IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 07.16

POLICE BODY CAMS

sometimes do more harm than good—but it's a problem that tech can solve

P. 30

IEEE

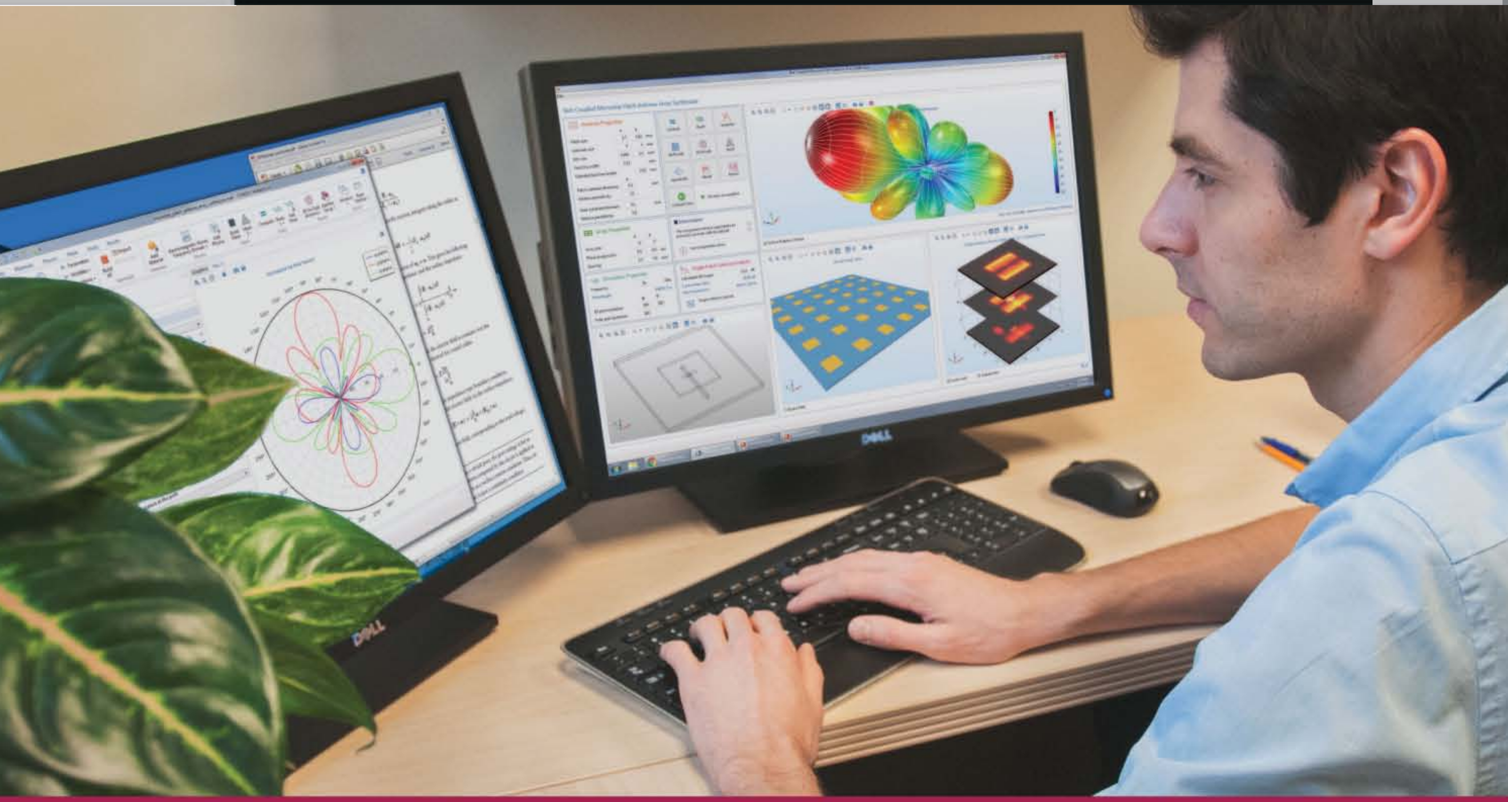

MULTIPHYSICS FOR EVERYONE

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

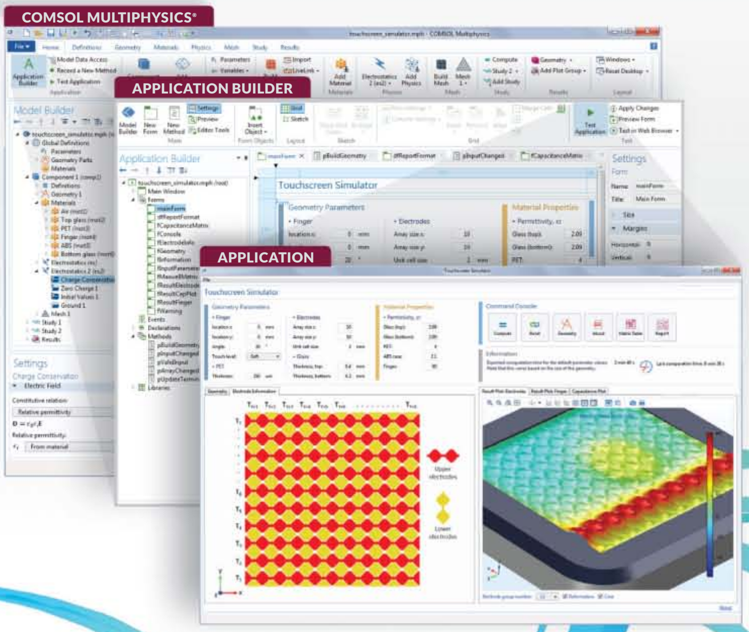

COMSOL

IEEE SPECTRUM

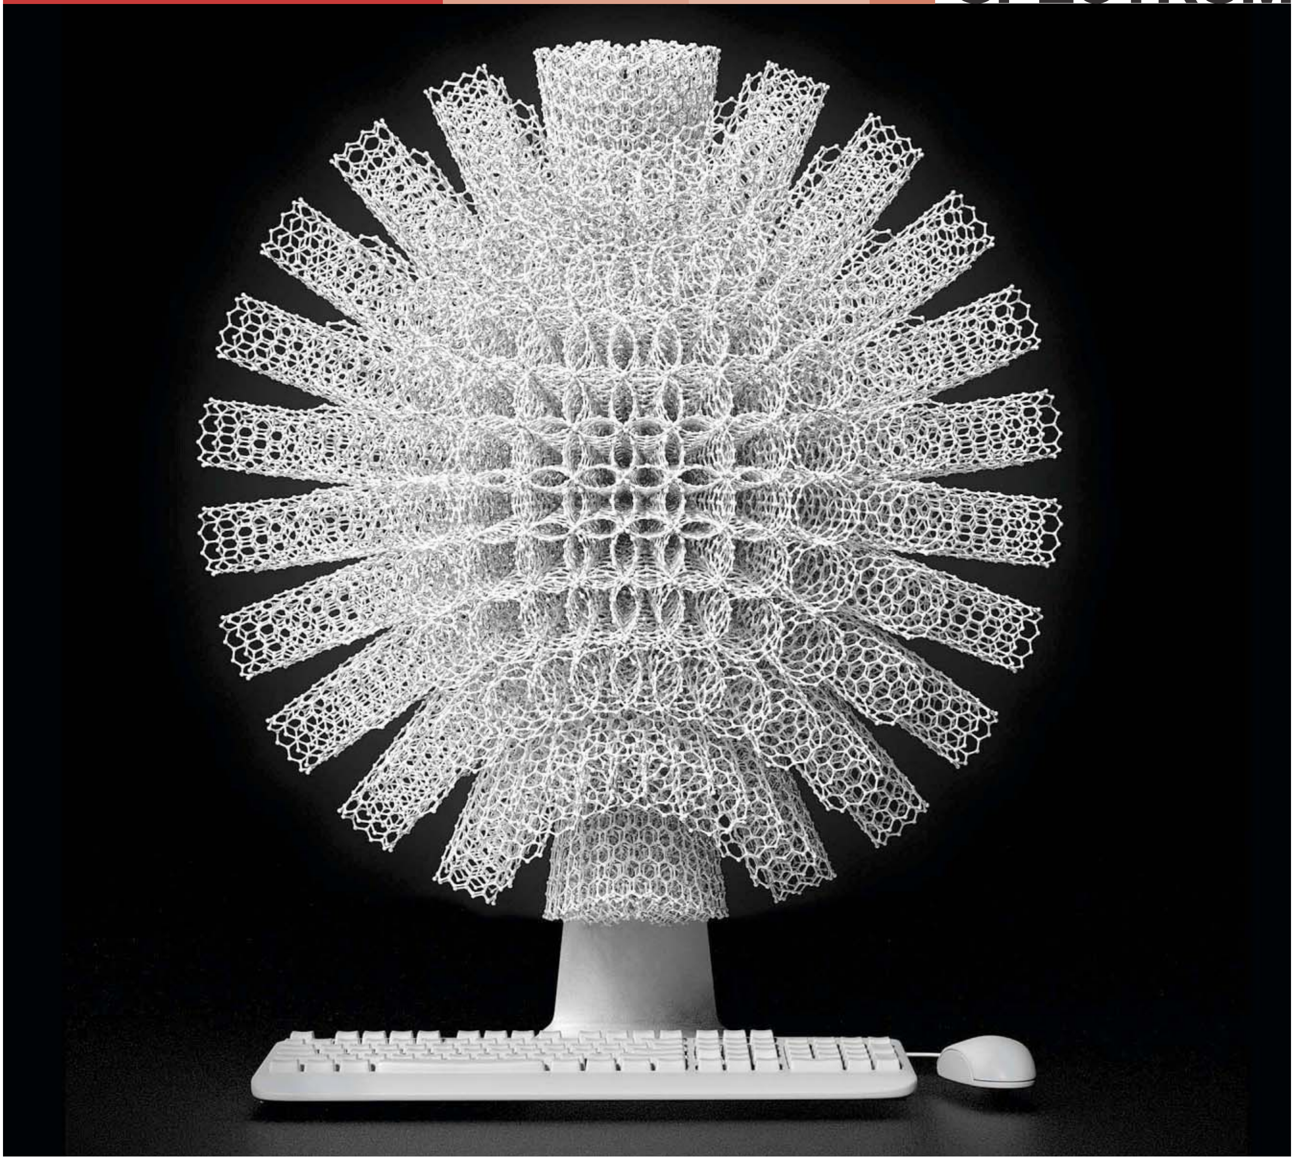

On the cover **Illustration for *IEEE Spectrum* by Flavio Montiel Cezaretti**

ILLUSTRATION BY **Eddie Guy**

DEPARTMENTS_07.16

IEEE SPECTRUM



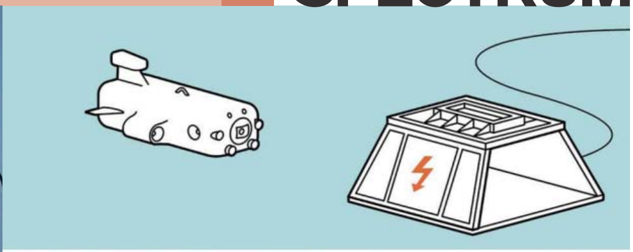
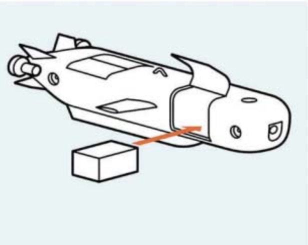
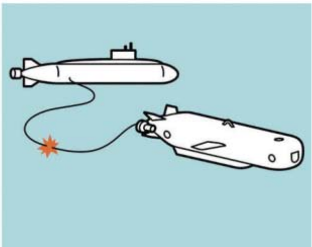

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Addressing Quantified Health With an Integrated Chip Solution—Available On Demand
- How Simulation Is Assisting With the Adoption of Internet of Things Through the Development of Smart Medical Implants—Available On Demand
- PTC and Scrum Inc. Present: PTC AgileWorx for Product Engineering Teams—Available On Demand

- **ANSYS ENGINEERING SIMULATION RESOURCE LIBRARY** http://spectrum.ieee.org/static/ansys-engineering-simulation-resource-library
- **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

MADE-TO-ORDER CLOTHES Nearly 10 percent of all unsold new clothing is thrown away. The founders of Unmade Studios aim to eliminate that waste with their suite of digital fabrication software for industrial knitting machines. The goal is to produce high-quality, made-to-order clothing in small batches for about the same price as mass-produced garments.

VIRTUAL BEAUTY Senior Member Parham Aarabi is founder and CEO of ModiFace, the world's largest and most advanced virtual-reality makeover platform. The company's technology is used in many of today's skin-analysis and facial-visualization systems.

CONNECTING CONTINENTS It's been 150 years since the first successful transatlantic cable was activated, reducing the time it took to send messages between Europe and North America from weeks to seconds.

FROM LEFT: GETTY IMAGES; RANDI KLETT; JAMES PROVOST (3)

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

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

BACK STORY_

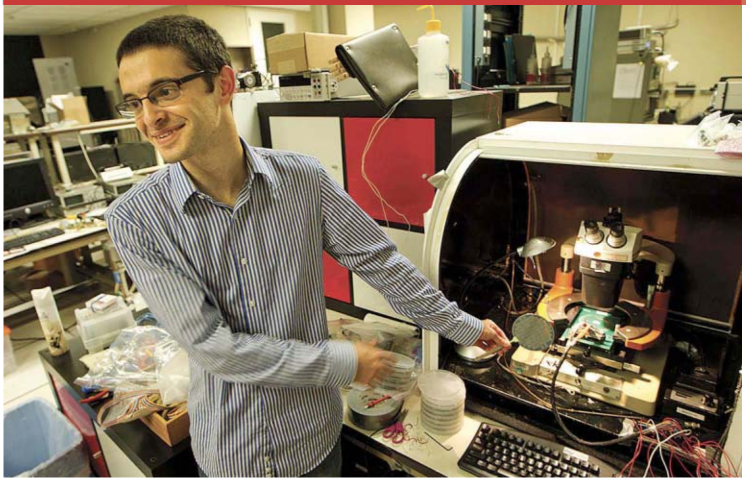

NORBERT VON DER GROEBEN

Carbon Country

MAX SHULAKER AND HIS COLLEAGUES made a splash in 2013 when they published a paper in *Nature* that described the first computer to use carbon-nanotube transistors. But, as with many research projects, that paper was the product of a lot of hard, unglamorous work. Months earlier, for example, team members pulled an all-nighter packing and moving test equipment, including a probe station [above], about 30 miles north, from their base at Stanford to the IEEE International Solid-State Circuits Conference in San Francisco. The team had been invited to perform a live demonstration of a handshaking robot that uses carbon-nanotube circuitry, and they needed the gear to connect those circuits to the outside world.

Since then, the station has largely remained in the lab, where Shulaker and his colleagues have worked on strategies to make the mass manufacture of carbon-nanotube circuitry commercially viable. The effort is described in this issue in "Computing With Carbon Nanotubes."

Shulaker, who recently completed his Ph.D., has been involved in the research since 2009. He was just finishing his freshman year at Stanford when he was drawn to the carbon-nanotube work being conducted in the labs of professors Subhasish Mitra and H.-S. Philip Wong, his coauthors on this month's article. The choice was a serendipitous one, Shulaker says. Silicon devices are approaching hard limits, and the semiconductor industry is hunting for strategies that will let it continue to make faster and cheaper circuits. "It's rare to have such a mature industry have so much uncertainty," he says.

Shulaker will continue work on carbon nanotubes in his laboratory at MIT, where he will start as an assistant professor this year. "I think my students will be in an even better position than I was in," he says, "because now we have enough building blocks to do some really amazing things." ■

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. 7 (INT), July 2016, p. 56, or in *IEEE Spectrum*, Vol. 53, no. 7 (NA), July 2016, p. 60.

07.16

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 EDITORS
Celia Gorman (Multimedia), celia.gorman@ieee.org
Amy Nordrum, a.nordrum@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.

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

PHOTO-ILLUSTRATIONS BY **Gluekit**

CONTRIBUTORS_

Barak Ariel

Ariel is a lecturer at the University of Cambridge's Institute of Criminology and an assistant professor at the Hebrew University of Jerusalem's Institute of Criminology. In 2012 and 2013, he helped to carry out an influential study of police body cameras, which showed them to be stunningly effective. "The very first month, we realized we were onto something," he says. But more recent results have been mixed, as he describes in "The Puzzle of Police Body Cams" [p. 30].

Theresa Sullivan Barger

Barger is a freelance business writer who contributes to *The New York Times, The Boston Globe,* and the National Society of Black Engineers' *Career Engineer* magazine. In this issue she looks at how companies can retain their female employees and the benefits they reap from doing so [p. 20]. "Women, overall, still shoulder more family and household responsibilities than do men," Barger notes. "They generally find greater job satisfaction in a collegial, collaborative work environment."

Alexander Hellemans

Hellemans, a physics reporter based near Naples, Italy, has been covering quantum computing and similar technology for *IEEE Spectrum* for more than 15 years. When he recently heard the European Union was going all in with a €1 billion effort, he says, "I thought it was a little bit late in coming." Still, this better-late-than-never project gave Hellemans the opportunity to cover the latest developments in Europe's quantum-computing research [see "Europe Bets €1 Billion on Quantum Technology," p. 10].

Harish Krishnaswamy

Krishnaswamy, an assistant professor of electrical engineering at Columbia University, specializes in getting semiconductors to do more with less. In 2013 he made CMOS transistors generate a record fivefold amplification of millimeter-wave signals. Now, with coauthor Gil Zussman, an associate professor at Columbia, he's made chips that direct radio signals differently depending on whether the waves are coming or going, producing a one-chip device that can talk and receive at the same time [p. 36].

Terry Ngo

Ngo is cofounder and CEO of Ignition Design Labs, where he has been rethinking the way Wi-Fi routers work [see "Why Wi-Fi Stinks—and How to Fix It," p. 42]. His résumé includes stints at Qualcomm and Atheros Communications, and he holds several patents on Wi-Fi and GPS communications. Ngo also collects vintage Wi-Fi routers, the oldest being an SMC Barricade 7404, which he bought in 2001; it has a maximum wireless data rate of 11 megabits per second.

Thwarting Poachers With Data

Remote, real-time monitoring of animals lets wildlife conservationists get the drop on poachers

OUR TOYOTA TRUCK burst through the dense jungle and into a clearing. Before us stood two startled men with rifles. Their wives, nearby, brandished machetes. Piles of dead animals were stacked high on two bicycles.

We had found our illegal hunters. In an instant, the soldier sitting next to me raised his pistol. I feared a shootout, but the bush hunters lowered their weapons and began to plead with the soldier in a shared language. Quickly a deal was struck: The hunters got to escape with their wives and pillage, while we got the prize we sought, a newborn chimpanzee.

We found the baby male chimp staked to the ground, a rope around his neck. The hunters had killed his parents. Once untied, the chimp clung tightly to the chest of the expert caregiver with us.

Our rescue mission happened 10 years ago in central Cameroon at a time when wildlife protection was becoming militarized. Park rangers, or soldiers on loan, like mine, routinely engaged in shootouts with poachers. Successes were rare. Too often, good guys paid with their lives.

They still do, which is why clever engineers are inventing new forms of "wild tech" that keep both rare animals, and their human protectors, safer. The revolution is based on remote, real-time monitoring of African animals, relying on attached collars or sensors, satellite tracking, embedded sensors in the landscape, and–in the near future–swarms of small drones. To analyze the vast and expanding pool of data about individual animals requires monitoring algorithms that rapidly create indicators about movement, proximity, and location, using "geofencing," or virtual perimeters of real-world geographic areas.

The goal, simply, is to engineer a "digital net" around every valued animal, so that it can be watched closely and continuously. Advocates of this technology hope that threat precursors can be identified and protective actions set in motion in advance. The ambitious program demands continued miniaturization of sensors, growing communication channels of various kinds, and novel software.

For large animals, such as elephants and giraffes, the digital tools are already revolutionizing conservation practices, as is demonstrated in a 2014 paper by elephant expert Jake Wall and his colleagues, based on their real-time monitoring of 94 elephants. As he and his team concluded, "Advancement of technology and the continued expansion of communications networks are allowing targeted, on-animal data collection and the economical and expedient distribution of such data."

For small animals, monitoring technologies offer even more dramatic advances. Many small animals are geniuses at hiding in the wild. Tracking them individually, even with "camera traps" embedded in the terrain, gives only a partial picture of their movements and the human threats they face. New techniques promise clear pictures of classes of wild animal life that receive only flickering documentation today.

The next phase of the environmentalism we want will be a kind of cyberwar between human poachers and human protectors. The battle lines will no longer exist only in the bush and savannas, along rivers and atop mountains. Soon, park rangers and amateur wildlife guardians will not draw guns on outlaw hunters. Instead, they will counterattack with sensor networks, drones, and dense computational studies managed remotely, with the touch of a keyboard.

