IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 02.16

It is the year 2023, and for the first time, a self-driving car strikes and kills a pedestrian. A lawsuit is sure to follow. But exactly what laws will apply? Nobody knows.

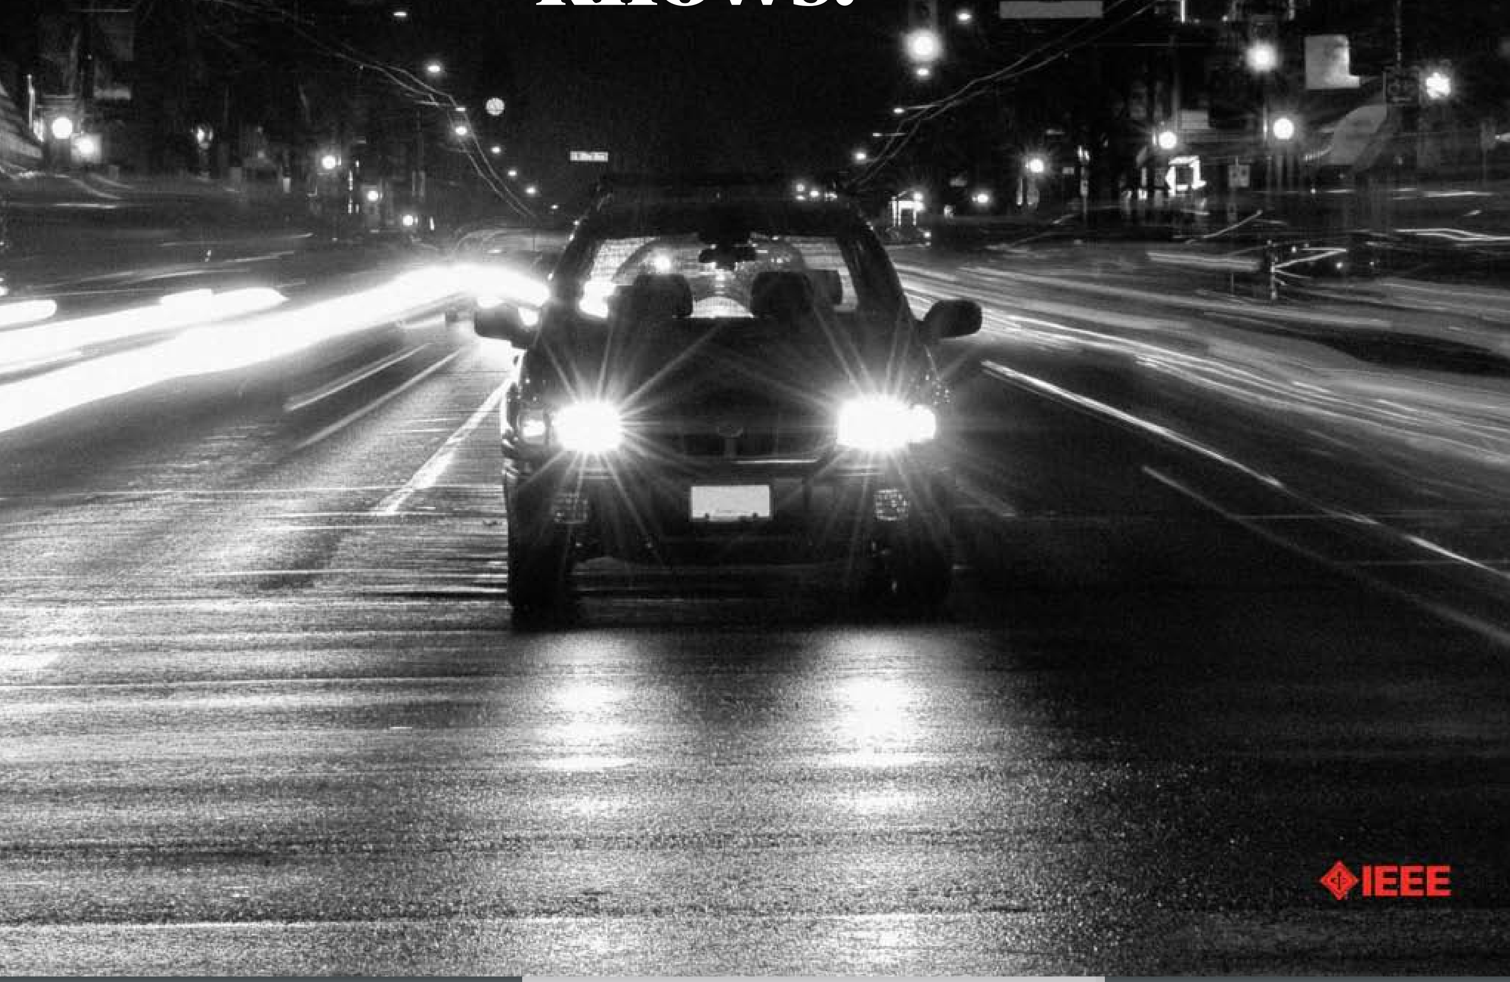

IEEE

MULTIPHYSICS FOR EVERYONE

The evolution of computational tools for numerical simulation of physics-based systems has reached a major milestone.

Custom applications are now being developed by simulation specialists using the Application Builder in COMSOL Multiphysics®.

With a local installation of COMSOL Server™, applications can be deployed within an entire organization and accessed worldwide.

Make your organization truly benefit from the power of analysis.

comsol.com/application-builder

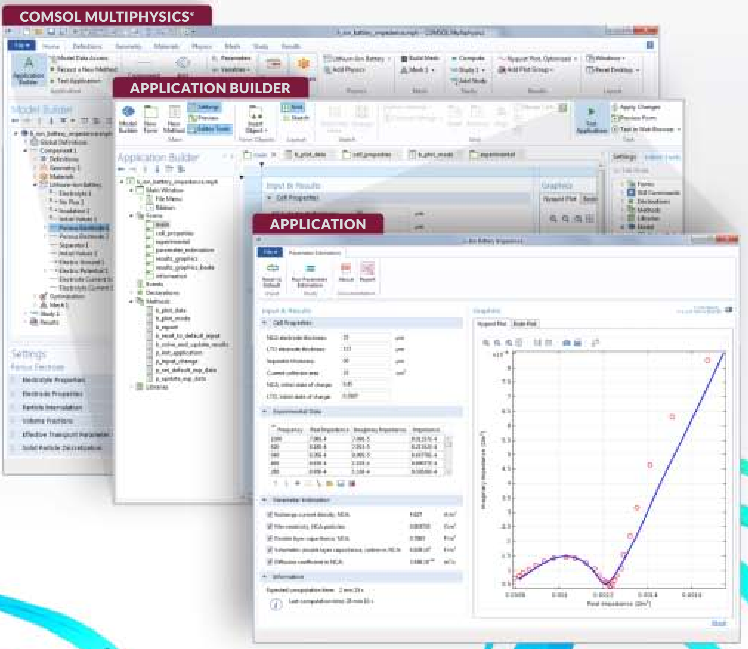

COMSOL

© Copyright 2016 COMSOL. COMSOL, COMSOL Multiphysics, Capture the Concept, COMSOL Desktop, COMSOL Server, LiveLink, and Simulation for Everyone are either registered trademarks or trademarks of COMSOL AB. All other trademarks are the property of their respective owners, and COMSOL AB and its subsidiaries and products are not affiliated with, endorsed by, sponsored by, or supported by those trademark owners. For a list of such trademark owners, see www.comsol.com/trademarks

FEATURES_02.16

IEEE SPECTRUM

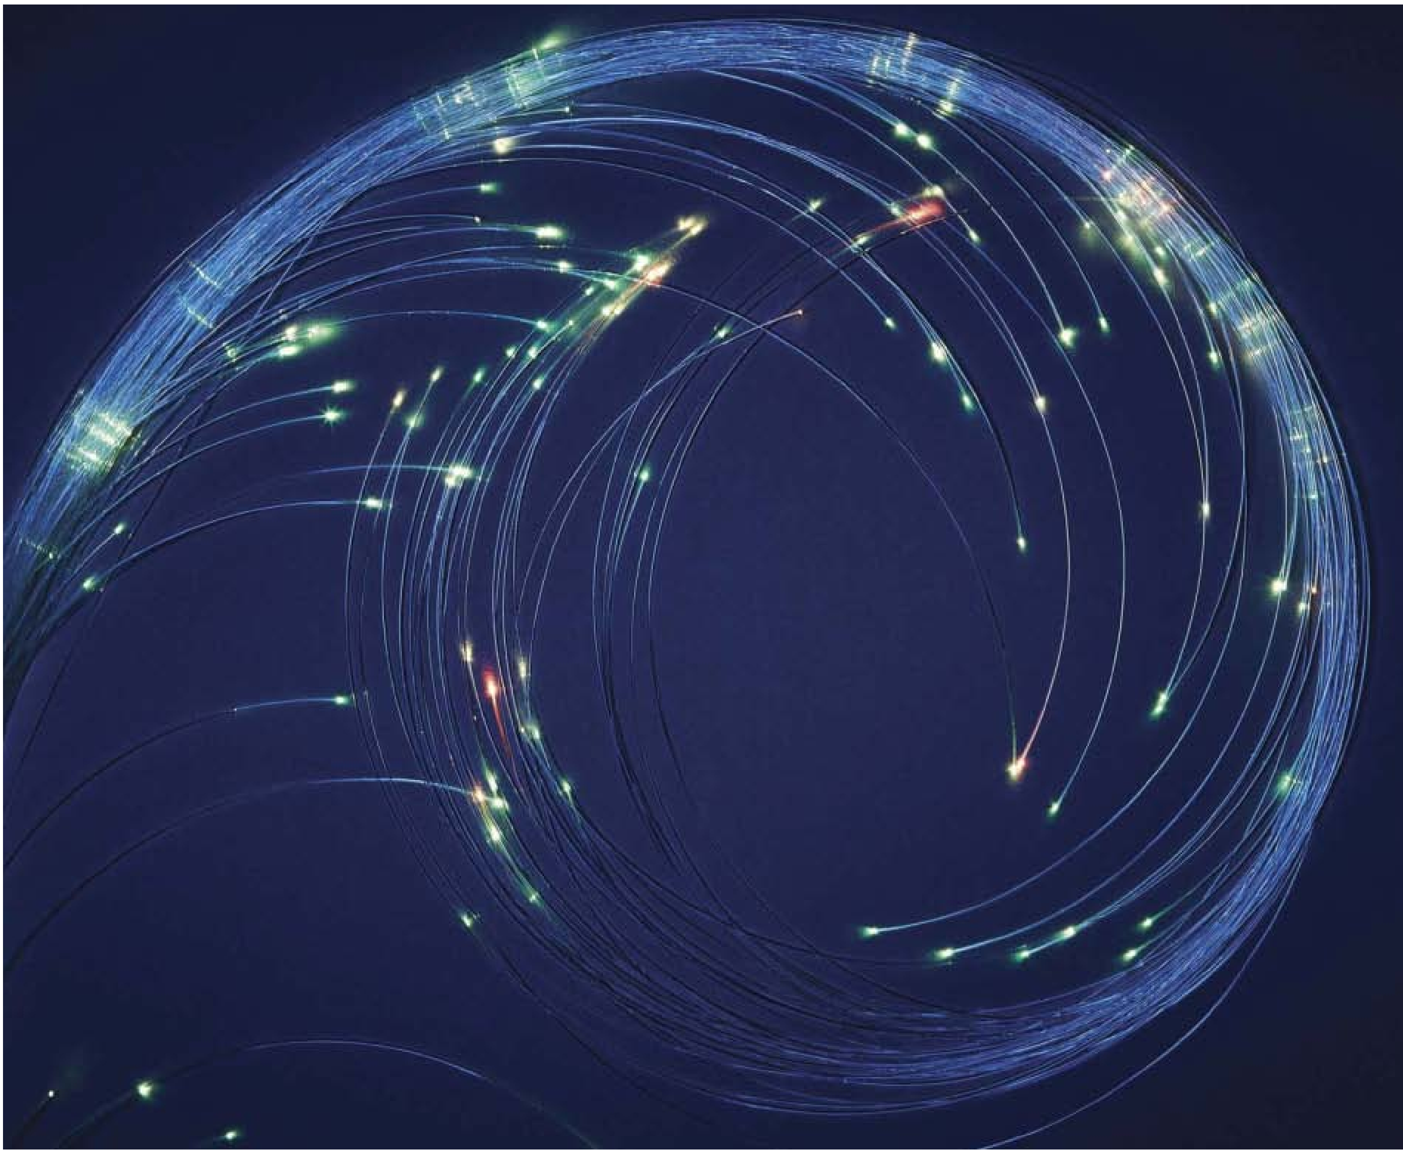

On the Cover **Photograph by Kevin Su/Alamy**

PHOTOGRAPH BY Sue Tallon

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

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

It starts with an integrated design approach for rapid prototyping.

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

>> See how at ni.com/5g

800 891 8841

NATIONAL INSTRUMENTS™

©2015 National Instruments. All rights reserved. LabVIEW, National Instruments, NI, and ni.com are trademarks of National Instruments. Other product and company names listed are trademarks or trade names of their respective companies. 24296

DEPARTMENTS_02.16

IEEE SPECTRUM

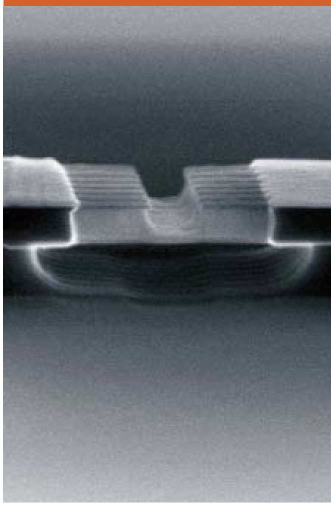

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Non Linear Magnetic Materials with COMSOL Multiphysics—11 February
- Grid Resiliency and Stability With Siemens Transformers—Available On Demand
- Signal Integrity: VNA Applications—Available On Demand
- **NEW PRODUCT RELEASE LIBRARY** http://spectrum.ieee.org/static/new-product-release-library
- **WHITE PAPERS** http://spectrum.ieee.org/whitepapers

The Institute

Available at theinstitute.ieee.org

- **BREAKTHROUGHS FROM SPACE** Research from the International Space Station has led to a robotic arm that lends a helping hand in delicate surgeries, an app that identifies clean water sources, and cameras that aid in disaster response efforts.
- **HOW CURIE SAVED THOUSANDS** You may know Marie Skłodowska Curie for her pioneering advancements in radiology, but did you know that she also developed technology that saved the lives of countless soldiers during World War I?
- **MEDICAL DEVICE INTEROPERABILITY** The IEEE Standards Association has unveiled new interoperability standards for medical devices, 3-D printers, and other e-health devices.

FROM LEFT: PURDUE UNIVERSITY; HEP SVADJA; JOHN LOCHER/AP PHOTO

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

BACK STORY_

El Camino Reality

FENDER BENDERS on the chronically congested stretch of El Camino Real that bisects Silicon Valley occur with some regularity. But the one Matthew N. Eisler came upon last summer caught his eye: a three-car pileup involving electric cars exclusively.

Talk about timing. Eisler, a visiting assistant professor in the department of integrated science and technology at James Madison University, in Harrisonburg, Va., was at Stanford University doing research on electric vehicles for a forthcoming book. He'd just left an interview with members of a local auto club when he came upon the accident. "There was a Prius and two Leafs," Eisler recalls. "They were all in a line: one, two, three. It was pretty minor–nobody seemed to be hurt. But it was unlike any accident I'd ever seen."

California has of course been ground zero for EV development for years. The state's stringent emission regulations have helped bring about such electrics as GM's ill-fated EV-1, Toyota's wildly popular Prius hybrid, and Tesla Motors' muscular Model S.

Eisler [shown with a Chevy Volt at his university] says the EV owners he's met have all been "to varying degrees evangelical–they're eager to explain why EVs are a good thing. But nobody sugarcoated it," he adds. "They're very forthcoming about the problems they've experienced."

"They consider themselves part of a broader experiment," he explains. In "A Tesla in Every Garage?," in this issue, Eisler analyzes the role of Tesla Motors in that experiment, concluding that the company's emphasis on luxury and performance has tended to obscure the real differences between EVs and their gasoline-fueled brethren. ■

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. 2 (INT), February 2016, p. 52, or in *IEEE Spectrum*, Vol. 53, no. 2 (NA), February 2016, p. 60.

02.16

IEEE SPECTRUM

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

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

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

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

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

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

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

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

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

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

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

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

EDITORIAL 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

RESPONSIBILITY FOR THE SUBSTANCE OF ARTICLES rests upon the authors, not IEEE, its organizational units, or its members. Articles do not represent official positions of IEEE. Reader comments will be posted online, or published in print as space allows, and may be excerpted for publication. The publisher reserves the right to reject any advertising.

REPRINT PERMISSION / LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Business Manager, IEEE Spectrum.

COPYRIGHTS AND TRADEMARKS *IEEE Spectrum* is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Careers, EEs' Tools & Toys, EV Watch, Progress, Reflections, Spectral Lines, and Technically Speaking are trademarks of IEEE.

LYNDA CHANDLER

IEEE

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

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

MANAGING DIRECTOR, PUBLICATIONS Michael B. Forster

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

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

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

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

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

DIRECTORS EMERITUS Eric Herz, Theodore W. Hissey

IEEE STAFF
EXECUTIVE DIRECTOR & COO James Prendergast
+1 732 502 5400, james.prendergast@ieee.org
CHIEF INFORMATION OFFICER Cherif Amirat
+1 732 562 6399, c.amirat@ieee.org
PUBLICATIONS Michael B. Forster
+1 732 562 3998, m.b.forster@ieee.org
CORPORATE ACTIVITIES Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
HUMAN RESOURCES Shannon Johnston, SPHR
+1 732 562 6343, s.johnston@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Eileen M. Lach, +1 212 705 8990, e.m.lach@ieee.org
CHIEF MARKETING OFFICER Patrick D. Mahoney
+1 732 562 5596, p.mahoney@ieee.org
EDUCATIONAL ACTIVITIES Jamie Moesch
+1 732 562 5514, j.moesch@ieee.org
CHIEF FINANCIAL OFFICER Thomas R. Siegert
+1 732 562 6843, t.siegert@ieee.org
TECHNICAL ACTIVITIES Mary Ward-Callan
+1 732 562 3850, m.ward-callan@ieee.org
MANAGING DIRECTOR, IEEE-USA Chris Brantley
+1 202 530 8349, c.brantley@ieee.org
IEEE PUBLICATION SERVICES & PRODUCTS BOARD
Sheila Hemami, *Chair*; John Baillieul, Babak D. Beheshti, Herbert S. Bennett, Jennifer T. Bernhard, Eddie Custovic, Jean-Luc Gaudiot, Ron B. Goldfarb, Glenn Gulak, Lawrence Hall, Christofer Hierold, James M. Irvine, Hulya Kirkici, Carmen S. Menoni, Paolo Montuschi, Thrasos Pappas, Michael Pecht, Michael Polis, Sorel Reisman, Tariq Samad, Fred Schindler, Gianluca Setti, Gaurav Sharma, Curtis A. Siller, David Tian, Ravi Todi, H. Joel Trussell, Leung Tsang, Stephanie M. White, Steve Yurkovich

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

CONTRIBUTORS_

Nathan A. Greenblatt

Greenblatt is an intellectual-property attorney at Sidley Austin in Palo Alto, Calif. In "Self-Driving Cars and the Law" [p. 42], he explains how the legal and policy implications of autonomous vehicles need to be developed in tandem with the technology. His interest in the topic sprang from "being frustrated by my daily commute and being fascinated by the potential of the technology," Greenblatt says. "I look forward to having vehicles on the road that will actually let others merge."

Jeff Hecht

In this issue, freelance writer Hecht takes a look at the fast-paced history of fiber-optic innovations and the ideas engineers are pursuing now to keep the momentum going [p. 24]. Hecht has covered fiber for roughly 40 years—"an embarrassingly long time," he says—and written several books on the topic along the way. In retrospect, he says, "It was sort of like covering a winning sports team on a roll." Now the question is what comes next.

Zachary C. Lipton

Lipton studies computer science at the University of California, San Diego, where he works closely with coauthor Charles Elkan, who is on the UCSD faculty. Lipton didn't have a computer science background when he started grad school. He was a jazz saxophonist in New York City. In 2011, he moved to California to join a health-tech startup and later enrolled at UCSD to study machine learning, a topic he and Elkan explore in "Playing the Imitation Game With Deep Learning" [p. 36].

Paul Perrault

In this issue, Perrault and coauthor Mike Teachman describe how they instrumented a set of Mason bee nests [p. 18]. The project was a novel challenge, says Perrault, who is a field applications engineer for Analog Devices. "While I've designed capacitive sensors for human user interfaces (smartphones) and for fluid instrumentation (industrial tanks), this was the first time I used them for a zoological purpose," he says.

Sue Tallon

A photo showing a uniform bouquet of glowing glass strands is standard for many fiber-optics articles. But Tallon, who shot the images for this month's fiber communication feature ["Great Leaps of Light," p. 24], was determined to do something different. She toyed with a variety of approaches over a three-day photo shoot in a darkened room. The winning subject: a 1.5-meter-long "whip" toy with a flashlight in its handle. It was surprisingly fun, she says, "seeing how much you could create with nothing there but light."

02.16

PHOTO-ILLUSTRATIONS BY **Gluekit**

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

GARY WATERS/GETTY IMAGES

The Trouble With Targets

The threat of failure hangs over efforts to solve big, multidimensional engineering challenges, but that shouldn't stop us from trying

AMONG TECHNOLOGY LEADERS burdened with unmet and ill-defined goals, setting engineering targets is the rage, promising specific results to people, corporations, and government.

Whether the demand is for capping the rise in Earth's temperature, creating therapies to halt memory loss in the elderly, or expanding farm output to meet the world's growing population, the answer is the same: Set targets.

The urge to target is varied and insistent—and contradicts the widespread view that technologies and their underlying physics and mathematics determine outcomes, not people. A movement to put humans at the center of engineering is fueled by the popularity of "effective altruism" and humanitarian engineering. Increasingly, politicians and the public talk about the technologies they want rather than settling for what Technology—with a capital *T*—can give them.

The sensibility informs a range of urgent questions for engineers. Does artificial intelligence pose threats to humanity? Target *good* outcomes of AI. Might robots destroy human employment? Create robots that only help workers. To counter Ebola, cancer, and other lethal diseases, invent vaccines or cures.

Targets (think putting humans on the moon) are a clever means of holding technologists to account, charting their progress, and insisting on results. In an age of limited resources, great inequality, and growing uncertainty, clear aims trump the value of free-wheeling inquiry.

But while appealing, targeting masks complexity and encourages overconfidence, even complacency. Existential threats, it turns out, are easy to identify but difficult to resolve.

Terrorism, chaotic climates, cyberweapons, mysterious diseases, vanishing species: The list of *fixable* perils grows longer with each year. Technologies of abundance altered our existence but came at a cost that is only now being more accurately counted. This tension—between the glories of our engineered lives and the price to be paid for them—is the essential drama of our times.

No one sits closer to the center of this gathering storm than the engineer. Only the engineer understands the contradictions of the human-built world and possesses the skills to craft solutions.

Yet engineering never occurs in isolation. Targets reflect the desires of masses of people. In a fragmented world, where cherished diversity spawns at times irreconcilable claims between factions, only existential threats generate the unity of purpose that in turn produces universally shared hopes for emerging technologies.

No surprise then that the technological landscape is littered with failed efforts at targeting, whether mounted by governments, corporations, or civil society.

Crafting targets is part art. Targeting seems least effective when goals are broad and fuzzy. Such targets as improving primary-school education, curing cancer, or preventing terrorists from using social media to win converts can seem impossible to reach. These challenges and others like them require pushing down multiple pathways toward many smaller targets, which then exponentially increases cost, complexity, and the chances of failure.

The multidimensionality of many goals has serious implications for targeters. Constraints on innovation aren't limited by ambition or even resources. The prospects for bending the physical world to humanity's wishes can never be fully tested in a lab or modeled on a computer. Because rising temperatures reflect many factors, halting warming will require many engineering projects, each with its own target. The interactions among the new targets raise the specter of an infinite regress.

Don't despair. Setting targets makes sense, especially if targets are concrete, feasible, and widely desired. Yet ambitious technological campaigns demand enormous humility. Grand aspirations for engineering must be matched with an awareness of the potential for choosing the wrong target or messing up in pursuit of the right one. All we can know for certain is that our good intentions are never enough. —G. PASCAL ZACHARY

G. Pascal Zachary, a professor of practice at Arizona State University's School for the Future of Innovation in Society, is the author of *Showstopper!* (1994), on the making of Windows NT.

CORRECTION: In our news article "Has Taiwan Given Up on Supercomputing?" [January], we misidentified "the Cabinet," Taiwan's executive branch of government, as the executive branch of the Chinese government.

