the team checks the bike's major systems and tops off the charge in each of the myriad battery cells. You have to do that just so, filling up one battery to a certain point, then catching up the others, then doing it again. It's the only way you can be sure not to overcharge any of them.

Well, that's not quite true: You could design an automatic charging system, as the Victory Motorcycles team has no doubt done. They're not far away, in the stretch that leads toward the burrito shack, and they are an actual company, based in Iowa. All they do is make motorcycles.

Hand charging is tedious. But then what would you do early on race day? The last few hours before the race seem to crawl, like in the old movies when the defense attorney and the defendant wait–and wait–as the jury deliberates. But for Barber, the guy whose life depends on sticking to that slick road surface, time is passing at a faster clip.

"It's wet up there, really wet," he says, returning from a little reconnaissance.

In another hour the sun would burn away that water, but the show must go on, and the electric motorcycles are the first category to climb the mountain. Barber thinks that sticking to the schedule is ridiculous: Surely, the famously waterproof gasoline-powered bikes should go first, he grouses. Surely, someone could go up there and sweep away the worst of the wetness. But there's no time.

After some hurry-up-and-wait delays, Barber gets the go-ahead. He turns on the clanging noise maker, scoots off to the starting line, and promptly enters radio silence as the pit team unexpectedly loses contact with the bike's transponder. Now the team must follow his progress fitfully, through a handful of visual sightings.

After an eternity, the news comes succinctly. "Our time is 11:16," says Brodsky, her ear pressed to a phone.

No 10-minute race record this time, or victory of any kind. Ringler is plainly crushed. Brodsky says nothing. Bonnell-Kangas conceals his disappointment. This was his last hurrah as team leader; now he'll take up his first job, a research position at the Detroit branch of Robert Bosch, the German auto supplier.

Hours later, after all the races are run and Barber can finally walk the bike down, he is ebullient. Or maybe he's just glad to be alive.

"You guys were great," he says. "It's as good as last year's time, and on a wet road. And the bike was great."

He's a pro. And so, literally, were the guys who won the category.

"The Victory bike was just so powerful," Barber says later, while sitting on a couch back at the rented house. He was referring to an electric version of a bike made by Victory and ridden by Dan Canet. It posted a 10:17.8 time, about a minute faster than the Buckeye Current.

"I don't know how they do it; Victory is very quiet about the design," says Barber. When he isn't racing, he works a day job in IT for Cisco Systems, in Manchester, England.

DAVID DID NOT BEAT GOLIATH, not this year. But maybe next year the team will build an automatic battery-charging system. Maybe they'll revamp the power management system, inverter included. Or they could follow the lead of vehicle manufacturers worldwide and pare weight from the bike.

"I can't say what we might do," says Brodsky, diplomatically. "We'll decide these things together, after looking at the data logs from today."

Brodsky, alone among the students, was present at the creation. That was six years ago, when as a senior in high school she wangled her way into a new and decidedly bare-bones operation. Now she has graduated. Next year she'll stay on as a grad student–on a mission.

"I'm the team leader next year," she says. "The rest of the team voted me in." ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/pikespeakrace1116

SOCIAL MEDIA'S

Bulletin-board systems built by hobbyists taught

PIONEERS OF CYBERSPACE: Welcome screens from various computer bulletin-board systems show their operators' wild creativity.