In a cruel paradox, when the digital net fails, we might watch an animal die in real time, as gripped by the deadly spectacle as we now are by the Super Bowl. What seems only the stuff of computer games will be woven into the fabric of our lives, and that of the animals we seek to save. And we will call this progress.
—G. PASCAL ZACHARY

G. Pascal Zachary teaches at Arizona State University's School for the Future of Innovation in Society. The chimpanzee he helped save 10 years ago is alive and well in the Sanaga-Yong Chimpanzee Rescue Center, in Cameroon's Mbargue Forest, thanks to the efforts of the center's founder, Sheri Speede.

CORRECTION: In "Doc Bot Preps for the E.R." [June], the names of two members of the research team at the University of California, Berkeley, were misspelled in a photo caption. The correct spellings are Animesh Garg and Aimee Goncalves.

CARL DE SOUZA/AFP/GETTY IMAGES

NEWS

50 NANOGRAMS

MASS OF DNA ENCODING A 5.27-MEGABIT BOOK IN AN EARLY EXPERIMENT

TECH COMPANIES MULL ARCHIVING DATA IN DNA

Could biological storage keep up with big data?

TEST TUBE BITS: Biology's data-storage method, DNA, might work for our data, too.

GETTY IMAGES

It was the looming sense of crisis that brought them together. In late April, technologists from IBM, Intel, and Microsoft joined an intimate gathering of computer scientists and geneticists to discuss the big problem with big data: Our data storage requirements are rapidly exceeding the capacity of today's best storage technologies: magnetic tape, disk drives, and flash memory.

The closed-door meeting in Arlington, Va., was convened to explore the potential of a new storage technology that is actually as old as life itself. The experts came together to weigh the merits of DNA data storage, which makes use of the marvelously compact and durable DNA molecules that encode genetic information inside living things. By converting digital files into biological material, warehouse-size storage facilities could theoretically be replaced by diminutive test tubes.

While this idea has been kicking around for many years, meeting attendee Victor Zhirnov says tech companies are now starting to consider DNA data storage as a real possibility. Zhirnov, » director of cross-disciplinary research and special projects for the Semiconductor Research Corp. (which cosponsored the meeting), says he was encouraged by the presence of "luminaries" from industry and academia who took active part in the two-day workshop. "The question was, can we demonstrate a prototype DNA storage machine within five to seven years?" explains Zhirnov. "It is a very ambitious goal, but we concluded that it is possible."

Here's how DNA data storage works. First you take any digital data that would normally be stored in a binary code of 0s and 1s and translate it into the genetic code of As, Cs, Gs, and Ts that represent the chemical building blocks of DNA. Then you give that DNA code (for example, GATTACA) to a synthetic biology company, which manufactures strings of DNA to your specifications. Next you stash the test tube in cold storage and walk away. When you want to retrieve the information, you take out the test tube and use a standard DNA sequencing machine to decode the material inside. That gives you once again the DNA sequence GATTACA, which you can translate back into binary to read your original file.

DNA is the densest storage medium in existence, able to store almost a zettabyte of data in a single gram of material. It's also extremely long lasting, as demonstrated by remarkable feats of paleontological derring-do. In 2013, for example, a team reconstructed the entire genome of an early horse species using DNA from a bone that was buried in the Arctic permafrost for some 700,000 years.

If DNA archives become a plausible method of data storage, it will be thanks to rapid advances in genetic technologies. The sequencing machines that "read out" DNA code have already become exponentially faster and cheaper; the National Institutes of Health shows costs for sequencing a 3-billion-letter genome plummeting from US $100 million in 2001 to a mere $1,000 today. However, DNA synthesis technologies required to "write" the code are much newer and less mature. Synthetic-biology companies like San Francisco's Twist Bioscience have begun manufacturing DNA to customers' specifications only in the last few years, primarily serving biotechnology companies that are tweaking the genomes of microbes to trick them into making some desirable product. Manufacturing DNA for data storage could be a profitable new market, says Twist CEO Emily Leproust.

Twist sent a representative to the April meeting, and the company is also working with Microsoft on a separate experiment in DNA storage, in which it synthesized 10 million strands of DNA to encode Microsoft's test file. Leproust says Microsoft and the other tech companies are currently trying to determine "what kind of R&D has to be done to make a viable commercial product." To make a product that's competitive with magnetic tape for long-term storage, Leproust estimates that the cost of DNA synthesis must fall to 1/10,000 of today's price. "That is hard," she says mildly. But, she adds, her industry can take inspiration from semiconductor manufacturing, where costs have dropped far more dramatically. And just last month, an influential group of geneticists proposed an international effort to reduce the cost of DNA synthesis, suggesting that $100 million could launch the project nicely.

The U.S. Intelligence Advanced Research Projects Activity (IARPA) cosponsored the meeting and may fund a research program to create a prototype "DNA hard drive," but Zhirnov says that hasn't been confirmed. "If such a program can be established, the teams are ready," he says. Research would likely focus first on the most obvious appli-

Digits to DNA and Back Again

1. Any digital file—a movie, medical records, the *Encyclopedia Britannica*—can be converted to a "genetic file" and stored as strands of DNA. First the digital file's binary code is translated into the four-letter genetic code, composed of the As, Cs, Gs, and Ts that represent the chemical building blocks of DNA strands.

2. Then a synthetic-biology company manufactures the strands to the customer's specifications.

3. A test tube containing the genetic file can be stashed away in cold storage until someone wants to retrieve the information.

4. A standard DNA sequencing machine reads out the genetic code.

5. The code is then translated back into binary.

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

cation for a DNA hard drive: using it for archival storage, in which the data remain unchanged until the entire file is retrieved for readout. However, an IARPA program could also fund researchers who have recently demonstrated that DNA can be used for random-access memory and can even be made rewritable.

Biotech consultant Rob Carlson attended the meeting, and says he expects that the intelligence agencies of several countries will fund work on DNA data storage to grapple with the onslaught of

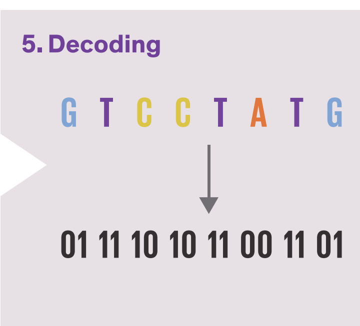

information now being gathered by surveillance technologies. "They're scratching their heads," he says, "and there's nothing else in the offing that can meet their storage needs." Carlson has written skeptically about the commercial market for synthetic DNA, yet he says that DNA data storage may be the application that makes the young industry viable. "We can imagine storing massive amounts of data in a very small volume, and we already know how to read and write it," he says. "Now the question is, can we read and write it at high throughput and at low cost?"
—ELIZA STRICKLAND

ISTOCKPHOTO

NEWS

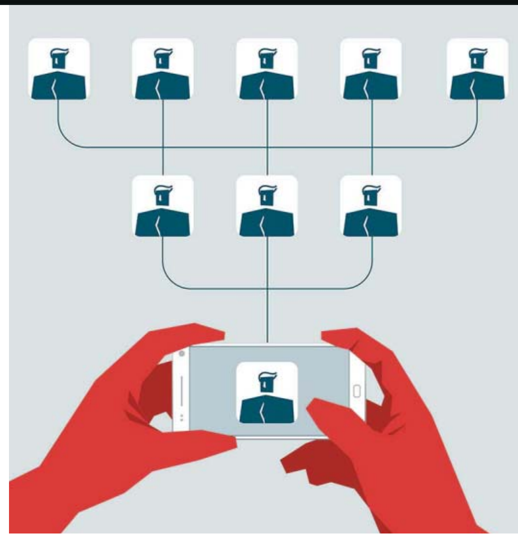

ANALYSIS SHOWS THAT U.S. INTELLIGENCE AGENCIES SEE LESS PHONE METADATA

A before-and-after glimpse into how Edward Snowden reshaped U.S. privacy policy

For years, the U.S. National Security Agency amassed metadata from the private phone lines of millions of people in the United States. This included information about whom a person called, when those calls were placed, and how long they lasted. When former NSA contractor Edward Snowden divulged this practice in 2013, his revelations undermined public trust and eventually led to nationwide privacy reforms. • At the time of his disclosure, the agency was operating under a "three hop" rule secured through the USA Patriot Act. This law permitted the agency to legally collect metadata from not only a terrorism suspect's phone but also from someone who called someone *who called someone else* who then called a suspect. Investigators could gather these records going back five years from the date of the request.

Soon after Snowden went public, legislators passed the USA Freedom Act to restrict the NSA to accessing metadata within two hops of a single "seed," or suspect, and to let it trace records back only 18 months. But Patrick Mutchler, a Ph.D. candidate at Stanford University, says citizens at the time had little information about the real-world impact of this policy shift.

That's why Mutchler and his Stanford colleagues recently published an analysis of phone metadata from 823 volunteers, mostly young men, showing the dramatic decrease in the agency's reach by removing a single hop and shortening the collection time frame [see illustration, "A Smaller Net," on the next page].

At the same time, their data also illustrate just how far the NSA's reach extends even today. Its agents can still request metadata from tens of thousands of private phones indirectly connected to the phone number of a suspected terrorist.

While they were at it, Mutchler's team also wanted to see if it could infer personal information about people based solely on calling patterns. It found that in at least some cases, investigators could surmise details about a person's location, health, relationship status, and religious beliefs by matching call records with publicly available databases of businesses

and organizations. The team also correctly identified the owners of a large percentage of random phone numbers using only Google searches and other public sources, casting doubt on claims that anonymous databases such as those used by the NSA protect the identities of those within.
—AMY NORDRUM

A SMALLER NET: Below are the numbers of people the NSA could legally collect phone metadata from in the course of a single investigation before and after the USA Freedom Act. A "hop" indicates a direct connection—that is, one person called another or sent him or her a text.

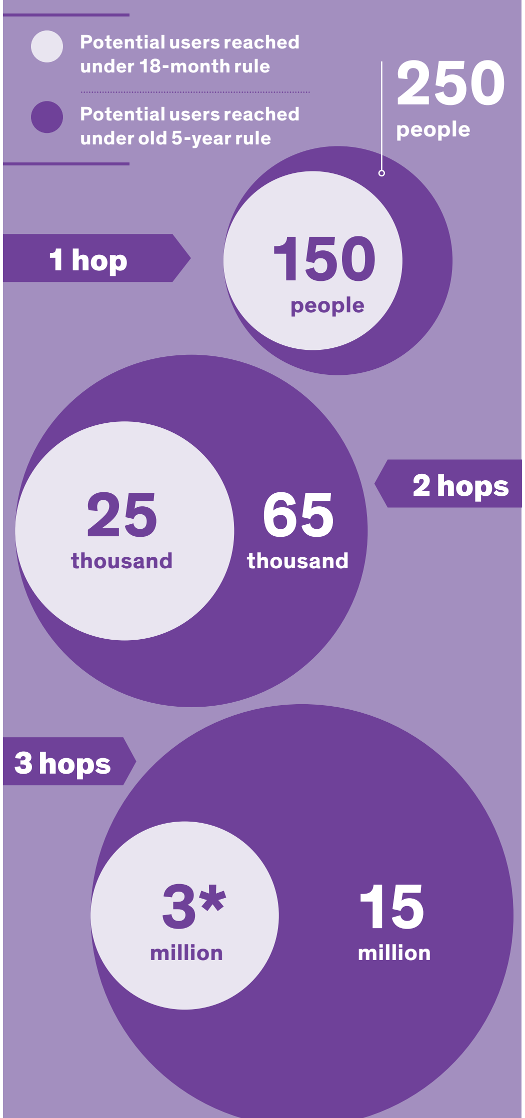

* For comparison only. The three-hop condition is not legal under the new rule.

EUROPE BETS €1 BILLION ON QUANTUM TECH

A 10-year-long megaproject aims to turn quantum physics into advanced tech

European quantum physicists have done some amazing things over the past few decades: sent single photons to Earth orbit and back, created quantum bits that will be at the heart of computers that can crack today's encryption, and "teleported" the quantum states of photons, electrons, and atoms. But they've had less success at turning the science into technology. At least that's the feeling of some 3,400 scientists who signed the "Quantum Manifesto," which calls for a big European project to support and coordinate quantum-tech R&D. The European Commission heard them, and answered in May with a €1 billion, 10-year-long megaproject called the Quantum Technology Flagship, to begin in 2018.

"Europe had two choices: either band together and compete, or forget the whole thing and let others capitalize on research done in Europe," says Anton Zeilinger, a physicist at the University of Vienna who did breakthrough work in quantum teleportation, which would be key to a future Internet secured by quantum physics.

The European Commission formally announced the funding for this project in May at a conference in Amsterdam, cohosted by QuTech, the most prominent quantum research lab in The Netherlands. There, about 350 scientists, representatives of European industry, and emissaries from U.S.-based global tech firms such as Google, Lockheed Martin, and Microsoft, hammered out the project's priorities. "The organization and shape of the flagship still has to be defined," says Anouschka Versleijen, a materials scientist and program director at QuTech. "But the ball has been set rolling."

In Amsterdam, the potential quantum technologies debated went far beyond the usual suspects: unbreakable encryption, an unhackable Internet, and computing that can crack problems it would take today's machines 100 years to solve. The new technologies included quantum simulation, quantum sensors, quantum imaging, quantum clocks, and quantum software and algorithms.

In quantum simulation, purpose-built quantum computers would perform quantum-mechanics-level modeling of

ILLUSTRATION BY Erik Vrielink

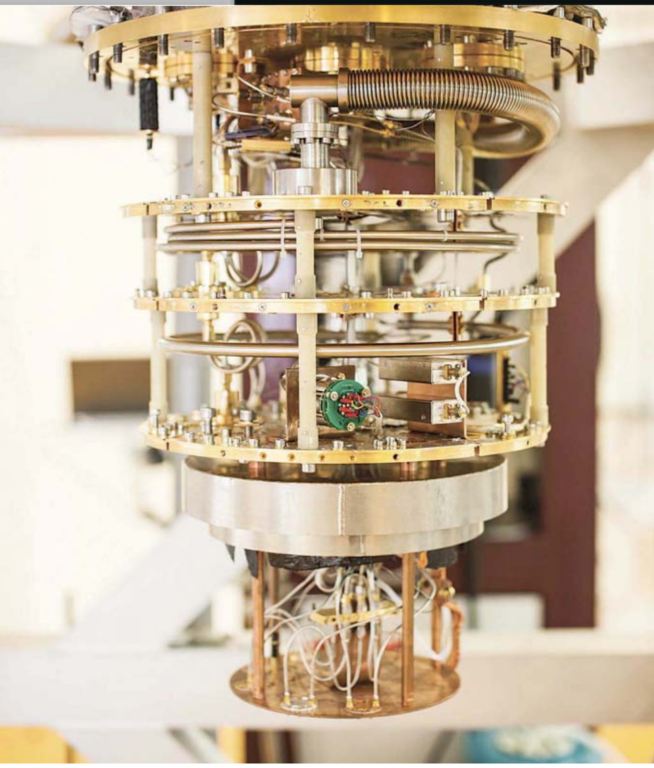

ERNST DE GROOT

COOL QUANTUM TECH: This dilution refrigerator can cool quantum dots to less than 5 millikelvins for experiments in quantum computing.

materials, which would be impractical on today's classical computers. The simulations would elucidate the fine structures of superconductors and map out complex chemical reactions to predict whether a newly engineered material would be stable.

Quantum sensors and quantum imaging will be especially useful in medicine. For example, they'll allow new ways to sense the heart's magnetic field, which could more accurately diagnose and distinguish heart diseases. "You will be able to obtain images of things that we were never able to see before," says Mark Everitt, a physicist focusing on quantum engineering at Loughborough University, in England.

Quantum clocks, which track the vibrations of a single atom to provide almost unimaginable accuracy, will serve a wide range of purposes, including accurate measurements of the local gravity potential and precision timing of financial transactions, says Helen Margolis, of the United Kingdom's National Physical Laboratory. She reports that the best of these quantum clocks could be made so accurate that they'd gain or lose no more than 1 second every 30 billion years.

New quantum algorithms could allow quantum computers to process data at a much higher speed, allowing for database searches, machine learning, and image recognition with unprecedented speed. Making use of such algorithms might be made easier for a broader range of coders because of quantum compilers that Microsoft and others are working on.

With such a potentially valuable set of technologies in the offing, why didn't Europe go all in sooner? According to Zeilinger, a portion of the problem was in the communication between scientists and engineers. Some of the weirdness of quantum theory puts off engineers, he says. "This is an old problem in Europe: Industry in Europe is more skeptical than in the United States, although it is now losing some of its skepticism," he says.

Everitt agrees. Many areas of quantum mechanics are no longer problems of physics; they are now engineering problems, he argues. "For these areas, we will see great progress that will lead to new products." –ALEXANDER HELLEMANS

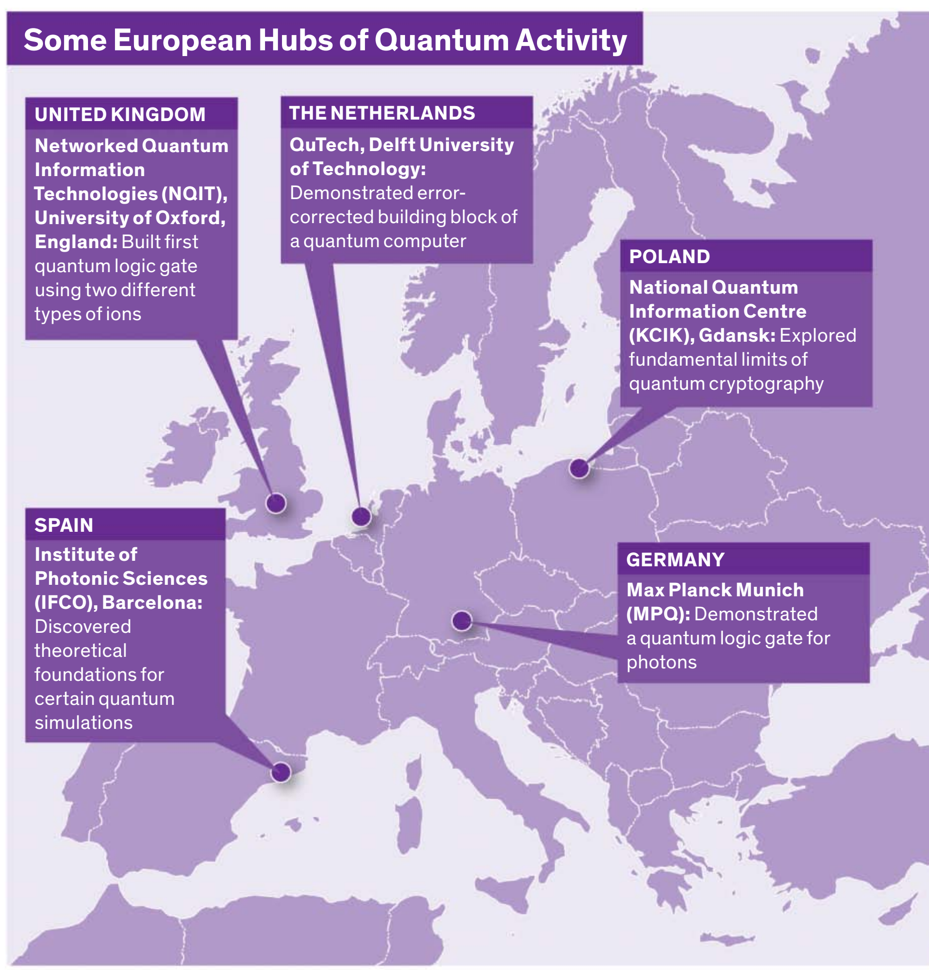

GOOGLE GLASS IS REVIVED IN THE ER

Experts can view patients from afar for quick emergency consults

Google Glass, despite its dramatic 2012 unveiling via a live skydiving demo, never became a consumer hit, and Google stopped selling the device in January 2015.

Now, a team of doctors at the University of Massachusetts Medical School (UMMS) may have discovered a killer app for the device: performing emergency-medicine consultations. It turns out that Glass enables off-site specialists to reliably and accurately observe and diagnose patients in real time. It may even help first responders triage victims in disaster scenarios.

Hospitals have tried using Glass before. Some efforts involved projecting medical records into a doctor's field of view, but physicians found the display to be too small to sort through charts and data, among other concerns.

Kavita Babu, Edward Boyer, and Peter Chai, medical toxicologists and emergency-medicine physicians at UMMS, hypothesized that it might be better to use Glass to send data out rather than for pulling it in. "As an emergency physician, you're really busy, and you end up making decisions with your specialists very quickly," says Chai. "A lot of those times you're talking to your specialist over the phone, and they're just hearing verbal descriptions. Everybody wants to be there to see the patient."

Telemedicine has been gaining popularity in health care, yet it typically consists of a computer and camera affixed to a bulky cart rolled from exam room to exam room. A simple head-mounted device with a camera and connectivity, on the other hand, is unobtrusive at the bedside and can transmit live, first-person images to a specialist anywhere.