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

NEWS

5-NM NODE

THE SEMICONDUCTOR MANUFACTURING GENERATION THAT MAY REQUIRE COMPOUND SEMICONDUCTORS INSTEAD OF SILICON

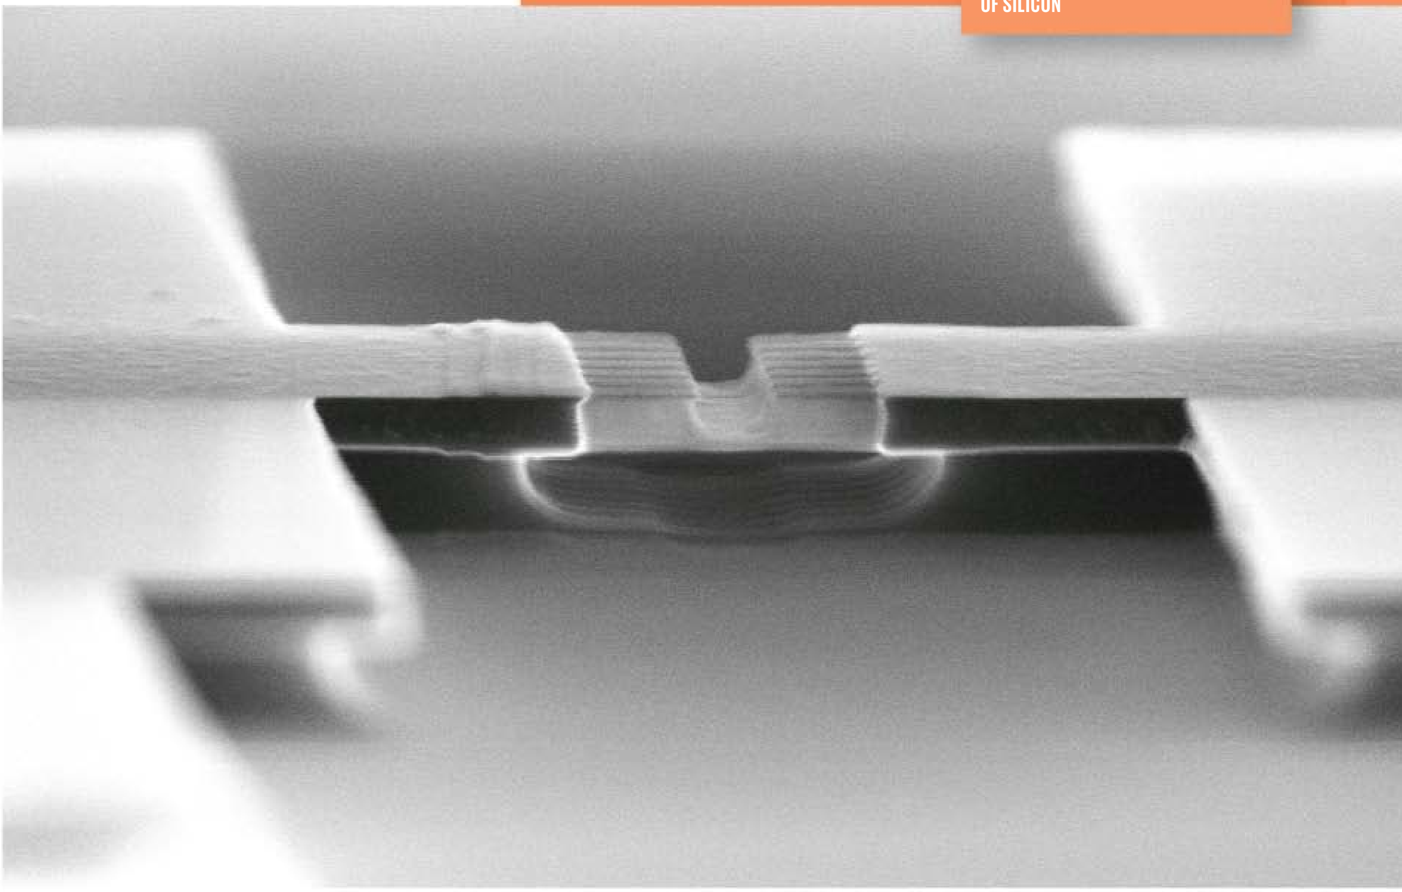

RISE OF THE NANOWIRE TRANSISTOR

Compound-semiconductor nanowires may keep smartphones charged longer

PURDUE UNIVERSITY

BRIDGING THE GAP: Germanium nanowires are suspended across a gap in this transistor.

We cherish our smartphones for delivering entertainment and information on the go, but their need for daily charging is a problem. Battery life can't get any shorter than it is today. (Well, it could, but consumers wouldn't be happy about it.) So when new smartphone models come on the market with microprocessors based on the latest foundry process, the increase in the number of transistors in the chips should be balanced by a reduction in the power that each transistor consumes.

For the remainder of the decade, this power reduction per transistor can be accomplished with today's workhorse device: the silicon FinFET. (It's so named because the channel through which current flows is shaped like a vertical fin.) But continuing progress further into the future will require an overhaul of the transistor's architecture: If the devices unveiled in December at the IEEE International Electron Devices Meeting (IEDM) are an indication, that overhaul will see the FinFET's silicon fin shrink vertically to become a nanometers-wide wire made from semiconductors other than silicon. »

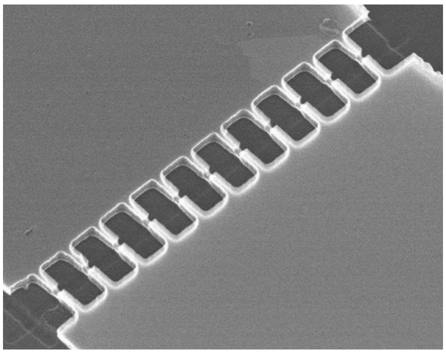

ODD WIRE OUT: Eleven nanowires form the channel to balance the germanium transistor's performance.

PURDUE UNIVERSITY

The superior semiconducting alternatives include germanium and III-V materials (so called because they combine a group III element such as gallium with a group V element such as arsenic). They transport charge faster and allow the production of transistors that can deliver the same amount of current as their silicon counterparts but at a lower voltage, saving power.

The change in geometry from fin to wire saves power in a different way: The gate, whose voltage controls the flow of current through the channel, can surround four sides of a nanowire channel but only three sides of a fin. So in the nanowire configuration, the gate is more effective at pinching off unwanted current that might otherwise leak through the channel, again saving power.

The building blocks of the CMOS circuits of today's processors require two types of transistors: one that transports electrons and another that carries holes, their positive counterparts. At IEDM, Peide Ye, a professor of electrical and computer engineering at Purdue University, in West Lafayette, Ind., championed germanium-based nanowire transistors for both types of devices. Meanwhile, a team of researchers from universities in Singapore showcased the promise of an all-III-V approach, and Niamh Waldron, principal engineer at Imec, the European microelectronics center, reported performance advances yielded by the pairing of an electron-transporting III-V transistor with a hole-transporting germanium transistor.

All three approaches have their pros and cons. Ye, who has spent a great deal of time over the past 15 years improving the performance of III-V transistors, argues that germanium devices are much easier to process. But the National University of Singapore's Xiao Gong points out that there are difficulties associated with making electron-transporting germanium transistors. Among the challenges: forming good electrical contacts.

Aaron Thean, vice president of logic at Imec, touts the European center's approach, which pairs a germanium hole-transporting transistor with a III-V electron-carrying device. He says it is best from a pure performance standpoint.

Fundamentally, the problem they all have to solve is how to marry their materials to the silicon that makes up the rest of the transistor and the chip substrate. That's a challenge because of the significant difference in the atomic spacing of crystalline silicon and that of germanium and III-V films. Consequently, depositing these materials onto silicon tends to lead to defects that destroy device performance.

But researchers think they've got the problem licked. "I think we have finally bridged that gap," says Waldron.

For Imec, the process for making these nanowire devices draws heavily on its in-house technology for making III-V-based FinFETs. Here, the III-V crystals are grown in V-shaped grooves in silicon, and the majority of the defects you'd ordinarily get are annihilated when the growing crystals meet the groove walls. To change a fin into a nanowire, one part of the fin, composed of indium phosphide, is etched away to leave indium gallium arsenide (InGaAs) nanowires.

Waldron and her colleagues first reported encouraging initial results for their InGaAs nanowires in 2014. They have now improved them by introducing a new process, developed in conjunction with ASM International, for adding an undisclosed material to the gate.

Electrical measurements at an operating voltage of 0.5 volt—which is about two-thirds of that used in circuits made with Intel's most modern process—highlight the superiority of the new gate. Transconductance, which reflects how quickly charge carriers move in the channel, has almost doubled to 2,200 microsiemens per micrometer.

The other key characteristic for assessing the suitability of these transistors is the subthreshold swing. This is related to the switching capability of the device and is governed by the quality of the interface between the transistor's channel and its gate dielectric.

For Imec's devices, the subthreshold swing is 110 millivolts per decade. (In other words, it takes 110 mV at the transistor gate to increase the current tenfold.) "We have to get it down for it to be a true contender" for next-generation CMOS, admits Waldron, who believes progress can be made by reducing the density of defects in the III-V layers. Still, she says, "I think this is a very big step forward for the technology. We are really confident for the future."

Ye's team has realized the far lower value of 64 mV per decade. "It's as good as you can wish," says Ye, who explains that FETs have a fundamental limit of 60 mV per decade.

This excellent result for subthreshold swing is the culmination of several years of hard work, which began with the development of a process for forming a good contact on electron-transporting germanium FETs and then doing so with germanium FinFETs. "Now we have moved even further," says Ye. "At this year's IEDM, we [presented] the first nanowire-based germanium CMOS."

The nanowire devices also yielded a transconductance of 1,057 μS/μm, which is a record for electron-transporting germanium nanowire FETs. These devices were combined with their hole-transporting cousins to form an inverter circuit with a gain approaching that of a silicon nanowire inverter. Although Ye's inverter operated at 1 V, he believes this could be trimmed to 0.5 V by further improving the contacts.

An advantage of an all-germanium approach is that it avoids many complications of crystal growth. Ye's team simply purchased a wafer from the French firm Soitec that had a silicon base bonded to the combination of an insulating oxide and a layer of germanium. Making the nanowires was then just a case of etching material away.

Like Ye and his colleagues, the Singapore team began with a germanium-on-insulator wafer. But they then had to grow several gallium-based semiconductors on it to produce nanowires of indium arsenide and gallium antimonide. In order to bridge the large difference in the atomic spacing between a gallium arsenide layer grown on the wafer and that of the gallium antimonide and indium arsenide layers, they developed a growth process that ensures that every 14th gallium atom at the interface does not bond with antimony.

"At around 10 nanometers from the interface, the quality of the material is very good, so we can grow the...channel layers," says Gong, of the Singapore team. The total thickness for the layers bridging the difference in lattice spacing is just 150 nm, compared with roughly 1 μm for traditional buffer layers, meaning less material, lower costs, and faster production.

The first devices created by the Singapore team produce low currents and high subthreshold swing, but they expect improvements to come along with the use of higher-quality material and a more suitable etching process for the gate.

There is much work to do before any of these technologies are ready for device production in the 2020s. Imec seems to be closest. But "there is work to be done," says Thean. –RICHARD STEVENSON

NEWS

LINKING CHIPS WITH LIGHT

Researchers integrated 70 million transistors and 850 optical components into a silicon processor using standard chipmaking tricks

Computer designers have long fantasized about using light rather than electrons to move data between microprocessors. Such optical interconnects would overcome the bandwidth bottleneck inherent in the wires and take full advantage of the leaps in processor speed. But combining two very different technologies–electronics and photonics–in the same silicon chip has been a high hurdle to overcome.

Now a group of researchers has come up with a way to build huge numbers of silicon transistors and optical components on the same chip, doing so for the first time, without a major overhaul of the chipmaking process. And they used it to build an IC containing 70 million transistors and 850 photonic components, which together provide all the logic, memory, and interconnection functions a processor needs.

In tests that linked two such chips–one acting as the processor, the other as memory–the optical interconnect could pass data at 2.5 gigabits per second in each direction along an optical fiber using a single wavelength of light, supplied by an external laser. Adding bandwidth is as simple as adding more wavelengths. Although they did not test it to this extreme, the chip should have been able to transfer 27.5 Gb/s in each direction through the same fiber.

The engineers involved–from MIT, the University of California, Berkeley, and the University of Colorado Boulder–invented what they call their "zero-change" approach to chipmaking. It relies on the standard CMOS processes used to make today's computer chips–specifically a high-performance process for the so-called 45-nanometer node, which debuted in 2007. "We didn't make any changes to the process," says MIT's Chen Sun.

They started with a silicon substrate; then they added a 200-nm-thick layer of silicon oxide, which acted as an insulator. Over that was the active layer–100 nm of crystalline silicon–plus a 100-nm layer of nitrides and a dielectric coating. The crystalline silicon included a small

LIGHTING THE PATH: Engineers integrated 850 photonic components with silicon transistors.

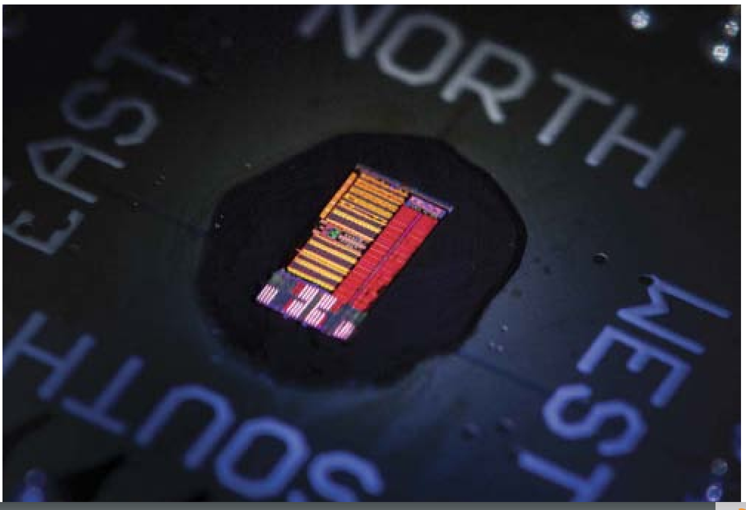

GLENN ASAKAWA/UNIVERSITY OF COLORADO

amount of germanium to produce strain on the silicon and speed up the circuits.

"We are able to use those existing layers to make our processors," Sun says. The processor they constructed was based on a dual-core RISC-V architecture—an open instruction set architecture originally developed at Berkeley. It also included 1 megabyte of static RAM.

Key to the process was that some of the silicon substrate was etched away. Because the oxide layer was so thin, the light would have passed through it to the silicon substrate, which would carry it away. Removing the substrate reduced that loss. The lack of silicon also allowed them to deliver light from an external laser to power the optical components, even though the chip was bonded face down to the printed circuit board. But etching couldn't be done across the entire chip. The team left the silicon in place under the microprocessor and memory, where no light is coming out anyway, so they could attach a heat sink to keep the processor cool.

The real heart of the photonics portion of the chip was the microring resonator, a loop 10 micrometers across that's coupled to a waveguide. They doped the structure with the same elements used to make *p-n* junctions in the transistors, and that action created a notch filter, which passes all incoming light except for a single wavelength. Putting a negative voltage across the junction pushed the charge carriers out of the ring, while a positive voltage returned them, creating a modulator that imprinted digital signals onto the light beam.

While the modulator is the answer to transmitting a light-encoded signal, receiving it and turning it into an electronic signal that the processor can work with requires a photodetector. Here the microring is also important. Normally, a photodetector made of the silicon germanium in the chip would have to be many millimeters to a centimeter long to have a chance of absorbing enough photons to actually detect the light. And that's way too big. But with the microring resonator it can be much smaller, because the light passes through it so many times that the SiGe can absorb it and generate an electronic signal.

Microring resonators have existed for a while, but "a lot of people in industry kind of ignored them," says Sun. That's because as they heat up, the index of refraction shifts and they drift away from the desired wavelength.

The solution was to develop active thermal stabilization. The stabilization system includes a separate photodetector and a digital controller. When the detector notices a change in the amount of photocurrent coming to it, the controller alters the voltage across the microring. This changes how much heat the structure dissipates, pushing its index of refraction back to normal.

Sun says his startup company, Ayar Labs, in Berkeley, hopes to be able to commercialize the technology within a couple of years, but at least one expert is skeptical. Anthony Levi, a professor of electrical engineering and physics in the Photonics Center at the University of Southern California, says the engineers involved in the optoelectronic processor are to be congratulated for making a working chip that integrates photonics and electronics, but he doubts the approach is practical. "The challenges of silicon photonics remain the same as they have always been: This includes too much optical loss, too much power dissipation, too much chip area, and so on," Levi says.

He says the U.S. Defense Advanced Research Projects Agency, which funded the work, has poured "massive amounts of precious research money" into silicon photonics, but the industry and customers, whose decisions come down to cost, haven't bought into it. "Even if the technology worked, there has to be a compelling reason to adopt a new and disruptive approach to building systems," Levi says. —NEIL SAVAGE

At its limits, the chip should be able to transfer 27.5 Gb/s in each direction through the same fiber

MIT'S FOOD COMPUTER: THE FUTURE OF URBAN AGRICULTURE?

Open-source, Internet-assisted farming aims for a new green revolution

Is urban agriculture on the verge of an Internet-enabled revolution? According to a team of technologists at MIT, this unexpected possibility may yet emerge from a series of recent technological breakthroughs. These include the development of high-efficiency blue LED lighting, whose inventors received the 2014 Nobel Prize in Physics.

The MIT researchers say that as technology enables new farming opportunities in indoor, warehouse-based settings, food production can be retooled to accommodate high-density urban living and maintain food security despite a future of increasing climate instabilities and vulnerabilities. The visionaries behind this new farming tech are working toward a networked agricultural system that looks to the open-source software movement for inspiration.

"The current state of [networked agriculture] is very one-off," says Caleb Harper, principal research scientist for the Open Agriculture (OpenAG) Initiative at MIT's Media Lab. "Everybody designs their own little unique 'data center' for plants, for a warehouse farm or for a vertical farm. They say it's super special, and they try to create intellectual property out of it. And they won't let people go in it, because they don't want other people to steal their brilliant ideas. That's the problem in my industry now. They haven't realized that there's a system that underlies this. And it won't scale until there is a common platform."

The common platform that Harper's OpenAG Initiative is developing is an open-source, digitized

THE FOOD COMPUTER: This tabletop machine lets you program the climate you want for your tasty experiments.

MIT MEDIA LAB

food-growing system that his group calls the Food Computer, or FC. As one of Harper's recent papers explains, the FC "creates a controlled environment using robotic control systems and actuated climate, energy, and plant sensing mechanisms. Not unlike climate-controlled data centers optimized for rows of servers, FCs are designed to optimize agricultural production by monitoring and actuating a desired climate inside of a growing chamber."

Sensors inside an FC monitor the growing conditions for the crop and fine-tune the light exposure, temperature, humidity, carbon-dioxide level, water cycle, and nutrient exposure according to a preset recipe for growing the plant. (FC plants do not grow in soil but instead are either hydroponic, with nutrients injected into the water surrounding the roots or misted onto the plant's dangling, open-air roots.) The recipes for each crop, as well as the controlling software and sensing data, will all be freely circulated among Food Computer users for tweaking and improvement—regardless of whether the FC is in a desert, the Antarctic, or on another planet.

Food Computers, as OpenAG conceives them, come in three sizes, with names that reveal Harper's data-center design background. On a recent visit to the Media Lab's fifth-floor OpenAG development space, Camille Richman, a mechanical engineer with the OpenAG Initiative, points out both a tabletop prototype called the Personal Food Computer and a plexiglass-walled walk-in growing module the size of a shipping container, filled with lights, pipes, and racks of plants. "That's the Food Server," she says. "And the warehouse-scale unit, which is the biggest we're looking at, we're calling the Food Data Center." The idea behind making FCs in three sizes is to standardize agricultural tech platforms for personal, small-scale, and large-scale use. What's more, growing conditions suitable to a particular crop that are developed inside a tabletop-size Food Computer can then be scaled up to the shipping-container-size Food Server or even the Food Data Center, which may occupy part of a warehouse, factory, or other large building.

To date, there are fewer than 10 Food Computers worldwide: six Personal Food Computers at MIT and Boston-area public schools, plus two Food Servers—one at MIT and one at Cinvestav (the Center for Research and Advanced Studies of the National Polytechnic Institute) Unidad Guadalajara, in Mexico. The first Food Data Center is scheduled to be built near MIT later this year.

But those numbers will undoubtedly be growing, says Dickson Despommier, emeritus professor of microbiology and public health at Columbia University, in New York City. The computerized vertical farming industry has been booming in recent years and will only get bigger. The expected explosion will come, he says, largely as a result of the phenomenal increase in efficiency that the latest generation of LED grow lights has introduced—from 28 percent efficiency five years ago to 68 percent today. Because grow lights' power usage represents the largest chunk of an indoor farm's operating expenses, Despommier says, the cost savings that the newest lights yield are enough to make vertical farming increasingly economically viable.

Japan, he says, has 145 vertical farms; Taiwan has 45. China just announced an order for 20. And the United States has more than a dozen scattered across California, Delaware, New Jersey, Texas, and Wyoming—and now, of course, in Cambridge, Mass. (Though, to be clear, only Cambridge has any of OpenAG's Food Computer-related hardware or software.)

The time it takes to see a full return on investment "has gone down from 10 or 12 years to between 6 and 8—and in some cases, even 5 years," Despommier says. "The other thing that's been remarkable is, I don't know of any [vertical farming] company in the United States that isn't in expansion mode. That tells you everything." —MARK ANDERSON

NEWS

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

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

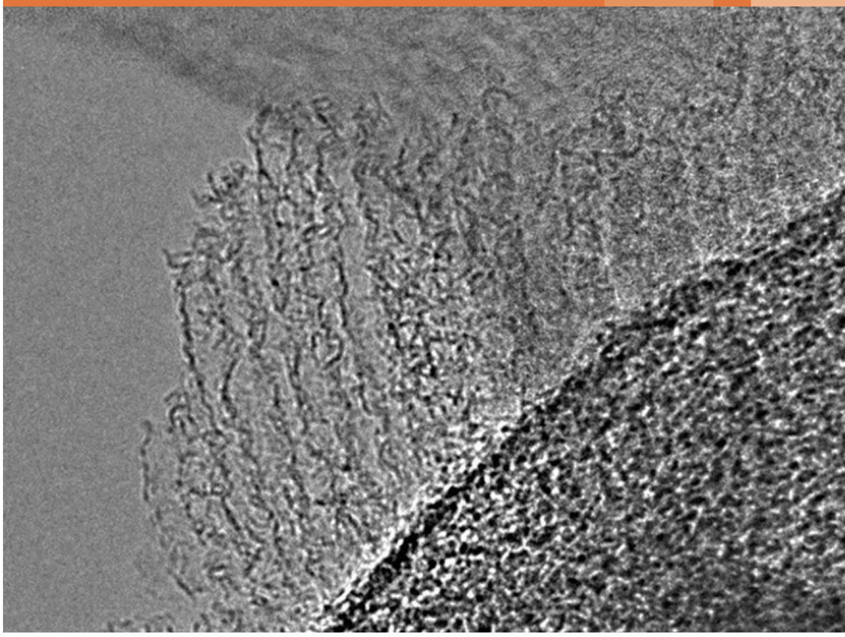

PORES OF POWER! Carbon tubes, laced with the occasional nitrogen atom, make for supercapacitors that store a lot of energy.

CHINESE ACADEMY OF SCIENCES/SCIENCE

NITROGEN SUPERCHARGES SUPERCAPACITORS

Quick-charging devices might finally match lead-acid batteries for energy storage

In lots of situations, the ideal energy storage device is not a battery, which stores lots of energy but can't deliver it particularly quickly. Nor is it a supercapacitor, which has limited storage but delivers what it's got quickly. Instead it would be something that could do both. Scientists in China and the United States recently took a big step toward that ideal component when they showed that nitrogen can triple the energy storage capacity of carbon-based supercapacitors, potentially making them competitive with some batteries in terms of the amount of energy stored.

Most supercapacitors in use today rely on carbon-based electrodes because their large surface area stores more charge. "We are able to make carbon a much better supercapacitor," says Fuqiang Huang, a material chemist at the Shanghai Institute of Ceramics.

Huang and his colleagues began with a framework of porous silica and lined the pores with carbon. They next etched away the silica, leaving porous tubes just 4 to 6 nanometers wide, each made of up to five layers of graphene-like carbon. They then doped the carbon with nitrogen atoms. The nitrogen altered the otherwise inert carbon, helping chemical reactions occur within the supercapacitor without affecting its electric conductivity.

These reactions enhanced the capacitor's ability to store energy roughly threefold without reducing its ability to quickly charge and discharge. Their devices could store 41 watt-hours per kilogram—comparable to what lead-acid batteries can store.

"It is as if we have broken the sound barrier," Huang says.

Other experts agree. "These results are a leap ahead in supercapacitor energy density," says physicist Nunzio Motta at Queensland University of Technology, in Australia, who did not participate in this study.

I-Wei Chen, a materials physicist at the University of Pennsylvania, in Philadelphia, who also worked on the breakthrough, put it in perspective by theorizing what the device could do in electric transportation. "A bus can run on an 8-watt-hours-per-kilogram supercapacitor for 5 kilometers, then recharge for 30 seconds at the depot to run the trip again," he says. "This works in a small city or an airport, but there is obviously a lot to be desired. Our battery has five times the energy, so it can run 25 kilometers and still charge at the same speed. We are then talking about serious applications in a serious way in transportation."

Materials chemist James Tour at Rice University, in Houston, who did not take part in this research, suggested that such advanced supercapacitors could also be useful for wearable devices. "Devices such as watches, glasses, and electronic skin do not favor big, heavy batteries, so a light, efficient supercapacitor could be a good alternative," he says.

The new supercapacitor does not store as much energy as lithium-ion batteries, which achieve 70 to 250 Wh/kg. However, the researchers say this supercapacitor beats them on power, cranking out 26 kilowatts per kilogram, compared with lithium-ion batteries' 0.2 to 1 kW/kg.

Huang and his colleagues are now investigating ways to create these supercapacitors in a scalable, robust, and inexpensive manner, he says. They are also experimenting with a variety of electrolytes to further improve these devices.

Future research could also aim to improve the porosity of the supercapacitors, and thus the amount of charge they can store, Queensland's Motta says. "However, a further increase in the porosity might lead to a very fragile material," he says, "and this could be a stumbling block."

—CHARLES Q. CHOI

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

There's a 5G moment of discovery out there.

We're here to help you find it.