DIAL-UP ROOTS

people how to interact online • By KEVIN DRISCOLL

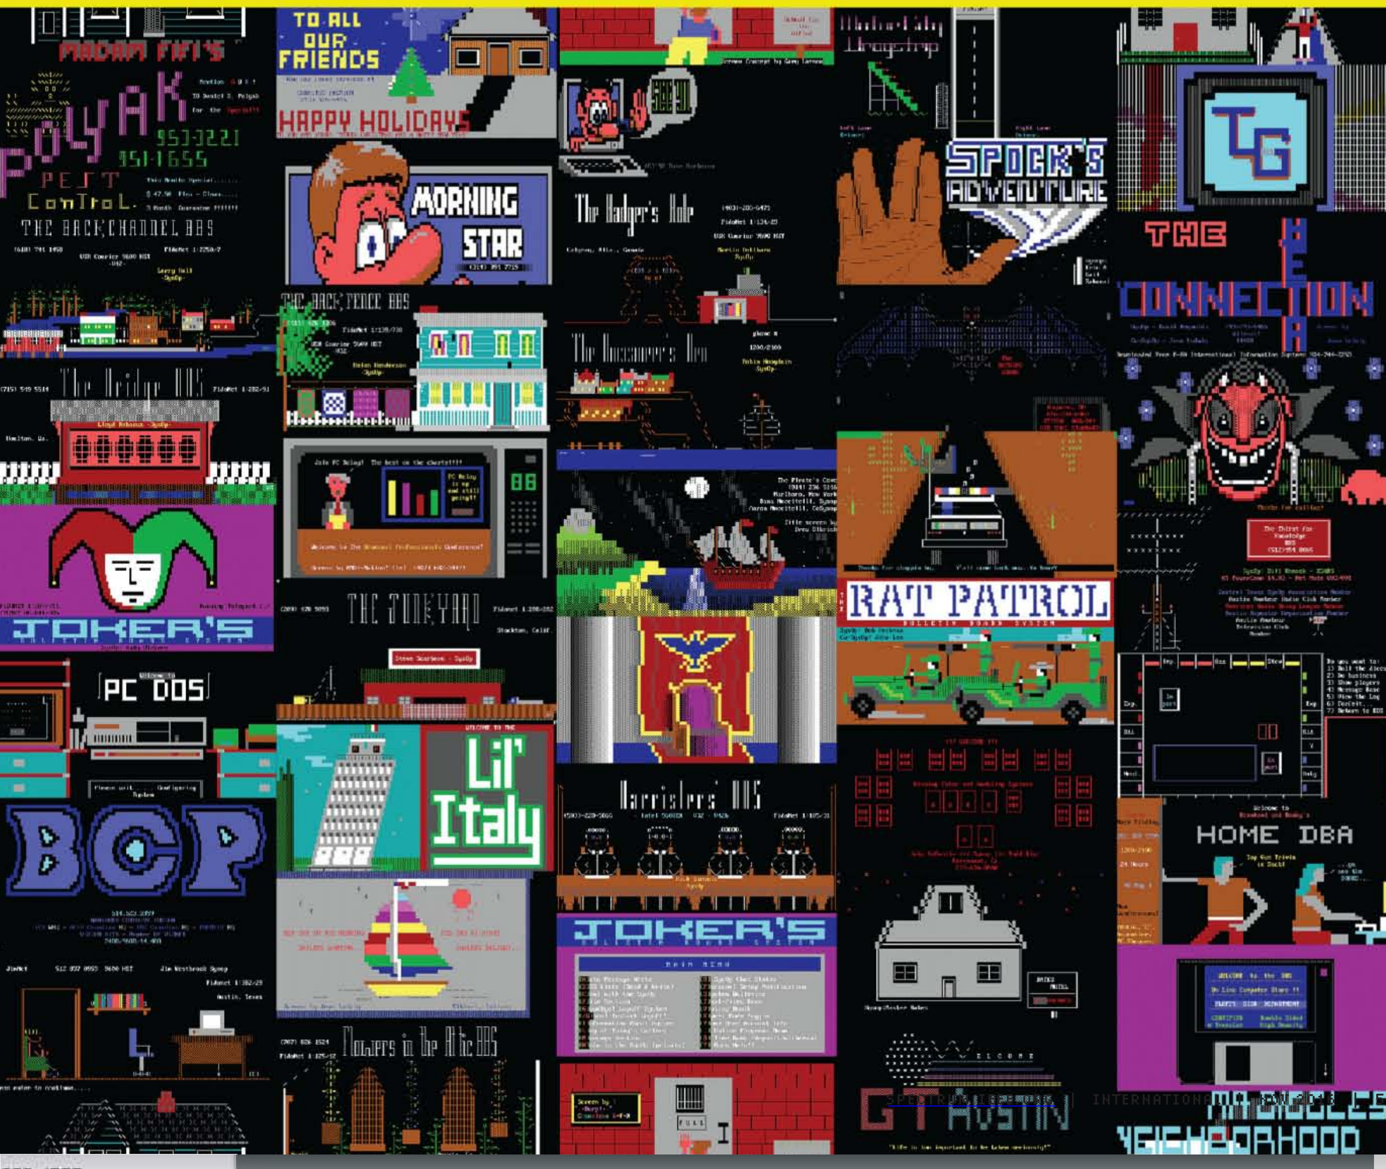

For millions of people around the globe, the Internet is a simple fact of life. We take for granted the invisible network that enables us to communicate, navigate, investigate, flirt, shop, and play. Early on, this network-of-networks connected only select companies and university campuses. Nowadays, it follows almost all of us into the most intimate areas of our lives. And yet, very few people know how the Internet became social. ¶ Perhaps that's because most histories of the Internet focus on technical innovations: packet switching, dynamic routing, addressing, and hypertext, for example. But when anyone other than a network engineer talks about the Internet, he or she is rarely thinking about such things. For most folks, the Internet is principally a medium through which we chat with friends, share pictures, read the news, and do our shopping. Indeed, for those who've been online only for the last decade or so, the Internet is just social media's plumbing–a vital infrastructure that we don't think much about, except perhaps when it breaks down. ¶ To understand how the Internet became a medium for social life, you have to widen your view beyond networking technology and peer into the makeshift laboratories of microcomputer hobbyists of the 1970s and 1980s. That's where many of the technical structures and cultural practices that we now recognize as social media were first developed by amateurs tinkering in their free time to build systems for computer-mediated collaboration and communication. ¶ For years before the Internet became accessible to the general public, these pioneering computer enthusiasts chatted and exchanged files with one another using homespun "bulletin-board systems" or BBSs, which later linked a more diverse group of people and covered a wider range of interests and communities. These BBS operators blazed trails that would later be paved over in the construction of today's information superhighway. So it takes some digging to reveal what came before.

How did it all start? During the snowy winter of 1978, Ward Christensen and Randy Suess, members of the Chicago Area Computer Hobbyist's Exchange (CACHE), began to assemble what would become the best known of the first small-scale BBSs. Members of CACHE were passionate about microcomputers, at the time an arcane endeavor, and so the club's newsletters were an invaluable source of information. Christensen and Suess's novel idea was to put together an online archive of club newsletters using a custom-built microcomputer and a hot new Hayes modem they had acquired.

This modem included an auto-answer feature, to which Christensen and Suess added a custom hardware interface between the modem and the hard-reset switch. Every time the telephone rang, the modem would detect the incoming call and then "cold boot" their system directly into a special host program written in Intel 8080 assembly language. Restarting the system with every call offered a blunt but effective means of recovering from hardware and software crashes–a common occurrence on home-brew hardware of the time.