Previously, Chai demonstrated the feasibility of using Google Glass in the emergency department of Rhode Island Hospital for dermatology consultations. Last year, the UMMS team successfully tested the use of Glass during toxicology consultations. An emergency-medicine resident wore Glass during a bedside evaluation of a patient and transmitted real-time video to a specialist to help with the diagnosis. The consulting toxicologists rated the audio-visual experience usable in 89 percent of the cases seen, and six patients received antidotes they would not have otherwise, based on recommendations from the consultants.

A QUICK CONSULT: Peter Chai, an emergency-medicine physician in Massachusetts, used Google Glass for consultations with specialists, such as dermatologists.

Most recently, Chai collaborated with Banner Health in Phoenix to evaluate the precision and reliability of a Google Glass-assisted consultation in a toxicology intensive care unit. Fifty patients at the Banner-University Medical Center Phoenix had an off-site toxicologist record physical examination findings through Google Glass during their bedside examinations by a toxicology fellow. The researchers compared the results of the two exams and found a high correlation between them. The study is soon to be published in the *Journal of Medical Toxicology*.

Aaron Skolnik, the assistant director of Banner Poison and Drug Information Center, who led the study, hopes to eventually use Glass to triage poisoned patients in rural parts of Arizona to determine who should be transferred to the medical center and who should stay at their local hospitals for treatment. Though Glass was a flop with consumers, "for medicine, and maybe industrial applications, it's actually a really great technology," says Skolnik. "It provides a huge amount of extra data reliably and accurately to a remote person at relatively low cost."

Despite its potential usability, there are still two big hurdles to overcome before we'll see ER docs regularly sporting Google Glass. First, it remains murky how physicians will be reimbursed for virtual consultations, a problem that all telemedicine is currently struggling with. Second, health-care providers are still figuring out how to include telemedicine visits in medical records, such as what type of data to save and how to store it.

As those issues continue to be sorted out, the UMMS team is now planning to test Glass in the field with first responders as they triage victims at disaster sites. First responders, such as firefighters, often don't have extensive medical training but are relied upon to evaluate patients in advance of the arrival of emergency medical technicians. With a wearable head-mounted computer, they could communicate live images and audio to a specialist back at headquarters to assist in a diagnosis. Even in an area without Wi-Fi, the responders could carry Wi-Fi hotspots in their pockets to enable communication.

In September and October, the UMMS team will participate in a mock active-shooter scenario in Massachusetts, equipping 50 to 60 first responders with Glass to communicate with about 15 disaster physicians and trauma surgeons.

"The end goal of all this is to develop a suite of sensors and technology that allows us to respond to patients remotely," says Chai.

—MEGAN SCUDELLARI

UNIVERSITY OF MASSACHUSETTS MEDICAL SCHOOL

PHOTOGRAPH BY Peti Kollanyi/Bloomberg/Getty Images

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

NO TIPPING REQUIRED

WHAT WILL TEENAGERS do to earn money after school when pizzas are delivered by self-driving robots? Starship Technologies, based in Tallinn, Estonia, is pointing us in that direction. The company has developed a six-wheeled delivery bot whose movements are tracked as it ferries packages or groceries. Its locked compartment can be opened only by the intended recipient with a unique code. The bot, which can accommodate two bags of groceries, travels on sidewalks like a mechanized pedestrian. Cameras and sensors help it avoid objects and people—such as the teens it'll put out of work.

RESOURCES

1943

THE FIRST YEAR DIGITAL AUDIO WAS USED, TO PROTECT SECRET PHONE CONVERSATIONS

REVIEW: DRAGONFLY DACS

UPGRADE YOUR SMART-PHONE'S AUDIO WITH THESE SLIM DIGITAL-TO-ANALOG CONVERTERS

RESOURCES_TOOLS

If you're a music lover, you probably know that plug-in devices are available to improve the audio quality of music files played back from your smartphone or laptop. The most intriguing of these devices combine a digital-to-analog converter (DAC) with a headphone amplifier. These USB DAC-headphone amps have been around for a few years, but some are as big as a phone themselves. ● Now a new generation are wonderfully compact—about the size of a thumb drive—and can handle high-definition music files consisting of 24-bit samples and a sample rate of 96 kilohertz. Of the half dozen such units on the market, two that stand out are the DragonFly Red (US $199) and DragonFly Black ($99) from AudioQuest in Irvine, Calif. ● Before I describe the sound from these little powerhouses, some background is in order. Starting with Android 5.0 (Lollipop), late in 2014, Google finally accommodated long-suffering audiophiles by enabling music streaming directly out of the USB port. This freed music lovers from the built-in DAC, which operates in the electromagnetically noisy environment inside the phone. (iOS users can get the audio using Apple's camera adapter, which costs $40.) ● Google's decision was a big deal for audiophiles because unlike Apple's mobile devices, many Android phones have a slot for a microSD card, which can give you an extra 128 gigabytes. This is important because a single high-definition song stored in the popular Free Lossless Audio Codec, or FLAC, format can occupy anywhere from 6 to 195 megabytes, so even a modest music library would swamp the internal storage of today's smartphones. ● To use the DragonFly devices with an Android phone, you need a USB On-the-Go cable, which costs about $4, and a music-player app such as ▸

PHOTOGRAPH BY Randi Klett

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

USB Audio Player Pro ($8) that will direct music to your phone's USB port.

I tried the DragonFly Red and Black with a variety of musical genres—orchestral, chamber music, jazz, and rock—and in formats ranging from standard MP3 to 24-bit/96-kHz FLAC. I listened with four headphones: the Audio-Technica IM70 earphones (retail price, $117); the NAD Viso HP50 ($249); the Audeze Sine ($449); and the Audeze EL8 Open-Back ($699). (Those last two were graciously made available by Andrew Lorberbaum of Park Avenue Audio in New York City.)

I'll get right to the point. The DragonFly DACs can turn your phone into a bona fide high-end audio machine. They will improve not just stored music but also streaming services such as Spotify and Google Play Music. But although the DACs themselves are inexpensive, you will not derive much benefit from them if you are using basic headphones: While listening through the relatively inexpensive Audio-Technica IM70s, I could not hear much difference between the DragonFly DACs and my Samsung Galaxy S5's normal sound.

But with better headphones, the differences were quite apparent, and the more so as I went up the cost spectrum. With the Audeze EL8s, the differences were stark: Bass notes were tighter and much more defined, percussion was sharper, and there were details and textures in vocals and strings that were simply not there without the DragonFly DACs. Listening with the Audeze EL8s to a recording of Mozart's clarinet quintet, I could close my eyes and visualize the cello at the right of the viola. Very rarely have I been able to perceive such a soundstage with mobile gear.

Much more subtle were the differences between formats, sampling rates, and sample sizes. I could generally hear a clear difference between MP3 and FLAC files. But my 54-year-old ears could hardly ever distinguish between FLAC files at, say, 16-bit/44.1-kHz and 24-bit/96-kHz resolution. So my suggestion is, if you own lots of CDs, load your microSD card with music ripped from them at 16 bits/44.1 kHz. If you're an audiophile, the DragonFly Red is a compelling proposition. The DragonFly Black is also well worth auditioning: It differs from the Red in its slightly lower power output and its analog, rather than digital, volume control. To my mind, at $99, it is one of the best bargains in high-end audio today.

—GLENN ZORPETTE

HI-FI: The author tested a DragonFly DAC with Audeze EL8 headphones. Good headphones are required for noticeable results.

DIY LIGHT-BASED 3D PRINTING

AN ENTICING IDEA TURNS OUT TO BE DEVILISHLY COMPLEX

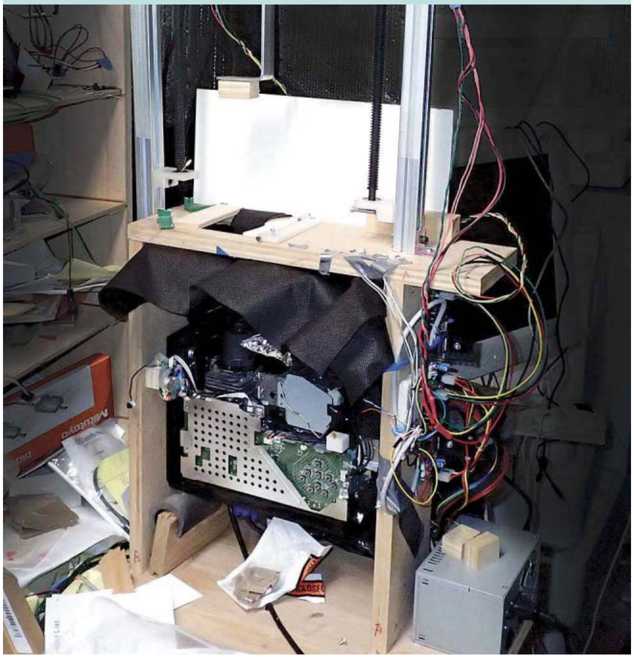

THE BASIC IDEA IS SO SEDUCTIVE. FORGET about forming an object from spools of plastic melted and laid down by a tiny nozzle as most consumer 3D printers do. Instead, use an optically reactive resin and modern display technology to zap an entire layer of solidified material into place in one shot. Then the next layer, and the next, until the object is complete. • After all, some professional stereolithography printers, such as the Form 1 and 2 and the Nobel 1.0, do exactly that. But could I, one person working at home in his spare time with mostly surplus parts, build such a machine? The answer is—spoiler alert!—*almost*. Like a tiny home-brew analogue of fusion power, my printer seems to stay just another month or two of concentrated work away from operation, no matter how many months go by. • My machine is based on a DLP printer, with the DLP standing for Digital Light Processing, a Texas Instruments technology that uses an array of individually controllable micromirrors. Ultraviolet light from a mercury bulb or an LED is directed at the array. Each micromirror sends light either to a beam dump or to the projector lens to focus an image

LEFT: RANDI KLETT; RIGHT AND OPPOSITE PAGE: PAUL WALLICH (5)

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

corresponding to a layer of the object. The image shines through the bottom of a vat to cure a microscopically thin layer of UV-sensitive resin. Instead of spreading the beam onto a wall, the system focuses the image onto an area maybe 10 centimeters across, so the physical resolution of the resulting 3D object is a small fraction of a millimeter. A movable stage lifts the cured layer up a fraction of a millimeter, more resin flows underneath, the projector displays the pattern for the next layer, and so on.

To construct this setup, I bought a DLP projector from an Internet overstock site for US $300. I tentatively dubbed my machine Lobachevsky (after the old Tom Lehrer song) because almost every part was lifted from some other project. The CPU for driving the projector is from an old digital picture frame, the aluminum extrusions from a disassembled ER1 robot, the stepper-motor drivers from my attempt at building a conventional 3D printer, and so forth.

After a few false starts, I figured out how to choreograph communications among the computer I was using to handle the 3D object data, the picture-frame CPU, and the microcontroller operating the stepper motors that raise and lower the stage. Then the real problems started.

Every step in the simple description of UV projection stereolithography I gave above hides a pitfall. For example, a standard consumer DLP projector uses a spinning filter wheel to subdivide each video frame into different time periods. During each period only red, green, or blue light passes, which also blocks UV light. So I had to get rid of the filter wheel. But I couldn't just unplug the wheel from the projector's motherboard, because the firmware will provide power only for the projection lamp after it's detected that the wheel is spinning at the right speed. Oh, and when you open up the projector to start messing with its internals, there's more firmware that reads a pressure switch that turns everything off, so you have to defeat that interlock too.

Meanwhile, the bottom of the vat where the printing happens must be transparent to UV and also completely unwilling to adhere to freshly cured resin. Other stereolithography home brewers have tried covering the bottom with a layer of PTFE film; a silicone gel used to encapsulate solar cells; various

LIGHT WORK

UV printing can allow objects designed on a computer [top] to be printed with high resolution. My printer uses a DLP projector [second from top] connected to a CPU taken from a digital picture frame to control stepper motors [second from bottom]. The motors raise the object as each layer is exposed to a pattern of light [bottom].

soaplike substances; and even carefully calibrated salt- or sugar-water solutions that are denser than the resins and so pool beneath them. When I tried the liquids, I was treated to an amazing display of surface chemistry: As long as I had just a tiny pool of resin on top of my solution, it floated in the middle of the vat just fine, but as soon as any resin touched the edge of my container, it would pull the rest of the pool over to the edge and then down to the bottom.

The encapsulant seemed to work better; objects formed, and sometimes the cured resin successfully detached from it as my stage lifted up. But that's when I discovered an entirely different problem: The resin wasn't actually curing as a result of the UV exposure from the carefully focused images coming from the projector. Instead, the pattern of light from the projector was heating small areas of the resin, increasing its sensitivity to ambient UV from light leaking out of the projector. As soon as I cut off those leaks, I stopped getting any curing at all. Further investigation showed that the projector optics included a previously unsuspected UV filter incorporated into a block of glass that protects the rest of the system in case the projector's bulb explodes. So I replaced it with a simple block of heat-resistant glass of precisely the same dimensions.

But while I was acquiring that block and installing it, time passed. My UV-sensitive resin, as complex chemical mixtures will, lost some of its sensitivity. It still curdles in response to a UV LED, but the light from the projector does nothing for it at all. I have more resin on order.

While I wait, there's another, potentially much simpler method that's attracting my interest. I took the backlight off an old LCD and now I can shine whatever I want through it. A UV LED produces a perfectly legible first layer at least, just like the contact prints 19th-century photographers made from their page-size glass negatives. The resolution of the old screen is miserable, and I know the ultraviolet will eventually damage the liquid crystals in the display, but it was due for recycling anyway. And if this version works, maybe I can convince someone to give me an out-of-date tablet or mobile phone, with 10 or 20 pixels to the millimeter. I'm sure it will be a simple matter of software to control the display.... —**PAUL WALLICH**

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

STRATEGIES FOR RETAINING FEMALE ENGINEERS

BOOSTING DIVERSITY CAN HELP THE BOTTOM LINE

IZABELA HABUR/GETTY IMAGES

Women earned 49 percent of all science and engineering bachelor's degrees, 43 percent of science and engineering master's degrees and 40 percent of science and engineering doctoral degrees in 2014, reports the National Student Clearinghouse. Yet women make up less than 25 percent of the STEM workforce and only 10.5 percent of employed engineers. Researchers have found that workplace culture and women's personal character traits play major roles in retention. So what are the things that make a difference?

Women prefer workplaces that are collaborative rather than hierarchical, explains Heather Metcalf, director of research and analysis at the Association for Women in Science. And they are more apt to stay in work environments that allow for creativity and flexibility, she says.

Conversely, women are fleeing companies that encourage employees to practically live at work, she says. While 71 percent of women with young children work outside the home, according to the Pew Research Center, women still shoulder more responsibility for child care and elder care than men. So living at the office to show they are committed to their jobs is not an option.

"Creating workplaces that have a lot of flexibility, that allow for people to work in a way that fits best with them, boosts creativity and job satisfaction," Metcalf says, and these are the settings where women stay and thrive.

Working in a cooperative setting, she says, creates a greater sense of belonging and connectedness, which all adds up to more innovation—and more patents. Studies show, she says, that for STEM (science, technology, engineering, and math) firms, women-owned, women-led businesses perform better economically than similar men-owned companies.

Natalie Babij, a technical-support specialist with Renewable NRG Systems in Hinesburg, Vt., says she loves working in a team-based culture. Her primary job is to help customers with problem solving and troubleshooting, but when she's not doing that, she works with her colleagues to improve existing processes.

Babij appreciates the fact that she doesn't have a rigid job description. Her coworkers know about her background as a data analyst and are happy to make use of her knowledge. Software engineer colleagues ask her to contribute to the design and planning process of creating new software features. "It makes me feel important and useful, and it's nice that I can use those skills," Babij says.

Research shows that women like Babij are likely to stay in their fields because engagement at work is "a very high predictor of career commitment," says Diana Bilimoria, a Case Western Reserve University professor who has researched women in STEM extensively.

No matter what type of organization women work for, large or small, public or private, their relationships with their immediate bosses are critical to whether they feel engaged and content. The ideal supervisor is committed to his or her subordinates' advancement and development, assigns stretch projects, and provides necessary support and feedback to help them be successful, Bilimoria says. And workplaces that employ women in higher levels are more apt to retain women at the lower levels.

"There need to be multilevel champions [of women] from the top as well as from the bottom and the middle, because women are more sensitive to dealing with gender bias," she says.

Workplace initiatives that offer leadership development, mentoring, and networking for women reap the benefits by retaining women, Bilimoria's research shows.

Software developer Christina Yakomin and a college friend were each recruited out of Villanova University, near Philadelphia, into leadership development programs at Fortune 500 companies. Both were promised opportunities for mentoring and advancement.

Yakomin, who chose the financial firm Vanguard, found that young women were a larger percentage of the leadership development program than she expected. There were also multiple women in leadership positions, and she was encouraged to present in meetings.

Her friend, also in a leadership program that looked comparable on paper, isn't allowed to think, is given menial tasks, and feels professional development opportunities are limited. She's leaving her company and joining Vanguard. **—THERESA SULLIVAN BARGER**

PROFILE: KEPLER COMMUNICATIONS
BUILDING A SPACE NETWORK TO HELP SATELLITES STAY IN TOUCH

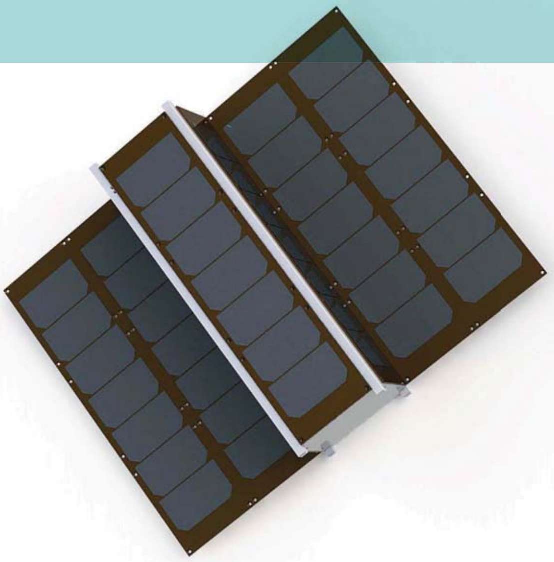

E.T. PHONE HOME: Kepler Communications plans to deploy small satellites to help other satellites stay in contact with their operators on Earth.

KEPLER COMMUNICATIONS

At the heart of the current wave of private space companies rests a single enabling technology: satellites. As these versatile vessels drop in price, there's no shortage of ideas for new satellite-based products and services. But there's a major stumbling block—as satellites orbit, they frequently lose contact with their owners. Any given satellite may fly over a ground station only once every few hours.

That's why Kepler Communications wants to build a grid of CubeSats (satellites roughly the size of a loaf of bread) to coordinate and relay messages sent by other satellites, acting as a service provider to fellow space startups. Kepler's founders aim to place their first two satellites into orbit by the end of 2017.

To assemble its complete network, Kepler's plans call for 50 satellites. They will be in a configuration that cofounder Mina Mitry says will guarantee that at least one Kepler satellite will always be within sight of satellites operating within the five most popular orbital planes.

At least one company-owned satellite will be passing above a ground station at any time. When clients send information to the nearest Kepler satellite, it can beam data around the network to find an available ground station.

As the private space industry heats up, there's growing enthusiasm for in-space infrastructure to serve companies and customers in this new, far-out economy. Richard M. Rocket, CEO at NewSpace Global, says his firm tracks more than 1,000 private space companies and expects that number to reach 10,000 over the next decade.

Kepler's debut also aligns nicely with the private space industry's shift toward standardization. Similar to the way satellites are now sold as modules that can easily be adapted to many missions, Kepler wants to provide a network that companies can simply plug into without the hassle of coordinating it on their own.

But to make this dream a reality, Kepler's founders must find a way to ensure the company's satellites can seamlessly communicate with other satellites while hurtling through space at roughly 28,000 kilometers per hour.

That's no easy task. The routing protocols written for the Earth-bound Internet rely on stable links to determine routes from computer to computer. These protocols don't work well in space, where links are constantly changing as satellites move toward and away from one another. Consequently, Mitry's team is designing a new data-transfer protocol for satellites.

A second challenge that Kepler will face is simply getting enough satellites into space in the first place. Though satellite and launch costs have dropped, few companies have managed to secure enough funding to send up a constellation of the size that Kepler intends to build. Kepler has earned investment of undisclosed amounts from four sources including Globalive Capital.

Despite the industry's poor track record, Joe Landon, chairman of Space Angels Network, says Kepler shouldn't be discouraged. He expects launch prices to fall further as companies including Blue Origin, SpaceX, and Virgin Galactic ramp up service.

"The entire industry is struggling to get enough capacity of satellites in space, but there's a lot of smart people working on that," he says. "I think that problem is going to get solved without Kepler having to do much."

Once its network is in place, Kepler will charge customers a monthly subscription fee based on the amount of data they transfer. Mitry hopes to achieve data rates of 1 to 40 megabits per second (8 Mb/s is roughly what is required to stream an HD video).

As Kepler prepares for its 2017 test run, NASA is leading its own effort to define space-based communications. In 2018, the agency plans to debut new standards and technologies to coordinate how commercial and public satellites interact with one another. Those developments may overlap or interfere with Kepler's system.