In just a few years, the fifth generation of wireless communications will be a reality. It will allow data to be transmitted up to 100 times faster than today's 4G networks. But getting to that speed won't be easy. You'll need genuine insights to overcome enormous technical challenges. We can help. We have the industry's first and most comprehensive 5G software library. It can significantly streamline design feasibility because it incorporates an iterative design sequence with every piece of Keysight 5G test equipment.

HARDWARE + SOFTWARE + PEOPLE = 5G INSIGHTS

- Designed for testing 5G simulation to verification
- Software platforms and applications that work seamlessly across our 5G instruments
- Incorporate iterative design and rapidly move between stages of your 5G development flow
- Industry's first and largest 5G library

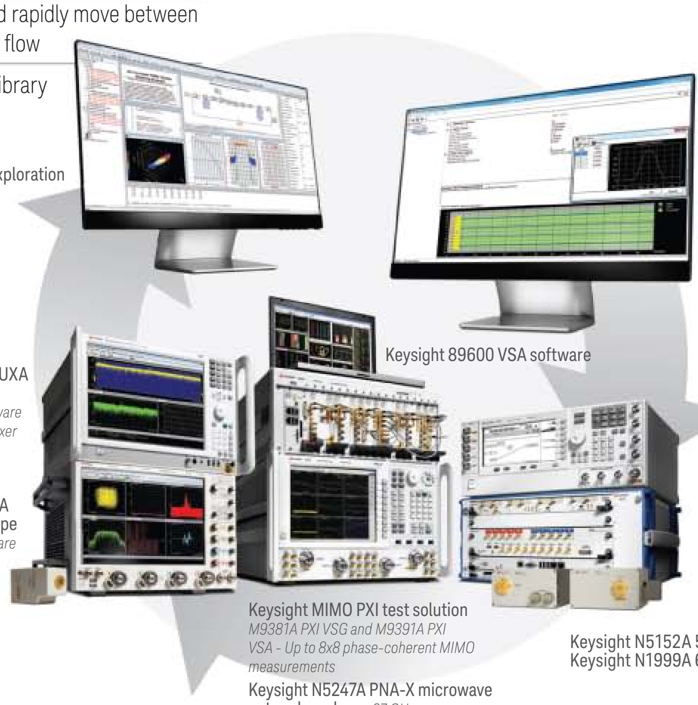

Download our white paper *Implementing a Flexible Testbed for 5G Waveform Generation and Analysis* at www.keysight.com/find/5G-Insight

USA: 800 829 4444   CAN: 877 894 4414

© Keysight Technologies, Inc. 2015

KEYSIGHT TECHNOLOGIES
Unlocking Measurement Insights

ILLUSTRATION BY Rex Features/AP Photo

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

FREIGHT FLYER

ON A TYPICAL DAY, millions of tons of goods are moved across the transportation network that crisscrosses the United States. But how do residents of countries lacking well-maintained roads, railways, and runways get the food, fuel, fabric, furniture, and footwear they desire? Expensively, if at all. London-based Reinhardt Technology Research says it has solved that problem with a novel way to haul freight: a vertical-takeoff-and-landing aircraft built to pick up and deliver a standard shipping container (artist's conception shown here). Electronic sensors and control algorithms allow the aircraft to identify and pick up containers, fly in formation with others of its type to save fuel, and drop off shipments with speed and precision.

THE BIG PICTURE NEWS

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

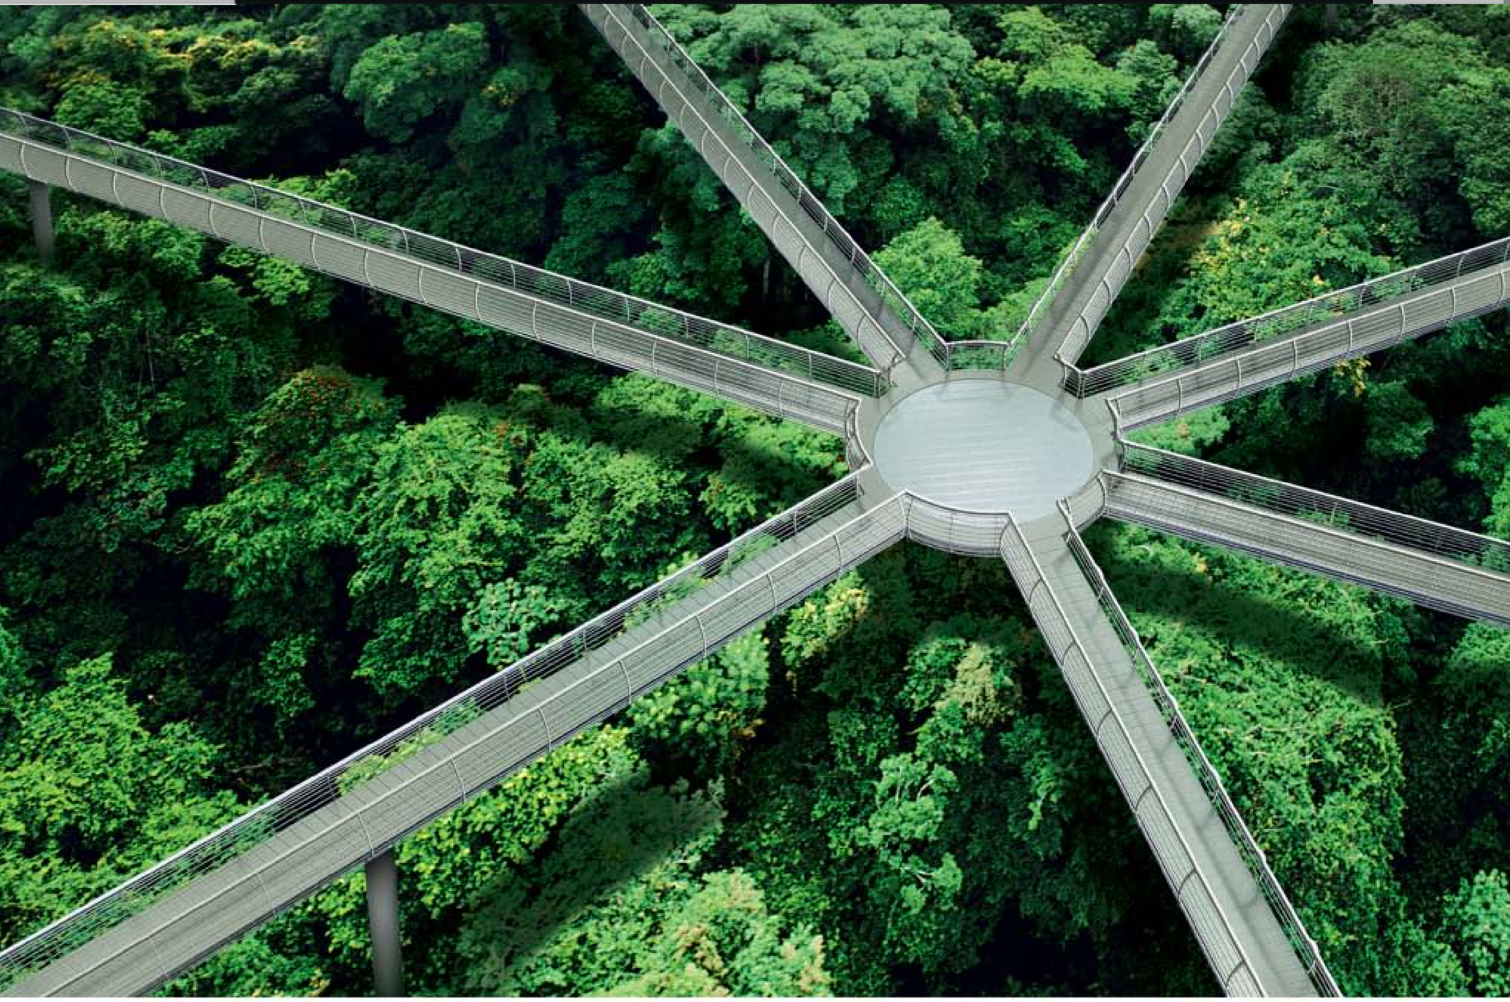

Make the Connection

Find the simple way through complex EM systems with CST STUDIO SUITE

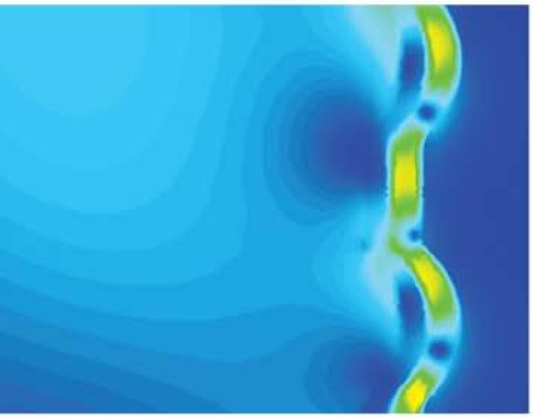

Components don't exist in electromagnetic isolation. They influence their neighbors' performance. They are affected by the enclosure or structure around them. They are susceptible to outside influences. With System Assembly and Modeling, CST STUDIO SUITE helps optimize component and system performance.

Involved in antenna development? You can read about how CST technology is used to simulate antenna performance at www.cst.com/antenna.

If you're more interested in filters, couplers, planar and multilayer structures, we've a wide variety of worked application examples live on our website at www.cst.com/apps.

Get the big picture of what's really going on. Ensure your product and components perform in the toughest of environments.

Choose CST STUDIO SUITE – Complete Technology for 3D EM.

CST

CST – COMPUTER SIMULATION TECHNOLOGY | www.cst.com | info@cst.com

RESOURCES

THE FIRST SYNTHETIC FIBERS WERE MADE FROM A PROCESS ORIGINALLY DEVELOPED TO MAKE LIGHTBULB FILAMENTS.

ANOUK WIPPRECHT
DYNAMIC DRESSES MERGE HIGH FASHION AND TECHNOLOGY

The current wave of wearable computing has mostly been about putting electronic smarts into fashion accessories of various forms: fobs, wristbands, gloves, shoes, and glasses. Meanwhile, the actual clothes we wear have been remarkably untouched by the digital revolution. But Anouk Wipprecht is leading the charge among high-tech designers to change that. She is creating clothes that can sense and respond to the person wearing them and the environment surrounding them. ▸

BOORS, BEWARE: Wipprecht in her Spider Dress 2.0, which extends to defend personal space.

HEP SVADJA

Like most high-end fashion designs, Wipprecht's projects are more about provoking discussion than trying to create something that can be purchased in a store. She is using her conceptual designs to explore what clothing could be, using increasingly elaborate technologies to push the boundaries of clothing's traditional dual roles of protection and display. For example, her Spider Dress 2.0 uses a 3-D-printed skeletal bodice equipped with sensors and motorized "limbs." Invade the wearer's personal space by standing too close too quickly and the limbs will extend to reestablish a comfort zone.

The Spider Dress 2.0 was built in partnership with Intel, which points to the dress as an eye-catching demonstration of the abilities of its new, embedded Edison chip. This is only one of several partnerships Wipprecht has formed with technology companies, as she seeks to expand the range of possibilities for what she calls "fashion-tech."

As a child, she was drawn to fashion, she says: "I found it fascinating how people used fashion to express and communicate their style and identity." By the time she was 14, she was already experimenting with fabrics and design, and just a few years later she was working in the fashion industry of her native Netherlands. "I didn't have rich parents, so I educated myself," says Wipprecht, and without access to expensive materials or tools, she "had to be creative." But she grew discontented: "After a few years, I got bored with the materials because they didn't naturally interact, they didn't do anything…that's why I got into robots."

In particular, Wipprecht was interested in robotics that could be intertwined with the human experience. "I didn't want robots like R2-D2," she explains. "I wanted robots to see like us and feel like us."

As part of her robot experiments, Wipprecht started working with microcontrollers such as the Arduino, and she saw how they could be used to create behaviors. Soon she was bringing her experience back to fashion: "I started to implement little motors into fabrics, which in return started to move and crawl." As her fashion experiments progressed, she says she "started to notice that system behavior is really important, so I studied interaction design." She went to Sweden and studied under David Cuartielles, one of the original creators of the Arduino, to learn more about electronics.

Her designs garnered attention on the international stage, not least when she was commissioned to create an interactive outfit for the lead singer of the Black Eyed Peas for the 2011 Super Bowl halftime show. Wipprecht began incorporating more and more sophisticated sensors and processing into her designs so that they could react to the wearer's brain waves, for example, as with her Synapse dress.

Indeed, Wipprecht has never been content with the technology available to her, which is what has led her to her tech-industry partnerships. She wants her clothes to be smarter, more perceptive, like a second skin acting on behalf of the wearer. She continues to expand her garments' circle of awareness, both in terms of physical distance and depth of perception. Wipprecht is now beginning to explore using machine learning so that her clothes can react in intelligent and subtle ways to the wearer's social experience, and a collaboration with Google is currently in the works.

Wipprecht finds that her keys to success are curiosity, being willing to work across disciplines, and having a clear mission. "Whatever triggers your attention, go for it," she says. "A lot of engineers ask me, 'How do I do design?' I say, just go out and make stuff. And a lot of designers ask me, 'How do I do technology?' Well, just buy an Arduino and play with it. Experiment. Explore. But if you don't have a mission in mind, it's hard to succeed." **—STEPHEN CASS**

FASHIONABLE THOUGHTS: The Synapse dress responds to a wearer's brain waves.

RESOURCES_HANDS ON

BEE COUNTERS
MEASURING A NEST'S OCCUPATION BY ITS CAPACITANCE

BEES ARE IMPORTANT: THEY POLLINATE **dozens of crops, including almonds, cacao, and coffee. While there has been a lot of attention paid to Western honeybees owing to colony collapse disorder, this specific disease and others like it are really measurable only once a colony collapses. And in any case, honeybees are not the only important bee pollinators. What we need is the ability to measure and monitor bee activity as it happens. ● Historically, such monitoring was the purview of undergraduates armed with clipboards. More recently, optical sensors have allowed for the automatic detection of bees entering and exiting the hive. But placing optical sensors in a habitat of pollen, mud, and other hive debris drastically degrades their effectiveness. What if there was a better way? ● A solution suggested itself when the two of us—a field applications engineer for Analog Devices and an amateur bee enthusiast—were working together on a previous project that involved capacitive sensing. Teachman (the bee enthusiast) commented to Perrault (the applications engineer) that the sensitivity of the AD7746 capacitance-to-digital conversion chip was**

LEFT: JASON PERRY; RIGHT: MIKE TEACHMAN

better than he had expected, and wondered, "Do you think we could measure bees with this?" All else being equal, the capacitance between two electrodes depends on the dielectric constant of the substance between them. Air has a dielectric constant of roughly 1, while water comes in at around 80. As living cells are mostly water, a bee should have a detectable dielectric signature. Intrigued by the idea, we developed a custom sensor setup to measure just that.

We concentrated on Mason bees, which are important pollinators (of fruit trees in particular). Unlike honeybees, these bees are solitary types: Every female works alone to build a nest and lay eggs. They build their nests in tubes, such as a reed or a hollow twig, and it's typically one bee to a tube. Starting from the back of the nest, the bee fills the tube with a series of cells, laying one egg in each cell. The egg is then sealed in, along with food in the form of pollen and nectar.

Previously, a group at the University of Prince Edward Island, in Canada, had used a capacitive sensor to measure bees transiting a hive entrance. But we realized that Mason bees' nesting preferences provided the opportunity to monitor activity *inside* their homes, as well as their entrances and exits. We would turn each bee's entire nest into a capacitive sensor.

We placed two 1.27-centimeter-wide strips of copper tape on either side of a 13.5-cm-long, 1.27-cm-diameter acrylic plastic tube and sealed it at one end. We had already connected the copper strips to shielded leads. By instrumenting the entire tube, not just the entrance, we could measure how active bees are within their nests, along with information about the material they bring in to construct cells.

Using some blocks of wood, we housed eight of these instrumented tubes together, and placed Mason bee cocoons on top of the blocks to ensure that the empty nests would be rapidly colonized once the bees emerged.

One AD7746 chip can handle two channels, so we mounted four on a custom breakout board and connected it to an Arduino microcontroller using the I2C serial protocol. We gathered information about local light levels, temperature, humidity, and pressure by using the Adafruit TSL2561, MCP9808, HTU21D-F, and BMP180 sensor boards, which we also connected to the Arduino via I2C. The Arduino logged data and relayed it to an SD memory card every 10 seconds, and we also sent data via a second, Wi-Fi–enabled Arduino to SparkFun's free cloud service. The total cost of the electronics was around US $200.

BEE AWARE

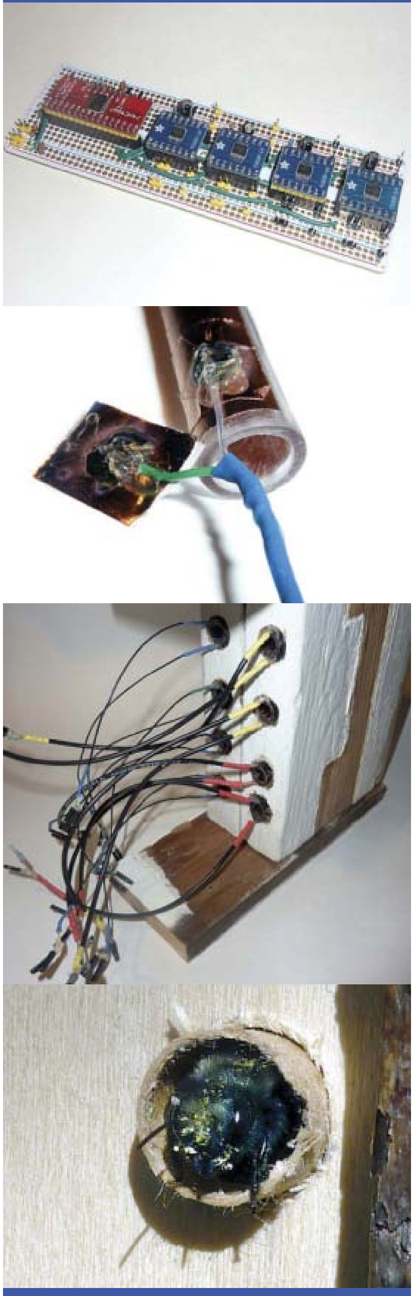

MIKE TEACHMAN (4)

We wired up four AD7746 chips [top] to strips of foil [second from top] in eight Mason bee nests [second from bottom], allowing us to gather data about the bees inside [bottom].

Converting the raw capacitance data into meaningful information about bee activity was the big challenge. For starters, over the course of 24 hours there are large swings in the baseline capacitance of the nests due to temperature and humidity shifts. Rather than trying to work out the exact impact of these shifts on the baseline from our separate temperature and humidity measurements, we simply blocked the entrance of one of the eight tubes to prevent its occupation. We then subtracted the baseline variations seen in the empty tube from the signal received from the other tubes. We also took video of the tubes and time-aligned it with the logged capacitance data to create a "learning" data set. This allowed us to be sure we were recognizing entrances and exits correctly.

As the bee brings material into the nest to build cells, it creates a permanent shift in the baseline capacitance of the cell. Although we've only just begun distilling the data, we believe it should be possible to determine not just the volume of material added to each cell but also details about the types of material involved (mud, eggs, pollen, and so on). As the bee moves about inside the tube, it also causes fluctuations in the capacitance. Together, this allows us to gauge how long the bees spend outside the nest, how active they are when they return, and how this changes depending on the time of day or other alterations in local conditions.

Through more rigorous analysis—perhaps by employing machine learning to process the data sets produced over the course of a year—it should be possible to predict how much pollen each bee is producing and the general health of the colony, and to better understand the dynamics of those important pollinators. During the next growing season, we plan to further develop the technology and see if it can be extended to other types of bees.

—PAUL PERRAULT & MIKE TEACHMAN

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

RESOURCES_TOYS

RETURN OF THE ASCII
FOUR TITLES THAT ESCHEW FANCY GRAPHICS

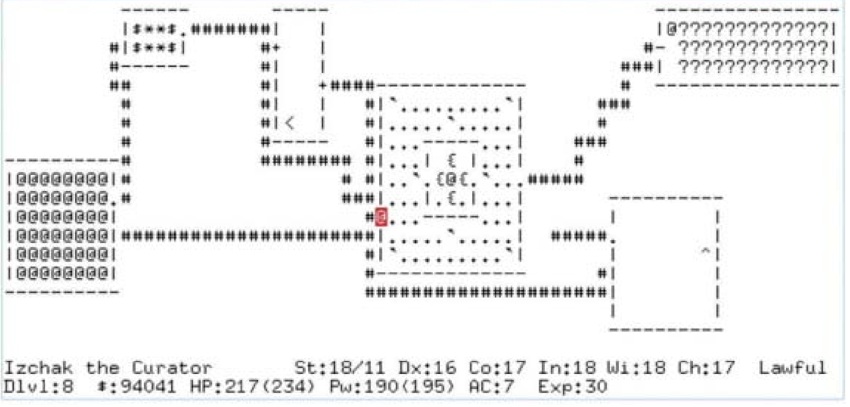

For the truest of nerds, last December brought a wondrous gift. No, not *Star Wars: The Force Awakens*. In fact, this gift is just about the furthest you can get from that movie's hyperrealistic computer-generated imagery and still have something that's considered screen-based entertainment. I am referring to ***NetHack* 3.6.0**, the first significant update in the *NetHack* game series in over a decade.

NetHack is free to download and available on many platforms. The goal is to plumb the depths of a procedurally generated dungeon and recover the Amulet of Yendor. Notably, you, the dungeon, and all the creatures and objects within it are represented by simple ASCII characters such as "%," ">," and "G" (denoting food, a descending stairway, and a gnome, respectively. Some people have created tile-based graphical interfaces to make things a little more intuitive).

The first version of *NetHack* was released in 1987 as a heavily modified descendant of a 1984 game called *Hack*. *Hack* in turn was based on *Rogue*, created around 1980 for Unix computers and the simple, character-based terminals of the day. *Rogue* and *NetHack* inspired the designers of many later games, in particular demonstrating how procedural generation could greatly extend the enjoyment of a game.

THE PURE QUILL: *NetHack* [above] hides complexity behind simple text. That spirit can be found in new games too, such as *Lifeline*.

By accreting ingenious and mischievous contributions from programmers over the years, *NetHack* maintained its popularity thanks to a richness and complexity that belied its simple presentation and hack-and-slash trappings. New versions were regularly released throughout the '80s and '90s. But eventually time seemed to have passed it by as even our phones became crowded with games sporting high-resolution color graphics. Now the series has been revived, officially incorporating many modifications created by die-hard enthusiasts. As a tribute to novelist, satirist, and noted *NetHack* fan Terry Pratchett, who passed away in 2015, the game has also been extensively seeded with quotes pulled from Pratchett's Discworld book series.

But *NetHack 3.6.0* is not the only recent game release to eschew modern 3-D or even 2-D graphics. Here are three other text-centered titles that prize thoughtfulness over pretty pixels.

***A Dark Room*:** Available in both a free online version and as a US $0.99 iOS app, this game has players alternate between reading snippets of text and navigating around a landscape depicted with text symbols, *NetHack*-style. You are given virtually no information at the outset, learning who your character is and what your purpose is only through exploration. The game is surprisingly good at evoking a subtly disquieting atmosphere that compels the player's attention. (A prequel called *The Ensign*, also available, delves more into the player character's back story.)

***Choice of Robots*:** Available in desktop and mobile versions for $4.99, this is not so much a game as an interactive novel, reminiscent of the old Choose Your Own Adventure books but definitely not written for children. You enter the story as a promising graduate student working on an advanced robot; initial choices center on what you value most in its design and how you are going to deal with your annoying Ph.D. supervisor. Soon, the stakes start rising as war looms. As with the best interactive fiction, *Choice of Robots* is very good about putting you on the horns of dilemmas, with downsides and upsides to every choice.

Lifeline... Optimized for use with the Apple Watch but playable with just an iPhone or iPad, this 99-cent game makes the unusual but effective choice of putting you at a considerable remove from the action. Your role is as an advisor for Taylor, an astronaut trainee who has crash-landed on an alien planet. Taylor's only link to the outside world is through text messages, which he uses to describe what's going on and ask questions about what he should do as he struggles to survive. Another unusual but effective game design choice is to have users play the game in real time, so once Taylor goes off to perform some action, he won't message you again till he's done, whether that takes seconds or hours, effectively building tension as you go about your real-life day, waiting for his signal. A sequel is now also available. **—STEPHEN CASS**

TOP: STICHTING MATHEMATISCH CENTRUM AND M. STEPHENSON; BOTTOM: 3 MINUTE GAMES

HEADS IN THE CLOUDS
TECH-INDUSTRY VETS FIGHT CLIMATE CHANGE

MARINE CLOUD BRIGHTENING PROJECT

LAB RATS: Marine Cloud Brightening Project engineers are working toward being ready for a field test within two years.

Silicon Valley is in the grip of the belief that only twenty-somethings are capable of coming up with the big, innovative technological ideas that will change the world. This belief persists, despite a landscape littered with startups that offer nothing more than minor variations on sharing or social apps that don't change the world so much as make certain parts of it slightly more convenient. So it's somewhat satisfying to discover that a Silicon Valley group that *does* have the biggest global problem of today firmly in its sights is composed of veteran engineers—none younger than 70, and the oldest, 80.

The goal is to engineer better clouds. Clouds, by reflecting sunlight, cool the land or sea below them. Whiter clouds reflect more sunlight, limiting the greenhouse effect. So what if you could intentionally make clouds whiter?

The physicist John Latham (later joined by the engineer Stephen Salter) came up with an idea of how that might be done in the early 1990s: by blasting tiny seawater droplets into the air. But could a machine to do this actually be built?

A small group of atmospheric scientists interested in testing the idea invited retired engineer (and IEEE Life Member) Armand Neukermans to meet with them about seven years ago. Neukermans knew a little something about spraying droplets: He helped develop the first generation of inkjet printers at HP. Perhaps a cloud-brightening system could be something like a big inkjet nozzle.