Once a connection was established, the host program welcomed users to the system, provided a list of articles to read, and invited them to leave messages. Christensen and Suess dubbed the system "Ward and Randy's Computerized Bulletin Board System," or CBBS. It was, as the name suggested, an electronic version of the community bulletin boards that you still see in libraries, supermarkets, cafés, and churches.

Anyone with access to a teletype or video terminal could dial into CBBS. And after a few months, a small but lively community began to form around the system. In the hobbyist tradition of sharing information, Christensen and Suess wrote up a report about their project titled "Hobbyist Computerized Bulletin Board," which appeared in the November 1978 issue of the influential computer magazine *Byte*.

The article provided details about the hardware they used and how they organized and implemented their software. The authors even included their phone numbers and invited readers to take CBBS for a spin. Acknowledging the experimental nature of the system, they encouraged readers to "feel free to hang up and try several times if you have problems." After the issue hit newsstands, calls to their computer started pouring in.

Over the next few years, hundreds of small-scale systems like CBBS popped up around the country. Perhaps inspired by the *Byte* article, many of these new systems were organized by local computer clubs. In 1983, TAB Books, publisher of numerous DIY electronics guides, published *How to Create Your Own Computer Bulletin Board*, by Lary L. Myers. In addition to explaining the concept and motivation behind online services, Myers's book included complete source code in the BASIC programming language for host software. The back of the book also listed the telephone numbers of more than 275 public bulletin-board systems in 43 U.S. states. Some charged a nominal membership fee, while most were free to use. The roots of social media were beginning to take hold.

PREVIOUS PAGES, WELCOME SCREENS: ORIGINAL BBSs/JASON SCOTT

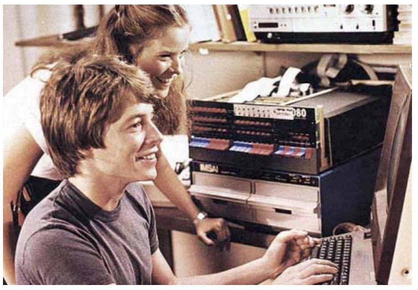

DIAL ME IN: [Clockwise from top left] Stacy Horn ran the Echo BBS in New York City; Sister Mary Elizabeth Clark ran the AIDS Education General Information System; CompuServe served up news; Matthew Broderick and Ally Sheedy starred in *WarGames*; Minitel provided online access in France.

CLOCKWISE FROM TOP LEFT: JAMES ESTRIN/THE NEW YORK TIMES/REDUX; BOB GRIESER/LOS ANGELES TIMES/GETTY IMAGES; AP PHOTO; UNITED ARCHIVE/ALAMY; OWEN FRANKEN/CORBIS/GETTY IMAGES

In retrospect, 1983 proved to be a critical year for popular computing. In France, the state-sponsored Minitel system completed its first full year of operation in Paris, making online news, shopping, and chat accessible to every citizen in that city. In the United States, novel commercial systems gained traction, with CompuServe reporting more than 50,000 paying subscribers.

Even Hollywood took interest in cyberspace. The 1983 movie *WarGames*, featuring a teenage hacker who explored remote computer networks from his bedroom, became an unlikely box-office smash. Although the IMSAI microcomputer and acoustic-coupler modem depicted in the movie once cost as much as a cheap used car, curious computer users inspired by the film could buy serviceable alternatives at the nearest Radio Shack for roughly the cost of a good-quality hi-fi stereo. And as the decade progressed, the online universe expanded rapidly from its original core of microcomputer hobbyists to encompass a much wider group.

Early dial-up BBSs were mostly local affairs. And no wonder: In the early 1980s, most Americans paid a flat monthly fee for unlimited local phone calls, but calls to another city or state were billed according to duration, distance, and time of day. Even calls within the same state could be quite costly. For tenderfoot BBS users, running up a monstrous phone bill was considered a rite of passage. To avoid long-distance calling altogether, more seasoned users restricted their online activity to nearby systems that could be reached toll free.

The local nature of BBSing meant that users could reasonably assume that the people they met online lived nearby. Even though many BBSs encouraged the use of pseudonyms, or “handles,” it was entirely possible that the person you were chatting with online tonight could be bagging your groceries or coaching your daughter’s soccer team tomorrow.

Many BBS administrators, or “sysops,” reinforced this sense of community by hosting regular in-person get-togethers, often at local parks or favorite watering holes. Online disagreements–flame wars–could be kept in check as well, because the cost of being a jerk escalated with the likelihood of later seeing your interlocutor face to face.

Another key distinction between the dial-up BBSs of the early 1980s and the social media services we use today is that, for the most part, each BBS was a world unto itself, whirring away in blithe isolation from all others. Over time, each system developed its own idiosyncratic personality. One might focus on trading shareware games, another on arguing politics, and a third on talking about TV shows.

Experienced BBS users would visit the various systems in their areas, maintaining a separate profile on each one. In big metro areas like Atlanta, Minneapolis, or Houston, there might be a dozen or more local BBSs to tap. But folks in less populated areas were seldom so lucky. Rural users with particularly narrow interests–say, collecting antique clocks–may have had only one or two other people nearby to chat with. They longed for more.