Mitry says Kepler wants to work with NASA and the rest of the private space industry to figure out what space-based communications and networks should look like in the future. Still, Badri Younes, NASA's deputy associate administrator for space communications and navigation, cautions companies against rushing ahead.

"Being there first does not guarantee that you will be successful," he says. "What you need to do is build it right." **—AMY NORDRUM**

Location: Toronto **Founded:** 2015
Employees: 6 **Funding:** Undisclosed

FLYING WITHOUT KEROSENE

ELIMINATING KEROSENE-BASED JET FUEL will be one of the greatest challenges to creating a world without carbon emissions. Aviation accounts for only about 2 percent of the global volume of such emissions and some 12 percent of the total released by the transportation sector, but converting to electric drive is much harder for airplanes than for cars and trains. ● Today's jet fuel–the most common formulation is called Jet A-1–has a number of advantages. It packs 42.8 megajoules into each kilogram, it can stay liquid down to -47 °C, and it beats gasoline on cost, evaporative losses at high altitude, and risk of fire during handling. No real rivals yet exist. Batteries capacious enough for intercontinental flights carrying hundreds of people are still the stuff of science fiction, and we will not see wide-bodied planes fueled by liquid hydrogen anytime soon. ● What we need is a fuel equivalent to kerosene that's derived from plant matter or organic waste. Such a biojet fuel would release no more CO_2 during combustion than the plants sequester during growth. The proof of principle has been demonstrated: Since 2007, test flights using blends of Jet A-1 and biojet have proved suitable as drop-in alternatives for modern aircraft. ● In March 2016, United began using biojet fuel at Los Angeles International Airport, taking deliveries from AltAir's Paramount refinery, where the fuel is made by hydroprocessing esters and fatty acids derived from beef tallow and pig fat and from agricultural wastes. Starting in 2018, Nevada's Fulcrum BioEnergy will produce biojet from municipal waste and supply it to United and Cathay Pacific. ● But United's biojet contract will provide only 2 percent of the airline's annual fuel consumption, an excellent example of the daunting scale of the required substitution. True, today's airliners are increasingly frugal: In 2015 they burned about 50 percent less fuel per passenger-kilometer than they did in 1960. But those savings have been swamped by the continuing expansion of aviation, which has raised annual consumption of jet fuel to more than 50 million metric tons by U.S. carriers and more than 250 million metric tons worldwide.

To meet this demand largely with biojet fuel, we would have to go beyond organic wastes and tap oil-rich seasonal or perennial oil crops whose cultivation would require large areas and create environmental problems. Temperate-climate oil crops have relatively low yields: With an average yield of 0.4 metric ton of biojet per hectare of soybeans, the United States would need to put 125 million hectares–an area bigger than Texas and California and Pennsylvania combined–under the plow to supply its own jet fuel needs. That's nearly four times the 34 million hectares that the country devoted to soybeans in 2016. Even the highest-yielding option, oil palm, which averages 4 metric tons of biojet per hectare, would still require more than 60 million hectares of tropical forest to supply the world's aviation fuel. That would necessitate quadrupling the area devoted to palm oil cultivation, leading to the release of carbon accumulated in natural growth.

But why take over huge tracts of land when you can derive biofuels from oil-rich algae? Intensive, large-scale cultivation of algae would require relatively little space and would offer very high productivity. However, Exxon Mobil's experience shows how demanding it will be to scale up to tens of millions of metric tons of biojet every year. Exxon, working with Craig Venter's Synthetic Genomics, began to pursue this option in 2009, but by 2013, after spending more than US $100 million, it concluded that the challenges were too great and decided to refocus on long-term basic research.

As always, the task of energy substitutions would be made easier if we wasted less, say, by flying less. But forecasts are for further substantial growth of air traffic, particularly in Asia. Get used to the unmistakable smell of aviation kerosene: It'll be here for a long time to come. ■

ILLUSTRATION BY Chad Hagen

REFLECTIONS_BY ROBERT W. LUCKY

OPINION

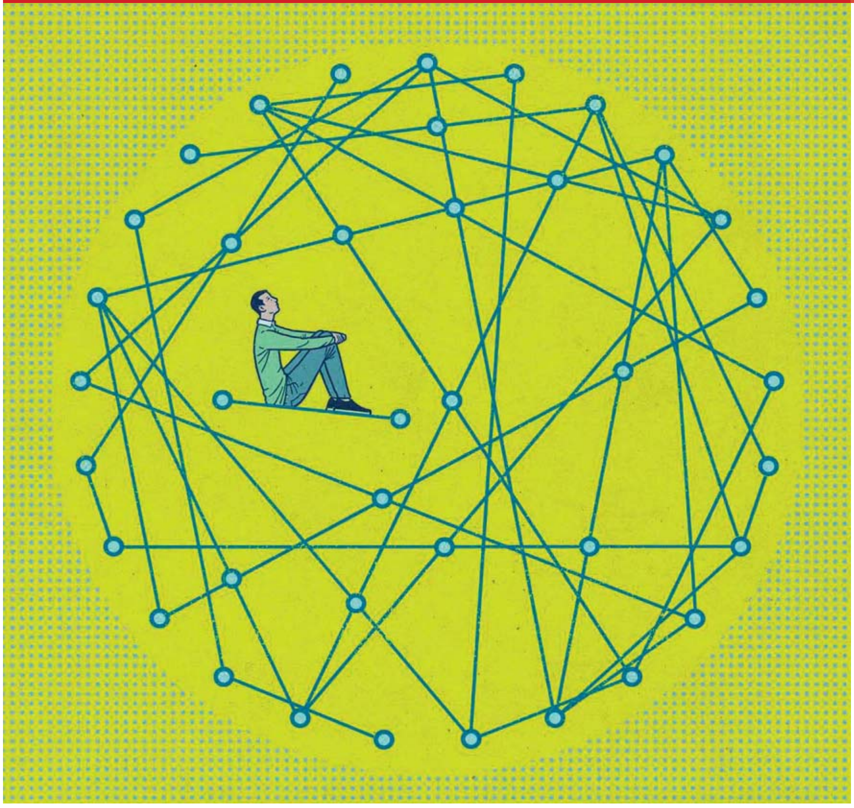

WAS THE INTERNET INEVITABLE?

MANY NOVELS HAVE PORTRAYED alternative versions of history, such as Robert Harris's *Fatherland: A Novel* (originally published by Hutchison, 1992) or Philip K. Dick's *The Man in the High Castle* (first published by Putnam, 1962). In both books the Axis powers are victorious in World War II, and the world today is dramatically different. However, in spite of technology's role in shaping history, there appear to be few books in which history is altered by imagined changes in technological innovation. An exception is the 1991 novel by William Gibson and Bruce Sterling, *The Difference Engine* (first edition by Bantam Books), in which Charles Babbage perfects a steam-driven analytical engine and the information age arrives a century earlier. • *The Difference Engine* launched an entire subgenre of speculative fiction and a corresponding aesthetic, both dubbed steampunk. But outside this subgenre's narrow confines (which often owe more to Jules Verne's conventional, forward-looking science fiction than to Gibson and Sterling's deliberately retrospective approach), there is a dearth of historical technology fiction. This may be because technology seems to evolve in a robust and inalterable manner, unchanged in the long run by outcomes of individual events. For example, suppose that the transistor had not been invented at Bell Telephone Laboratories in 1947. In all likelihood, that invention would have occurred soon thereafter at some other place. Before long someone else would have fabricated the first integrated circuit, and ultimately technology might be little changed today. It's as if there are preordained stepping-stones along a pathway of innovation that are simply uncovered when the time is right. The paradox, of course, is that those stepping-stones are evident only in retrospect.

I tried to imagine a substantial revision of technology history: Could I envision a singular event that, when changed, would have resulted in there being no Internet as we know it today? This seemed difficult because so many different people and organizations were involved in its evolution. Nonetheless, I wondered what might have happened if in the 1960s and '70s the Advanced Research Projects Agency had had other priorities and decided not to fund work in computer networking.

Without ARPA's funding, vision, and project management, there would have been less R&D in computer networking, but even so, there would have still been many pockets of work in the field. What would have been absent is the role of government as a neutral steward of the evolving network. So in my imagined scenario, information networks, instead of being designed by the users themselves, empowered by the open TCP/IP platform, are designed by the telecommunications industry. Now, instead of an *Inter*net, there is a balkanized tapestry of many competing proprietary systems largely controlled by telco service providers.

Each country has its own system, and the browser and the World Wide Web never evolve as such. Telephones have built-in displays and log in automatically to the local service provider, where users immediately encounter an enormous tree of menus. Fees are charged by the bit and for selected interactions, so the service is relatively expensive and usage is sparse. With the low participation, regionalization, and tight control of information services, national brands do not emerge—no Google, Amazon, or Facebook.

Well, all this seems like a bad dream, but in truth such a scenario would have been very unlikely. My own belief is that something akin to today's Internet would have been so compellingly attractive that it would have emerged from some alternative pathway through the swirling chaos of actions and interactions.

But we'll never know. ■

ILLUSTRATION BY Dan Page

Computing With Carbon Nanotubes

The tiny carbon nanocylinder has returned. Now it's ready to give silicon some competition

By Max Shulaker, H.-S. Philip Wong & Subhasish Mitra

THE SILICON SEMICONDUCTOR INDUSTRY has chugged along for more than 50 years. Like a steamroller, it has trundled over bumps and holes, while defying repeated warnings that it was running out of fuel or was about to be overtaken by flashier competitors. • So we technologists are understandably reluctant to speculate about the end of silicon. And yet, speculate we must. After decades of steady improvements to the efficiency and speed of our computer chips, brought about by physically shrinking the dimensions of silicon transistors, we've reached a point where the massive effort to miniaturize those switches ekes out only very modest gains in performance. The steamroller still rolls, but it's slowing down, and the maintenance and upkeep on it are fast becoming unsustainable. • Researchers are pursuing many options to keep integrated electronics on its exponential arc. One possible approach is to simply give up on improving the silicon transistors themselves, and instead focus on changes to the architecture or design of computers. This strategy has had some success in the past, when, for example, we moved from single cores to multicore processors, or began incorporating specialized accelerators into systems. Yet there are only a limited number of design tricks that can be played, and we have used many of them already. »

ILLUSTRATION BY Eddie Guy

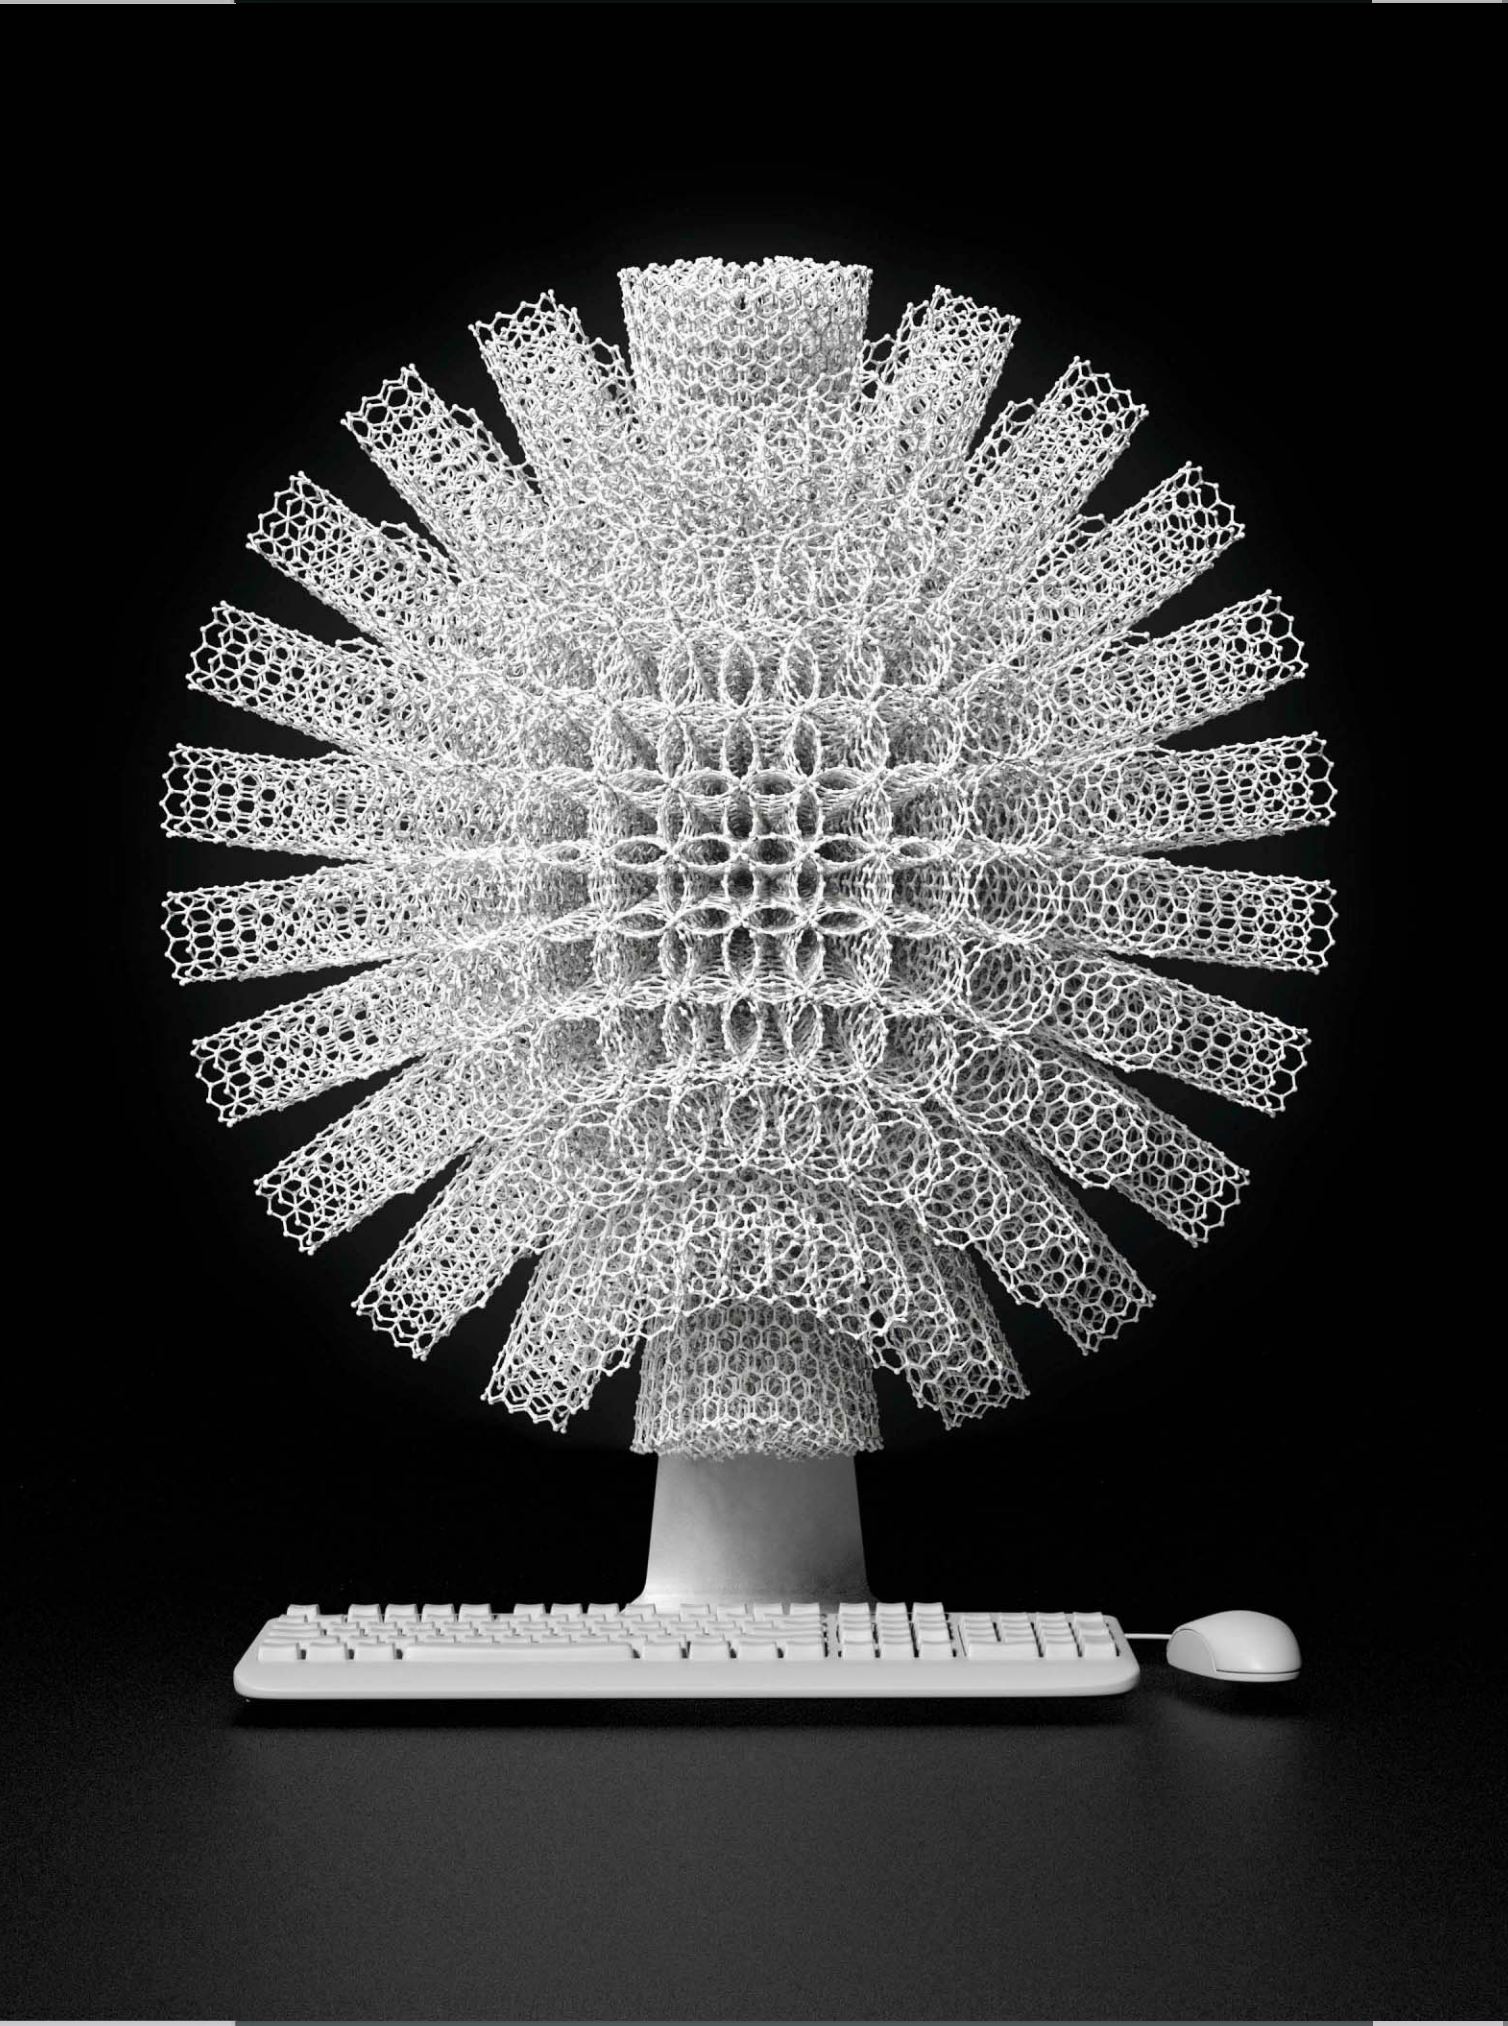

So future progress must still involve improving the underlying switches. Modifying materials and device geometries can create transistors with better electronic properties than silicon transistors have today. But such evolutionary approaches will yield only small benefits. To get bigger gains, we're left with one option: use novel nanotechnologies to supplant silicon altogether.

There are many candidates for this role. But the most exciting and mature contender by far is the single-walled carbon nanotube, a rolled-up sheet of linked carbon atoms. If you're experiencing some déjà vu now, it's no wonder. There was great enthusiasm for carbon nanotubes in the 1990s, and a number of big-name semiconductor companies began investigating them. But the carbon nanotube fell out of fashion when researchers ran into obstacles in circuit fabrication that seemed too hard to overcome. When the first measurements of the electronic properties of graphene–flat lattices of carbon atoms–were reported in 2004, many were all too eager to abandon carbon nanotubes and take up that shiny new technology.

Now the carbon nanotube is back and better than ever. Unlike its much-hyped two-dimensional cousin, the carbon nanotube comes in a natural semiconducting form, which means it can be turned on and off to make a binary switch. And recent research into the material has moved well beyond the demonstration of simple logic gates. In 2013, for example, our group at Stanford University built the first complete digital system–a basic computer–entirely out of carbon-nanotube circuits. What's exciting is that the fabrication and design techniques we employed to build that computer can be seamlessly integrated into a semiconductor fabrication facility today and used to manufacture chips with the billions of transistors needed to compete with state-of-the-art silicon.