That meeting was the de facto end of Neukermans's retirement. In 2008, he brought what came to be called the Marine Cloud Brightening Project to Silicon Valley, assembled a group of fellow semiretirees, sublet some lab space in Sunnyvale, and started working. The core developers include Gary Cooper, formerly of Syntex and Roche; Jack Foster, an early laser engineer who worked at Sandia, Sylvania, HP, and Xros; semiconductor-manufacturing guru Lee Galbraith, who worked at Sandia and Tencor; Suds Jain, formerly of Synoptics, Bay Networks, Broadcom, and Acterra; and Bob Ormond, formerly of Metara and now with Aqua Metrology Systems. That's just the engineering team. The science part of the project involves collaborators from the University of Washington, the Pacific Northwest National Laboratory, and various collaborators in England.

These engineers are able to tap into long-standing networks of former colleagues and friends. "This is Silicon Valley," Neukermans says, "so when we need to make small holes, we go to the Nanolab at Stanford. We needed a diamond; we went to sp3 Diamond Technologies, and they gave us flat pieces of diamond. We needed laser drilling; we knew where to go for that: DPSS Lasers."

Being senior citizens, though, means they are past the stay-at-the-lab-all-night-drinking-energy-shots stage. "This is a geriatric crew, so we only work four days a week," Neukermans says. "And not everyone always shows up, but they do if they can."

Still, they've managed to get a lot done. The key, they know, is making the droplets extremely tiny, because that way they'll drift up into clouds instead of falling, and they'll be more reflective. Initially, the group did design something like a continuous inkjet printer, with holes one-thousandth the size of a traditional inkjet. But they found that if salty water clogged just one hole, the fluid flowed irregularly into the other holes and caused problems. Now the group is developing a system that works more like a snowmaking machine, but with droplets one-thousandth the size of snowflakes.

They are planning to build five full-size machines for a test run conducted by the University of Washington's Joint Institute for the Study of the Atmosphere and Ocean. It will take US $6 million—including the cost of observational aircraft and the data analysis—to complete the testing. The group aims to raise that money from both private investment and government grants and is looking forward to testing the results of their work outdoors, with actual clouds. "We know this project can't be just old guys in a garage," says Neukermans. "We are going to deliver a prototype."

The engineers don't necessarily expect it to work. Neukermans is at pains to point out that the group is not currently advocating for large-scale cloud-based geoengineering: "We are advocating that research be done." He reasons that even if things go badly, it's valuable to know what doesn't work, as it often leads to what might work. But the team considers it worth a shot. **—TEKLA S. PERRY**

FOOD WASTE

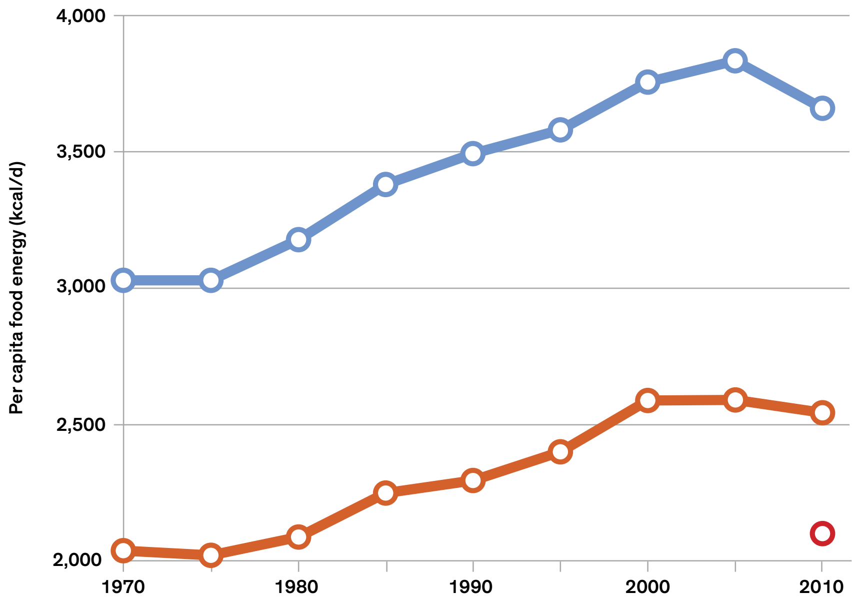

U.S. FOOD WASTE is shown by average per capita food supply [blue], food supply adjusted for waste and spoilage [orange], and actual average daily intake [red].

SOURCES: FOOD AND AGRICULTURE ORGANIZATION OF THE UNITED NATIONS; U.S. DEPARTMENT OF AGRICULTURE; U.S. NATIONAL HEALTH AND NUTRITION EXAMINATION SURVEY

THE FOOD SUPPLY OF THE UNITED STATES COMES TO A little under 3,600 kilocalories per person per day. That's supply, not consumption–and a good thing, too. ● Consider that if you omit babies and housebound octogenarians, whose daily requirements are less than 1,500 kcal, that would leave more than 4,000 kcal available for adults–a quantity appropriate only for a lumberjack. Americans may eat too much, but not that much. The U.S. Department of Agriculture (USDA) adjusts these figures for "spoilage and other waste" and puts the actual daily average at about 2,600 kcal per person. But even that isn't quite right. Both surveys of self-reported food consumption (done by the National Health and Nutrition Examination Survey) and calculations based on expected metabolic requirements indicate that the actual average daily intake in the United States comes to about 2,100 kcal per person. Subtract 2,100 kcal in intake from 3,600 kcal in supply and you get a loss of 1,500 kcal, which means about 40 percent of our food goes to waste. ● This was not always the case. In the early 1970s, the USDA put the average food availability adjusted for preretail waste at less than 2,100 kcal per day, nearly 25 percent less than what it is now. ● The National Institute of Diabetes and Digestive and Kidney Diseases estimates that the United States' per capita food waste increased by 50 percent between 1974 and 2005 and that the problem has gotten worse since then. ● But even if the average daily loss had remained at 1,400 kcal per capita, a simple calculation shows that in 2014 (with 319 million people) this wasted food could have provided adequate nutrition (2,200 kcal per capita) to about 200 million people, or almost exactly the entire population of Brazil, Latin America's largest nation and the world's sixth most populous country.

This food loss inevitably entails a significant waste of energy that is used directly to operate field machinery and irrigation pumps and indirectly to produce the steel, aluminum, and plastics needed to make those mechanical inputs and to synthesize fertilizers and pesticides. The extra agricultural effort also ends up hurting the environment by causing soil erosion, nitrate leaching, the loss of biodiversity, and the growth of antibiotic-resistant bacteria.

Yet even as they waste food, Americans are still eating far more of it than is good for them. The prevalence of obesity–defined as a body mass index of 30 or greater–more than doubled between 1962 and 2010, rising to 35.7 percent from 13.4 percent among adults over age 20. Add to this number the merely overweight (a BMI between 25 and 30) and you find that among adults 74 percent of males and 64 percent of females have an excessively high weight. Most worrisome, as obesity is usually a lifelong condition, that proportion is now above 50 percent for children above the age of 6 as well.

In short, the United States needs to produce considerably less food and consume it with considerably less waste. And yet the mantra of higher food production is chanted as loudly as ever. Its most recent permutation is to produce more by eventually flooding the markets with fake meat made from altered legume proteins. Instead, why not try to find clever ways to halve food waste to a more acceptable 20 percent? Excessive food waste has become the norm in other high-income countries as well (only Japan keeps it at a moderate level), and for different reasons such waste is unacceptably high in many low-income countries with a barely adequate food supply. So cutting that waste in half would lead the way to a more rational use of food worldwide. When will less have more appeal than more? ■

INFOGRAPHIC BY Erik Vrielink

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

WHEN TECHNOLOGY HATES US

Go into any of the little cafés or horlogeries on Paris's Left Bank…and sooner or later you will hear someone say, "*Les choses sont contre nous.*" "Things are against us." —*Paul Jennings*

OUR LIVES ARE IMMEASURABLY EASIER and more connected than we could ever have imagined just 20 years ago. But to reap these benefits, we've had to cozy up to our tech. So close, in fact, that it has become personal, almost a part of who we are. But every now and then, that closeness enables us to see a darker side of the relationship: Sometimes, it seems, our devices are actively working against us. The phone that spontaneously reboots just when you need it most; the computer that crashes only when you have unsaved changes; a bug that suddenly appears when you demo your code. • This tendency for things not just to fail but to seemingly fail *out of spite* was labeled *Resistentialism* by the British humorist Paul Jennings back in the 1960s, but similar terms have been coined in the years since. We are often **FOBIO** (frequently outwitted by inanimate objects) or **FOILED** (frequently outwitted by inanimate, lame electronic devices). The editor and writer Edward Tenner coined the phrase **revenge effect** to refer to an unintended and negative consequence of some new or modified technology: Get everyone on e-mail and the result isn't a paperless utopia but soaring consumption of office paper; give everyone faster Internet connections and the result isn't a boom in leisure time but a dearth of time away from work. Programmers talk of **Heisenbugs**, software errors that disappear or change their behavior when the programmer tries to trace them or examine them (incorrectly named after the physicist Werner Heisenberg, whose Uncertainty Principle states that we can never know perfectly the position or momentum of a particle beyond a certain minimum level, and which is often confused with the *Observer Effect*, where the act of observing a particle alters its properties).

It's nice to have a world of information in your pocket, but make devices too small and you get the **fat-finger problem**: the tendency to make errors on a device where screen elements are too small for human digits. (Scammers are taking advantage of the related problem of **fat-finger dialing** by setting up toll numbers that differ by one digit from popular numbers, so as to glean money from misdials.) The solution is phone software that checks our spelling, but this in turn has created the **autofail**, an indecorous or nonsensical error introduced into a message by such software. A synonym is the **Cupertino effect**, so-named because early versions of Apple's spell-checkers would often change the word *cooperation* to *Cupertino*, the location of Apple's headquarters. The equivalent for an error created by voice-recognition software is the **speako** (a play on *typo*).

The profusion of options offered by a digital technology can make it harder to use. In particular, when it becomes difficult to discern the exact state or mode of the device, the result is called **mode confusion**, which has been observed in situations both serious (airline accidents) and trivial (trying to enter a password with the Caps Lock key inadvertently on).

We live in a world increasingly run by complex systems in which multiple technologies not only interact but work properly only if all the other technologies are working correctly. (Such a system is said to be **tightly coupled**.) Most complex systems have redundancies and fail-safe mechanisms to prevent failures, but the interactions between technologies in a complex system are so, well, *complex* that it isn't possible to predict all the ways any one failure will affect the system. Therefore, problems in these systems are more or less inevitable, a phenomenon called, oxymoronically, the **normal accident**. This isn't strictly Murphy's Law. Instead, it's a variation on the theme: "If something can go wrong, it usually won't, but eventually it will." Things, after all, are against us. ■

ILLUSTRATION BY Richard Mia

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

GREAT LEAPS OF LIGHT

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

Fiber-optic bandwidth has increased exponentially for decades. Can the trend continue?

BY **JEFF HECHT** PHOTOGRAPHY BY **SUE TALLON**

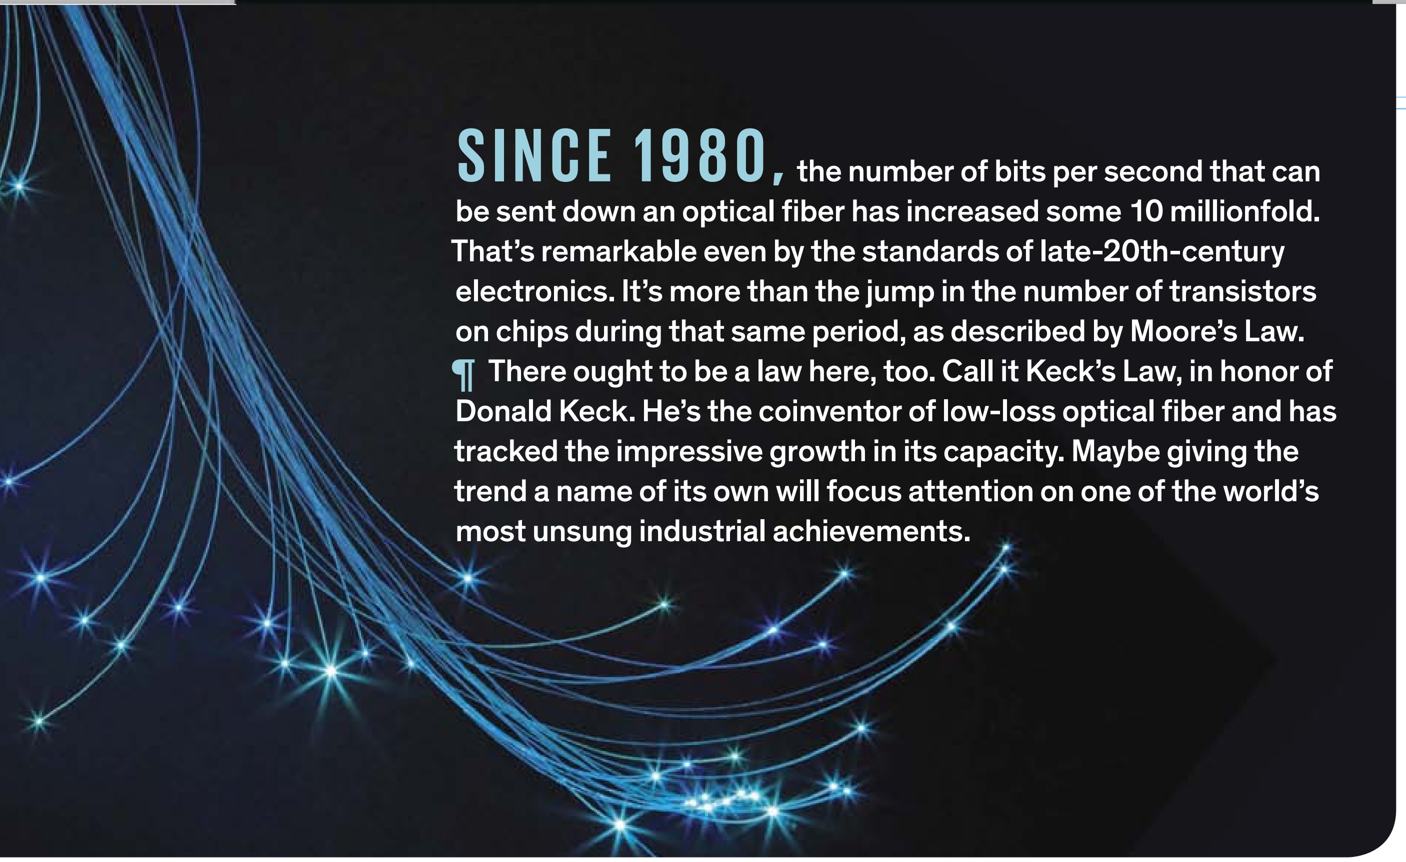

SINCE 1980, the number of bits per second that can be sent down an optical fiber has increased some 10 millionfold. That's remarkable even by the standards of late-20th-century electronics. It's more than the jump in the number of transistors on chips during that same period, as described by Moore's Law. ¶ There ought to be a law here, too. Call it Keck's Law, in honor of Donald Keck. He's the coinventor of low-loss optical fiber and has tracked the impressive growth in its capacity. Maybe giving the trend a name of its own will focus attention on one of the world's most unsung industrial achievements.

Moore's Law may get all the attention. But it's the combination of fast electronics and fiber-optic communications that has created "the magic of the network we have today," according to Pradeep Sindhu, chief technical officer at Juniper Networks. The strongly interacting electron is ideal for speedy switches that can be used in logic and memory. The weakly interacting photon is perfect for carrying signals over long distances. Together they have fomented the technological revolution that continues to shape and define our times.

Now, as electronics faces enormous challenges to keep Moore's Law alive, fiber optics is also struggling to sustain the momentum. For the past few decades, a series of new developments have allowed communications engineers to keep pushing more and more bits down fiber-optic networks. But the easy gains are behind them. To keep moving forward, they'll need to conjure up some fairly spectacular innovations.

The heart of today's fiber-optic connections is the core: a 9-micrometer-wide strand of glass that's almost perfectly transparent to 1.55-μm, infrared light. This core is surrounded by more than 50 μm of cladding made of glass with a lower refractive index. Laser signals sent through the core are trapped inside by the cladding and guided along by internal reflection.

Those light pulses zip down the fiber at a rate of about 200,000 kilometers per second, two-thirds the speed of light in vacuum. The fiber is almost perfectly clear, but every now and then a photon will bounce off an atom inside the core. The longer the light travels, the more photons will scatter off atoms and leak into the surrounding layers of cladding and protective coating. After 50 km, about 90 percent of the light will be lost, mostly due to this scattering.

Communications engineers therefore need to boost the intensity of the light at regular intervals, but this approach has limitations of its own. The interaction between a powerful, freshly boosted signal and the glass in a fiber can cause distortions in the signal that build up with distance, a bit like a haze in the air that obscures distant objects more than nearby ones. These distortions are called nonlinear because they don't double if the intensity of the light is doubled. Instead they increase at a faster rate. When the light is intense enough the distortions will drown the signal in noise. The story of fiber is a saga of finding ways to boost the data rate and the length of transmission despite the scattering and distortion problems.

The very first fiber-optic messages were encoded by simply switching the laser source on and off. Engineers made steady

MODERN FIBER OPTICS: THE BASICS

Electronic signals are converted into optical signals at the source [left] and sent into fiber. Optical amplifiers made from specially doped fiber boost the signal. In

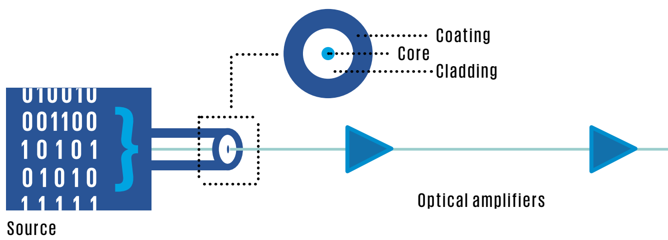

improvements in how quickly that switching could be done. By the mid-1980s, a few years into the dawn of commercial fiber networks, the strategy could be used to send several hundred megabits per second through a few tens of kilometers of glass.

To keep the signal going after the first 50 km, a repeater was then used to convert light pulses into electronic signals, clean them up, amplify them, and then retransmit them with another laser down the next length of fiber.

This electro-optical regeneration process was cumbersome and costly. Fortunately, a better approach soon emerged. In 1986, David Payne of the University of Southampton, in England, showed that it is possible to amplify light directly inside an optical fiber instead of using external electronics.

Payne added a dash of a rare-earth element called erbium to a fiber core. He found that by exciting erbium atoms with a laser, he could prime them to amplify incoming light with a wavelength of 1.55 μm—just the point where optical fibers are most transparent. By the mid-1990s, amplifiers containing erbium-doped fibers were already being installed to stretch fiber transmission distances. Depending on their spacing, a series of amplifiers could relay signals over a distance of 500 to several thousand kilometers before the signals had to be converted to electronic signals for cleaning up and regeneration with more expensive gear. Today, chains of erbium-fiber amplifiers can extend fiber connections across continents or oceans.

The emergence of the erbium-fiber amplifier opened the door to another way to boost data rates: multiple-wavelength communication. Erbium atoms actually amplify light across a range of wavelengths, and can be made to do so quite uniformly from 1.53 to 1.57 μm. That's a wide enough band to accommodate multiple signals in the same fiber, each with its own much narrower band of wavelengths.

THE LIGHT EXPONENTIAL

Fiber-optic capacity has made exponential gains over the years. The data in this chart, compiled by Donald Keck, tracks the record-breaking "hero experiments" that typically precede commercial adoption. It shows the improvement in fiber capacity before and after the introduction of wavelength-division multiplexing [light blue section], mentioned below.

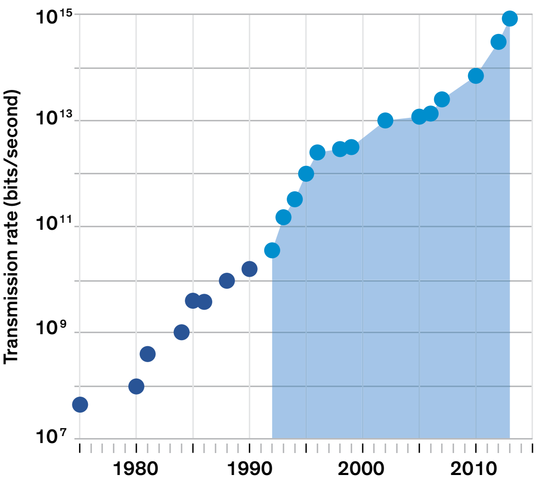

DATA SOURCE: DONALD KECK

This multiwavelength approach, dubbed wavelength-division multiplexing, along with further improvements in how quickly laser signals could be turned on and off, led to an explosion in capacity in the mid- and late-1990s. By 2000, fiber-optic transmission systems were commercially available that could amplify as many as 80 separate signals, each carrying 10 gigabits per second. In reality, nobody needed all that transmission capacity at the time, so transmission systems were installed with only a few wavelengths and the option to add more channels later.

Network operators added more wavelengths to existing fibers as the Internet took off in the early 2000s. But about a decade ago, it became clear that the traditional way of encoding signals was reaching its limit, and that some routes would soon run out of capacity without new technology or more fibers. An on-off signal carries only one bit at a time (if light in a given interval exceeds a threshold power, it generally represents 1; light below the threshold represents 0). The only way to pack more bits per second with this approach is to do what engineers had traditionally done to such signals: Shorten how long each pulse—or lack of pulse—lasts.

Unfortunately, the shorter the pulse, the more vulnerable it becomes to an optical effect called dispersion. This is the same phenomenon that causes prisms to spread light into a rainbow of colors. It arises because the speed of light in glass varies with wavelength. Even a pulse of laser light, whose spectrum is as close to a single wavelength as you can get, will stretch out as it travels through a fiber. And as pulses stretch out, they interfere with one another. The problem gets worse as the data rate increases and the interval between successive pulses gets shorter. The upshot is that a fiber that could carry 10 Gb/s for 1,000 km would carry 100 Gb/s for only 10 km before the signal would need to be cleaned up and regenerated.

Fibers with improved designs were devised to cut down on pulse dispersion, but replacing the existing fiber network would have been prohibitively expensive. And by 2001, overbuilding during the Internet bubble had left behind vast

a switching center, a receiver extracts the signal from its carrier wave and converts it to electronic form. A switch divides and directs that signal into transmitters, which convert their input into an optical signal that is sent along fiber to its next destination.

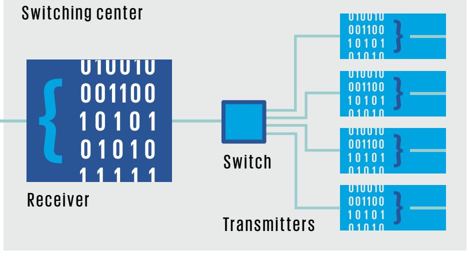

ILLUSTRATION BY Erik Vrielink

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

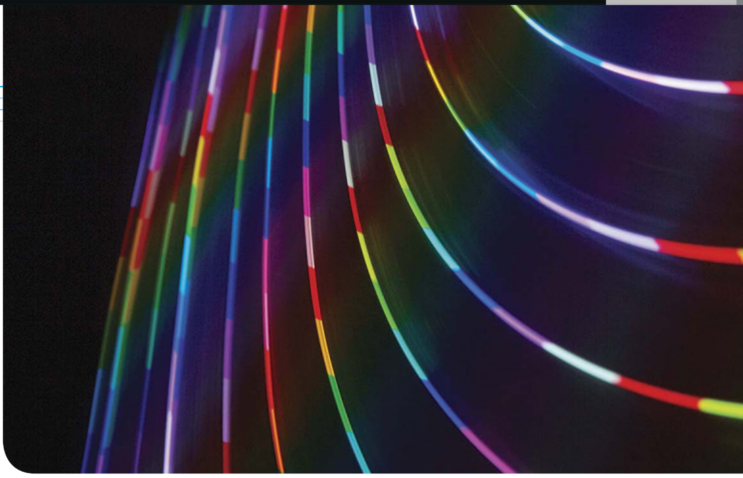

amounts of unused, "dark" fiber. Fortunately, engineers had other tricks, including two techniques that were previously used to squeeze more wireless and radio signals into narrow slices of the radio spectrum.

One was a change to the way signals are encoded. Instead of turning the laser on and off, leave it on all the time and modulate its phase—the timing of the arrival of its peaks and troughs. The simplest such digital phase modulation shifts the peak of a wave by a quarter wavelength, or 90 degrees, ahead or behind the wave's natural arrival time. The wave representing a 1 will then be at its peak when the wave representing a 0 is at its trough. This approach will still yield two bits, but the capacity of the signal can be doubled by combining two waves. Together, they shift the phase in smaller increments, by +135, +45, -45, or -135 degrees. The four resulting states are used to represent the four possible two-bit combinations: 00, 01, 10, and 11.

In 2007, Bell Labs and Verizon used this approach, which is called differential quadrature phase-shift keying, to send 100 Gb/s through some 500 km of fiber in Verizon's Florida network. This was a big deal but still not quite good enough for Verizon, which, like other long-haul carriers, wants signals to be able to travel 1,000 to 1,500 km along the cables in its workhorse backbone system before requiring an expensive regenerator.

Fortunately, a second technique was able to bridge that distance. This one exploits coherence, an intrinsic property of laser light. Coherence means that if you cut across the beam at any point, you'll find that all its waves will have the same phase. The peaks and troughs all move in concert, like soldiers marching on parade.