To overcome the economic cost of long-distance calling and the social cost of isolation, BBS operators needed to find a way to interconnect their systems, to create a network of BBS networks–a people's Internet. Christensen and Suess had floated the idea of a network of BBSs in the conclusion of their 1978 *Byte* article, but it was fellow amateur Tom Jennings who designed a real inter-network of BBSs.

In 1984, Jennings was distributing free of charge from his home in the San Francisco Bay Area a BBS host program for Microsoft's operating system, MS-DOS. He called the program Fido. The popularity of the Fido BBS grew alongside that of MS-DOS, and soon there were a dozen or more BBS systems running the program.

On a lark, Jennings added an experimental feature to his software that enabled two Fido BBSs to call each other automatically and exchange data. After a coast-to-coast test with John Madill, a Fido BBS operator in Baltimore, Jennings organized the first network of BBSs and called it FidoNet. Within a decade, FidoNet grew into a massive 20,000-node network reaching users as far away as South Africa and New Zealand. Unlike with today's Internet, machines on this network would typically store data for many hours before forwarding it on to its destination. But FidoNet nevertheless functioned as a valuable global data network.

The early versions of FidoNet included a number of clever design decisions that facilitated the rapid growth of the network. From the start, for example, Jennings encouraged collaboration by publishing technical documentation on FidoNet's protocols and file formats. As a result, support for FidoNet was added to other BBS-software packages and soon became a de facto standard for building such networks.

Because FidoNet was run primarily by hobbyists, it was designed to keep costs down. Exchanges between nodes normally took place in the middle of the night, when long-distance phone rates were lowest. This interval, during which many BBSs refused to answer calls from users, became known to insiders as "national mail hour."

Initially, FidoNet provided only inter-BBS email services. But FidoNet allowed BBS users to imagine themselves contributing to a vast conversation with people from all over the world. Newcomers later added other features and capabilities.

HALL OF FAMERS: In the 1970s and 1980s, Ward Christensen [top], Randy Suess [middle], and Tom Jennings [bottom] helped to pioneer the then-new online world. Christensen and Suess developed "Ward and Randy's Computerized Bulletin Board System," and Jennings later created FidoNet.

In 1986, Jeff Rush, a BBS operator in Dallas, created a conferencing mechanism for FidoNet. Rush's system, dubbed Echomail, worked similarly to the forums on CompuServe or the newsgroups on Usenet: You could post a comment, and anyone on the network could reply to it. But participation in those earlier forums required either a costly monthly subscription or access to a public Unix system, whereas anyone with a PC and a modem could access FidoNet's Echomail. As a result, Echomail became hugely popular and was particularly useful for people living in rural areas or those with niche interests. For the first time, the whole population of BBS users could engage in discussion together.

Beyond the United States, BBSs seemed to spring up anywhere that a microcomputer could be connected to a telephone line. In 1987, Pablo Kleinman, a sysop in Buenos Aires, helped to connect the first four FidoNet nodes in Argentina, and a year later, Juan Dávila, a sysop in Puerto Rico, announced the creation of Latino Net, a Spanish-language conference. Following the dissolution of the Soviet Union in 1991, a FidoNet user based in the United States circulated technical information about reading and writing Cyrillic messages and encouraged fellow FidoNet enthusiasts to reach out to the "overwhelming" number of Russian users joining the network.

By 1993, Randy Bush, a long-time advocate of FidoNet in the developing world, estimated that 59 percent of the nodes were located in North America, 30 percent in Europe, 4 percent in Australia and New Zealand, with the remaining 7 percent split among Asia, Latin America, and Africa. The African systems may have been small relative to the whole network, but these small networks could have an outsize impact for local users. In some regions of Africa, noted Bush, a FidoNet gateway provided an important means for poorly funded academics and NGO staffers to keep up to date with the latest research and news from abroad.

The growth of FidoNet during the 1980s was part of a larger movement toward greater connectivity among computer networks. Those interconnections were made using special gateways that translated messages between otherwise incompatible networks. In 1986, the administrators of several institutional Unix systems opened gateways to nearby BBSs, thereby enabling the exchange of messages among users of FidoNet, Usenet, and the nascent Internet. Indeed, several Usenet newsgroups began as automatically created copies of material posted in popular FidoNet Echomail conferences.

ORIGINAL PHOTOGRAPHS: JASON SCOTT

ILLUSTRATIONS BY QuickHoney

It's tempting to imagine that these gateways introduced BBS users to the Internet. But the Internet at that point held no particular appeal for most BBS users, who were already enmeshed in their own lively online worlds. The early Internet was limited to people with access to a large university or research center, whereas BBSs were open to anyone. As a result, they could be wacky and weird–independent free agents in cyberspace. Indeed, it may be more accurate to say that gateways exposed the relatively staid Internet to the wild ways of BBSers.

As the 1980s passed into the 1990s, falling prices for computers as well as the availability of used PCs allowed many more people to start using the technology. Modems also shifted from an expensive option to a standard component on new PCs. As these barriers to participation fell, a wave of new recruits came online, and the culture of BBSing began to reflect their more diverse interests. Teens, educators, artists, and entrepreneurs joined the growing population of BBS hosts and users.