Not only can carbon nanotubes make traditional chips more powerful, they can also be used to create entirely new computing systems–ultrafast and efficient three-dimensional chips that can be employed in data centers as well as the wide world of mobile computers used in cars, smartphones, and the sensor-filled Internet of Things. Such 3D systems could mix carbon nanotubes and emerging memory technologies, and they could also be built directly on top of conventional silicon circuitry. So while carbon nanotubes may eventually leap ahead of silicon, they could also share the road with it. In both cases, the resulting massive advances in computational capability would affect our lives profoundly.

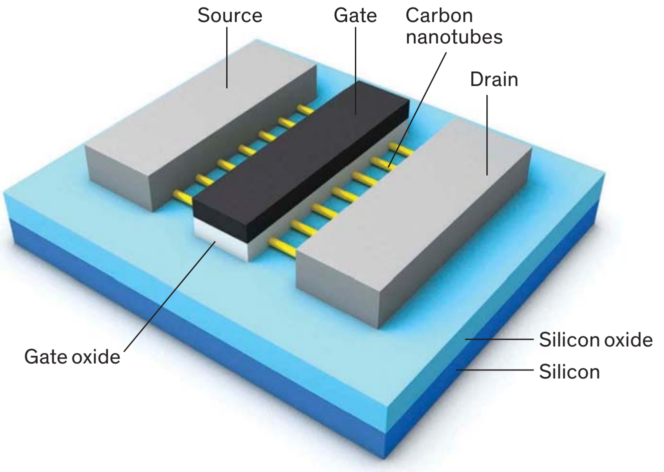

SWITCHING WITH CARBON: **A carbon-nanotube field-effect transistor (FET) resembles a conventional silicon FET. The main difference is the current-carrying channel, which is made from carbon nanotubes [yellow] instead of silicon.**

THE INITIAL DISCOVERY of carbon nanotubes is hard to pinpoint in time. But many cite a landmark 1991 paper by physicist Sumio Iijima, who grew "graphitic carbon needles" on an electrode in a gas-filled vessel. Iijima's work inspired a flurry of interest, and soon electrical engineers began suspecting that carbon nanotubes might be an ideal transistor technology.

To build a transistor, you need several things: a source and drain–the origin point and destination for current; a channel between the two where the current can flow; and a gate near the channel to control that flow. Semiconductor-device physics is complex, particularly when it comes to today's nanoscale switches, but as a general rule of thumb, there are two basic properties that determine a transistor's performance. First, the more current that flows through the channel, the faster a circuit can be. That's because greater current allows the gate of the next device in the circuit to be charged faster. Second, the thinner the channel, the more energy efficient the device will be, because it will be easier for the gate to exert the control needed to turn the transistor on and off. Think of the channel as a water hose and the gate as your foot. The thinner the hose, the less effort you'll need to make in order to stem the flow of water.

Unfortunately, in today's ultrasmall silicon devices, these two factors–the thinness of the channel and the speed at which charge can flow through it–are at odds with each other. Like other three-dimensional bulk materials, as silicon gets thinner, the sides and edges of the crystal have a stronger and stronger effect on its electrical properties. These so-called surface effects impede the movement of charge through the material.

But the carbon nanotube naturally and simultaneously fulfills both of those transistor channel requirements. The nanotube's smooth-walled structure is naturally devoid of dangling bonds, permitting charges to flow through it quite rapidly, many times faster than in silicon. And because carbon nanotubes are also very small–in the realm of about a nanometer in diameter–a gate can easily control the flow of current along its length.

A carbon-nanotube transistor looks much the same as a silicon transistor. The main difference is that the channel is made of carbon nanotubes instead of silicon. Just as in a silicon transistor, the electric field created by a

ILLUSTRATION BY Jean-Luc Fortier

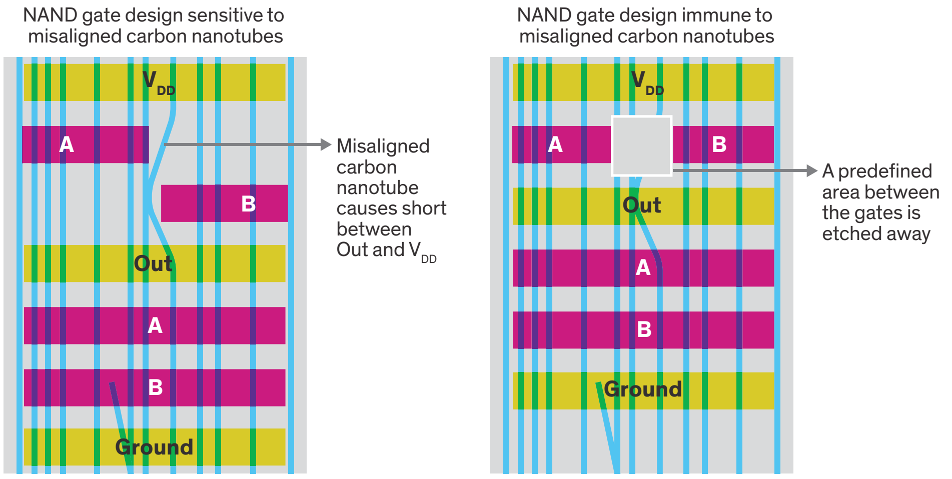

TACKLING IMPERFECTIONS: Carbon nanotubes don't always grow in perfectly parallel lines. Some circuit designs are more sensitive to such imperfections than others. In the NAND gate design at left, for example, a wayward carbon nanotube can weave between gates, creating a short. Other designs, such as the one for the NAND gate at right, will function properly regardless of imperfections in the arrangement of the underlying carbon nanotubes.

gate controls the flow of current. And just as with silicon chips, circuits can be formed by connecting the carbon-nanotube transistors to one another with tiny metal wires built into additional layers of the chip.

Two independent teams, based at the Delft University of Technology, in the Netherlands, and at IBM, reported the first carbon-nanotube field-effect transistors (CNFETs) in 1998. And in the years after, researchers devised a number of experiments to probe the fundamental physical properties of these devices. They also began building basic circuits; a notable advance came in 2006, when an IBM group built a ring oscillator–a key logic demonstration circuit–by constructing 12 CNFETs from a single 18-micrometer-long carbon nanotube.

Impressive though those achievements were, there was a big gap between a small test circuit built from a single carbon nanotube and a chip that could compete with state-of-the-art silicon. To make a chip, researchers would have to find a way to construct billions of transistors from a sea of carbon nanotubes. A few requirements for high-performance circuits emerged early on. Each transistor channel would need to contain multiple, densely packed carbon nanotubes in order to drive enough charge to compete with the speed of silicon circuitry. What's more, these dense carbon nanotubes should be spaced uniformly, minimizing electrical variation between transistors.

The natural way to reach this level of density and consistency is to build the sources, drains, and gates of the transistors on top of uniform arrays of closely spaced carbon nanotubes. But, as you might imagine, the process of creating a neat arrangement of billions of identical curled-up carbon sheets, distributed evenly on a silicon substrate, isn't easy.

The problem begins with carbon-nanotube synthesis. Any process that is used to grow carbon nanotubes almost inevitably results in an assortment of tubes with different diameters and atomic arrangements. Variations in the underlying atomic structure yield carbon nanotubes with a mix of electrical properties. Some are semiconducting, while others are metallic. The semiconducting carbon nanotubes are exactly what you want for digital logic, while the metallic ones basically act as wires because they can't be turned off by the gate. As a result, they always conduct current, even when the transistor should be turned off, which wastes power and can cause errors in logic operations.

To make matters worse, it's hard to grow carbon nanotubes in nice parallel arrays on a surface. If some of them wind up in divergent orientations, they can extend beyond the confines of their own transistor channel and create random circuit connections. The density of nanotubes can also vary from spot to spot, which will result in transistors with a wide variety of current-carrying capacities. In fact, some could easily have no carbon nanotubes in their channel at all, resulting in device failure.

IN 2004, OUR DIVERSE GROUP of students and postdocs at Stanford University, led by professors Subhasish Mitra and H.-S. Philip Wong, set out to tackle these problems. (One of us, Max Shulaker, began with the group in 2009 as an undergraduate and will soon join the faculty at MIT.) The field of carbon-nanotube research had shrunk significantly by the time we began our work. But the hope was that we, along with others in that small but innovative community, could finally figure out a way to transform the nanomaterial from a laboratory curiosity into the stuff of tomorrow's chips.

To start, we focused on tackling two key obstacles to high-performance circuits mentioned earlier: the presence of misaligned carbon nanotubes and also of metallic ones. Our first step was to find a way to grow carbon nanotubes that were as aligned as possible. The approach we settled on, inspired by work done at the University of Illinois at Urbana-Champaign and the University of Southern California, was to grow them on a crystalline quartz substrate.

When surrounded by a carbon-rich gas, carbon nanotubes will grow out from metal nanoparticles stuck to the surface of a quartz wafer in nice, orderly parallel lines. We can then transfer the carbon nanotubes onto a more traditional chip substrate–an insulating layer of silicon dioxide on top of silicon–by coating the carbon nanotubes with a thin layer of metal.

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

The metal acts as an adhesive, which we can then physically peel off the quartz and transfer to the silicon. This metal is then chemically etched away, leaving the array of highly aligned carbon nanotubes behind. (A similar transfer process is already used in commercial facilities to make silicon-on-insulator wafers for traditional chips.)

But even with quartz as a guide, we found that a small fraction–about 0.5 percent–of carbon nanotubes still do not grow properly. They might begin growing straight only to develop a kink and end up on a diagonal. To deal with these misaligned nanotubes, we devised a way to effectively design circuits around them. An automated design algorithm, based on graph theory, decides which areas of a circuit should be etched away so that it's mathematically impossible for a stray nanotube to connect two electrodes that it isn't supposed to. This scrubbing ensures that no misaligned carbon nanotube can cause a short circuit that interferes with the logic. Importantly, this process requires no knowledge of the exact positions of the aberrant carbon nanotubes–the same design works for any starting configuration. Without this strategy, we'd need to inspect every wafer and every carbon nanotube in order to find the misaligned ones, an impractical and costly process.

We found that we could tackle the metallic nanotubes with a similar mix of processing and design. In 2001, an IBM team showed that, with sufficient voltage, it's possible to destroy metallic carbon nanotubes by burning them away. We adapted this technique to operate across an entire integrated circuit at once, by building temporary wiring that effectively turns a whole chip into one giant transistor. After the circuits are made, we apply a voltage to the gates of the transistors to turn off all the semiconducting carbon nanotubes. Then we apply a voltage across each source and drain region. Because the semiconducting carbon nanotubes have already been turned off, only the metallic nanotubes in the channels will conduct electricity. With a sufficiently strong voltage pulse, even the relatively low resistance of these bad nanotubes will create a brief heat spike that fries the nanotubes, turning them into minuscule puffs of CO_2. Today, we use a variation on this approach that can remove the entire length of a metallic carbon nanotube and eliminate more than 99.99 percent of them (a small and manageable fraction–now less than 1 percent–of the good, semiconducting carbon nanotubes are sacrificed in the process as well).

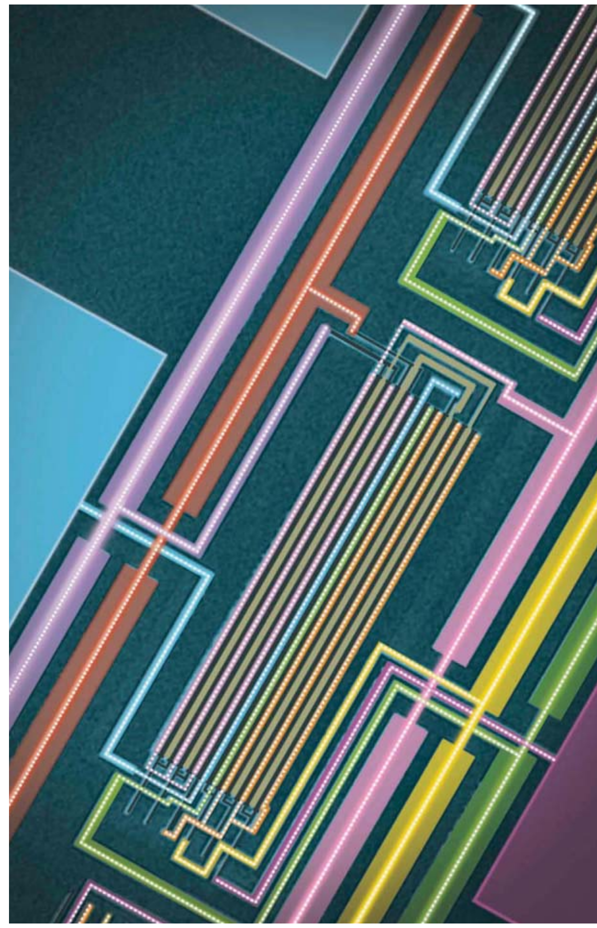

CARBON LOGIC: A scanning-electron-microscope image shows a portion of the circuitry in the first carbon-nanotube computer, which the authors reported in 2013. The 1-kilohertz processor was primitive by modern standards, but it demonstrated that any digital system could be made using the devices.

Combining metallic-nanotube removal with our strategy for addressing misaligned nanotubes allowed us to build the first digital systems out of carbon nanotubes using only practical, large-scale manufacturing approaches (we refer to this combination of processing and design as the "imperfection-immune paradigm"). Our first proof of principle was Sacha, a handshaking robot that we demonstrated live at the 2013 IEEE International Solid-State Circuits Conference, in San Francisco. When a conference attendee grabbed the hand of the robot, a carbon-nanotube-based subsystem would sense the change in the value of an embedded capacitive sensor and trigger the robot to shake hands. While this was admittedly a simple show-and-tell demonstration, Sacha did make a very important point: CNFETs had reached a level of maturity where a live, system-level demonstration could be made.

To drive the viability of carbon-nanotube circuitry home, we also created a complete digital

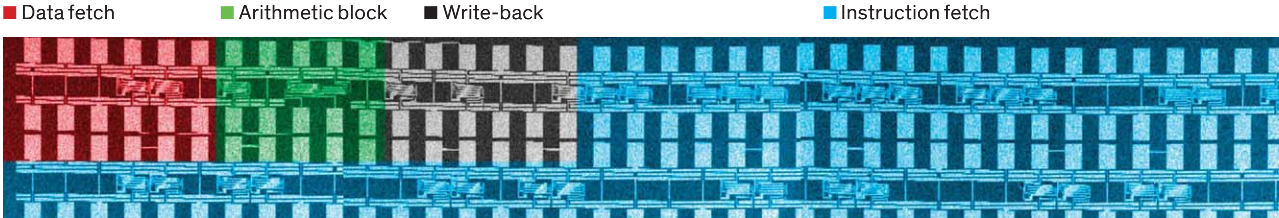

METAL AND CARBON: A view of the 2013 carbon-nanotube computer's circuitry shows various functional areas. More than 30,000 carbon nanotubes were used to create the computer's 178 transistors.

TOP: MAX SHULAKER/NATURE; BOTTOM: MAX SHULAKER

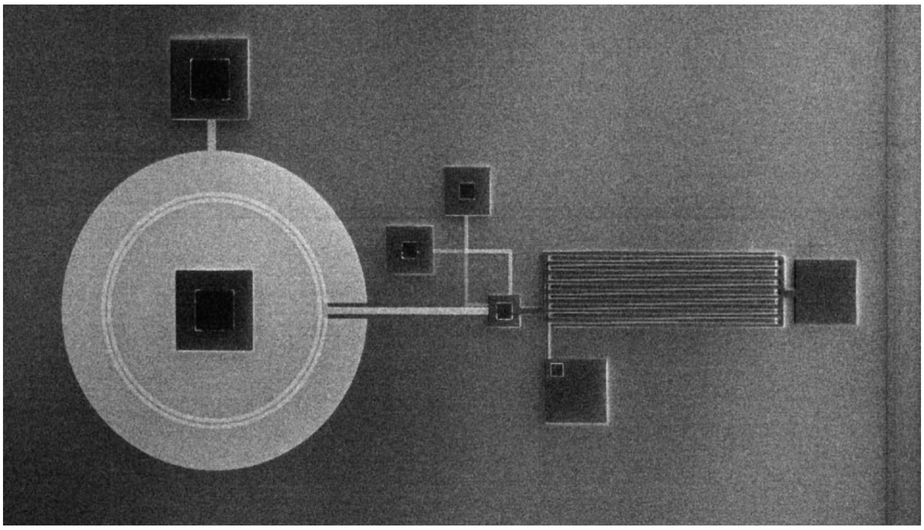

BUILDING UP: Circuits built one on top of another, connected by dense wiring, could dramatically improve speed and energy efficiency. This FPGA switching element [left] contains a carbon-nanotube transistor, a silicon transistor, and two resistive RAM memory cells. Each element sits on its own layer.

system: the first computer processor built entirely from CNFETs. The processor–which was built from only 178 CNFETs containing more than 30,000 carbon nanotubes–was certainly modest by modern standards. It used a 1-bit architecture and had a clock rate of 1 kilohertz, less than a millionth the clock rate of today's CPUs. But we weren't optimizing for speed; the chip was slow in part because we put probes at each node to monitor every possible transition inside the computer. Still, the processor was a complete computer; it could perform multitasking and emulate instructions from the commercial MIPS instruction set. More important, it showed that CNFETs could be used to build any kind of digital system, a first for any of today's emerging nanotechnologies.

OF COURSE, INTEGRATION alone doesn't make a technology. The performance of CNFET circuits was still far behind silicon. A core challenge here was carbon-nanotube density. At the time we reported our carbon-nanotube computer in 2013, we were able to grow arrays of carbon nanotubes with a linear density of 1 to 10 carbon nanotubes per micrometer, along the line perpendicular to the transistor channel. To create enough current-carrying capacity for a high-performance CNFET, we'd need 10 to 100 times that value. We'd also need to make sure that density would be uniform from spot to spot to keep transistor variations to a minimum.

In 2014, we reported a new technique that could simultaneously address those challenges. The approach builds on the strategy we used to move carbon nanotubes from the quartz substrate they're grown on to the substrate that will be used to build a chip. We showed that this method could be used multiple times to layer one array of carbon nanotubes on top of another. Because it's generally easier to make uniform arrays that are fairly sparse, this multiple-transfer approach also lets us decouple the density problem from the uniformity one. In other words, we can grow multiple sparse arrays and then layer them, one on top of the other, to achieve high density.

Unlike other potential nanoelectronic devices, the CNFETs we've made with this approach can deliver currents that are more or less on par with a silicon transistor of the same size. This was unprecedented performance for a carbon-based nanotechnology, and CNFETs are expected to get better still.

Ours isn't the only approach to making carbon-nanotube circuitry. For example, IBM and several other research groups have done work with semiconducting and metallic nanotubes suspended in solution. The metallic nanotubes are separated and removed. Then the remaining nanotubes are dried on a substrate that has been patterned so that they land, as designed, in particular locations and orientations. Researchers are still working to drive down the fraction of metallic nanotubes that remain and boost the density of the semiconducting ones, but this method shows significant promise.

Despite advances like these, we still encounter skepticism when we talk about carbon-nanotube computing. One question that often arises is whether it's possible to build complementary logic. Today's processors use a semiconductor technology called CMOS, which stands for complementary metal-oxide semiconductor. The "complementary" refers to the fact that it employs two different kinds of transistors: *n*-type transistors, which carry electrons across the channel, and *p*-type ones, which use holes–positive charges that are actually an absence of electrons. There are many advantages to having both types of devices, but for years researchers primarily made only *p*-type carbon-nanotube transistors. That was because the first CNFETs were *p*-type and the materials used to make them are somewhat easier to work with. Indeed, our carbon-nanotube computer was made entirely from *p*-type CNFETs.