Coherence can be used to drastically improve a receiver's ability to extract information. The scheme works by combining an incoming fiber signal with light of the same frequency generated inside a receiver. With its clean phase, the locally generated light can be used to help determine the phase of the noisier incoming signal. The carrier wave can then be filtered out, leaving the signal that was added to it. The receiver converts that remaining signal into an electronic form carrying the 1s and 0s of the information that was sent.

QUADRATURE PHASE-SHIFT KEYING

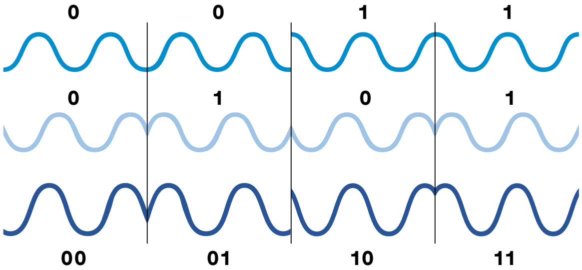

A light wave's phase—the arrival times of peaks and troughs—can be used to encode information. In quadrature coding, four phases [top and middle waves] can be combined to create four different two-bit combinations [dark blue waves at bottom].

Achieving such coherent reception with infrared light was trickier than with radio waves; it was difficult for optical receivers to match the frequency of the input light signal. That changed with the development of advanced digital signal processors in the early 2000s. They allowed the receiver to deal with the mismatch between the local light and the incoming signal, reconstruct the signals' phase and timing, and correct for pulse spreading that occurred en route.

Together, quadrature coding and coherent detection—along with the ability to transmit using two different polarizations of light—have carried optical fibers to their present limit. Today, new transmitter and receiver systems allow a single optical channel—a single wavelength—to carry 100 Gb/s over long distances, in fibers that were designed to carry only 10 Gb/s. And because a typical fiber can accommodate roughly 100 channels, the total capacity of the fiber can approach 10 terabits per second.

Most of the optical fibers installed since 1990 in regional, national, and international cable systems are compatible with this technology, and in the past six years, many backbone networks have been upgraded to carry signals at that rate. "It's out there in pretty robust quantities in long-haul terrestrial transport, and most if not all transoceanic submarine cable upgrades are being done at 100 gig," says Erik Kreifeldt, a senior analyst at the research firm TeleGeography.

For a sense of the numbers, consider a recent fiber system by Ciena Corp., a Hanover, Md.-based company. The system can transmit 96 channels, each carrying 100 Gb/s, across hundreds to thousands of kilometers. All together that amounts to 9.6 Tb/s—enough for 384,000 people to stream Ultra HD from Netflix. And that's just one fiber; today's fiber-optic cables can carry anywhere from about a dozen to several hundred fibers.

But, other than a brief period after the tech bubble collapsed in the early 2000s, the world has never had enough bandwidth. Global Internet traffic increased fivefold from 2010 to 2015, according to a recent report from Cisco. The trend is likely to continue with the growth of streaming video and the Internet of Things.

So developers are considering their options.

One idea is to adopt even more advanced signal-coding techniques. The quadrature phase shifting used today encodes

SOURCE: WIKIPEDIA

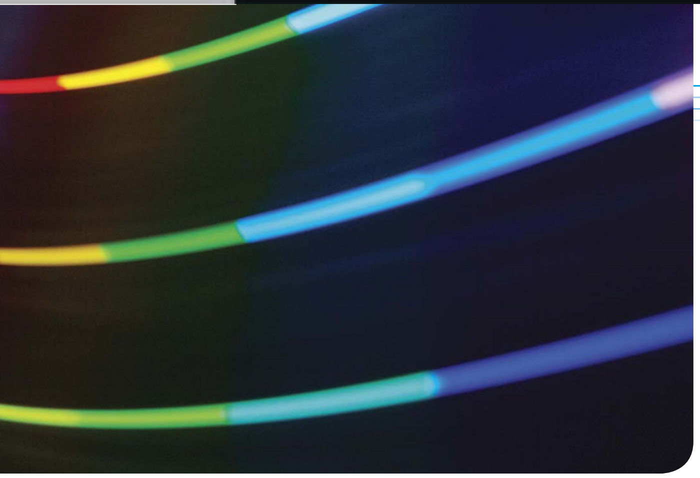

two bits per signal interval, but Wi-Fi and other wireless systems use even more complex coding. The widely used 16-QAM code, for example, can carry all 16 possible combinations of four bits, from 0000 to 1111. Some cable television equipment uses 256-QAM.

Such advanced coding schemes do work in fiber, but as you might expect, there's a trade-off. The more complex the encoding, the closer the information is packed together. The signal can tolerate fewer perturbations before parts of it wind up in the wrong place. Turning up the power can help, but it also creates nonlinear distortion, which worsens with distance. As a result, system makers are generally considering 16-QAM only for relatively short links–of up to a few hundred kilometers.

For longer-haul fibers, engineers have instead devised a way to squeeze channels closer together. There's room to work with: Today's advanced long-haul fibers may contain dozens of channels, but they leave chunks of wavelength unused between adjacent channels to prevent cross talk. If those buffer zones are removed, more channels could be packed into each fiber, creating what system engineers call a superchannel, which transmits at every wavelength inside a fiber-optic band. The change can increase transmission efficiency up to 30 percent, says Helen Xenos, director of product and technology marketing for Ciena.

The trick is to find a way to encode signals so they don't interfere with one another, and at least a few companies have found ways to make it work. In 2013, Ciena and British telecommunications group BT packed multiple channels together without buffers to create an 800-Gb superchannel along a 410-km stretch running between London and Ipswich. At least one Ciena customer, the company says, is in the process of deploying a superchannel system on fibers in a transoceanic cable.

Ciena says it uses individual chips to generate each laser signal. But they can also be combined onto one chip, a more compact and potentially cheaper approach. "Our secret sauce is our photonic integrated circuit technology," says Geoff Bennett, director of solutions and technology at Infinera. In 2014, Bennett says, the company demonstrated a short, 1-Tb superchannel made using 10 laser transmitters incorporated into a single photonic integrated circuit. He says future systems should be capable of taking fiber capacity to 12 Tb/s along long-haul networks–and twice that rate for shorter systems used in metropolitan areas.

Such 12-Tb/s superchannels are still a few years away. But when they arrive, they'll likely be the last capacity boost that we can give to the current generation of installed fibers. That's because those fibers will be approaching a fundamental barrier called the nonlinear Shannon limit. It's an extension of a limit, described in 1948 by information theorist Claude Shannon, that says a transmission channel can carry only so much information without error given its bandwidth and signal-to-noise ratio. The nonlinear version includes an additional factor: a limit on how much the power of a signal can be turned up before nonlinear effects that can arise in glass–but not in air–generate enough noise to drown out the signal.

There is no getting around the nonlinear Shannon limit. But when the time comes to install more fibers, carriers will have other options. "The change that is best established and best understood," says Infinera's Bennett, is to simply introduce fiber with a larger core. Early fibers were designed with small cores, which strongly limited the number of paths light could take. Using a smaller core helped prevent photons in the signal from bouncing off the core-cladding interface at different angles. If the photons in a pulse did that, they'd take different paths–some longer, some shorter–spreading out the pulse so that it would interfere with the next one.

New fiber designs use novel core microstructures, such as photonic crystals, to constrain the light to follow the same path through cores with up to about twice the cross-sectional area as the standard 9-μm fiber. Because the signal has more space, cross-section-wise, to pass through, its energy density is lower. This decrease in energy density cuts down on those same nonlinear distortions that limit transmission distances and speeds. The end result is an increase in data rates; future versions might boost | CONTINUED ON PAGE 47

À LA MODE

Signals with different "modes"—different spatial profiles—can be sent into the same fiber to boost capacity. The intensity of light in a mode varies from point to point in its cross section. Here is a sampling [top] of modes as they might appear when introduced into a fiber. The bar at the bottom shows how one mode might oscillate between different states as it moves along the fiber.

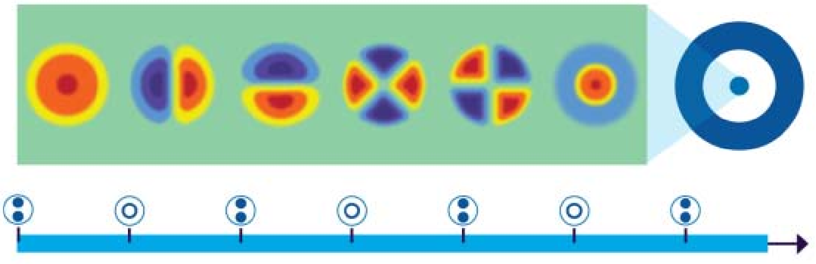

ILLUSTRATION ADAPTED FROM P.J. WINZER, BELL LABS TECHNICAL JOURNAL, VOL. 19, PP. 22–30 (2014)

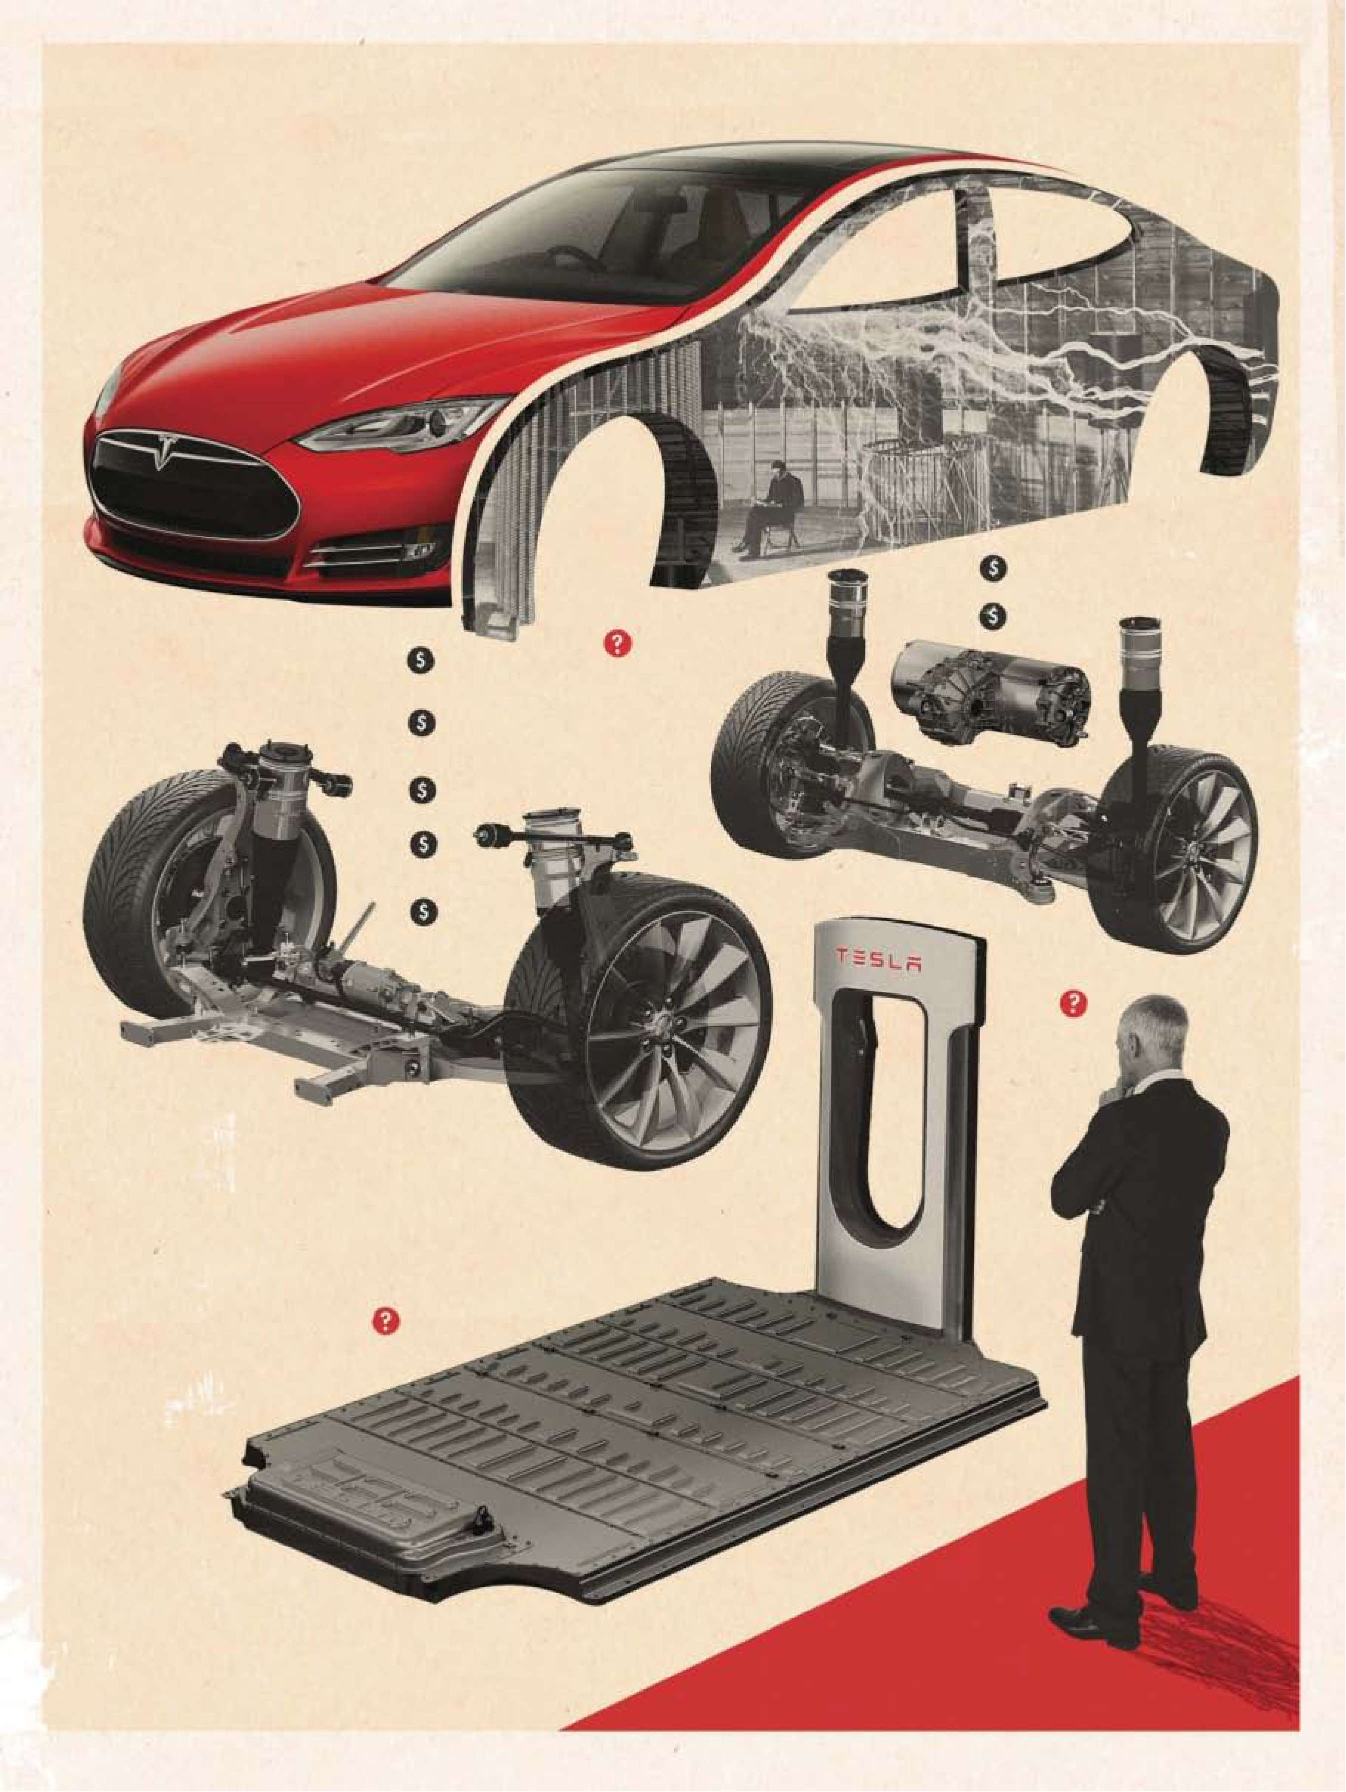

ORIGINAL PHOTOGRAPHS: NIKOLA TESLA: DICKENSON V. ALLEY/WELLCOME IMAGES; TESLA PARTS: TESLA MOTORS

PHOTO-ILLUSTRATION BY Lincoln Agnew

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

?

?

a TESLA *in* EVERY GARAGE?

$

$

$

HOW NOT TO SPARK AN ELECTRIC-VEHICLE REVOLUTION

By Matthew N. Eisler

LAST SUMMER, AS I DROVE AROUND the San Francisco Peninsula, I caught glimpses of a sea change in American automobile culture. Plug-in electric vehicles and charging stations seemed to be everywhere. Near the entrance to Stanford University, I witnessed a three-car fender bender involving only electric cars. And perhaps most remarkable: the prevalence of the Tesla Motors Model S, a luxury electric sedan that's become the new status symbol among the tech-savvy. With more than 90,000 built to date, the Model S is now a common sight on Bay Area highways and byways. ¶ The Model S and Tesla's newer Model X sport-utility vehicle are unquestionably engineering marvels. And in the vision of Tesla CEO Elon Musk, they are nothing less than agents of change. Rivaling the best-performing internal-combustion automobiles, these green supercars are designed to reconcile comfort, power, and convenience with environmental sustainability, an ethos encapsulated in the current marketing slogan for the Model S: "Zero Emissions. Zero Compromises." ¶ Sales of the high-end models S and X—versions of which sell for well over US $100,000—are supposed to provide the revenue needed to produce the Model 3, an affordable electric supercar for the masses, slated to be unveiled in March and in production in 2017. Eventually, as they and other electric vehicles proliferate and get tied into

a SHORT HISTORY of EVs

A combination of regulation and innovation has revived prospects for the electric car

1832–1839: Robert Anderson invents an electric carriage powered by nonrechargeable primary cells.

1900: Of 4,192 cars produced in the United States, more than a quarter are powered by electricity.

1901: Thomas Edison patents rechargeable batteries made from nickel-zinc and nickel-iron.

1920s–1930s: Most electric vehicles vanish from U.S. roads, replaced by tens of millions of cheap internal combustion engine vehicles.

1997: Toyota introduces Prius hybrid electric vehicle, which uses a NiMH battery and a gasoline engine.

2000: GM begins recalling its EV-1s and destroys nearly all of them.

2003: Silicon Valley engineers Martin Eberhard and Marc Tarpenning found Tesla Motors to produce high-end battery electric cars.

2004: Elon Musk invests in and becomes chairman of Tesla Motors.

2006: Tesla unveils its Roadster, with a base price of US $98,000, equipped with a battery pack using repurposed lithium cobalt oxide cells.

2008: China's BYD Auto Co. releases the F3DM, the world's first commercial plug-in hybrid.

2009: Nissan unveils all-electric Leaf, equipped with a lithium manganese oxide battery pack.

2010: Chevrolet releases the Volt plug-in hybrid, equipped with a lithium manganese oxide battery pack from LG Chem.

2012: Tesla begins selling the Model S full-size luxury electric vehicle, equipped with a battery pack using lithium nickel cobalt aluminum cells produced by Panasonic.

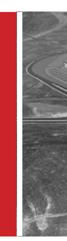

rooftop photovoltaic and energy storage systems, they will disrupt Rust Belt car manufacturing as well as the fossil-fuel industry, accruing environmental benefits in the process. As Musk sees it, Tesla Motors is not simply an automaker. It is an "energy innovation company," a critical element in its broader quest for "zero emission power generation."

Skeptics in financial and environmental circles have questioned the ability of the Bay Area startup to deliver this future, as well as the whole idea that the car market can be substantially "greened." I won't rehash their arguments here. I'm a historian of science and technology, and I do believe that Tesla is making history, of a sort. I spent last summer talking with electric vehicle owners about cars, class, and politics. And what I learned suggests that the battery electric vehicle, or BEV, represents a more thorough upsetting of the existing order of things than Musk and his acolytes might like to admit.

ELECTRIC CARS HAVE BEEN AROUND as long as gasoline-powered ones. But their strong resurgence in recent decades is due almost entirely to California's air-quality politics, which led to the creation of the Zero Emission Vehicle (ZEV) mandate in 1990.

Promulgated by the California Air Resources Board (CARB), the state's "clean air" agency, the mandate was the first such initiative in the world, and it had international repercussions. It compelled the seven major automakers then doing business in California (Chrysler, Ford, General Motors, Honda, Mazda, Nissan, and Toyota) to start producing vehicles that emitted no harmful tail-pipe pollutants, with a rolling quota for each manufacturer based on its market share of conventional vehicles. The larger the share, the larger the ZEV quota. Specifically, CARB mandated that a carmaker's California fleet consist of "at least 2 percent ZEVs [in] each of the model years from 1998 through 2000, 5 percent ZEVs in 2001 and 2002, and 10 percent ZEVs in 2003 and subsequent model years."

The board was forbidden from specifying the kinds of technology that carmakers could use to achieve the desired air-quality outcomes, but the only technology that then met the ZEV criteria was the battery electric car. The car industry, uninterested in shaking up existing product lines, responded with something like this: *Consumers want an*

CLOCKWISE FROM TOP: SMITHSONIAN (2); TOYOTA; GM; NISSAN; TESLA

1959: Allis-Chalmers tests a fuel-cell electric tractor, first vehicle of any kind to employ such a power source.

1964: General Motors tests the Electrovair, a battery electric concept vehicle. Later, it tests the Electrovan fuel-cell concept vehicle.

1966: Britain's Electricity Council deploys the Enfield 8000, an experimental battery EV equipped with lead-acid batteries.

1979: John B. Goodenough and a team at the University of Oxford invent lithium cobalt oxide cathode.

1982: Stanford R. Ovshinsky develops nickel metal hydride (NiMH) battery and founds Ovonic Battery Co.

1990: GM unveils the Impact, an all-electric concept car; California's Zero Emission Vehicle (ZEV) mandate is implemented.

1994: Daimler-Benz unveils the NECAR 1 hydrogen fuel-cell electric concept vehicle.

1996: To satisfy ZEV quotas, GM introduces all-electric EV-1, for lease in California and Arizona.

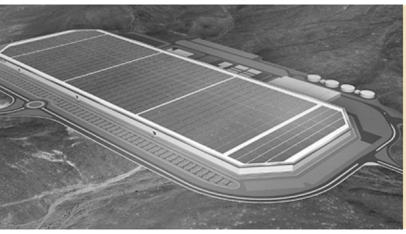

2014: Tesla selects site near Reno, Nev., for Gigafactory 1, where it will produce lithium-ion battery packs with Panasonic.

2015: Tesla begins selling Model X crossover electric SUV.

First quarter 2015: Major car companies purchase $51 million in ZEV credits from Tesla.

March 2016: Planned unveiling of Tesla's moderately priced Model 3 electric car.

CLOCKWISE FROM TOP: GM; GEORGIOS MICHAEL; TESLA (2); GM; DAIMLER; GM

automobile that does it all—one that is clean but also as fast, long-ranged, stylish, and convenient as a good gasoline vehicle. Given the state of the art of electric vehicles, this is impossible. We will roll out advanced zero-emission technology eventually, but it will take time.

This was a straw argument. Automakers had never before let the quest for perfection get in the way of marketing a merely good or just plain awful gasoline automobile. Nevertheless, the argument worked: In 1996 CARB weakened the mandate, eliminating the quotas for the years 1998 to 2002 (which it later extended by a year). In return, the industry agreed to deploy some BEVs ahead of schedule, but in very small numbers and only for lease.

Even as the industry dragged its collective heels on the battery EV, it continued to push for alternatives to satisfy the ZEV category—in particular, hybrid electric and fuel-cell electric vehicles, which the board agreed to recognize in 1998. The inclusion of hybrids enabled portions of an automaker's ZEV quota to run on gasoline. Eventually, CARB created the oxymoronic ZEV subcategory of "partial" zero-emission vehicles (PZEVs), which implied hybrid technology.

Hydrogen fuel-cell cars, by contrast, would emit no tail-pipe pollutants. Ever since its invention in the mid-19th century, the fuel cell had impressed technologists as a kind of miracle engine that combined the best features of the battery and the internal combustion motor. By electro-oxidizing gaseous and liquid fuels stored in conventional tanks, fuel cells would enable electric vehicles to use the vast and existing fossil-fuel infrastructure and drive farther than battery-powered EVs.

At least that was the theory. In fact, researchers were—and still are—struggling to make the technology practicable. Yet the promise of a fuel-cell super EV was enough for CARB to further revise the mandate. In a series of increasingly complex equivalence formulas, the board allowed automakers to gain ZEV credits through a mixture of battery, fuel-cell, and hybrid electric vehicles. It also allowed carmakers to trade ZEV credits and bank them.

The ZEV mandate set the stage for General Motors' EV-1. The company had unveiled an all-electric concept car called the Impact at the January 1990 Los Angeles Auto Show. But the Impact had been commissioned by then CEO Roger Smith for internal political reasons rather than out of any serious desire for large-scale production, according to Michael Shnayerson in his classic 1996 book, *The Car That Could: The Inside Story of GM's Revolutionary Electric Vehicle*. Smith quickly realized he'd let a genie out of a bottle and went into damage control, publicly pointing out the limitations of existing rechargeable batteries. But with ZEV quotas to fulfill, GM introduced the Impact to the market in 1996 as the EV-1, for lease in California and Arizona.

Now, CARB didn't require GM or any other carmaker to maintain its leased fleets of such first-generation "compliance cars." And so the automakers didn't, instead choosing to repossess and decommission them. GM sent most of its 1,117 EV-1s to the crusher, an event that outraged popular sentiment, as documented in the film *Who Killed the Electric Car?*

BIG BATTERY BET: The Gigafactory 1, now under construction near Reno, Nev., will produce lithium-ion battery packs for Tesla Motors' mass-market electric car, the Model 3, as well as Powerwalls, the company's energy-storage units for the home.