Hundreds of BBSs appeared in the late 1980s and early 1990s to serve particular communities and interests. The Batboard in Columbia, Mo., was dedicated to all things Batman, for example, while the Complete Baseball BBS in Cambridge, Mass., featured a different sort of bat. Followers of the Grateful Dead ("deadheads") arranged tape trades and carpooling on the WELL in Sausalito, Calif., and hardcore gamers exchanged homemade *Doom* levels on Software Creations in Clinton, Mass. The Back Door in San Francisco supported a lively LGBT community, while the Backdraft in Key Largo, Fla., provided a meeting place for firefighters. The Christian-themed Winplus in Kent, Wash., and the pagan-oriented Brewer's Witch in Houston were both well known for their friendly off-line get-togethers, while DharmaNet linked up hundreds of Buddhist BBSs.

Free from the staggering size and profit imperative of commercial services, such small-scale community-oriented BBSs were sites of experimentation. Given the preponderance of men in the early BBS culture, many sysops endeavored to create welcoming environments for the relatively few women who called in. Stacy Horn, proprietor of the long-running Echo BBS in New York City, was especially tenacious in her appeals to women, offering new female users 12 months of free service and women-only online conferences. "I thought it would be a cooler place if there were women there," she remarked in an interview from 1996. "Who cares to talk to 20-year-old white guys all the time?"

Unlike with social-media services today, it was common for small-scale BBSs to vet new users before granting them full access. The TARDIS BBS in Indiana, for example, verified all of its users with a voice telephone call. Users claiming to be women were invited into the "Ladies Only" sections of the board only after being cleared by one of its female moderators. Even the board's male founder, Tom O'Nan, was shut out. "To this day," he recently joked, "I don't know what went on in that room!"

For communities in crisis, a BBS could be an important hub for sharing information. At the outset of the AIDS epidemic, as thousands were dying amid media coverage that was mostly suspicious or hostile, BBSs provided an invaluable source of health information and social support. Between 1985 and 1993, more than 100 computer bulletin boards were set up to share information regarding HIV and AIDS.

One cornerstone in this network was the AIDS Education General Information System (AEGIS) operated by Sister Mary Elizabeth Clark out of her home in San Juan Capistrano, Calif. Each day, Clark searched medical databases for the latest research about treatment and prevention and then used FidoNet to circulate this information to BBSs in more than 40 countries throughout North America, Europe, and Africa. Meanwhile, the United Methodist Church's Computerized AIDS Ministries Network (CAM) in New York City helped combat the isolation faced by many people affected by the crisis, by answering visitors' questions and generally providing support.

For two decades, bulletin-board systems provided countless people with an accessible platform for online community. Jason Scott, director of *BBS: The Documentary*, estimates that there were at least 106,418 BBSs in operation at one time or another between 1978 and 2004. And yet, in spite of their wide-ranging scale, geographic reach, and cultural influence, BBSs seem to be all but forgotten today. Why?

Part of the reason is their swift demise. With the commercialization of the Internet in the mid-1990s, thousands of BBSs seemed to vanish almost overnight. Some attempted to move, wholesale, to Usenet or the World Wide Web, while others simply pulled the plug and turned off the lights.

But a curious thing happened along the way to oblivion: Thousands of BBSs quietly metamorphosed into Internet Service Providers, or ISPs. In December 1995, Jack Rickard, editor of *Boardwatch* magazine, noted that the available statistics on Internet usage didn't add up: The major ISPs, like UUNET, Netcom, PSINet, and InternetMCI, were not reporting enough subscribers to account for the total number of active users online.

Using a database compiled by his magazine, Rickard estimated that more than 95 percent of the 3,240 ISPs created in the previous two years were former BBSs operating under new names. The same telephone lines and modems that once connected local callers to one another were now being used to provide direct connections to the global Internet. The dial-up BBSs were bulldozed, it seems, to build Internet on-ramps.

For Internet users today, there is more to the BBS story than mere nostalgia. As our online lives are channeled through an ever-smaller number of service providers, reflecting on the long-forgotten dial-up era can open our eyes to the value of diversity and local control.

After all, people often get frustrated that today's social-media platforms can't deal with small-scale disputes or apply different policies to different constituencies. BBS users of the past could complain directly to the sysop–maybe even with a phone call–and he or she could act immediately and unilaterally on their behalf. Try that with Twitter or Facebook. You'll gain a much better appreciation for the systems often running on little more than solder, duct tape, and a dream that first enabled regular people to connect with one another through their computers. ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/bbshistory1116

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

FACULTY POSITIONS

Electrical and Computer Engineering

NYU SHANGHAI

NYU Shanghai is currently inviting applications for tenured/tenure track positions at the rank of assistant, associate, or full professor. We will consider applicants in all areas of electrical and computer engineering, and have a particular interest in applicants in the fields of Multimedia Analytics, Neural Engineering, Systems Neuroscience, Robotics, Smart Cities and Smart Vehicles. The applicants must have demonstrated abilities in both research and teaching. 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 ***January 15, 2017*** *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), electronic copies of up to five recent, relevant publications and the names and addresses of three or more individuals willing to provide letters of reference. Please visit our website at **http://shanghai.nyu.edu/en/about/work-here/open-positions-faculty** for instructions and other information on how to apply. If you have any questions, please e-mail **shanghai.faculty.recruitment@nyu.edu**.

About NYU Shanghai:
NYU Shanghai is the newest degree-granting campus within New York University's global network. It is the first higher education joint venture in China authorized to grant degrees that are 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 of the highest caliber who are committed to NYU's global vision of transformative teaching and innovative research and who embody the global society in which we live.