But carbon nanotubes don't intrinsically have a type, and it's fairly easy to dope them to carry either charge carrier. Instead of inserting other types of atoms into the material, as is done in silicon, this doping can be done through the choice of materials used to make the insulating layer between the gate and the channel or the metal source and drain regions of the transistor. In the case of the source and drain, those materials essentially act as filters, permitting either electrons or holes into the carbon-nanotube channel. Many research | CONTINUED ON PAGE 49

MAX SHULAKER

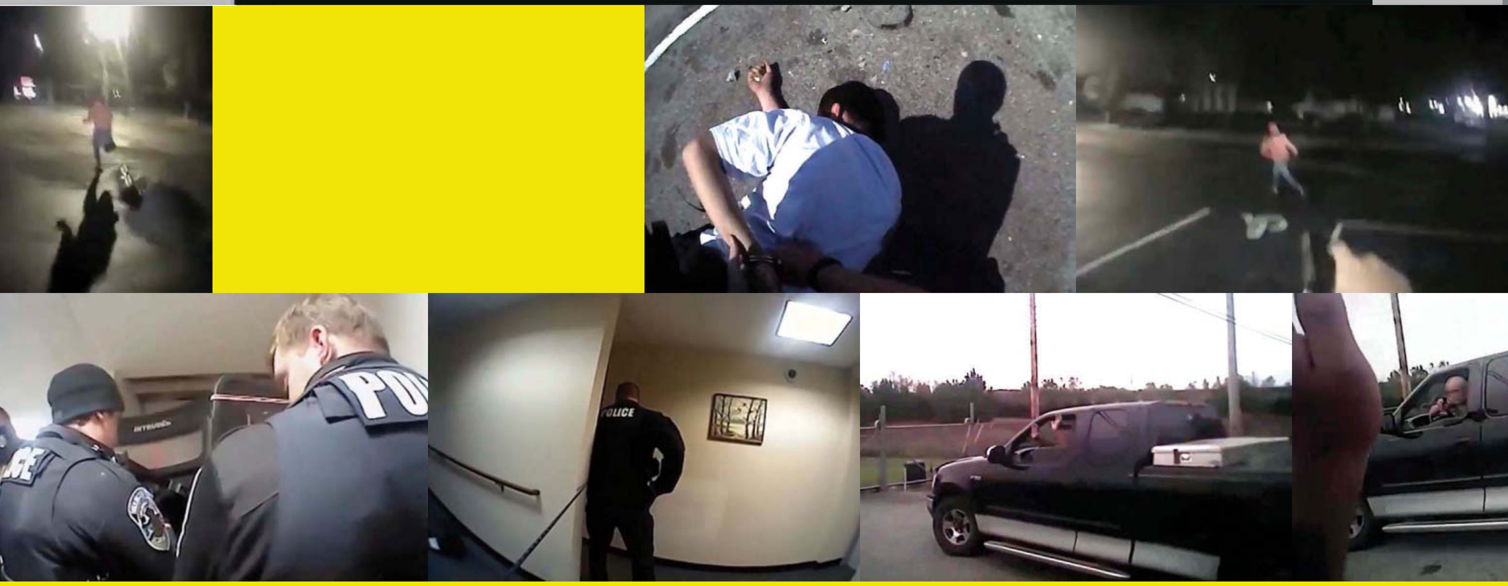

The Puzzle of Police

P **OLICE BODY CAMERAS are popping up everywhere, often to good effect because both police and suspects normally behave better in their presence. No wonder these small devices, enthusiastically endorsed by police, politicians, and civil-rights advocates, have generated a burgeoning industry. Yet people know very little about how and why they work, so the intended and unintended consequences of using them remain nebulous. ● That's not for lack of effort. There have been nearly 40 studies on the use of body cameras, including a dozen randomized controlled trials on the magnitude of their effect on policing. Despite all this work, it's still not entirely apparent why these cameras are helpful, under what conditions, or for whom. ● Here I'd like to offer my interpretation of all that research and to delve into what sets police body cameras apart from other video-recording equipment, such as closed-circuit television, dashcams, and everyday smartphone cameras.**

SEEING IS BELIEVING: Body cameras, in widespread use by police departments around the world, document a huge variety of interactions between officers and civilians, a few of which can be seen in these video frames.

BODY-WORN CAMERAS GIVE MIXED RESULTS, AND WE DON'T KNOW WHY

By **BARAK ARIEL**

Body Cams

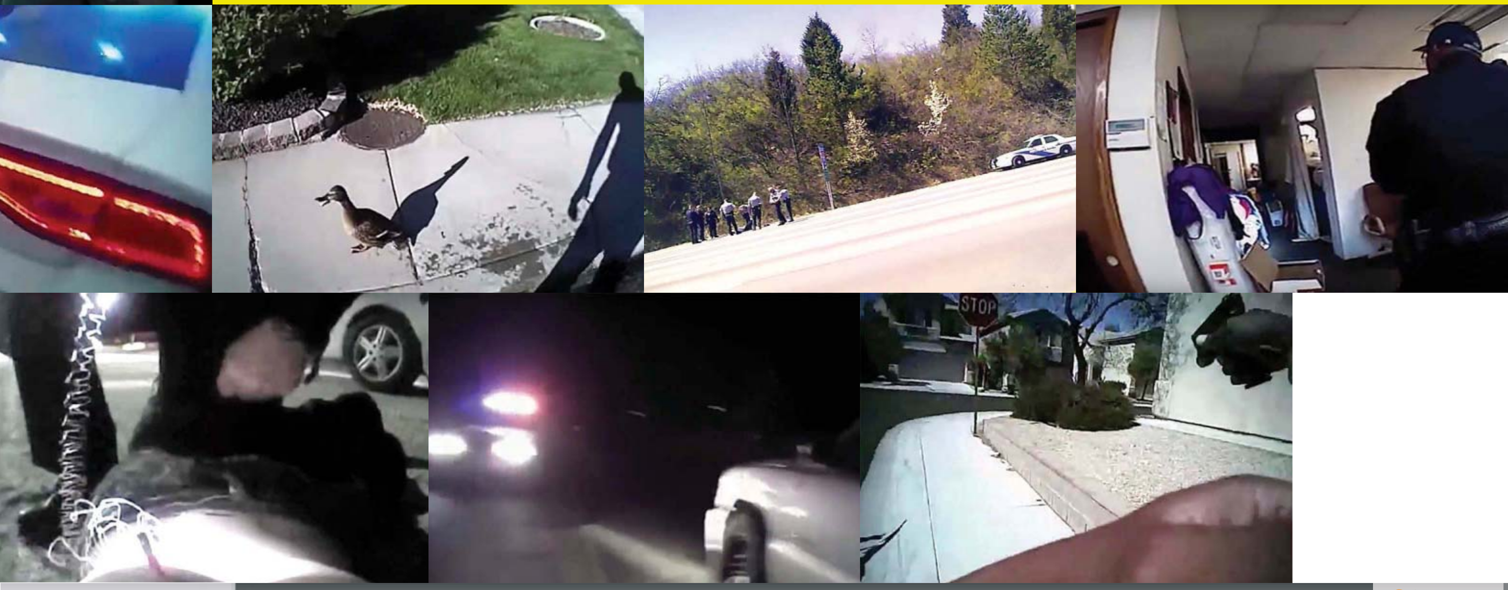

Ultimately, body cameras are just video cameras, albeit specialized ones for capturing evidence in a reliable way. Although body cams are relatively straightforward technology, their effectiveness has proven highly variable, for reasons that remain puzzling. Why, for example, does their use reduce by more than half the number of times officers apply force during their encounters with the public in some places, while in other places it nearly doubles the reported use of force?

To address these questions, it's helpful first to consider other surveillance devices used in policing. These have a longer track record and could offer some insight into the new phenomenon of body cameras.

Over the past 25 or so years, surveillance cameras have increasingly become an integral part of law enforcement as technological advancements made video cameras better, more reliable, and substantially cheaper. Most, if not all, police agencies in developed nations use closed-circuit television (CCTV) with the aim of deterring criminal activity, investigating crimes that have taken place, and prosecuting those responsible for them.

London, for instance, is under such heavy police surveillance that it's hard to find streets not covered by CCTV. And many of the blind spots that remain are targeted by private video cameras. So when you're in London, as well as many other cities, nearly every move you make in public is videotaped, tagged, and filed away somewhere.

But do these cameras really prevent crime and disorder? Common sense says, of course they do. If potential offenders are rational actors, aware that their misdeeds will be caught on video, they will surely be deterred from wrongdoing. This is the rationale for concluding that bad guys will not victimize us when there is CCTV, an ever-more-common part of life for which we all pay the price of infringed privacy.

Despite such intuition, and despite the huge investments made, overwhelming evidence indicates that CCTV equipment has, in fact, little deterrent effect. At least 44 studies illustrate that CCTV reduces the overall level of crime by only about 16 percent, with half of that reduction concentrated in parking lots. There is no effect at all on assaults, robberies, and similar against-person violent crimes. The evidence also tends to show that what little criminal behavior CCTV prevents is just displaced to other areas.

In short, CCTV cameras are pretty much a failure. Perhaps that's because offenders know where the blind spots are. Or maybe it's because CCTV does not work well in the dark or when the perpetrator wears a hoodie. Another possibility is simply that police lack the resources to assign somebody to sit through endless hours of recorded video hoping to find clues in not-so-serious crimes. Or it could also be that CCTV has become such an integral part of everyday life that its presence escapes people's conscious attention. Whatever the reason, it's clear that although CCTV may make you feel more secure, it does not really make you any safer, although it makes it modestly less likely you'll get your car broken into or stolen.

Another common type of video surveillance, the one that has accompanied the proliferation of smartphones, might actually be more important. Video cameras are ubiquitous, and the video recording of engagements between police and public is incredibly influential, especially when misconduct is caught, be it the infamous beating of Rodney King in Los Angeles in 1991 or the killing of Eric Garner in New York City in 2014. These recordings certainly demonstrate the effect of cameras on public reactions to the police, having sparked the Los Angeles riots of 1992 and the Black Lives Matter movement of 2015.

Given the notoriety of such videos, a camera at the scene of a police-public encounter ought, logically, to

Not-So-Candid Camera

LINE OFFICERS in West Valley City, Utah, wear body cameras while on duty, a practice the town began in March 2015. They use Taser International's Axon system, which allows the camera to be affixed to the side of special Oakley eyeglasses or mounted to the officer's collar. The camera is attached by cable to a main unit on which the officer triggers recording by pushing a button. The equipment is stored in a special docking station used to upload video to a cloud-based service while recharging the unit's batteries.

PREVIOUS PAGES: POLICEACTIVITY/YOUTUBE AND RIALTO POLICE DEPARTMENT send out an accountability cue, eliminating feelings of anonymity on either side. There is no strong evidence, however, to support the conclusion that mobile-phone cameras deter officers from misconduct. The Garner incident is particularly telling in this regard: As clearly shown in the recording released to the media, officers were well aware of the cameras filming them (some of them looked directly at the cameraman), yet they still used a prohibited choke hold. Why?

There are probably two reasons. First, the officers might have been aware of a camera, but it still didn't really register. Here, as in many other highly charged encounters, people just don't think much about all the cameras around them before making the decision to throw a punch or otherwise misbehave. Second, even if they are aware of being recorded, the parties involved might not perceive a strong possibility that the footage captured can and will be used to hold them accountable for any type of misconduct. Because it is not official evidence, the video from a civilian's smartphone camera doesn't seem to inspire concern that it will surface later.

I should point out that one particular type of camera is effective in preventing unwanted behavior: road or speed cameras. A systematic review of 35 rigorous tests of road cameras has shown that they reduce serious and fatal accidents by as much as 44 percent. The evidence is unequivocal. They do exactly what they are intended to do. Even a speed-camera sign, with no actual camera anywhere, prompts drivers to slow down. Unlike CCTVs or smartphones, speed cameras work because punishment for illegal actions is virtually guaranteed.

THESE PAGES: GEORGE FREY/GETTY IMAGES (5)

WHERE DO BODY CAMERAS sit on this spectrum of effectiveness? There have been 12 randomized controlled trials, with another 30 ongoing research projects seeking the answer. My students in the police executive program at the University of Cambridge, in England, and I have conducted most of the randomized controlled trials that have been published so far, including one that is now commonly referred to as the "Rialto experiment."

This study, which really gave rise to today's heated debate about police body cameras, took place in Rialto, Calif., a town of about 100,000 located some 100 kilometers east of Los Angeles. I worked on this experiment with Tony Farrar, then Rialto's police chief, while he was completing a master's thesis, for which I served as his academic advisor.

My Cambridge colleague Alex Sutherland helped Tony and me give this evaluation of body cameras a high degree of statistical rigor. Indeed, we devised this test just as if we were investigating the effectiveness of some new drug therapy. The study involved all of the town's 54 frontline police officers, who for an entire year, starting in February 2012, would be assigned either to treatment (camera-wearing) or control (not camera-wearing) conditions when they went out on patrol.

During treatment shifts, officers were asked to take video of all their interactions with the public, to announce that the encounter was being recorded, and subsequently to store the footage on a secure cloud-based server. In control shifts, the officers were told not to use body cameras at all. Outcomes were then measured in terms of officially recorded use-of-force incidents and complaints lodged against Rialto police officers. At the end of a year, we were able to compare nearly 500 patrol shifts during which all police-public encounters were assigned to treatment conditions with a roughly equal number of shifts assigned to control conditions.

The results were stunning. There were roughly 50 percent fewer incidents of force being used while the officers were wearing body cameras compared with control conditions (8 incidents as compared with 17). And after reviewing the footage, we discovered that all eight times the camera-using officers resorted to force, they did so in response to violent behavior on the part of the people they were engaged with. The evidence suggests that in 4 of the 17 instances in which officers not wearing cameras resorted to force, it was the officer who initiated physical contact. This seems a key finding, because it really points to cameras making police officers less likely to use force without ample justification.

What's more, there was a 90 percent reduction in citizens' complaints against police officers compared with the

12 months prior to the experiment. That's particularly remarkable, because it was a 90 percent reduction in the total number of complaints filed, not just those filed against officers wearing cameras.

Perhaps because all officers wore cameras some of the time, their behavior even when not wearing cameras changed. The statistics suggest as much, because the officers not wearing cameras resorted to force only half as often as they did in the year before the experiment. The moderating effect of the cameras seems to have been contagious.

A Burgeoning Industry

THERE IS NO SHORTAGE of companies manufacturing body cameras for the law-enforcement market, including those shown here. So police departments can choose from among many models and capabilities.

TASER INTERNATIONAL
Scottsdale, Ariz.

EDESIX
Edinburgh

REVEAL MEDIA
Surrey, England

CLOCKWISE FROM TOP: TASER; REVEAL MEDIA; EDESIX

The results of the Rialto experiment would lead you to think that police body cameras are an unequivocal success. In Rialto they were, but studies my colleagues and I have since conducted elsewhere require that I add a rather large note of caution. You see, if you consider the 10 other places where we have now completed tests of such cameras, you would conclude that their overall effect on police use of force is a wash: In some instances they help, in some they don't appear to change police behavior, and in other situations they actually backfire, seemingly increasing the use of force.

That wearing a camera would ever cause an officer to use force more than he or she would otherwise do is puzzling, to say the least. But some hints of what is happening come from looking at how well officers complied with the experimental protocol.

In places where they closely followed the instructions (use the camera during each encounter if you're in a treatment group; don't use it if you're in a control group), the results were positive—a 37 percent reduction in use of force on average. But if you allow the treatment group discretion to choose when to turn it on, the result is 71 percent greater use of force. Thus the problem seems to arise mainly when officers are allowed to turn cameras on at times of their own choosing.

A **LTHOUGH THERE ARE EXCEPTIONS,** especially where officers are allowed too much discretion, it nevertheless seems clear to me that police body cameras properly used can be very helpful because officers and suspects alike become more certain that they'll be punished for bad behavior when a camera is rolling.

For this reason, body cameras are sometimes seen as a panacea: They record everything and therefore tell the unmediated story of what took place. They increase transparency, heighten accountability, and keep the actions of all parties in check. Because no rational person wants to get into trouble (or into more trouble than he or she is already in), police-public interactions become less heated.

Very often that's the case, but it doesn't always play out that way. It's common for officers to interact with people who are mentally ill, drunk, high, raging with anger, or otherwise emotionally disturbed—people who are not likely to be aware of a camera even if told about it. Similarly, officers in emotionally heightened situations—such as during a high-speed pursuit or while subduing a resisting offender—might easily fail to turn the camera on or just ignore it even if it's running.

Here's where better technology could help. Body cameras could be automatically activated immediately when certain cues are triggered, such as when the officer enters a crime hot spot, leaves a police vehicle,

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

Perplexing Inconsistencies

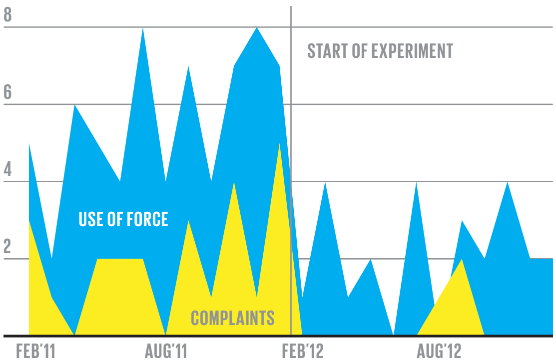

Results of the Rialto Experiment (incidents)

The number of times officers from Rialto, Calif., resorted to using force and the number of complaints received both diminished markedly after just half the patrols began using body cameras in February 2012.

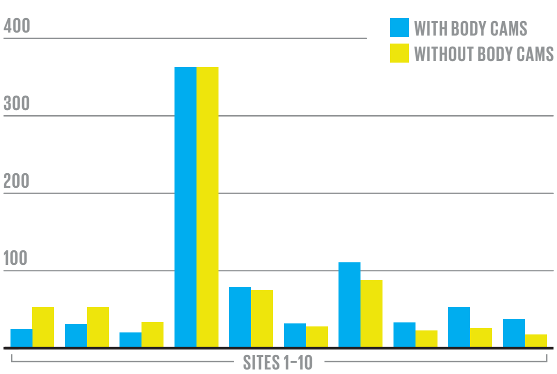

Use of Force in Multisite Trial (incidents)

The results from the author's trials of body cameras at 10 undisclosed sites showed the use of force sometimes diminishes when cameras are worn, sometimes stays the same, and sometimes even increases.

takes out handcuffs or a weapon, turns on the siren, or makes a call for assistance on the radio. This way, the officer doesn't have to think about activating the body camera in a tense situation.

One manufacturer, Taser International of Scottsdale, Ariz., is already offering body cameras that activate automatically when an officer turns on a vehicle's flashing lights or draws a "smart" weapon. And others are investigating the use of cameras that stream live video back to headquarters. Right now, power and bandwidth limitations prevent that general practice, but one day in the not-too-distant future it will likely be the norm.

WHILE THE VARIABILITY in the effectiveness of these cameras is disappointing, perhaps it shouldn't be so surprising. After all, what could work for a sheriff's department in Iowa may not necessarily apply to a national police force in the Middle East. The results are going to hinge on the way official records are stored, the prevalence of police misconduct, and how legitimate the force is in the eyes of the people they serve.

So even though I'm a great supporter of body cameras, I'd advise police departments to move slowly and thoughtfully in their adoption. Test them in ways that allow these cameras to fail as much as to succeed. Try out different procedures and pick the ones that work best. After all, who's to say that the policies practiced in the United States or the United Kingdom are appropriate elsewhere? Police administrators need to consider, in particular, whether rank-and-file officers will embrace the cameras or lash back at "big brother" and the threat to their autonomy.

There are certainly a lot of issues to address: Should superior officers be allowed to view all of the patrol officers' footage, or ought those rights be limited to specific cases? Should officers be able to review their own recordings before filing their written reports? And who should be in charge of curating all this video, police departments or some independent agency subject to public-records laws? The different people involved will no doubt answer those questions very differently based on their hunches and past experiences, so to find what really works best will require much experimentation.

I would also add that there are still plenty of technological hurdles. The most pressing is the need to integrate the evidence captured by body cameras with other information-technology systems used in law enforcement. These are already troublingly fragmented. Typically, police forces have one system for taking emergency calls, another for handling investigations, and another yet for recording their reactions to crimes. Offender databases do not talk to victim databases. And there is often no link whatsoever with databases held by prosecutors, so court outcomes are never then fed back to the police. The list goes on.

These kinds of problems will grow only more acute as police departments look to incorporate additional layers of information with their body-cam video, including location data and the results from facial-recognition technologies. To fully explore the possibilities and rewards, many more tests are needed.

There's another reason to test these cameras and their associated technologies carefully before handing them out. Imagine a police force is looking to use body cameras as a way to boost public trust. If police administrators simply announced, "We're using cameras now, and we find that police misconduct is extremely rare," most people would take such statements with a grain of salt. If, however, the police force had first partnered with independent researchers to test the effectiveness of the cameras, positive results would more likely be believed. That, in my view, is the best prescription for any police agency that wants to move forward with body cameras and the technology that is swiftly evolving around them. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/bodycams0716

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

1 Chip, 2x the Bandwidth

Today's radio technology has almost nothing in common with its counterpart from a century ago. Countless advances have produced gear that is fantastically smaller, more reliable, and more efficient in its use of power and bandwidth. And yet, one enormous limitation remains from long ago: Radios still can't transmit and receive on the same frequency at the same time. • Such a capability, called full duplex, would be a great advance for wireless. It would in one fell swoop double a network's capacity, the physical ability to carry data. At a time when unused radio-frequency spectrum has already been mined to near extinction and demands for data are increasing—and with emerging 5G networks targeting a 1,000-fold increase in data flows—full-duplex wireless has become a holy grail in the search for a way to ease the spectrum crunch. • Now, at last, we and others have demonstrated that a full-duplex wireless system can be practical and reliable. Research in our own labs at Columbia University under the FlexICoN project and in Europe under the DUPLO initiative has demonstrated full-duplex operation within CMOS integrated circuits, the kind that are ubiquitous in today's computing and communications gear.