DAVID CALVERT/WASHINGTON POST/GETTY IMAGES

It was as a direct response to this fiasco that Silicon Valley engineers Martin Eberhard and Marc Tarpenning founded Tesla Motors in 2003. In crafting their business model, they embraced the auto industry's rationale for *not* producing the BEV: that existing batteries were inadequate and that consumers would not buy electric cars that did not deliver pulse-elevating performance. Accordingly, they decided to produce a high-end sports car—the Roadster—equipped with a lithium-ion battery, the most powerful available chemistry and also the most promising.

In this way, the green supercar myth, invented by the automaking establishment to maintain the status quo, became the basis of a business plan to undermine that establishment.

JOINING TESLA MOTORS IN 2004, Musk kept the supercar philosophy alive. As a business model, it made sense. For an upstart company trying to overthrow the status quo in an industry with enormous barriers to entry, the hype associated with supercars was and is essential to Tesla's ability to attract investment capital.

In important ways, however, the supercar mythos has obscured the technological and industrial implications of battery-powered electric vehicles. Observers tend to focus on battery cost per kilowatt-hour as the most important factor in commercializing the electric vehicle. Indeed, Tesla is devoting a vast effort to drive costs from upwards of $400 per kilowatt-hour to below $200. With plenty of help from the state of Nevada and in partnership with Panasonic, it is now building an enormous $5 billion battery-pack assembly plant (the so-called Gigafactory 1) near Reno.

Large battery packs do represent a significant share of the cost of an electric vehicle. In the highest-performance versions of the Model S, these packs account for up to 30 percent of the suggested retail price of $95,000 to $138,000.

But this obsession with cost per unit energy obscures the real challenge with EV batteries: battery aging. While increasing the capacity and power of batteries has been relatively straightforward, making batteries last longer has proved a hard and incremental slog. Battery lifetime in cellphones, tablets, and other portables is less of an issue because portable consumer electronics are designed to last only a few years. Where electric vehicles are concerned, however, lifetime is crucial.

Evaluating the useful life span of a battery is an imperfect science. The basic metric is the number of charge-discharge cycles, which can mean very different things depending on how and where the battery is used. Designers of most contemporary EVs favor lithium-ion chemistry, and in particular the variants that are optimized for power and energy. Such batteries tend to have a relatively short shelf life, age rapidly when exposed to high temperature and voltage, and get worked harder than the batteries in cellphones or laptops. Some types of lithium-ion batteries are also highly flammable and require expensive control systems to be used safely.

EV batteries are, of course, rigorously tested in the lab. But real-world driving presents a much more complex set of conditions, and assumptions about battery lifetimes are often disproved in the field. Some Nissan Leaf owners, for instance, were dismayed to learn that driving their cars a lot caused the battery capacity to decline far more rapidly than they expected.

What is certain is that even the best batteries will wear out long before the electric motors they serve. Which means that battery-powered EVs have hefty replacement costs that consumers may or may not even realize, let alone be willing to pay. Consider as well that carmakers typically do not produce their own advanced battery cells (although a few, like Tesla and BMW, do assemble such cells into battery packs). Auto executives no doubt loathe the idea of creating an electric fleet whose chief component is made not by them but by dedicated manufacturers like LG Chem or Panasonic. In the world of vertically disintegrated manufacturing, parts suppliers get paid first.

Why don't automakers simply build their own EV batteries? For one thing, established battery makers have economies of scale, often in multiple lines of business; trying to catch up would require Herculean effort. Only Nissan has taken this route, and then only in collaboration with the Japanese electronics giant NEC. Even there, as demand for the Nissan Leaf has waned, the automaker has been looking to cut its costs by switching to South Korea's LG Chem, which makes the battery pack for the Chevrolet Volt plug-in hybrid electric.

Indeed, most new EV users are choosing to lease rather than buy, in part because automakers are promising cheaper and more capable batteries in future versions. But that tactic may backfire: Subsidizing the cost of battery aging in this way may keep consumers happy, but it will doubtless erode the carmakers' ability to profit from EVs.

IT'S NO SURPRISE, THEN, that most mainstream automakers have been reluctant to develop battery-powered all-electric vehicles. An alternative, less risky approach to EV technology is the hybrid, which combines a small internal combustion engine with a small battery pack.

The only major automaker thus far to bet big on hybrids is Toyota. It selected a relatively small nickel metal hydride battery, a reliable, safe, and reasonably powerful chemistry pioneered by the U.S. inventor Stanford R. Ovshinsky. In the first-generation hybrids, the battery was used in a power-assist role; unlike the battery in an all-electric car, it was rarely if ever deeply discharged, a practice that shortens battery life span. A hybrid battery typically doesn't last as long as the automobile itself, but replacement costs average less than $3,500 including labor. As a result, the original Prius, which Toyota introduced in 1997, had a life cycle resembling that of a conventional gasoline automobile.

You can't say the same for the battery-powered all-electric car. Power packs for the Model S, which currently come in 70-, 85-, and 90-kilowatt-hour versions, are technological wonders, bolstered by an eight-year unlimited-mileage warranty. The battery may indeed still work after eight years, but it will do so at a significantly reduced capacity. Most drivers won't be too happy to have their driving range erode by 25 percent or more. And the sustainability of such warranties is also an open question, given Tesla's dependence on generous but temporary public subsidies. These include the $7,500 U.S. federal tax credit to buyers (available for all makes of electric vehicles, up to 200,000 cumulative unit sales), a host of state subsidies, and the ability to earn cash by selling ZEV credits to its competitors.

All this uncertainty over battery lifetimes has had an interesting effect on the used EV market. Lower-end models depreciate faster than gasoline vehicles and are now available at bargain prices. A 2012 Nissan Leaf that sold for around $36,000 retail was worth a little over $8,000 on average in 2015.

The Model S, on the other hand, is retaining most of its value, thanks to Tesla's resale guarantee, which encourages customers to trade in after three years. The success of this approach has had the unintended consequence of generating a large number of very expensive used EVs. When I visited the Tesla plant in Fremont, Calif., last June, I was struck by the huge lot devoted to surplus supercars. It made me wonder: Why would a reasonably well-informed consumer pay $65,000 to $100,000 for a used Model S with a limited warranty when she could plunk down a few thousand more for a brand-new one with a full warranty?

WHILE BATTERY-LIFETIME ECONOMICS tends to be glossed over in public discussions of electric cars, the question of charging infrastructure is robustly engaged. Here, too, however, confusion and paradoxes abound, in good measure because of the supercar ideology. Just like the assumption about EVs themselves, the assumption here is that users must have an experience that replicates the convenience of gas stations as much as possible. So charging means fast charging–the ability to replenish a battery to 80 percent of capacity in 30 minutes or less. This is the approach embraced by Tesla, which has built a proprietary network of supercharging stations capable of adding 270 kilometers of range to a vehicle in a half hour, for free.

Because the large-battery Model S and Model X are so long-legged, this network is located primarily along intra- and interstate corridors in the United States (plus a few in Canada). And because it can be used only by Tesla cars, EV boosters and builders complain that it has stymied the popularization of electric cars. Indeed, Tesla's supercharging is one of three largely incompatible fast-charging standards in the United States (the others are CHAdeMO and SAE). Some critics have called out the federal government for its failure to support fast charging while simultaneously subsidizing the production of electric vehicles.

Yet the fact is that all EV charging systems require changes in consumer behavior. EV drivers have to learn to plug in their cars at home overnight or at other points of access during the day; they may also have to hire an electrician to reconfigure the receptacles in their garages. Even a high-amperage Supercharger requires the user to wait 30 minutes or more for a "refueling."

California currently has more nonresidential charging stations than any U.S. state–more than 5,700 as of August 2014. But many of these stations are clustered at a few high-profile workplaces and are virtually nonexistent in apartment complexes and commercial parking lots. This was a frequent complaint among the EV owners I interviewed. Pacific Gas & Electric Co., the major electric utility in northern California, plans to install an additional 7,500 level 2 (that is, 240-volt AC) chargers, at a cost of $222 million to taxpayers.

But the reality is that in some places, charging infrastructure for the masses is hiding in plain sight. Back before direct-ignition and fuel-injection devices improved the starting capabilities of gasoline engines, drivers in colder climates often had to warm the engine blocks of their cars. As a result, 120-volt AC sockets were widely installed in public parking lots in places like Alaska, Canada, and Scandinavia. Much of this Frostbelt infrastructure still exists, and similar plug-in spots could be built quickly and cheaply elsewhere as a cost-effective way of stimulating affordable | CONTINUED ON PAGE 49

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

Playing the Imitation Game With Deep Learning ⟷ By Mimicking People, Recurrent Neural Networks Gain Some Amazing Abilities ⟷ By Zachary C. Lipton & Charles Elkan

"On tap at the brewpub. A nice dark red color with a nice head that left a lot of lace on the glass. Aroma is of raspberries and chocolate. Not much depth to speak of despite consisting of raspberries. The bourbon is pretty subtle as well. I really don't know that find a flavor this beer tastes like. I would prefer a little more carbonization to come through. It's pretty drinkable, but I wouldn't mind if this beer was available."

Besides the overpowering bouquet of raspberries in this guy's beer, this review is remarkable for another reason. It was produced by a computer program instructed to hallucinate a review for a "fruit/vegetable beer." Using a powerful artificial-intelligence tool called a recurrent neural network, the software that produced this passage isn't even programmed to know what words are, much less to obey the rules of English syntax. Yet, by mining the patterns in reviews from the barflies at BeerAdvocate.com, the program learns how to generate similarly coherent (or incoherent) reviews. ⟶

ILLUSTRATION BY John Hersey

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

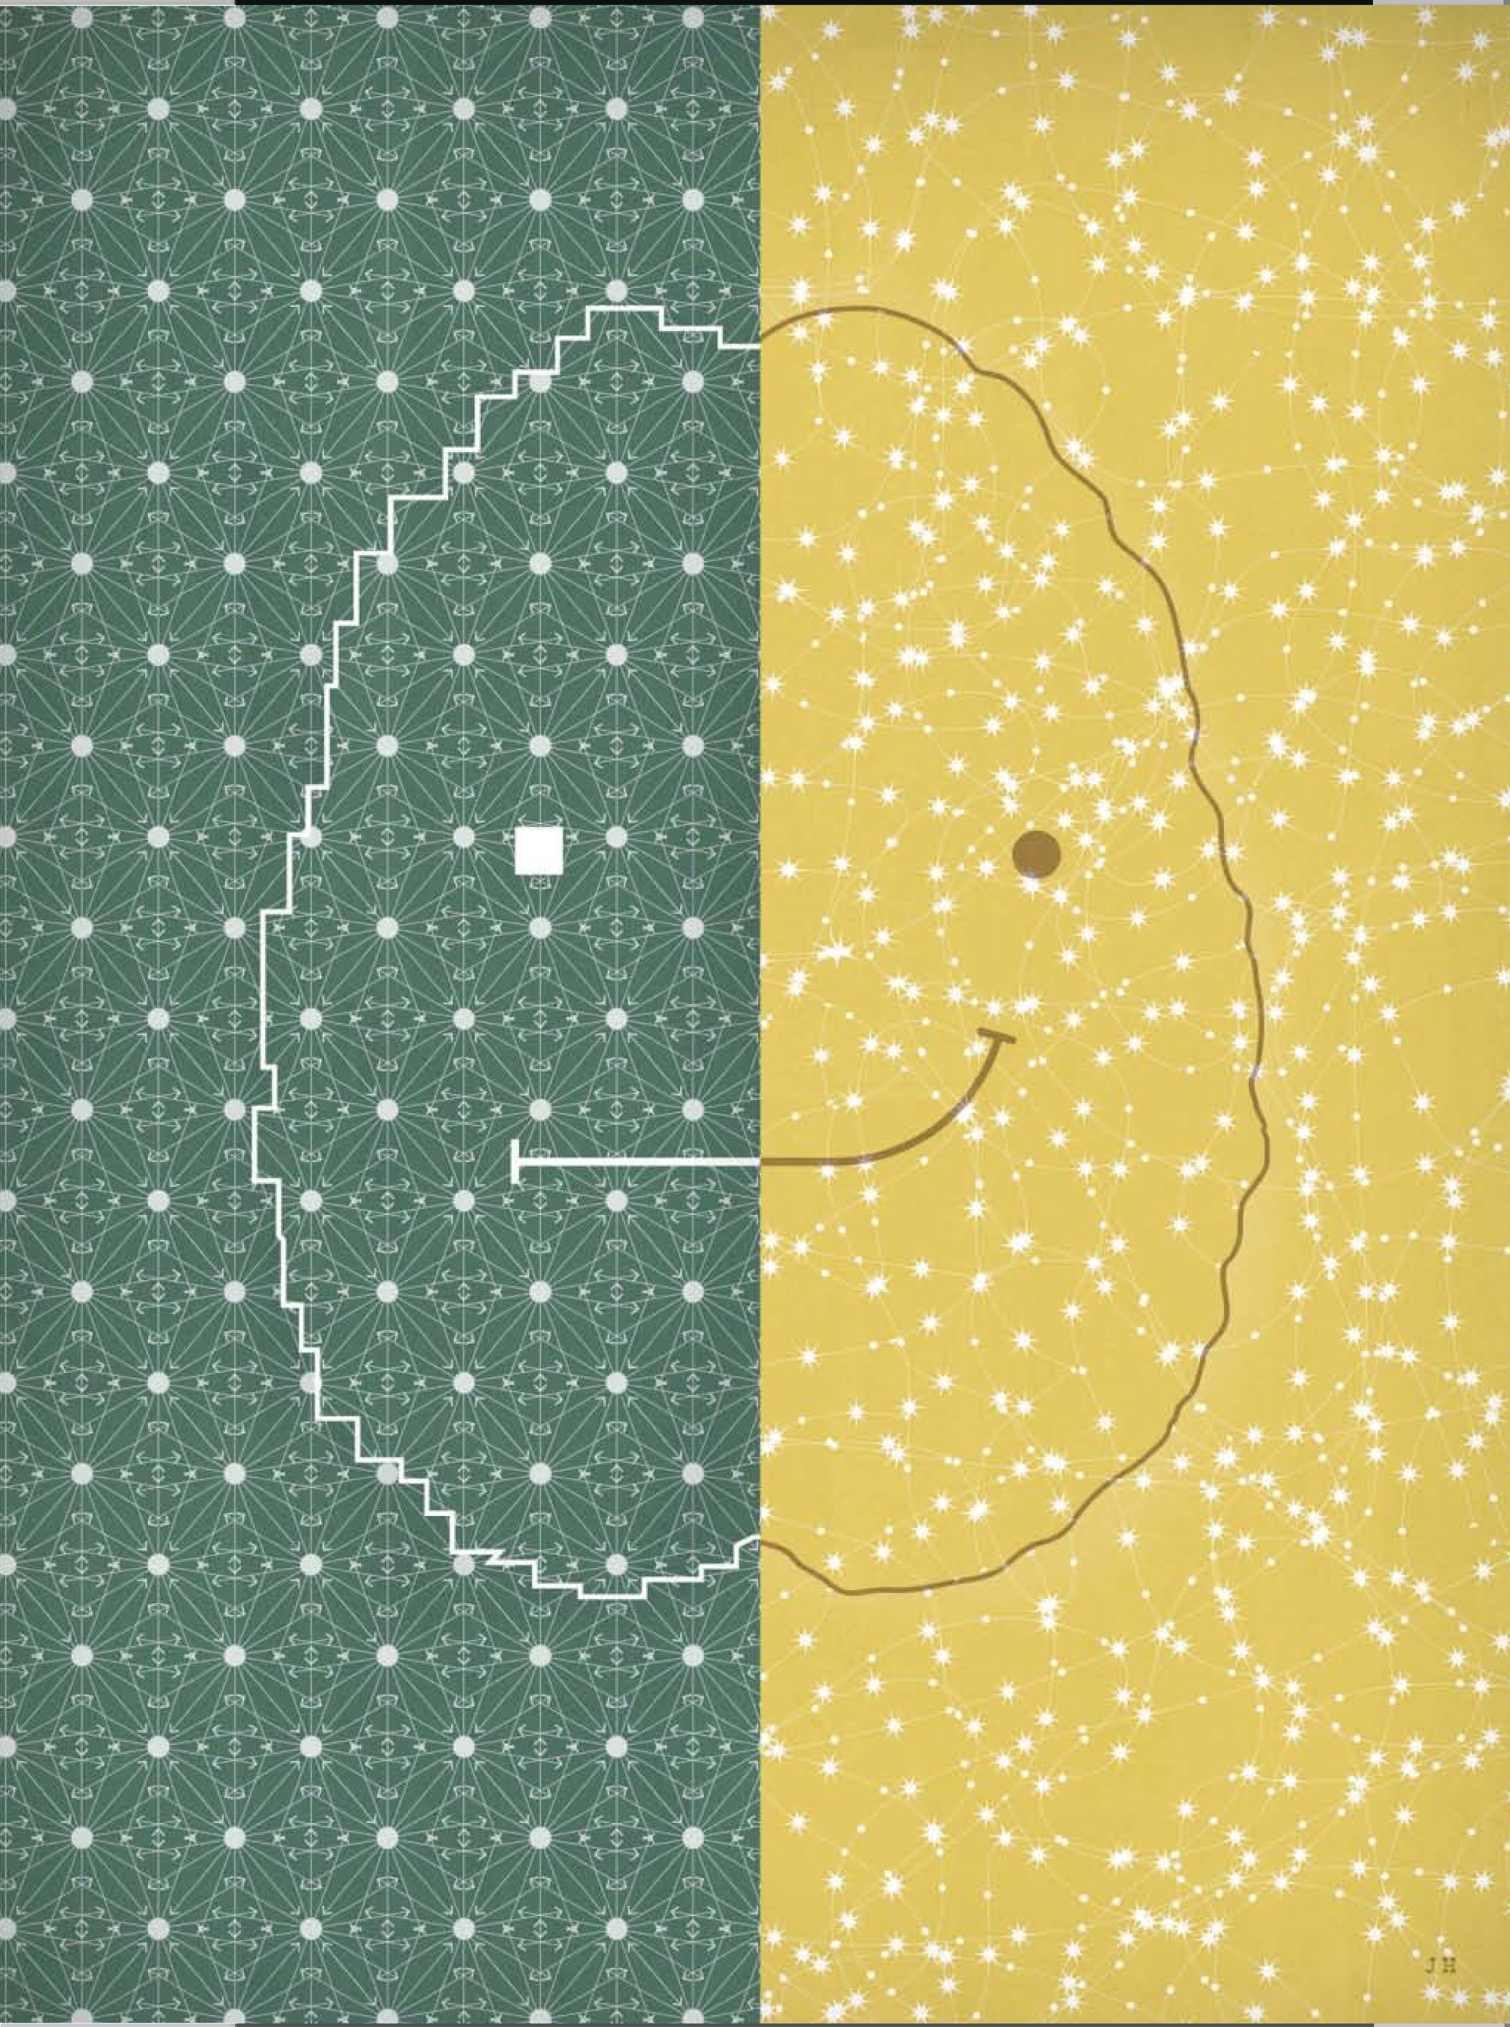

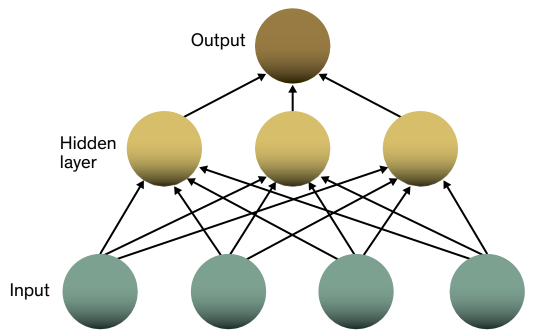

BOTTOM'S UP: A standard feed-forward network has the input at the bottom. The base layer feeds into a hidden layer, which in turn feeds into the output.

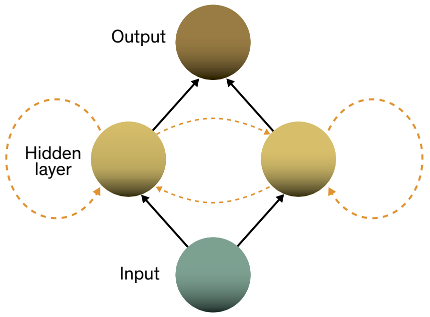

LOOP THE LOOP: A recurrent neural network includes connections between neurons in the hidden layer [yellow arrows], some of which feed back on themselves.

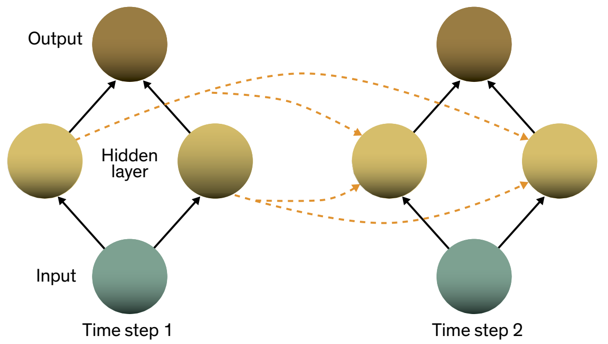

TIME AFTER TIME: The added connections in the hidden layer link one time step with the next, which is seen more clearly when the network is "unfolded" in time.

The neural network learns proper nouns like "Coors Light" and beer jargon like "lacing" and "snifter." It learns to spell and to misspell, and to ramble just the right amount. Most important, the neural network generates reviews that are contextually relevant. For example, you can say, "Give me a 5-star review of a Russian imperial stout," and the software will oblige. It knows to describe India pale ales as "hoppy," stouts as "chocolatey," and American lagers as "watery." The neural network also learns more colorful words for lagers that we can't put in print.

This particular neural network can also run in reverse, taking any review and recognizing the sentiment (star rating) and subject (type of beer). This work, done by one of us (Lipton) in collaboration with his colleagues Sharad Vikram and Julian McAuley at the University of California, San Diego, is part of a growing body of research demonstrating the language-processing capabilities of recurrent networks. Other related feats include captioning images, translating foreign languages, and even answering e-mail messages. It might make you wonder whether computers are finally able to think.

That's a goal computer scientists have pursued for a long time. Indeed, since the earliest days of this field, they have dreamed of developing truly intelligent machines. In his 1950 paper, "Computing Machinery and Intelligence," Alan Turing imagined conversing with such a computer via a teleprinter. Envisioning what has since become known as a Turing test, he proposed that if the computer could imitate a person so convincingly as to fool a human judge, you could reasonably deem it to be intelligent.

The very year that Turing's paper went to print, Gnome Press published *I, Robot*, a collection of Isaac Asimov's short stories about intelligent humanoids. Asimov's tales, written before the phrase "artificial intelligence" existed, feature cunning robots engaging in conversations, piloting vehicles, and even helping to govern society.

And yet, for most of the last 65 years, AI's successes have resembled neither Turing's conversationalists nor Asimov's humanoids. After alternating periods of overenthusiasm and subsequent retrenchment, modern AI research has largely split into two camps. On one side, theorists work on the fundamental mathematical and statistical problems related to algorithms that learn. On the other side, more practically oriented researchers apply machine learning to various real-world tasks, guided more by experimentation than by mathematical theory.

Until recently, both sides of this divide focused on simple prediction problems. For example: Is an e-mail message spam or not spam? Or: What's the probability that a loan will default? A cynic might say that we dreamed of creating humanlike intelligence and got spam filters instead. However, breakthroughs in neural-network research have revolutionized computer vision and natural-language processing, rekindling the imaginations of the public, researchers, and industry.

The modern incarnation of neural networks, commonly termed "deep learning," has also widened the gap between theory and practice. That's because, until recently, machine learning was dominated by methods with well-understood theoretical properties, whereas neural-network research relies more on experimentation.

ZACHARY C. LIPTON/UNIVERSITY OF CALIFORNIA, SAN DIEGO (2)

As high-quality data and computing resources have grown, the pace of that experimentation has similarly taken off. Of course, not everyone is thrilled by the success of deep learning. It's somehow unsatisfying that the methods that perform best might be those least amenable to theoretical understanding. Nevertheless, the capabilities of recurrent neural networks are undeniable and potentially open the door to the kinds of deeply interactive systems people have hoped for—or feared—for generations.

YOU MIGHT THINK that the study of artificial neural networks requires a sophisticated understanding of neuroscience. But you'd be wrong. Most press coverage overstates the connection between real brains and artificial neural nets. Biology has inspired much foundational work in computer science. But truthfully, most deep-learning researchers, ourselves included, actually know little about the brain.

So here's how neural networks actually work. Neural networks consist of a large number of artificial neurons, building blocks that are analogous to biological neurons. Neurons, as you might recall from high school biology class, are cells that fire off electrical signals or refrain from doing so depending on signals received from the other neurons attached to them. In the brain, a neuron connects to other neurons along structures called synapses. Through such connections, one neuron's firing can either stimulate or inhibit the firing of others, giving rise to complex behavior. That's about all the brain science you'll need to know to get going with deep learning.

Like biological neurons, an artificial neural network's simulated neurons work together. To each connection between one artificial neuron and another, we assign a value, called a weight, that represents the strength of the linkage. A positive weight could be thought of as an excitatory connection, while a negative weight could be thought of as inhibitory. To determine the intensity of an artificial neuron's firing or, more properly, its *activation*, we calculate a weighted sum of the activations of all the neurons that feed into it. We then run this sum through an aptly named *activation function*, which outputs the desired activation. This value in turn contributes to the weighted sums for calculating other neurons' activations.