NYU's global network includes degree-granting campuses in New York, Shanghai, and Abu Dhabi, 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.

NYU Shanghai is an equal opportunity employer committed to equity, diversity and social inclusion. We strongly encourage applications from individuals who are under-represented in the profession, across color, creed, race, ethnic and national origin, physical ability, and gender and sexual identity. NYU Shanghai affirms the value of differing perspectives on the world as we strive to build the strongest possible university with the widest reach.

上海纽约大学
NYU SHANGHAI

EOE/AA/Minorities/Females/Vet/Disabled/Sexual Orientation/Gender Identity Employer

---

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) robotics, including human/robot interactions, (ii) mobile health sensing and health analytics, (iii) senior and junior faculty in cybersecurity (including hardware-enabled cybersecurity) and mobility, (iv) electric machines and variable frequency drive systems, and (v) Director of the ElectroScience Laboratory and Professor of Electrical and Computer Engineering. For additional information on these positions, see **https://ece.osu.edu/about/employment**. All positions may involve joint appointments with other engineering departments. Applicants must have a Ph.D. and outstanding academic credentials. Successful candidates are 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 eng-ece-search@osu.edu. Some 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.

Queen's UNIVERSITY

Faculty Position: Applied Mathematics & Math & Engineering

The Department of Mathematics and Statistics, Faculty of Arts and Science at Queen's University, Kingston, Canada, invites applications for a Tenure-track faculty position in Applied Mathematics and Mathematics and Engineering at the rank of Assistant Professor with a starting date of **July 1, 2017**.

In areas related to Mathematics and Engineering, there are presently prominent research groups in Information and Communication Theory, Control Theory, and Geometric Mechanics. For more information about the Mathematics and Engineering program, please see **http://www.mast.queensu.ca/meng/**. A successful candidate will be expected to work in any of these or complementary research areas, and to contribute to both the graduate and undergraduate programs. A candidate who joins the Mathematics and Engineering group will be expected to obtain a license as a Professional Engineer; an undergraduate degree in Engineering is a strong asset towards obtaining the license.

For the full position announcement and information on how to apply, please, visit: **http://www.mast.queensu.ca/positions/**.

NATIONAL TAIWAN UNVERSITY
Department of Electrical Engineering
Faculty Openings

The top choice for the students in natural sciences and engineering in Taiwan is inviting you to join us to establish a world-class department. We are seeking candidates with enthusiasm in teaching and potential in research, and invite applications for tenure-track/ tenured positions of Professor, Associate Professor and Assistant Professor. For more information: **http://www.ee.ntu.edu.tw/en**

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

BAYLOR UNIVERSITY 1845

Faculty Position

The Electrical and Computer Engineering Department of Baylor University seeks faculty applicants for a Lecturer or a Senior Lecturer Position. Any area of expertise in ECE will be considered but applicants in computer engineering will be given special consideration. Applicants must demonstrate potential for excellent teaching; applicants for Senior Lecturer must additionally present evidence of achievement in 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, 2017 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 interests,

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

Additional information is available at **www.ecs.baylor.edu**. Should you have any questions on the position, feel free to contact the search chair, Dr. Keith Schubert at keith_schubert@baylor.edu. Please submit materials to **https://apply.interfolio.com/38070**.

Baylor University is a private not-for-profit university affiliated with the Baptist General Convention of Texas. As an Affirmative Action/Equal Opportunity employer, Baylor is committed to compliance with all applicable anti-discrimination laws, including those regarding age, race, color, sex, national origin, marital status, pregnancy status, military service, genetic information, and disability. As a religious educational institution, Baylor is lawfully permitted to consider an applicant's religion as a selection criterion. Baylor encourages women, minorities, veterans and individuals with disabilities to apply.

Two Tenure-Track, Tenured Faculty Positions
Department of Electrical and Computer Engineering
Queen's University at Kingston

Queen's UNIVERSITY

The Department of Electrical and Computer Engineering in the Faculty of Engineering and Applied Science at Queen's University invites applications for two Tenure-track or Tenured faculty positions with a specialization in the field of Computer Engineering Software, with priority areas in Software Engineering, Machine Learning, Autonomous and Intelligent Systems, Data Analytics and Big Data, and Mechatronics and will also consider exceptional candidates in other fields of Computer Engineering Software. Preference will be given to candidates at the rank of Assistant professor; however, exceptional candidates at the rank of Associate Professor may also be considered

The successful candidates must have a PhD in Electrical and Computer Engineering or a related discipline completed at the start date of the appointment. Licensure as a Professional Engineer in Canada, or eligibility to acquire licensure in Canada, is an essential requirement. Postdoctoral or industrial experience will be considered an asset.

Queen's University is one of Canada's leading research-intensive universities with a global reputation and a recognized leader in Canadian higher education. The Department of Electrical and Computer Engineering is medium-sized with 23 full-time and 12 cross-appointed faculty, 470 undergraduate students, and 170 graduate students. The department offers undergraduate, master's, and doctoral programs in electrical and computer engineering and has recently launched a unique entrepreneurial engineering program referred to as ECE Innovation Stream, ECEi.

The Department of Electrical and Computer Engineering is home to Queen's Centre for Energy and Power Electronics Research (ePOWER) and has connections to a number of multi-disciplinary centres such as: Canadian Microelectronics Corporation (CMC), Human Mobility Research Centre, Green Centre Canada, Innovation Park, and Queen's Innovation Connector.