ILLUSTRATION BY Greg Mably

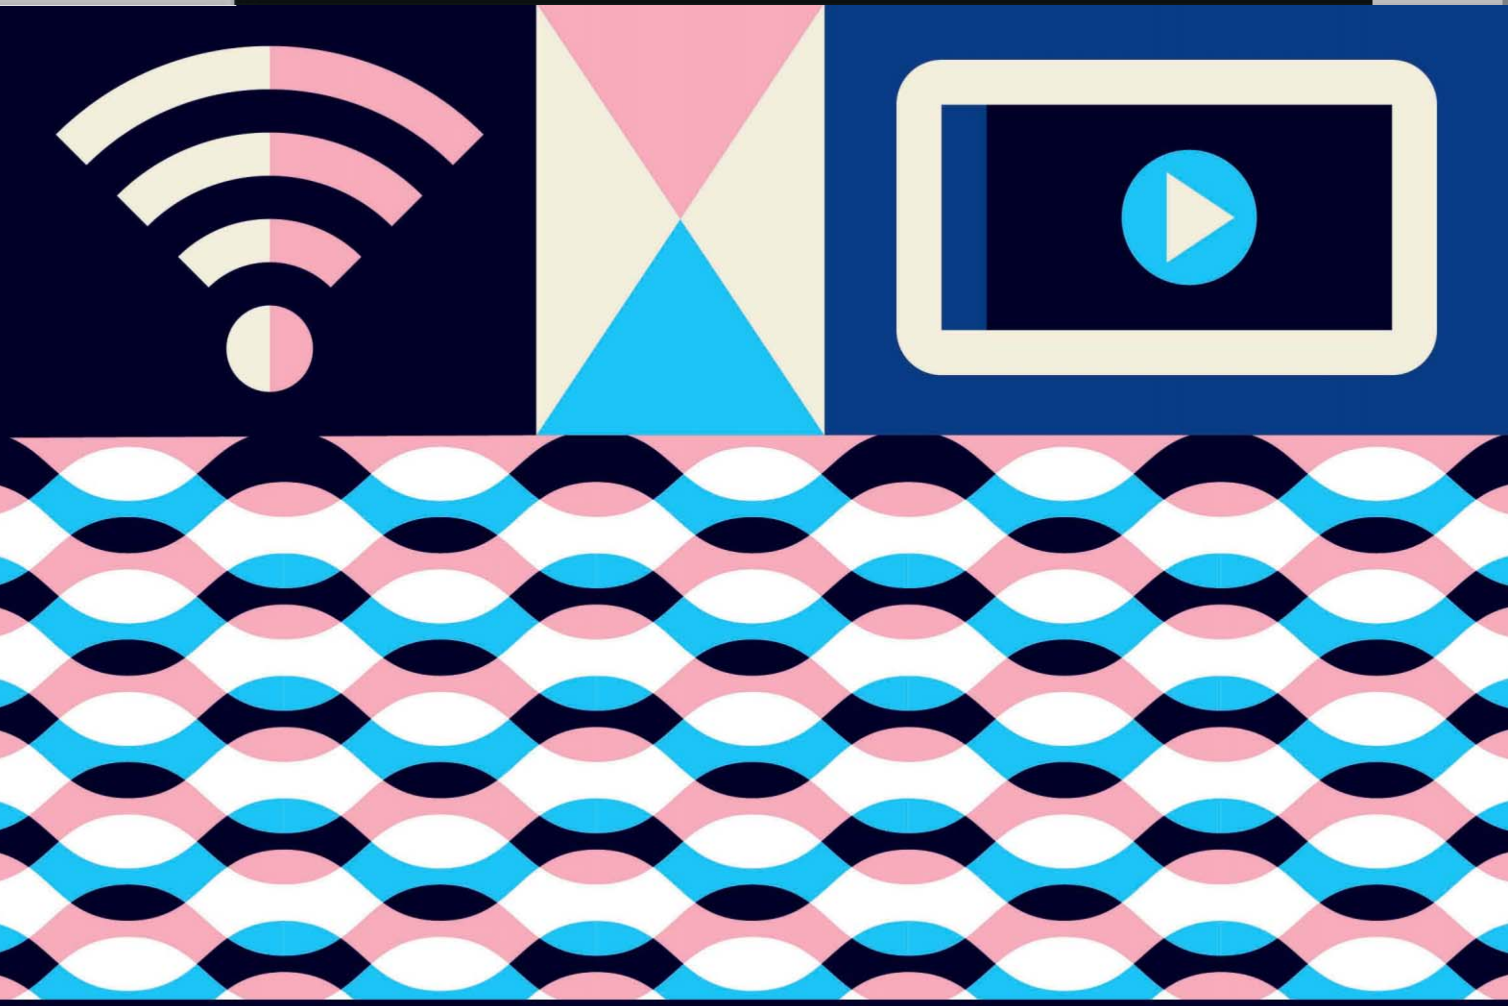

A ONE-ANTENNA CHIP THAT CAN SEND AND RECEIVE SIMULTANEOUSLY COULD DOUBLE THE DATA CAPACITY OF PHONE NETWORKS

By **Harish Krishnaswamy** *&* **Gil Zussman**

This work follows the first demonstrations of this technique, which were accomplished only a few years ago in separate projects at Rice and Stanford universities using benchtop laboratory equipment. The Stanford research has since been spun off into a startup company, Kumu Networks, that is using discrete components to bring full-duplex capability to base stations and infrastructure, where cost and size constraints are more relaxed than in handsets.

In the wired world, duplex circuits are old hat. Early pre-electronic telephone handsets were able to use a single channel to send and receive simultaneously by isolating the earpiece from the microphone using a hybrid transformer circuit; that way, the outgoing and returning signals could each pass on a pair of twisted wires without interfering with each other.

In the wireless realm, the idea dates to the 1970s, when the Plessey Groundsat system gave soldiers full-duplex radio over channels within the 30- to 76-megahertz VHF band. But in those days the feat could be managed only where money was no object and where it was possible to put some distance between the transmitting and receiving antennas. In today's military systems, simultaneous transmit and receive capability on the same frequency channel is being pursued using photonics techniques to isolate the receiver from the transmitter.

But getting full duplex into civilian applications involving compact units, for example in cellular communications and Wi-Fi, is harder because the transmission signal creates at the receiver a tremendous self-interference, or echo. Because that echo can be anywhere from a billion to a trillion times as strong as the signal that needs to be detected, the system can be made to work only by canceling the echo very, very accurately.

Military systems pursue full-duplex capability by accepting more weight, bulk, and cost than their civilian counterparts could tolerate. To make the technique feasible for consumer products like smartphones, researchers had to push the state of the art in such disciplines as antennas, circuit design, and algorithms.

That's why today's wireless networks are only half duplex. Transmitters and receivers either transmit and receive in different time slots (which is called time-division duplexing, or TDD) or at the same time but at different frequencies (frequency-division duplexing, or FDD). Because time or frequency resources are being used only part of the time, such

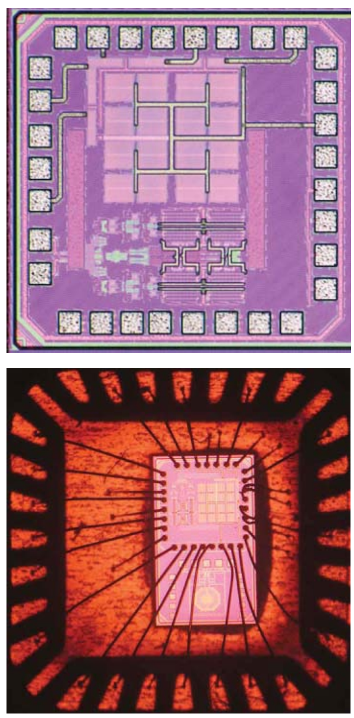

DO THE TWIST: This IC "twists" the polarity of an electromagnetic wave, making it the first nonmagnetic, on-chip circulator [chip, above; commercial package, below].

networks achieve only half the basic network capacity that is possible in an ideal full-duplex network.

To go from half to full duplex requires solving a basic problem: A wireless device has trouble transmitting and receiving at the same time for the same reason you can't hear a whisper while screaming at the top of your lungs: The interference is louder than the signal. Canceling that interference at the receiver involves subtracting the known transmitter signal. But that pithy description hardly does the task justice. Remember, the echo is a billion to a trillion times as loud as the signal you're straining to hear, so you have to cancel reverberations extremely accurately. That means performing cancellation across several domains: radio frequency, analog, digital, and even within the antenna interface, as we will describe later. The cancellation in each domain must be coordinated with the cancellation in all the others.

In addition, the transmitter signal is not entirely known. As the transmitter signal leaks into the receiver, it undergoes frequency distortion through the antenna interface. Further complicating matters is the fact that the transmitter signal can reflect off nearby objects and follow several paths that arrive at the receiver at different times. Therefore, to achieve near-perfect cancellation—in which a constructed replica of the echo signal is subtracted from the echo so that no more than, say, one part per billion of the echo remains—the wireless self-interference channel from the transmitter to the receiver must be determined and reproduced very accurately.

The interfering signals are of course received in analog circuits, and in an ideal world they'd be immediately handed off to digital circuits, which can process them with far greater flexibility. But in full duplex the echo is so powerful that the analog circuitry would hand off a badly distorted version of the echo. So we have to do some cancellation in the analog side first.

One method is called time-domain cancellation. Let's say the interference signal from the transmitter arrives at the receiver along with time-delayed copies of itself that have been reflected off nearby objects. Here, synchronizing the cancellation signal with the interference requires slowing down the cancellation signal by sending it through prepared pathways, perhaps centimeters long, that mimic the delays experienced by the interference signal. The problem is that an integrated circuit is itself much smaller than a centimeter square, so such a time-domain cancellation scheme couldn't physically fit onto the chip.

Therefore, one of us (Krishnaswamy), along with his Ph.D. student Jin Zhou

NEGAR REISKARIMIAN (2)

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

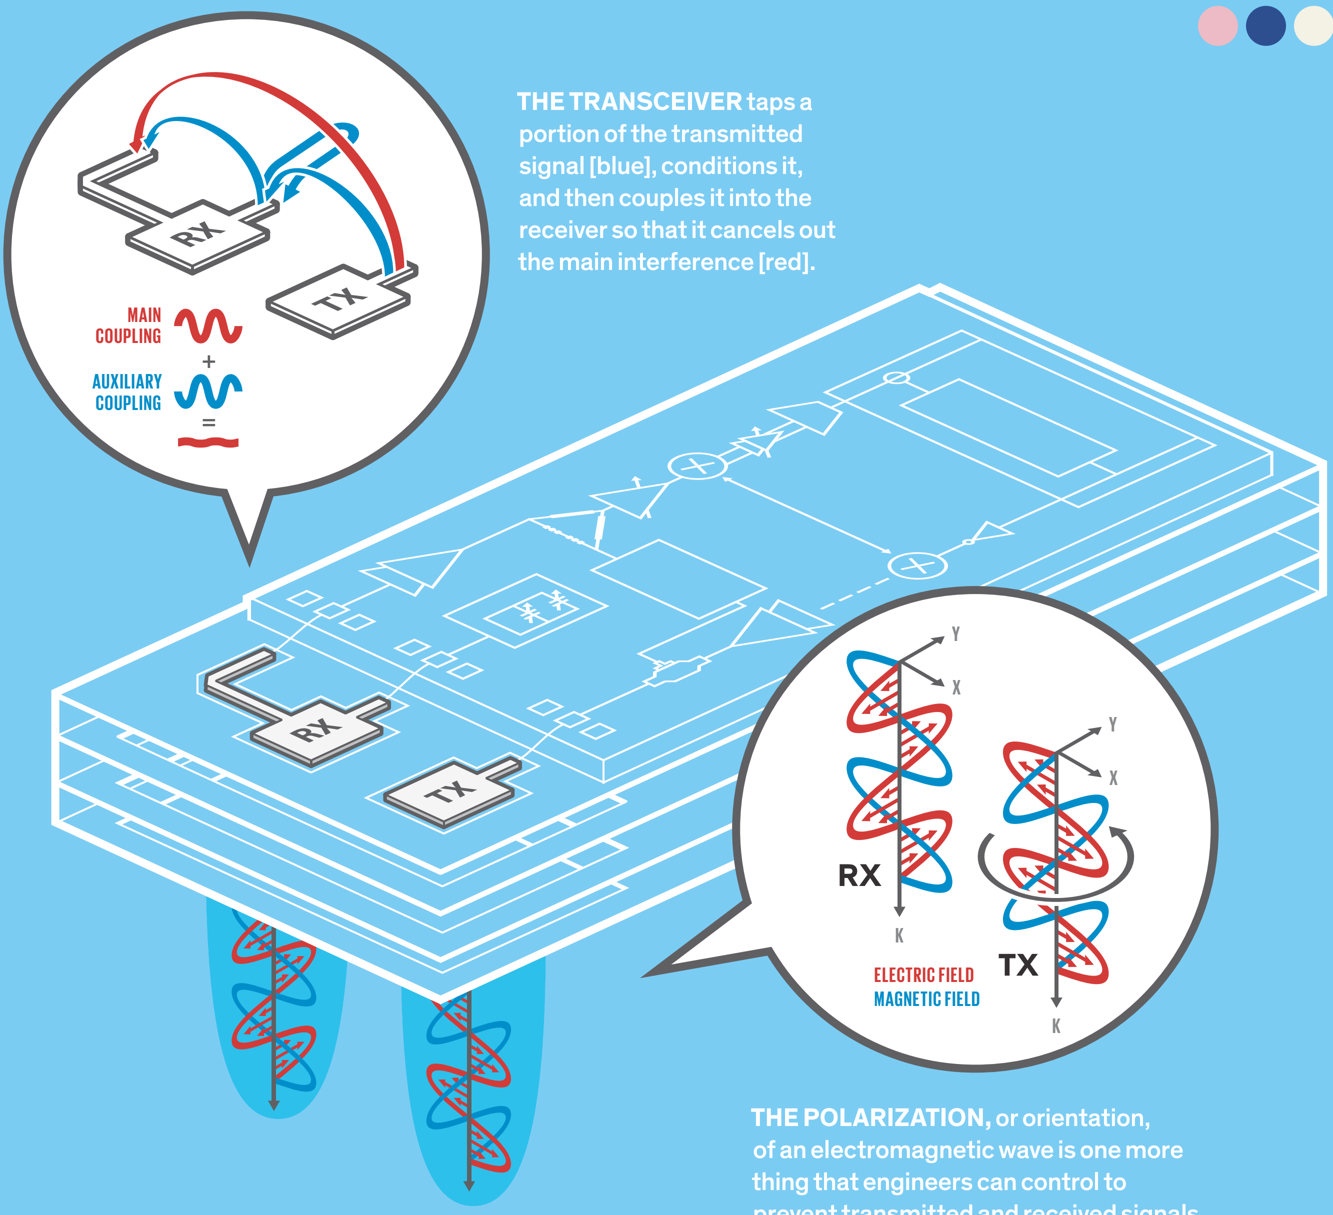

THE TRANSCEIVER taps a portion of the transmitted signal [blue], conditions it, and then couples it into the receiver so that it cancels out the main interference [red].

THE POLARIZATION, or orientation, of an electromagnetic wave is one more thing that engineers can control to prevent transmitted and received signals from interfering with each other. This 60-gigahertz full-duplex transceiver developed at Columbia University uses this polarization technique to separate the transmitted from the received signals.

Polarization-Based Duplexing

proposed another approach—frequency-domain equalization. It works something like the graphic equalizer in a stereo system, which corrects sound by adjusting the power of audio signals in particular bands of frequencies. To divide the signal into multiple frequency bands that can be individually manipulated, we use filters that each have a very sharp frequency response (or high quality factor); they take the incoming signals and let only a very narrow range of frequencies through. We use a number of filters centered at different frequencies spread out across the signal's full spectrum.

Historically, it had been impossible to build sharp radio-frequency filters on a chip. We did it in nanoscale CMOS chips by using a circuit-design trick, called the N-path filter. A conventional filter uses inductors and capacitors, but inductors are notoriously hard to put on a chip. The N-path filter instead uses switches, which is to say transistors—the essence of integrated-circuit technology.

Other groups had already put N-path filters on a chip, but we were the first to apply N-path filters for frequency-domain equalization. The result was

ILLUSTRATION BY James Provost

echo cancellation across a very wide band of frequencies for a full-duplex wireless radio.

In our system, a bank of N-path filters taps a small portion of the transmitter signal. Next it divides that RF signal into two frequency bands (though more than two is also quite feasible). Then it conditions the signal in each of those bands to mimic the self-interference that's arriving at the receiver.

This multiband approach divides the bandwidth into bite-size chunks, a divide-and-conquer strategy that makes it easier to condition each chunk of bandwidth—that is, to adjust it for power and for phase. Circuitry performs the conditioning for each band according to the weights assigned to that band. Again, to use the audio analogy, it's like dialing up the bass and dialing down the treble, and doing whatever else it takes to get the output to match the input signal.

The next step is to automate this weighting process so that the output changes accordingly as the environment changes. Of course, these environmental changes are dynamic, fluctuating from second to second, so the process has to be automated. Though we have shown some initial and promising demonstrations of such automation, there is more work to be done here.

We designed a prototype of the receiver-cum-canceler and fabricated it using 65-nanometer CMOS technology. Our full-duplex receiver can operate at any frequency between 0.8 and 1.4 gigahertz, and the RF self-interference canceler suppresses the transmitter interference, for a variety of antenna types, over a bandwidth that is about 10 times as great as what you can get with existing, conventional cancellation techniques. We achieved this 10x performance advantage with just two N-path filters in the bank. That's good enough to make it compatible with many advanced wireless standards, including LTE and Wi-Fi. More filters would enable even wider cancellation bandwidths.

Another advantage of our frequency-based cancellation scheme is its compatibility with existing wireless systems, which use what's called multiband frequency-division duplexing. As mentioned earlier, FDD is a half-duplex strategy in which the transmitter and the receiver operate at the same time but at different frequencies. It requires duplexer filters to separate the transmitter and the receiver signal at the common antenna. Because these off-chip duplexer filters cannot be tuned, today's smartphones use a separate duplexer filter to support each of the FDD bands—and 4G LTE supports 25 bands, therefore requiring 25 separate duplexer filters!

You can reduce the bulk and cost of the radio component of a cellphone by replacing those filters with just a few tunable duplexers, but such duplexers typically are less effective in isolating the transmitter from the receiver than their fixed-frequency counterparts. Consequently, the receiver is particularly prone to transmitter self-interference. And that's where self-interference cancellation comes in.

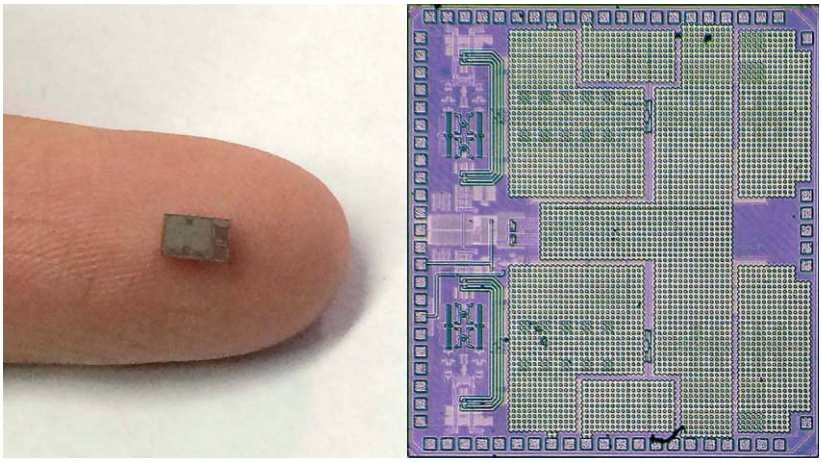

ONE SMALL PACKAGE: This is the world's first full-duplex transceiver on a chip. It cancels interference with the familiar technique of frequency-domain equalization.

JIN ZHOU (2)

The entry point for self-interference is right at the antenna, and it would be wonderful to suppress the interference there, before it has a chance to leak into the receiver. The main challenges are to keep the antenna compact—say, for use in a cellphone—and make sure that the self-interference doesn't come back every time the electromagnetic environment changes. In other words, we need the antenna to be smart.

Such a smart antenna can manipulate the radio wave's obvious electronic characteristics—amplitude, phase, and frequency—but also the extra dimension: wave polarization. A radio wave is really two fields joined at the hip, one electrical, the other magnetic—hence the word "electromagnetic." Each field oscillates at a given frequency, and the oscillation of the electric field induces the magnetic field, and vice versa. The two fields are perpendicular, and the way the pair of them are oriented in space is called their polarization. Electromagnetic waves of different polarization can pass through each other without interference.