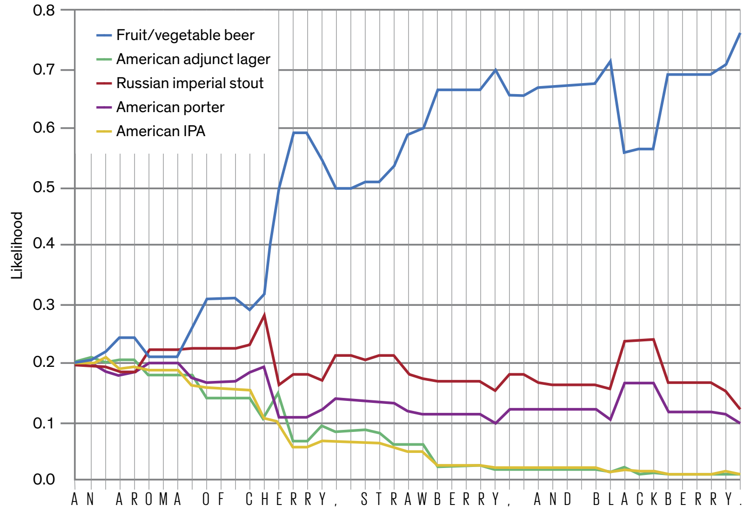

ONE BY ONE: A recurrent neural network is used here to classify the subject of a beer review, tackling that task one character at a time in sequence. The likelihood that the beer being described belongs to one of the five possible categories shifts with each new character presented to the network.

So you might wonder: Which neuron's activation should be calculated first? In a densely connected network, it's not clear. Because each neuron's activation can depend on every other neuron's activation, changing the order in which we process them can produce radically different outcomes. To dodge this problem entirely and to simplify the computations, we typically arrange the neurons in layers, with each neuron in a layer connected to the neurons in the layer above, making for many more connections than neurons. But these connections go only from lower to upper, meaning that the output of a lower layer influences the one above, but not vice versa.

For obvious reasons, we call this a feed-forward architecture. To provide input to this kind of neural network, we simply assign activations to the neurons of the lowest layer. Each value could represent, for example, the brightness of one pixel in an image.

We calculate the activations in each higher layer successively based on input from the layer below. The ultimate output of the network—say, a categorization of the input image as depicting a cat, dog, or person—is read from the activations of the artificial neurons in the very top layer. The layers between input and output are called *hidden layers*.

Writing a program to produce outputs from inputs in this way might seem awfully easy—and it is. The hard part is training your neural network to produce something useful, which is to say, tinkering with the (perhaps millions of) weights corresponding to the connections between the artificial neurons.

Training requires a large set of inputs for which the correct outputs are already known. You also need some way to measure how much the generated output deviates from the desired output. In machine learning parlance, this is called a *loss function*. Once you have those things, you present randomly selected training examples to the network. Then you update the weights to reduce the value of the loss function. That is, you repeatedly adjust the connection weights a small amount, bringing the output of the network incrementally closer to the ground truth.

While determining the correct updates to each of the weights can be tricky, we can calculate them efficiently with a well-known technique called backpropagation, which was developed roughly 30 years by David Rumelhart, Geoff Hinton, and Ronald Williams. When the loss function has a convenient mathematical property called convexity, this procedure is guaranteed to find the optimal solution. Unfortunately, neural networks are nonconvex and offer no such guarantee. Computer scientists apply this strategy anyway, often successfully.

Early on, computer scientists built neural networks with just three tiers: the input layer, a single hidden layer, and the output layer. Training networks with many hidden layers (known as deep nets) seemed computationally infeasible. This perception has been overturned, thanks in part to the power of graphical processing units (GPUs). It's now common to train networks with many hidden layers in just hours or days.

ONCE YOU'VE TRAINED a neural network, it can finally do useful things. Maybe you've trained it to categorize images or to gauge the likelihood that an applicant for a loan will default. Or perhaps you have it tuned to identify spam e-mails. These are impressive capabilities, to be sure, but they are limited by the amnesia of feed-forward neural networks. Such networks have absolutely no memory. Each prediction starts afresh, as if it were the only prediction the network ever made. But for many real-world tasks, inputs consist of sequences of contextually related data–say, a set of consecutive video frames. To determine what is happening at any frame in a video, it's clearly advantageous to exploit the information provided by the preceding frames. Similarly, it's hard to make sense of a single word in a sentence without the context provided by the surrounding words.

Recurrent neural networks deftly handle this problem by adding another set of connections between the artificial neurons. These links allow the activations from the neurons in a hidden layer to feed back into themselves at the next step in the sequence. Simply put, at every step, a hidden layer receives both activation from the layer below it and also its own activation from the previous step in the sequence. This property gives recurrent neural networks a kind of memory.

Training networks with many hidden layers seemed computationally infeasible

Let's walk through things slowly as the network processes a sequence of inputs, one step at a time. The first input (at step 1) influences the activations of the hidden layers of the network and therefore affects the output of the network. Next, the network processes the second input, again influencing its hidden layers. At step 2, however, the hidden layers also receive activation flowing across time from the corresponding hidden layers from step 1. Computation continues in this fashion for the length of the sequence. Thus the very first input can affect the very last output, as the signal bounces around the hidden layers.

Computer scientists have known for decades that recurrent neural networks are powerful tools, capable of performing any well-defined computational procedure. But that's a little like saying that the C programming language is a powerful tool, capable of performing any well-defined computational procedure. It's true, but there's a huge gap between knowing that your tool can in theory be used to write some desired program and knowing exactly how to construct it.

Working out the correct architecture for a recurrent network and training it to learn the best values for all the weights are hard problems. As with all neural networks, there's no mathematical guarantee that the training procedure will ever find the best set of weights. Worse yet, calculations to determine the weight updates can easily become unstable with recurrent nets. This can stall the learning process, even blowing up the network when the magnitude of the update is so large that it wipes out all that the model has learned.

Fortunately, over the years computer scientists have overcome many of these difficulties in practice, and recurrent neural networks are proving exceptionally powerful. One particular variety, called the Long Short-Term Memory model, developed in 1997 by Sepp Hochreiter and Jürgen Schmidhuber, has seen the most success. It replaces simple artificial neurons with specialized structures called memory cells. We won't go too far into the weeds describing memory cells here, but the basic idea is to provide the network with memory that persists longer than the immediately forgotten activations of simple artificial neurons. Memory cells give the network a form of medium-term memory, in contrast to the ephemeral activations of a feed-forward net or the long-term knowledge recorded in the settings of the weights.

ONE KILLER APPLICATION for recurrent nets is language translation. Here, the training data consists of pairs of sentences, one in the source language and one in the target language. Amazingly enough, the sentences don't need to be the same length or share the same grammatical construction.

In one such approach, by researchers Ilya Sutskever, Oriol Vinyals, and Quoc V. Le, the source sentence is first passed as input, one word at a time. Then, the recurrent network generates a translation, also one word at a time. This program matched the accuracy of many state-of-the-art machine-translation programs, despite lacking any hard-coded knowledge of either language. Other impressive examples demonstrating what can be done with such networks include recognizing and even generating handwriting.

In our own research, we recently showed that recurrent neural networks can recognize many medical conditions. In collaboration with David Kale of the University of Southern California and Randall Wetzell of Children's Hospital Los Angeles, we devised a recurrent neural network that could make diagnoses after processing sequences of observations

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

A baseball game in progress with the batter up to plate.

A brown bear standing on top of a lush green field.

A woman lying on a bed in a bedroom.

WHAT'S MY LINE: A recurrent neural network developed by Jeff Donahue of the University of California, Berkeley, and others wrote these image captions. The computer-generated descriptions are remarkably accurate (with some perplexing exceptions).

A black and white cat is sitting on a chair.

A close up of a person brushing his teeth.

A person holding a cell phone in their hand.

taken in the hospital's pediatric intensive-care unit. The sequences consisted of 13 frequently but irregularly sampled clinical measurements, including heart rate, blood pressure, blood glucose levels, and measures of respiratory function. Our objective was to determine which of 128 common diagnoses applied to each patient.

After training the network, we evaluated the model by presenting it with a new set of patient data. The network proved able to recognize diverse conditions such as brain cancer, status asthmaticus (unrelenting asthma attacks), and diabetic ketoacidosis (a serious complication of diabetes where the body produces excess blood acids) with remarkable accuracy. Surprisingly, it could also reliably recognize scoliosis (abnormal curvature of the spine), possibly because scoliosis frequently produces respiratory symptoms.

THE PROMISING RESULTS from our medical application demonstrate the power of recurrent neural networks to capture the meaningful signal in sequential data. But for other applications, such as generating beer reviews, image captions, or sentence translations, evaluation can be difficult. It's hard to say objectively in these cases what constitutes good performance, because the ground truth that you use for the correct answer may be only one of many reasonable alternatives.

Computer scientists have come up with various metrics for these circumstances, particularly to evaluate translation systems. But these metrics are not quite as satisfying as simple measures, like accuracy, which work for straightforward prediction tasks. In many cases, though, the true measure of success is whether someone inspecting the network's output might mistake it for human-generated text. The most impressive image-captioning research papers tend to contain eye-popping examples of what the neural net can do. Success in this context really means getting someone to declare, "There's no way a computer wrote that!"

In this sense, the computer-science community is evaluating recurrent neural networks via a kind of Turing test. We try to teach a computer to act intelligently by training it to imitate what people produce when faced with the same task. Then we evaluate our thinking machine by seeing whether a human judge can distinguish between its output and what a human being might come up with.

While the very fact that we've come this far is exciting, this approach may have some fundamental limitations. For instance, it's unclear how such a system could ever outstrip the capabilities of the people who provide the training data. Teaching a machine to learn through imitation might never produce more intelligence than was present collectively in those people.

One promising way forward might be an approach called reinforcement learning. Here, the computer explores the possible actions it can take, guided only by some sort of reward signal. Recently, researchers at Google DeepMind combined reinforcement learning with feed-forward neural networks to create a system that can beat human players at 31 different video games. The system never got to imitate human gamers. Instead it learned to play games by trial and error, using its score in the video game as a reward signal.

We expect reinforcement learning to rise in prominence as computers become more powerful and as imitation-based learning approaches its limits. And already, some pioneering work is combining recurrent neural networks with reinforcement learning.

As research at this intersection of approaches develops, we might see computers that not only imitate our abilities but eclipse them. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/recurrent0216

UNIVERSITY OF CALIFORNIA, BERKELEY (6)

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

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

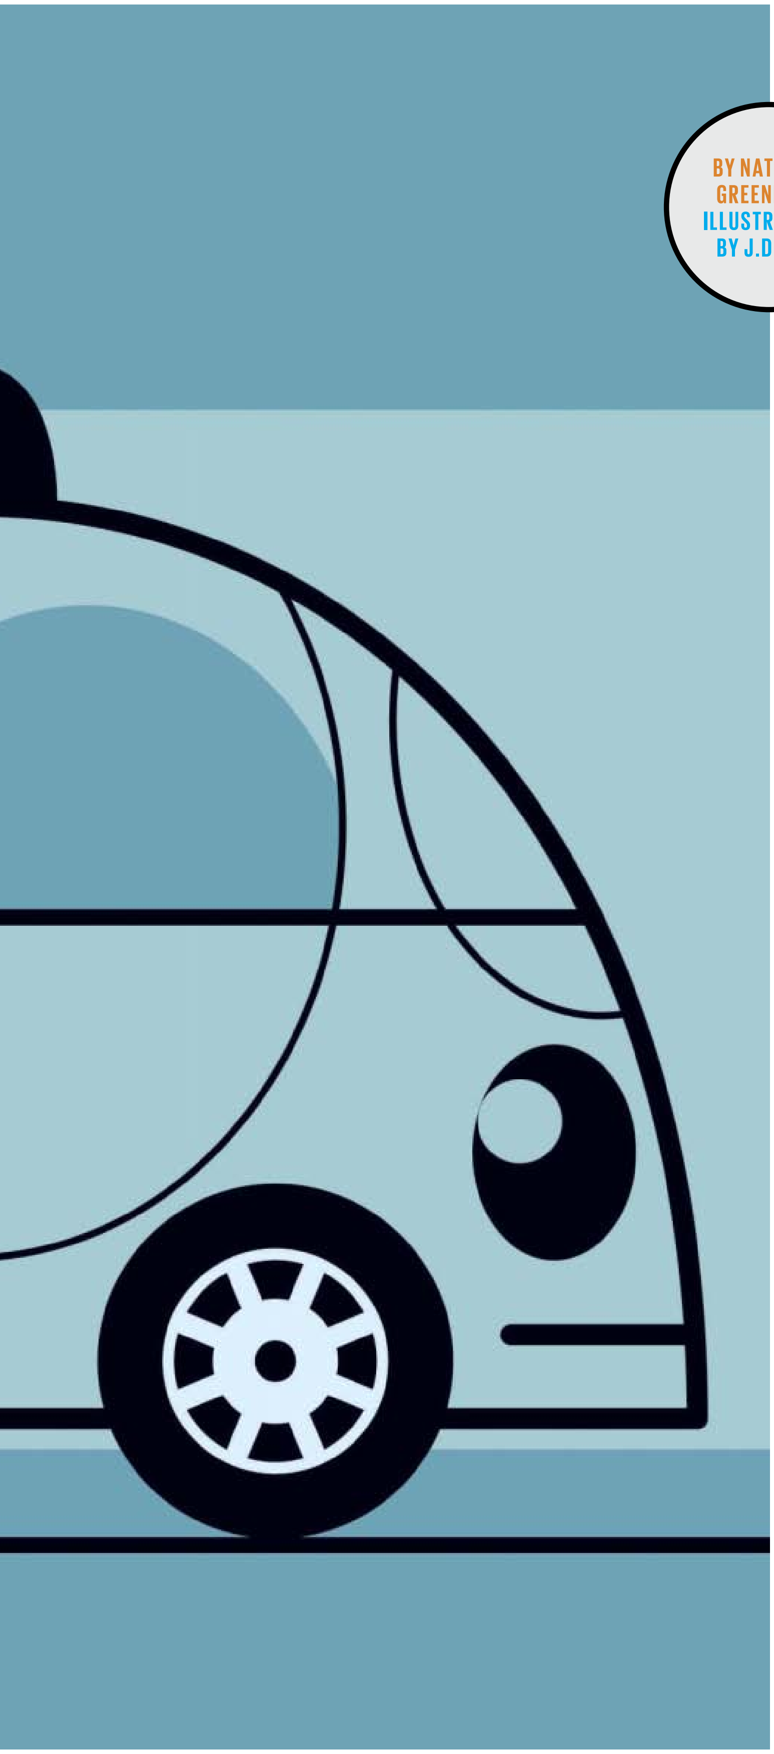

BY NATHAN A. GREENBLATT
ILLUSTRATIONS BY J.D. KING

SELF-DRIVING CARS AND THE LAW

Putting autonomous vehicles on the road isn't just a matter of fine-tuning the technology

IT IS THE YEAR 2023, and for the first time, a self-driving car navigating city streets strikes and kills a pedestrian. A lawsuit is sure to follow. But exactly what laws will apply? Nobody knows. Today, the law is scrambling to keep up with the technology, which is moving forward at a breakneck pace, thanks to efforts by Apple, Audi, BMW, Ford, General Motors, Google, Honda, Mercedes, Nissan, Nvidia, Tesla, Toyota, and Volkswagen. Google's prototype self-driving cars, with test drivers always ready to take control, are already on city streets in Mountain View, Calif., and Austin, Texas. In the second half of 2015, Tesla Motors began allowing owners (not just test drivers) to switch on its Autopilot mode.

The law now assumes that a human being is in the driver's seat, which is why Google's professional drivers and Tesla owners are supposed to keep their hands near the wheel and their eyes on the road. (Tesla's cars use beeps and other warnings to make sure they do so.) That makes the vehicles street legal for now, but it doesn't help speed the rollout of fully autonomous vehicles. ▶

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

US $16,000

Median damages in automobile accident cases compared with $748,000 for product liability cases

It's not only the law that's playing catch-up but also the road system. We've invested billions of dollars in a transportation infrastructure designed for human vision, not at all for computers. But it's possible to make changes to the laws that govern the roads and the infrastructure, and those could go a long way toward making driverless cars the rule instead of the rare exception.

No matter how the laws and infrastructure evolve and how smart the cars become, bad things will still happen and manufacturers will end up in court. So far, we have no strictly applicable case law, for although Google cars have been involved in 17 accidents to date, the robot was at fault in none of them.

But accidents caused by autonomous machines are hardly theoretical. Consider manufacturing robots, surgical robots, and subway trains—smart machines that have led to injury, death, and lawsuits. Just one 2009 crash caused by a malfunction of a train's automatic train-control system, for example, led to 21 lawsuits and 84 out-of-court claims. Attorneys don't even have to wait for actual accidents to sue: A 2015 lawsuit against Ford, GM, and Toyota accused the companies of "hawking vehicles that are vulnerable to hackers who could hypothetically wrest control of essential functions such as brakes and steering."

FOR NOW, the legal landscape is a hodgepodge. Laws in California and Nevada, for example, allow self-driving cars on public roads so long as a human driver is sitting behind the wheel on alert, and other states are allowing testing on designated roadways. European regulators have allowed limited tests of self-driving cars and even tractor-trailers. The United Kingdom authorized testing starting last year and has begun reviewing road regulations to figure out how to eventually allow a fully autonomous shuttle. Japan allowed its first road test of an autonomous car in 2013, although much of the research being done by Japanese car companies is happening in the United States.

The United States' National Highway Traffic Safety Administration has been carefully watching the technology and is generally endorsing it, stating, for example, that the agency is "eager" to "support the development of vehicle automation" due to the "exciting opportunity [automated vehicles] represent to the American public." At a minimum, regulations need to continue smoothing the path for testing and rolling out at least limited versions of the technology.

We can't put off changing the laws until the advent of robotic driving, because today's laws leave a lot of room for uncertainty, and uncertainty stalls progress. A car company can't be expected to invest in putting out a new fleet of autonomous cars when it could be forced to pull them all off the road after the first accident. We won't have truly autonomous cars on the road until this gets sorted out.

Volvo recently announced that it would take the blame if "any of its self-driving cars crashes in autonomous mode." Although that may sound like a big deal, it doesn't represent progress. Under current U.S. law, Volvo would most likely take the blame anyway. The real questions—which terrify Volvo and other manufacturers—are: Exactly when will they be held responsible, and how much will they have to pay?

Most legal scholars think that an accident will lead to a major design-defect lawsuit. That worries the car companies for several reasons.

First, it's expensive no matter who wins. A multimillion-dollar legal case is nearly a certainty when new, complex driving systems involving millions of lines of source code are involved.

Second, the outcome of that case is hard to predict. Generally, the key question in a product liability lawsuit is whether the product had a "defective condition" that was "unreasonably dangerous." This often involves determining whether the product designer could have made the product safer at an acceptable cost. But what's "reasonable" for a new technology? Is "reasonably safe" defined by the average human driver, the perfect human driver, or the perfect computer driver?

Third, a lawsuit can lead to a recall. A legal determination that a design is defective, caused an accident, and will likely cause another can be a powerful incentive for a recall. Recalls and mistakes can be expensive: GM's ignition-switch recall cost the company US $4.1 billion in 2014, and Volkswagen's diesel emissions scandal will likely cost the company over $7 billion. For at least some autonomous-car defects, however, the recall could take the form of a software patch delivered wirelessly and inexpensively.

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

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

Finally, punitive damages can come into play. Punitive damages are generally available in the United States for outrageous conduct in designing or manufacturing a defective product. The 1994 case *Liebeck v. McDonald's*, which involved a punitive damage award of $2.7 million for burns from hot coffee, is a famous example.

Because of the risk of such a lawsuit, the potential legal costs faced by manufacturers of autonomous vehicles are higher than the costs faced by human drivers. Most auto accidents usually result in pretrial settlements, and the 4 percent of cases that go to trial have relatively low legal costs and low potential damages, compared with those of a design-defect lawsuit. The U.S. Department of Justice reported that median damages in automobile accident cases were $16,000, compared with $748,000 for product liability cases. These legal costs create unfair disincentives for autonomous vehicles.

THE SOLUTION to the lawsuit problem is actually pretty simple. To level the playing field between human drivers and computer drivers, we should simply treat them equally. Instead of applying design-defect laws to computer drivers, use ordinary negligence laws. That is, a computer driver should be held liable only if a human driver who took the same actions in the same circumstances would be held liable. The circumstances include the position and velocity of the vehicles, weather conditions, and so on. The "mind" of the computer driver need not be examined any more than a human's mind should be. The robo-driver's private "thoughts" (in the form of computer code) need not be parsed. Only its conduct need be considered.

That approach follows basic principles of negligence law. As Dobbs's *Law of Torts* (2nd ed.) explains: "A bad state of mind is neither necessary nor sufficient to show negligence; conduct is everything. One who drives at a dangerous speed is negligent even if he is not aware of his speed and is using his best efforts to drive carefully. Conversely, a person who drives without the slightest care for the safety of others is not negligent unless he drives in some way that is unreasonably risky. State of mind, including knowledge and belief, may motivate or shape conduct, but it is not in itself an actionable tort"–that is, wrongful conduct.

For example, a computer driver that runs a red light and causes an accident would be found liable. Damages imposed on the carmaker (which is responsible for the computer driver's actions) would be equal to the damages that would be imposed on a human driver. Litigation costs would be similar, and the high costs of a design-defect suit could be avoided. The carmaker would still have a financial incentive to improve safety. In fact, the manufacturer would have greater incentives than with a human-driven vehicle, because of publicity concerns. Correction of systemic problems could be implemented via a predictable mechanism, such as a mandatory crash-review program with government oversight, without excessive risk to the manufacturer.

As the safety of autonomous vehicles improves and as legal costs become more predictable, stricter safety standards could be imposed to encourage further progress. This scheme would help encourage development of the technology without undermining marginal incentives for safety. Insurers have a century of experience in predicting accident costs for human-driven cars. Courts have a century's worth of benchmarks on which to draw to ensure that a fair comparison is made. Making it just as easy for the courts to judge cases involving self-driving cars would shield manufacturers from excessive financial risk while compensating accident victims no less than they are today. With such predictability, it is likely that self-driving car manufacturers would pay about the same for insurance per vehicle as an average human driver does. Insurance costs could even be lower because the self-driving car would qualify for all the "good driver" discounts.

30,000

Estimated number of lives autonomous driving technology can save each year in the United States alone

Implementation of any of these policies in places like the United States and Canada will have to happen on a state-by-state or province-by-province basis, as the rules of the road in these countries aren't set nationally; in Europe and many other areas, however, it will evolve country by country.

PUBLIC POLICY is holding back self-driving cars in another way–it influences the design of the roads and the way they are governed based on the needs of drivers that "see." The rules require that we stop on red, yield when we see a triangle-shaped sign, and obey metering lights at freeway entrances. That's easy and intuitive for humans, not so easy for machines. Today's autonomous vehicles recognize objects with a combination of object tracking using distance and velocity (it doesn't really matter if an object cruising down the road is a car, a rolling boulder, or a flying saucer; a computer driver can avoid hitting it without knowing what it is), and object recognition (it makes sense to slow down if a small child or a deer is near the curb, but zipping past a fire hydrant is fine). This technology still has years of development ahead.

A more costly but potentially simpler approach would be to make the infrastructure friendlier to autonomous vehicles. These changes wouldn't eliminate all the challenges–a car would still have to "see" the child approaching that intersection–but it would simplify much of the burden. Radio frequency transmitters in traffic lights, for example, could tell a computer driver if a light is green or red more quickly and with greater accuracy than a machine vision system struggling with shadows and glare. These kinds of changes will have to happen on national levels, with international coordination if possible, through both regulation and standardization of technology.

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

THERE'S A REASON to speed the rollout of autonomous vehicles. By replacing error-prone human drivers, autonomous driving technology can potentially save 30,000 lives each year in the United States alone. It can annually prevent 5 million accidents and 2 million injuries, conserve 7 billion liters of fuel, and save so many hundreds of billions of dollars in lost productivity and accident-related costs that the figures are beyond comprehension.

That's because computer drivers are in principle fundamentally safer drivers. They never text, do their makeup, or fall asleep at the wheel. (Human error, in contrast, causes roughly 93 percent of crashes.) Robo-drivers can have 360-degree vision, and thanks to lidar, radar, and ultrasonic sensors, they can see through fog and in the dark.

Computer drivers can have "telepathy": A computer driver could let another computer driver know that it is considering changing lanes before making the decision to do so. It could communicate with traffic lights to minimize wait times at intersections and optimize traffic flow.

Computer drivers react faster. Humans rely on chemical signals, with reaction speed limited to about 1.5 seconds—or 37 meters (about 120 feet) at highway speeds. Computers rely on much faster electrical signals and gigahertz-scale processors to react.

Computer drivers can take far more rigorous driver tests than the 20-minute road tests offered by U.S. departments of motor vehicles today. Recorded or virtual information could test a computer driver's ability to safely drive, say, a million miles before handing over a license.

POST YOUR COMMENTS at http://spectrum.ieee.org/robocarlaw0216

And, finally, computer drivers have the potential to accumulate far more wisdom than any human. It is said that wise men learn from the mistakes of others; only fools learn from their own. Every autonomous vehicle can learn from thousands of others, through incremental and permanent engineering improvements. Humans, unfortunately, often repeat mistakes others have made.

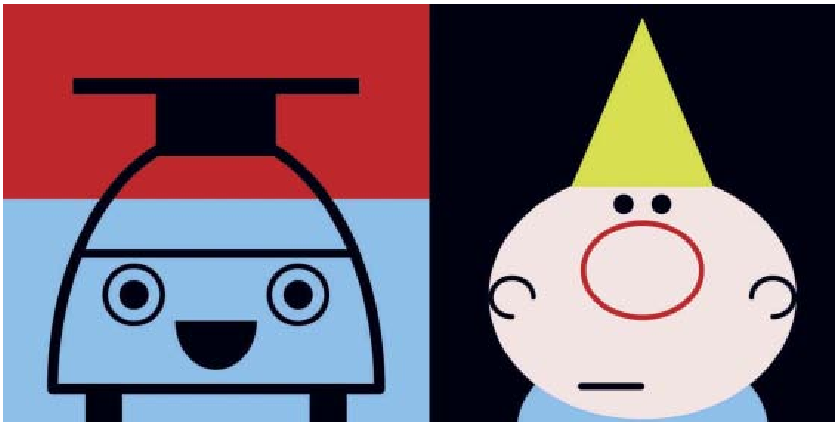

WHEN SELF-DRIVING CARS do succeed, the effect will be widespread. And they will succeed, despite having two giant thumbs on the human side of the scale—the stacked deck of liability rules and the transportation infrastructure that relies on vision rather than other means of communication. Then a host of new social and legal issues will emerge.