Queen's historic campus is located in the heart of the vibrant Kingston community in the Thousand Islands region of South Eastern Ontario. Queen's is positioned centrally with respect to three major metropolitan areas: Toronto, Montreal, and Ottawa.

The full posting, including additional information and the requirements for submitting an application is available at **http://my.ece.queensu.ca/About-Us/Employment.html**. Inquiries can be sent to ece-search@queensu.ca . The deadline for applications is December 31, 2016; however, applications will continue to be reviewed until the position has been filled.

A PEN FOR THE ATOMIC AGE

In 1958, the Parker Pen Company developed the "Atomic Pen," which later made a cameo appearance in Stanley Kubrick's *2001: A Space Odyssey*. The design never went into production, but it called for a tiny packet of radioactive isotopes that would heat the ink to produce a selectable range of line densities. The idea of a radioactive pen may sound crazy now, but in an era promising atomic cars and atomic planes, it no doubt seemed perfectly reasonable. ■

ALFRED EISENSTAEDT/THE LIFE PICTURE COLLECTION/GETTY IMAGES

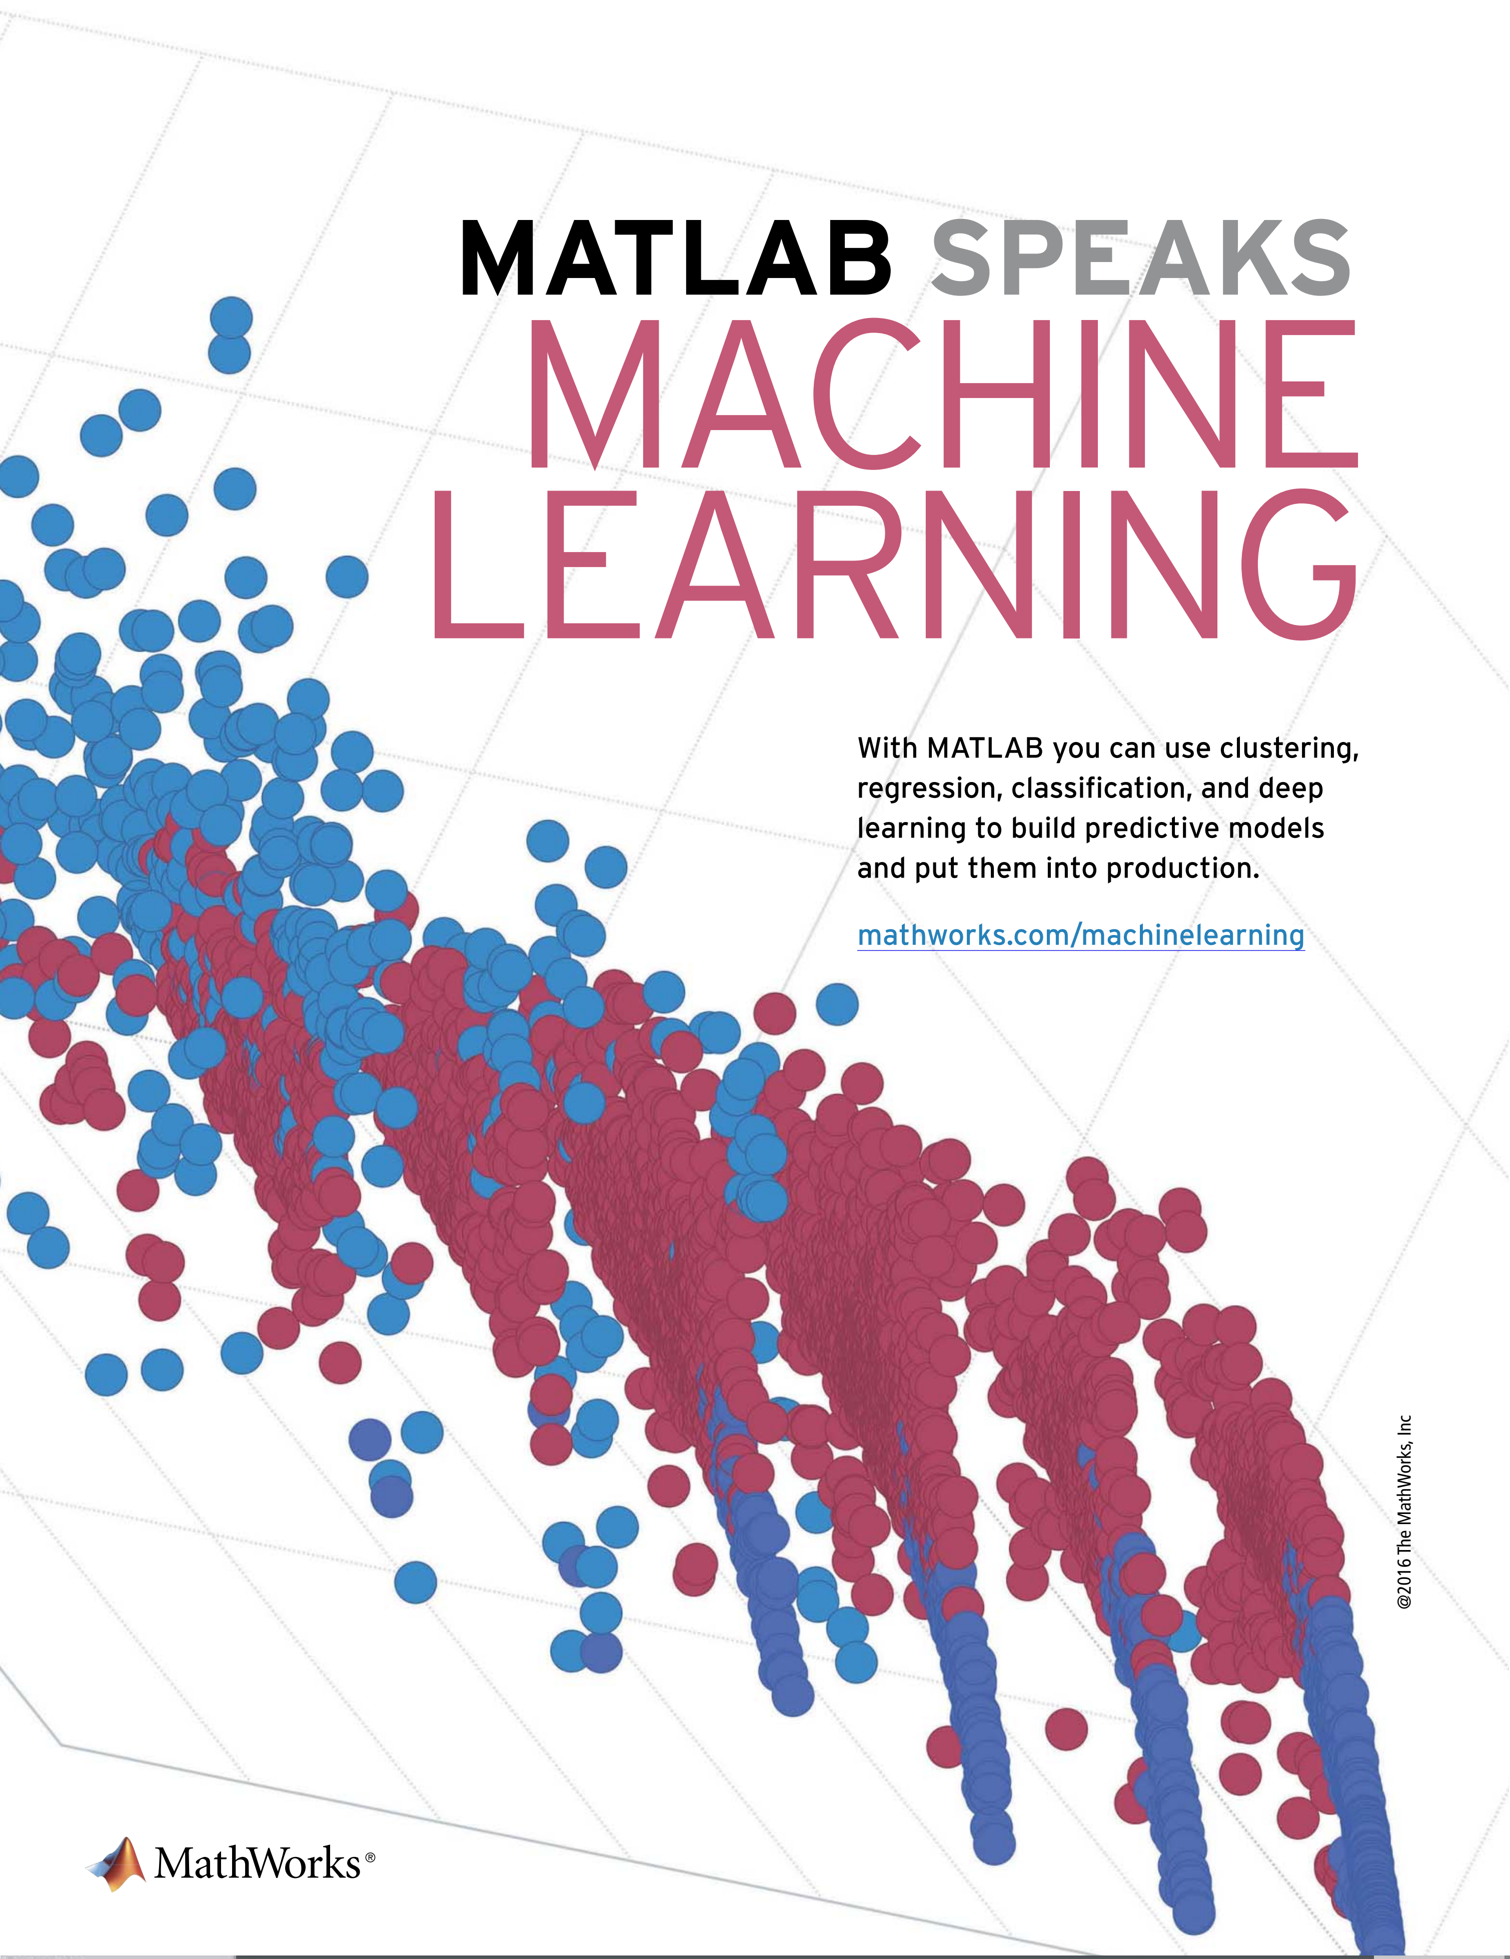
MATLAB SPEAKS
MACHINE
LEARNING
With MATLAB you can use clustering, regression, classification, and deep learning to build predictive models and put them into production.
mathworks.com/machinelearning
@2016 The MathWorks, Inc
MathWorks®