The most immediate impact will likely be on transportation for hire—the sandbox of Uber, Lyft, and taxis. Uber is already bullish on replacing its human drivers with computers, and it has hired 50 Carnegie Mellon University scientists to develop the technology. Uber may be able to count on computers not filing a class-action lawsuit against the company, but it should plan for angry human former drivers accusing it of "economic terrorism" and lengthy negotiations with regulators. And Uber will not be the only robo-taxi startup.

Self-driving cars will also disrupt the standard model of car ownership and use. Currently, cars typically are parked 95 percent of the time. If people ordered a self-driving car only when needed, utilization rates would rise, ownership costs would decline, and as an added benefit, we would typically ride in newer-model cars with a smaller environmental footprint. But what will that do to car sales?

Privacy will be a concern. Self-driving cars are the ultimate connected, on-the-grid machines. Not only would they know your exact location and route, but in a robo-taxi or shared-ownership model, the cars might have video monitoring or other means of preventing vandalism (or passenger failure to clean up all those breakfast-sandwich wrappers). In addition, self-driving cars continuously monitor other drivers on the road. Whether the gigabytes of generated information can be permanently stored—and how they can be used later—is not settled.

And self-driving cars will have a profound effect on city design. Parking spaces take up, on average, about 31 percent of city central business districts. Self-driving cars can park themselves in peripheral areas, or, in a shared-ownership/taxi model, they could pick up the next passenger. In either case, more land could be devoted to pedestrian zones, shopping, parks, and other valuable uses.

These issues will be worked out because ultimately we want to choose the best technology in terms of costs and benefits to society. So 50 years from now, in a world with no traffic accidents, people will look back and conclude that human drivers were a design defect. ■

The author is an intellectual-property lawyer. Complex product liability issues are beyond the scope of this article.

GREAT LEAPS OF LIGHT

CONTINUED FROM PAGE 29

capacity by as much as a factor of 10, says Bennett.

These larger-core fibers are already being deployed, mostly in submarine cables where transmission capacity is most valuable. And they're generally a good option for a new connection, says Bennett: "If somebody is planning to deploy new terrestrial fiber, they might as well deploy large-area fiber." But as attractive as they may be, large-core fibers don't completely eliminate the nonlinear distortion problem.

A potentially more promising approach is to create multiple parallel paths along which separate light signals can travel. Developers call it spatial-division multiplexing because the strategy splits the transmitted data into different physical paths.

The term actually refers to three very different kinds of parallel transmission. The simplest and most obvious approach is to add more physical paths by adding more fibers to a cable. Multifiber cables are already in wide use, but boosting capacity can be costly and complex because each fiber in a cable needs its own transmitters, receivers, and amplifiers.

Much bigger rewards may be possible if engineers can find a way to integrate separate light paths within the same fiber in a compact fashion. One way to do that is to construct fibers that contain several light-guiding cores running along their length. Like an ordinary fiber, a multicore fiber is made by first assembling the required materials in a cylindrical "preform," which is then heated so the glass can be pulled into a long thin fiber.

Unlike multifiber cables, which need a separate fiber amplifier for each fiber, a multicore fiber could be paired with a multicore amplifier. An eight-core amplifier could potentially cost much less than eight single-fiber amplifiers.

An alternative is to make a core that can guide light in a few distinct ways, called modes. Light signals in two different modes pass through each other along the fiber, but they can be isolated from each other when they emerge at the end of the fiber.

To create multiple modes in a fiber, the mode for each signal must be shaped to have the right cross section as it goes into the fiber. Each mode would need to be generated by its own laser, and optics and electronics at the receiving end must be able to separate out the modes. This separation is already done in radio systems using multiple-input/multiple-output antennas.

So far, both multimode and multicore fiber transmission are in the early stages of development. There have been multiple laboratory tests, dubbed "hero experiments" because they're out to set records that impress reporters or super-

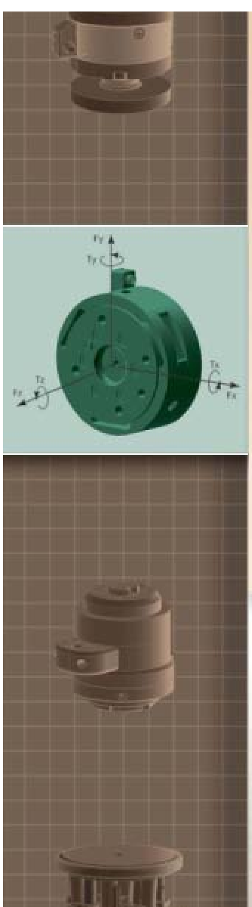

ROBOTIC END-EFFECTORS

Measure all six components of force and torque in a compact, rugged sensor.

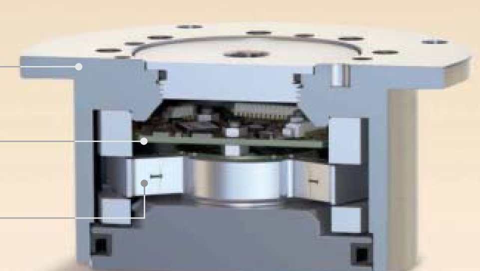

The F/T Sensor outperforms traditional load cells, instantly providing all loading data in every axis. Engineered for high overload protection and low noise, it's the ultimate force/torque sensor. Only from ATI.

ATI INDUSTRIAL AUTOMATION
Engineered Products for Robotic Productivity
www.ati-ia.com/ift
919.772.0115

Singapore NRF Fellowship 2017

The Singapore National Research Foundation (NRF) invites outstanding researchers in their early stage of research careers to apply for the Singapore NRF Fellowship.

The Singapore NRF Fellowship offers:

- A five-year research grant of up to SGD 3 million (approximately USD 2.4 million).
- Freedom to lead ground-breaking research in any discipline of science and technology in a host institution of choice in Singapore.

The NRF Fellowship is open to all researchers (PhD holders) who are ready to lead independent research in Singapore.

The Call is now open!
Apply online at **https://rita.nrf.gov.sg**

Closing Date: **23 February 2015, 3pm (GMT+8)**

For further queries, email us at **NRF_Fellowship@nrf.gov.sg**

About the National Research Foundation, Prime Minister's Office, Singapore

The National Research Foundation (NRF), set up on 1 January 2006, is a department within the Prime Minister's Office. The NRF sets the national direction for research and development (R&D) by developing policies, plans and strategies for research, innovation and enterprise. It also funds strategic initiatives and builds up R&D capabilities by nurturing research talent. The NRF aims to transform Singapore into a vibrant R&D hub that contributes towards a knowledge-intensive, innovative and entrepreneurial economy; and make Singapore a magnet for excellence in science and innovation. For more information, please visit **www.research.gov.sg**.

NATIONAL RESEARCH FOUNDATION
PRIME MINISTER'S OFFICE
SINGAPORE

visors. Such demonstrations suggest that each approach has the potential to multiply fiber capacity significantly. Together they might push capacity up by a factor of perhaps a few hundred.

But the systems needed to exploit these approaches aren't practical yet, and a host of questions remain. "Basically all spatial-division multiplexing techniques have their own showstopping problems today," says Bennett. For example, for multicore and multimode fibers, simply connecting the ends of the fibers to transmitters and receivers is far more complex than for standard fibers. In both cases, much more mechanical precision is needed. Great care has to be taken to make sure light goes in exactly as intended. And for multicore fibers with multicore amplifiers, the cores in each system have to match up with extreme precision.

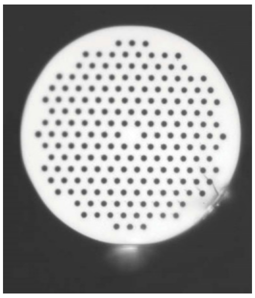

Larger-core fibers, such as this one, employ novel structures to help confine signals while boosting capacity.

NKT PHOTONICS

Barring an engineering breakthrough, "it's almost always easier to just light up another fiber," Bennett says. "This is what service providers are telling us."

Peter Winzer, a distinguished member of technical staff at Bell Labs and a leader in high-speed fiber systems, agrees that installing new cables with even more fibers is the simplest approach. But in a recent article, he warned that this approach, which will add to the cost of a cable, might not be popular among telecommunications companies. It wouldn't reduce the cost per transmitted bit as much as they had come to expect from earlier technological improvements.

New ideas continue to emerge. In June 2015, Nikola Alic of the University of California, San Diego, and colleagues reported a way of increasing fiber transmission distance by using optical frequency combs, which naturally lock laser wavelengths relative to one another, eliminating jitter and improving signal quality. "We can at least double the data rate of any system" by using a frequency comb, says Alic. "It is very nice and solid work," says Winzer, but he doubts it would have much practical impact, because developers want a bigger increase.

What will come next? Today telecommunications carriers have their hands full installing 100-Gb coherent systems. Superchannels will boost maximum capacity by 30 percent or so, and spatial-division multiplexing looks like the best candidate for the next big jump in capacity. But beyond that, who knows?

Perhaps some new twist on an old idea might come along. Coherent transmission, which was finally adopted around 2010, was actually a hot topic in the 1980s, but it lost out then to other technologies that were ready to deploy. Something totally new might emerge from the fertile ground of photonics research. And we could always lay more fibers. In any case, the global thirst for data will keep engineers working very hard to keep pumping up the bandwidth. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/fiber0216

A TESLA IN EVERY GARAGE?

CONTINUED FROM PAGE 35

short-range battery electric fleets. And yet such a system is not being considered seriously in California or anywhere else in the Sunbelt, which speaks to the pervasive influence of big-battery, supercar thinking in current EV planning.

GIVEN THE HIGHLY COMPETITIVE nature of the auto industry, why did Tesla receive so much support from other automakers? For example, Daimler and Toyota invested $50 million apiece in Tesla in 2009 and 2010, respectively. And in 2010, the company was able to buy its Fremont plant, which had been built for the Nummi partnership between GM and Toyota, for the rock-bottom price of $42 million.

The answer is that in various ways the major car companies saw the Bay Area startup as a means of hedging their risk during a period of forced technological change. For them, producing small numbers of unprofitable compliance cars was, and continues to be, a short-term annoyance and also raises the disconcerting possibility that these all-electric vehicles will eventually become popular.

But with Tesla, they had an easy way out: The establishment players purchased ZEV credits from Tesla, thereby decreasing the number of compliance cars they had to produce. Such sales raised $51 million for Tesla in the first quarter of 2015.

Daimler and Toyota also contracted with Tesla to supply power trains and battery packs for their compliance cars, notably the Mercedes-Benz B-class hatchback and the RAV4 sport-utility vehicle. Besides outsourcing their ZEV responsibilities, the auto giants were likely motivated by the opportunity to learn what the startup was doing, in the same way that the Nummi partnership allowed GM to learn about Toyota's lean manufacturing. With all of this industry and government assistance, Tesla's stock skyrocketed, from around $26 per share in 2010 to more than $250 in mid-2015. In October 2014, amid plummeting oil prices, Daimler and Toyota cashed out, reportedly at a tidy profit. Analysts offered varying opinions as to why, but I think it's likely that by then their experience with Tesla had proved beyond a doubt that in a world of cheap gasoline, BEVs made no economic sense once the federal rebate money ran out.

And for Toyota, the hybrid is the electric automobile of the future. In a speech in 2013, Takeshi Uchiyamada, team leader of the original Prius project and the current chair of Toyota, dismissed the idea that hybrids are just an interim solution on the way to an all-electric vehicle, adding that his company saw the technology as a "long bridge and a very sturdy one."

NOBODY SHOULD WRITE OFF Musk and Tesla just yet. The fundamental conditions that gave rise to the company persist. Blessed with a mild and dry but fragile ecology as well as rich engineering and capital resources, California has an ecofriendly entrepreneurial culture. It was the ideal setting

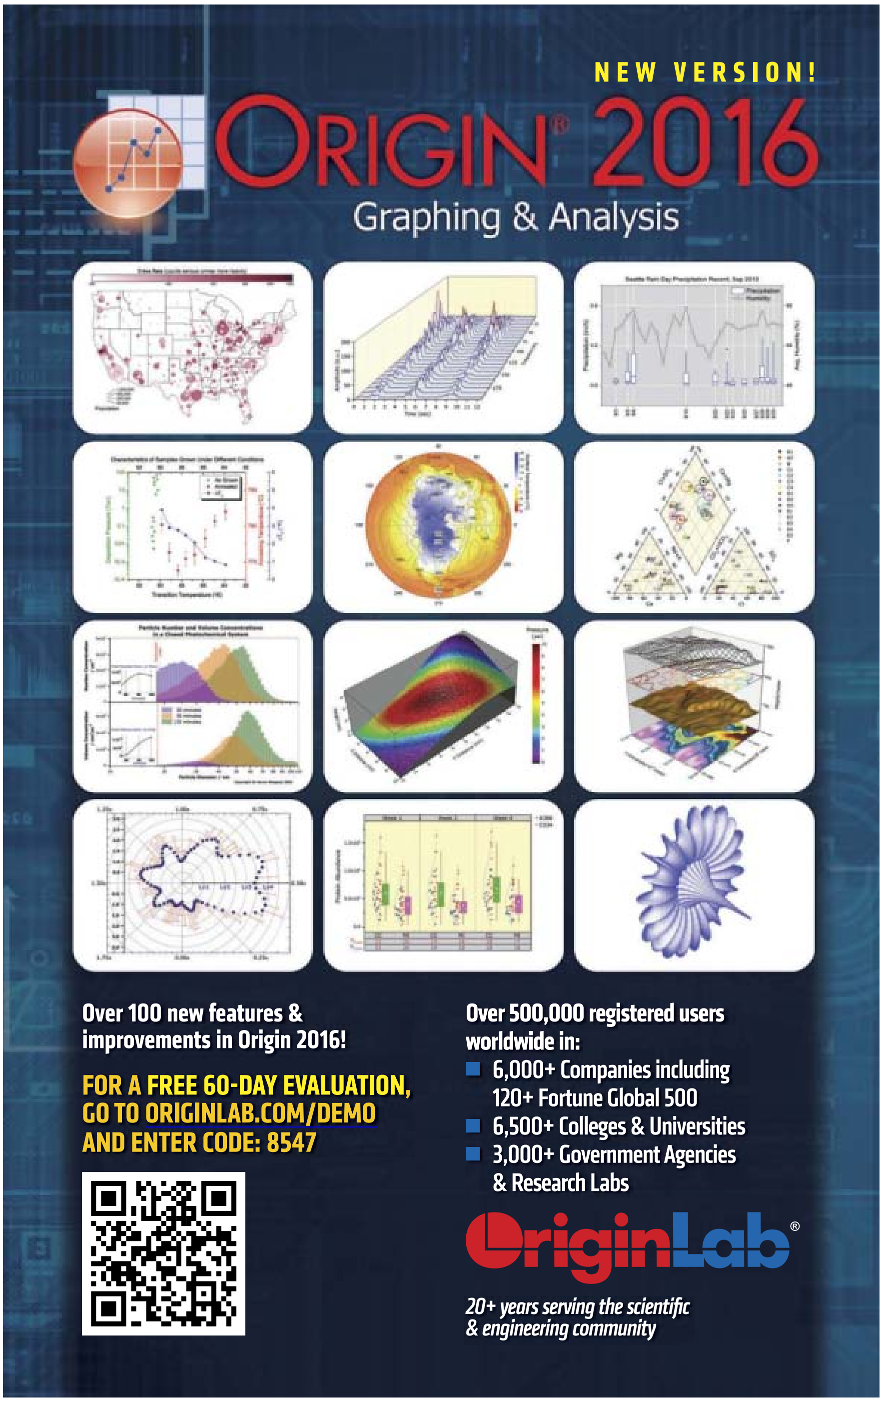

IEEE.tv gets a mobile makeover

Bring an award-winning network of technology programs with you:

- Generations of industry leaders.
- The newest innovations.
- Trends shaping our future.

Go mobile or get the app.
www.ieee.tv

Available on the App Store
Available in Android Market

IEEE

Toyota Technological Institute (TTI) has an opening for a professor (tenured or tenure-track), or an associate professor (tenured or tenure-track) in the Department of Advanced Science and Technology, Faculty of Engineering.

The successful candidate is expected to conduct not only quality research and education in the field of information science and technology specified below, but also to promote cooperative research and educational programs in collaboration with our sister institute, Toyota Technological Institute at Chicago that focuses on fundamental computer science.

Research field: Intelligent information processing including artificial intelligence and its application, communication and network systems, computer vision, cyber-physical systems, human machine interface.

Qualifications: The successful candidate must have a Ph.D degree (or equivalent), a record of outstanding research achievements, and the ability to conduct strong research programs in the specified area. The candidate is expected to teach mathematics and programming at the introductory level and machine learning, information theory and signal processing at the advanced level. The supervision of undergraduate and graduate students in their research programs is also required.

Starting date: October 2016, or at the earliest convenience
Documents: (1) Curriculum vitae
(2) List of publications
(3) Copies of 5 representative publications
(4) Description of major accomplishments and future plans for research activities and education (3 pages)
(5) Names of two references with e-mail addresses and phone numbers
(6) Application form available from our website (http://www.toyota-ti.ac.jp/english/employment/index.html)

Deadline: **April 15, 2016**
Inquiry: Search committee chair Professor Tatsuo Narikiyo
(Tel) +81-52-809-1816, (E-mail) n-tatsuo@toyota-ti.ac.jp

The above should be sent to:
Mr. Takashi Hirato
Administration Division
Toyota Technological Institute
2-12-1, Hisakata, Tempaku-ku
Nagoya, 468-8511 Japan

(Please write "Application for Intelligent Information Processing Position" in red on the envelope.)

for Eberhard, Tarpenning, and Musk to cultivate their electric supercar ideology. Nowhere else does the belief in better living through advanced technology run as deep.

Tesla cars, and modern BEVs in general, have largely been shaped within this unique context and are now being exported to the world. But when the technology has been transferred to other contexts, things have not always proceeded smoothly: Recall *New York Times* reporter John M. Broder's abortive test drive of a Model S, in which he tried to go from Delaware to Connecticut. He became stranded after he exhausted the battery in cold weather.

The ensuing row between him and Musk (who accused Broder of sabotage) overshadowed the fact that Broder acted precisely as Tesla implied its users should. That is, he took at face value the suggestion that the all-electric vehicle requires no changes in consumer behavior or expectations. Doubtless Broder would have been better off researching how cold as well as hot weather can detract from BEV performance.

Several inconvenient truths thus come into focus. The first is obvious but nevertheless bears repeating: Tesla sloganeering aside, electric vehicles are significantly unlike internal combustion vehicles in fundamental ways. Most important, they age in ways that are not yet fully understood and that have huge implications for ownership and manufacturing.

Less obvious but equally pertinent is that the user's experience of an electric vehicle will vary significantly depending on climate, culture, and class. Many people identify strongly with their gasoline-fueled vehicles. To them, a silent, clean, superefficient EV is just not a car.

With the Model 3, expected next month, Musk will be targeting not his loyal and wealthy fans but rather the non-true believers: drivers socialized to the very high levels of cost-effectiveness, convenience, and sophistication of modern gasoline-engine technology. Musk wants the Model 3 to compete with BMW's 3 Series, but the kind of people who are considering a BMW are unlikely to be awed by electric supercar mythology. And on a practical level, there's the problem of how to profitably build and market the Model 3 in large volumes, especially outside California. A recent Morgan Stanley analysis forecasts an average cost that's nearly double the Model 3's planned $35,000 retail price.

The reality is that an affordable electric vehicle for the masses already exists in the form of the hybrid electric car. Toyota has sold more than 8 million of them to date, mostly Priuses, an order of magnitude greater than the cumulative sales of all BEVs. Hybrid electrics are a formidable market segment that will prove difficult to challenge, especially as more data on the economics of large EV batteries come to light.

No doubt the technology of BEVs is here to stay in one form or another, and for that alone, Tesla has accomplished something significant. But we can't ignore the fact that the technology inevitably demands enormous changes in the ways automobiles are manufactured, used, and maintained. The question is not whether the battery electric vehicle can be made to adapt to society but whether society is willing to adapt to it. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/teslastrategy0216

Faculty positions in Electrical and Computer Engineering in Africa

The College of Engineering at Carnegie Mellon University, a world leader in information and communication technology, has extended its global reach into Africa. In 2012 we became the first U.S.-based research university offering on-site master's degrees in Africa at our base in Kigali, Rwanda.

Carnegie Mellon University in Rwanda is educating future leaders who will use their hands-on, experiential learning to advance technology innovation and grow the businesses that will transform Africa.

We are seeking highly qualified candidates to join our world-class faculty, who share in our vision of developing creative and technically strong engineers that will impact society. Faculty members are expected to collaborate with industry and deliver innovative, interdisciplinary graduate teaching and research programs.

Carnegie Mellon is seeking exceptional candidates who can deliver innovative, interdisciplinary graduate programs in these areas:

- Software engineering
- Mobile and cloud computing
- Communications and wireless networking
- Cybersecurity and privacy
- Embedded systems
- Energy systems
- Image and signal processing
- Data analytics
- Applications in healthcare, agriculture, finance and infrastructure
- Innovation and technology management

Candidates should possess a Ph.D. in a related discipline and an outstanding record in research, teaching and leadership.

- Please contact us at **info@rwanda.cmu.edu** for full application requirements.
- Further information about CMU in Rwanda can be found at **www.cmu.edu/rwanda.**
- Applications should be submitted by email to **director@rwanda.cmu.edu**.

Carnegie Mellon University
Rwanda

PAST FORWARD_ BY EVAN ACKERMAN

CAN YOU HEAR ME NOW?

From the lab notebook of X., RCA engineer, 1 June 1951: In retrospect, testing ultrahigh-frequency antenna designs by hand with thunderclouds overhead may not have been the best idea.... In other news, the company is now seeking qualified applicants to fill four open positions in the ultrahigh-frequency antenna design group....The antennas we're working on will be able to receive a staggering 70 television channels—people will no longer be limited to just 12 VHF channels. We're confident that 82 channels is more than any rational human could ever possibly need. ■

PETER STACKPOLE/THE LIFE PICTURE COLLECTION/GETTY IMAGES

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

Portable Data Acquisition that's really...PORTABLE

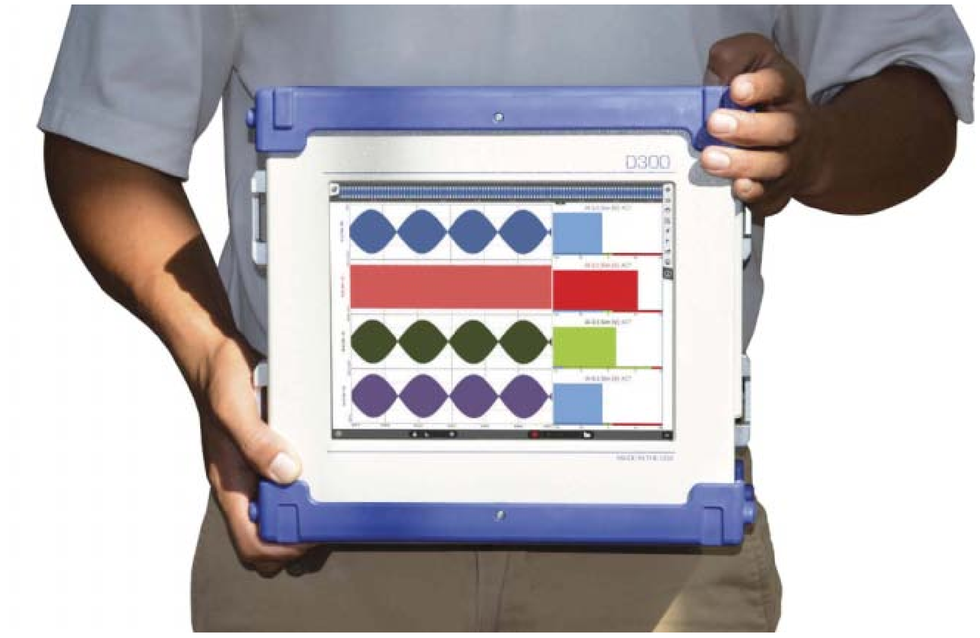

The D300 from Dewetron is so small and light that you can carry it wherever you go. And yet it won't come up short in data acquisition capabilities.

At just 11.5 x 11" (and 4.5" thick), you can carry it around as easily as a football. The bright touchscreen display makes setup and operation fast and intuitive. It runs from a very wide DC power range, and includes a worldwide compatible AC/DC adapter. We offer signal conditioners for every kind of sensor, including strain gauges, accelerometers, thermocouples, low and high voltages, and more.

Need more channels? Plug in up to 16 of our 8-ch EPAD2 series "bricks" for thermocouples, RTDs, 50 VDC, and 4-20 mA inputs, and add 128 more channels!

Call today! +1 (401) 284-3750
www.dewamerica.com/D300

Please contact us to set up a demonstration today!

- 16 FAST analog input channels
- Up to 128 SLOW analog inputs
- Up to 1 MS/s/ch sample rate
- Choice of signal conditioners: voltage to ±200V, IEPE accelerometers, and multi-function strain gage inputs
- Two or more counters, plus DIO
- Optional CAN BUS, ARINC, 1553...
- Turn-key software - plug and play
- Made in USA
- Fast Delivery Time!

DEWETRON


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