Krishnaswamy and his Ph.D. student Tolga Dinc were able to use polarization for duplexing within a pair of compact antennas (configured for 4.6 GHz), one for the transmitter and one for the receiver. We were able to place them right next to each other because the waves that were coming and going to the two antennas were orthogonally polarized with respect to each other, which effectively isolated them. But though this isolation minimizes self-interference, it does not eliminate it entirely. That's why we also installed a port in the receiving antenna that's copolarized with the transmitting antenna. The port samples a small portion of the transmitted signal, conditions the signal through a filter, and then passes it on to the receiver port. Result: near-perfect cancellation. Because the filter | CONTINUED ON PAGE 51

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


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

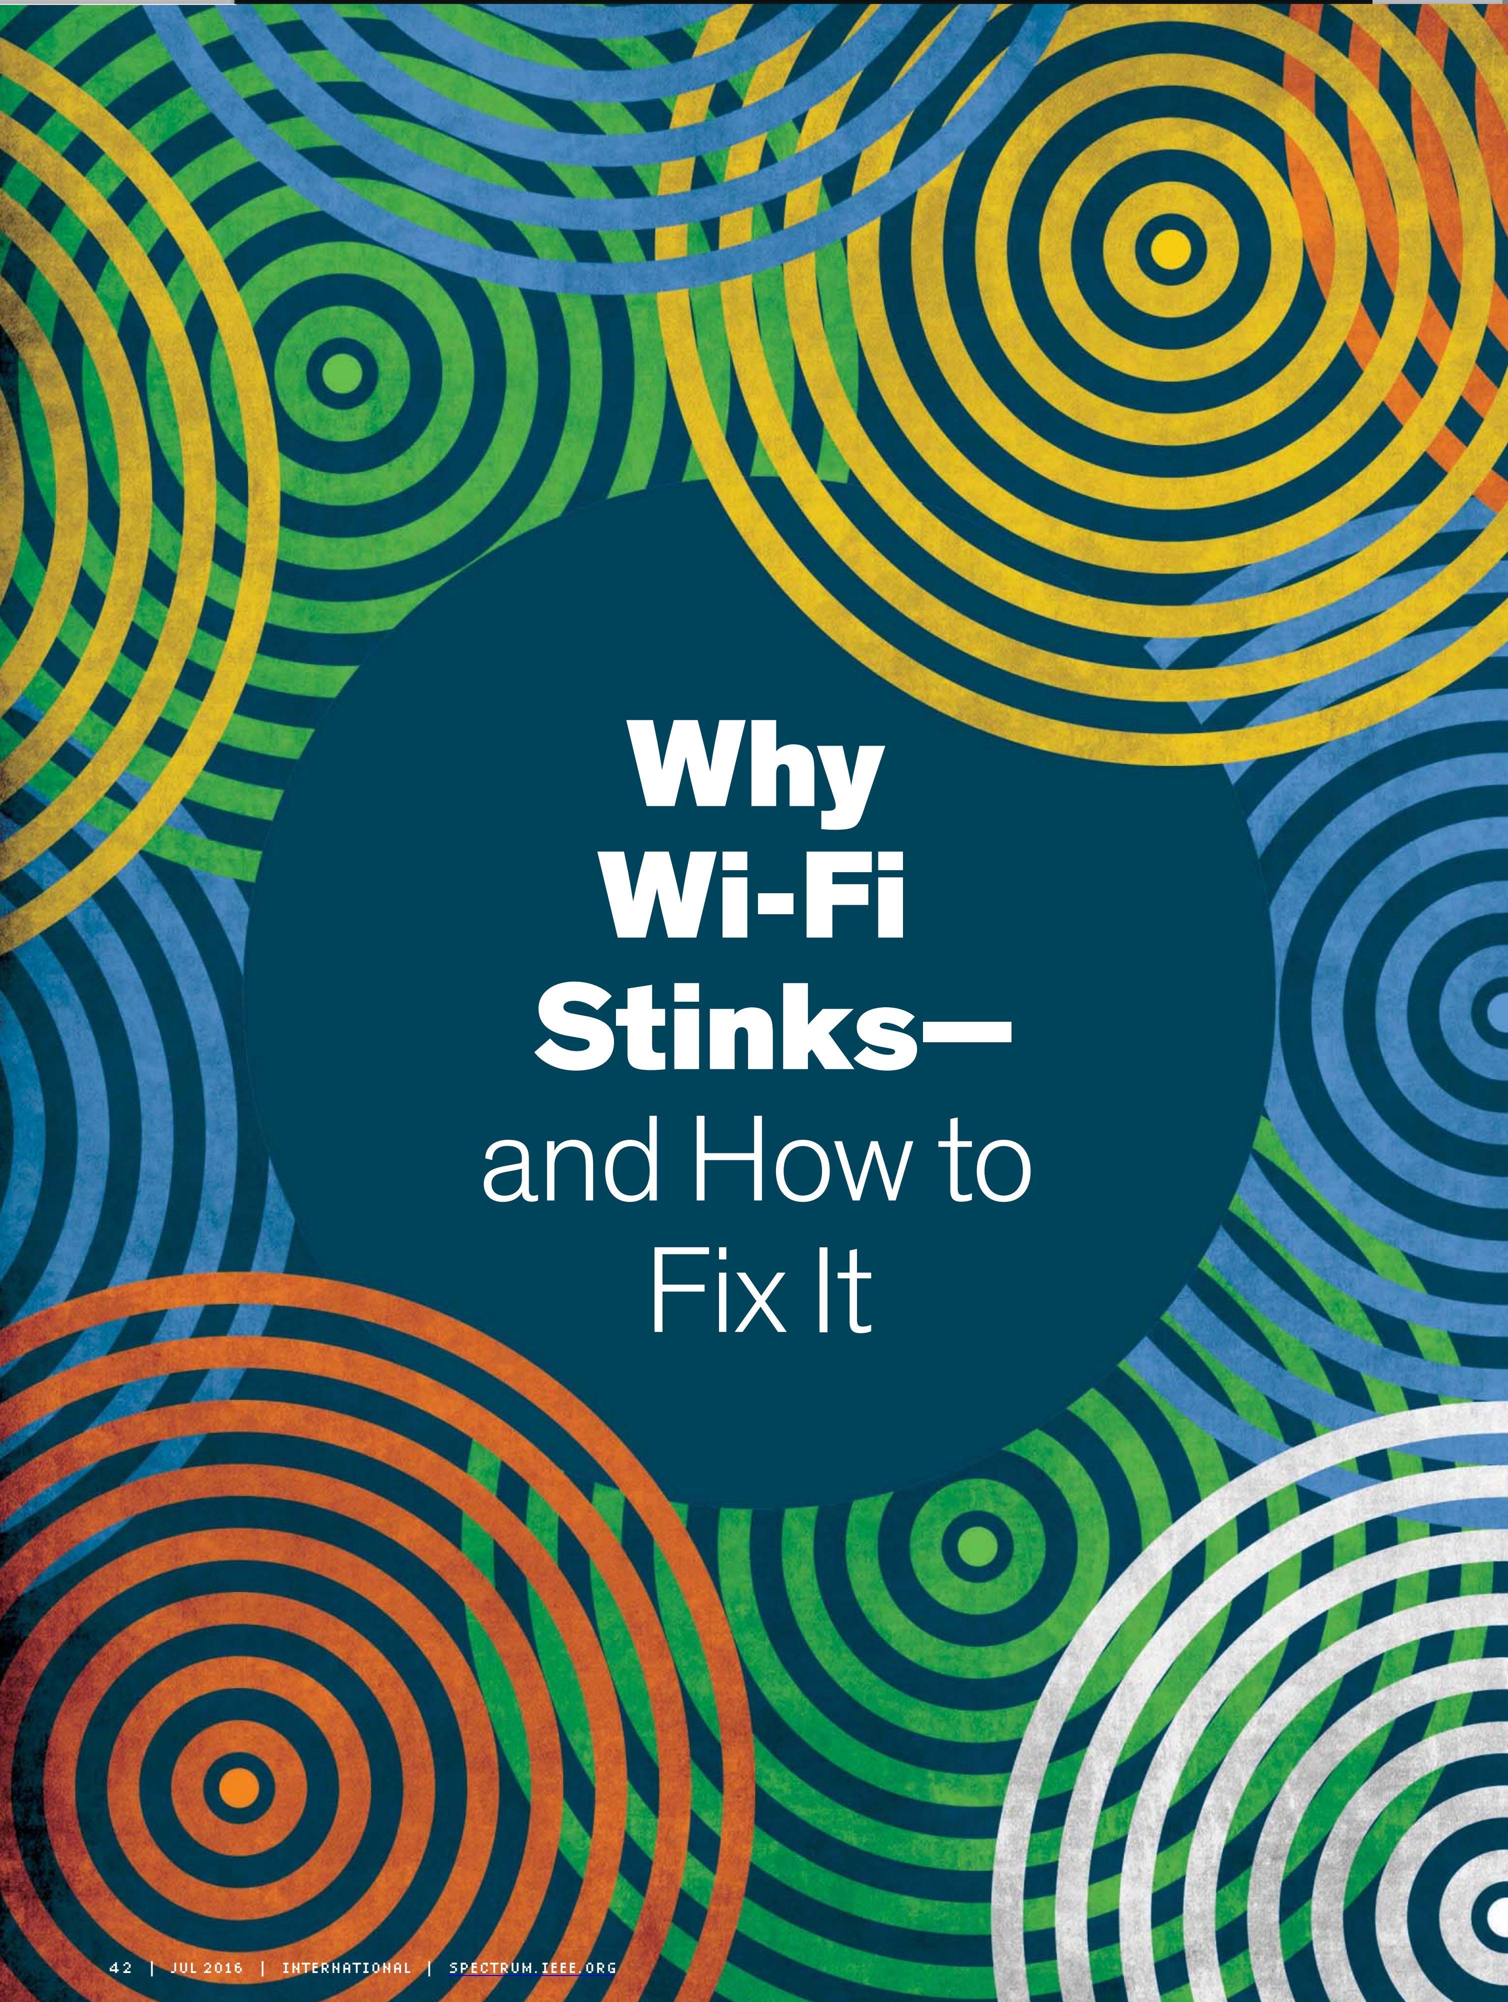

Why Wi-Fi Stinks—and How to Fix It

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

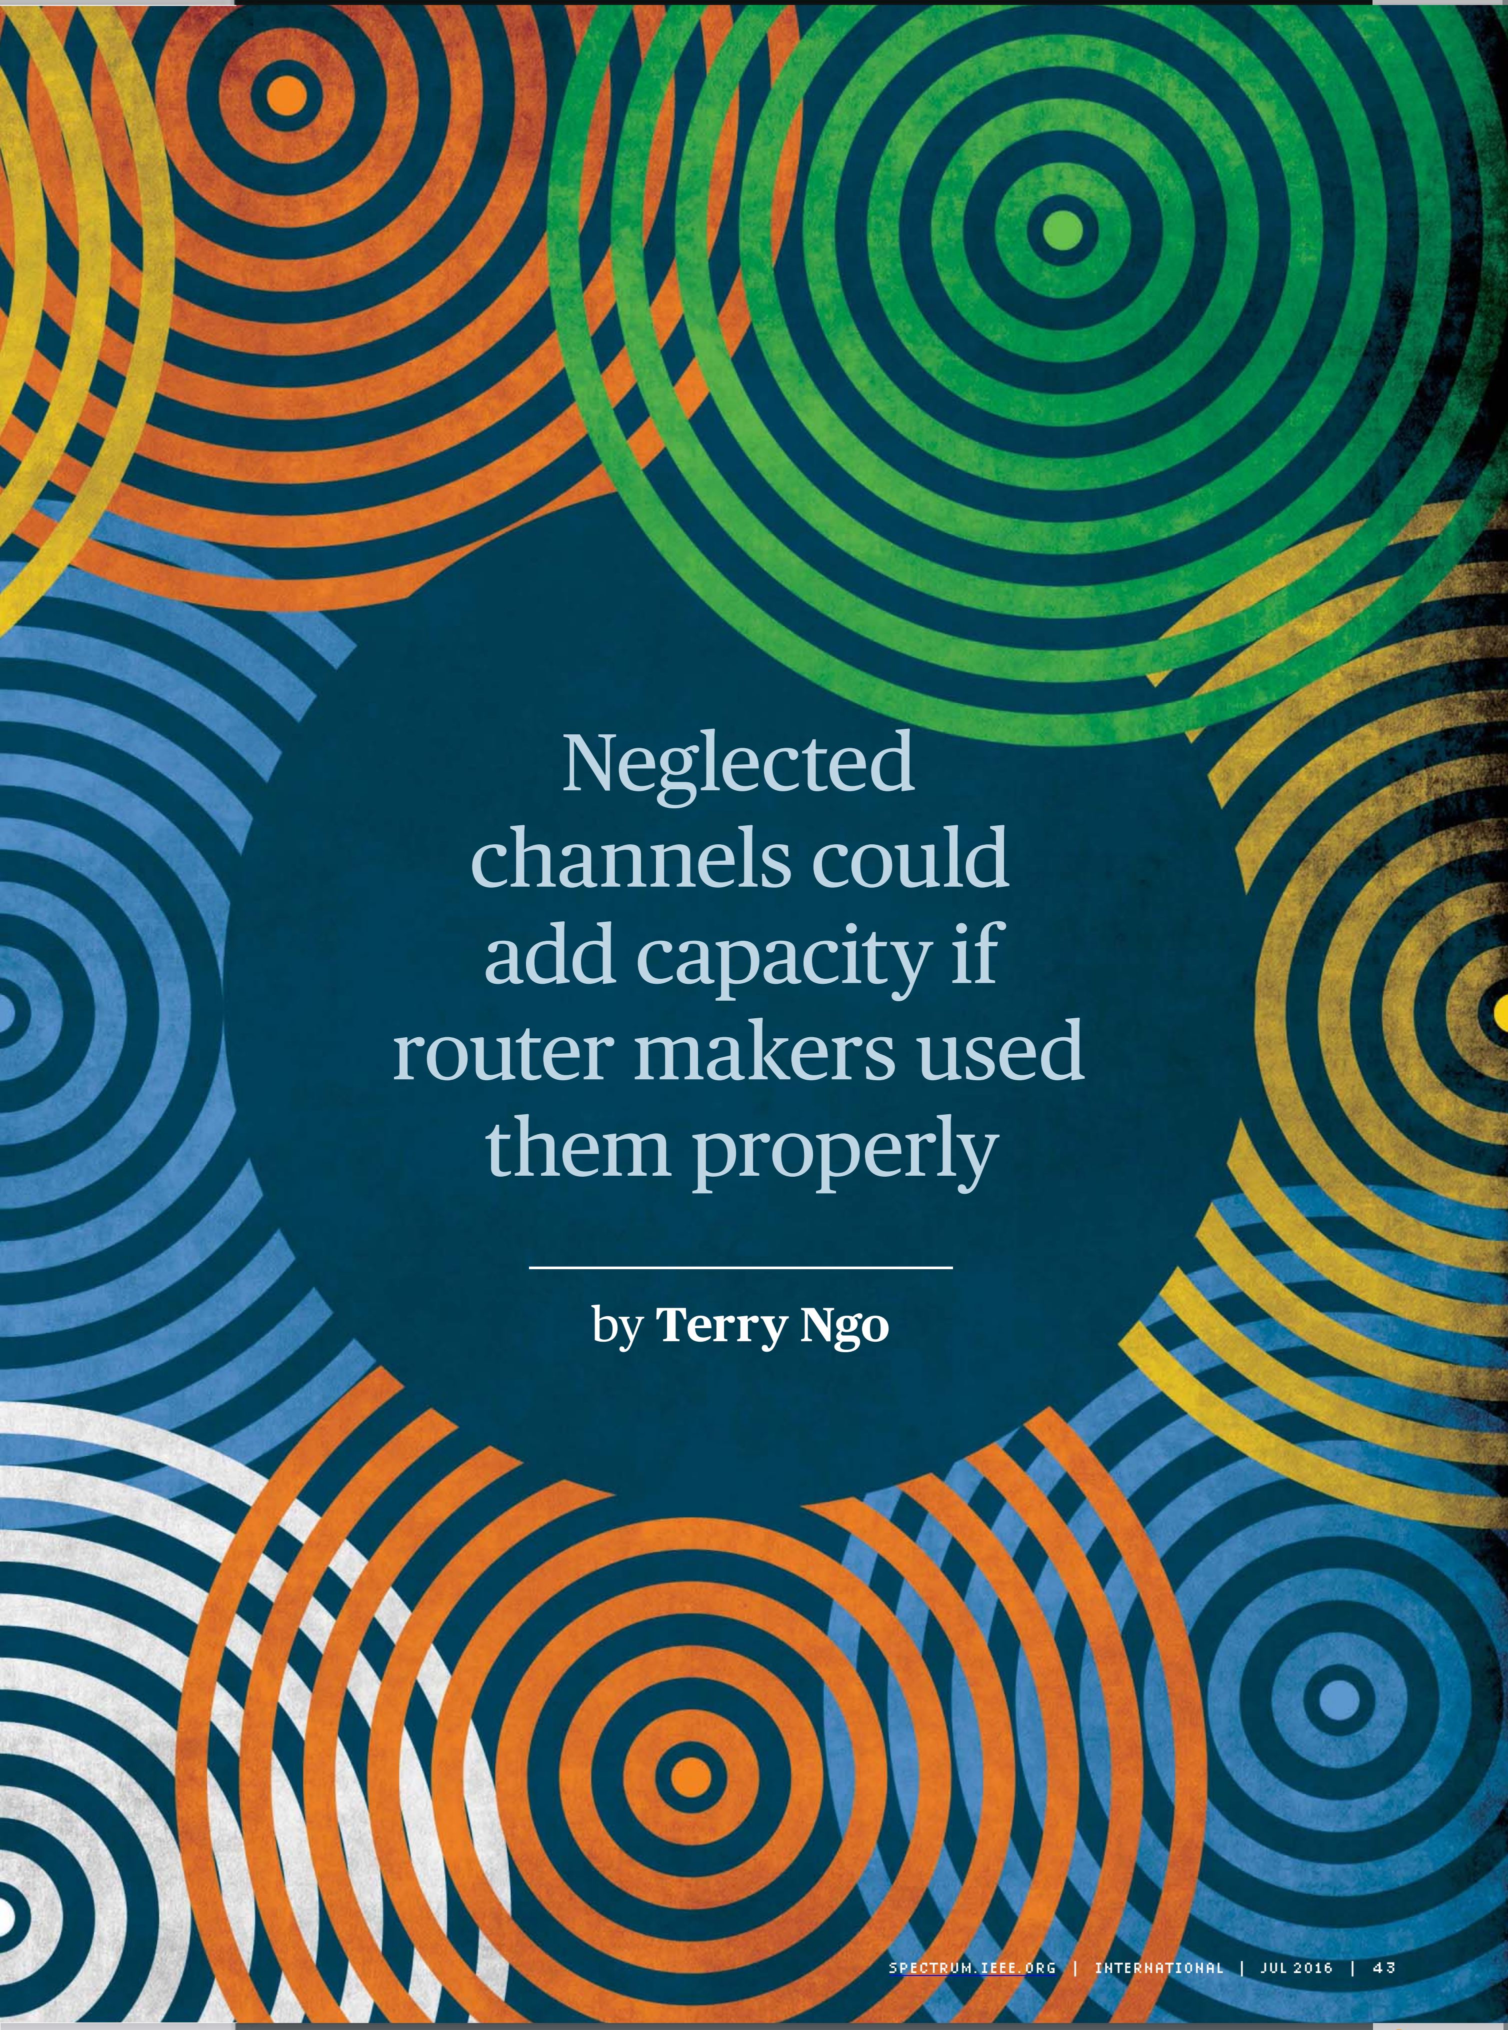

Neglected channels could add capacity if router makers used them properly

by **Terry Ngo**

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

Ask consumers in the developed world about their household Wi-Fi connection and they'll likely tell you that lately it seems to be getting worse, not better. Some might even say, "It stinks!" Even the residents of the White House have Wi-Fi problems. In an interview with the BBC just before Super Bowl 50, First Lady Michelle Obama complained, "It can be a little sketchy. The girls are just irritated by it."

The White House, along with homes inhabited by more than 80 percent of the United States and 50 percent of the worldwide population, are in urban areas where Wi-Fi connections are steadily getting worse. The reason would appear to be obvious: There are many more people–and things–using Wi-Fi than a decade ago, and the numbers continue to grow. Today, 6.4 billion connected devices are in use around the globe. By 2020, that will mushroom to 20.8 billion–that's 2.8 mobile devices for every person on Earth. So certainly the wireless highways through which Wi-Fi traffic moves have gotten and will continue to get more crowded.

But the reason for the virtual traffic jam is not as straightforward as simply having more vehicles to accommodate: The roads themselves are causing conflicts. And the situation has been made worse by three changes in the market.

First, not only is every house on your block likely to have a router but quite a few will have more than one, and many communities are also served by public Wi-Fi networks. Second, an increased demand for speed demands wider virtual lanes on the Wi-Fi highway, which means there will be fewer of them. And finally, cellular operators are dumping traffic into the Wi-Fi spectrum–to stay with the roadway analogy, imagine all the people who had been commuting by train suddenly taking to private cars.

Here's how Wi-Fi has become a victim of its own success–and what engineers can do to improve things.

Wi-Fi operates in what is known as the unlicensed spectrum. That is, though the Federal Communications Commission (FCC) generally demands licenses to use the airwaves in the United States, and national spectrum regulators like Japan's Ministry of Internal Affairs and Communications do the same elsewhere, the regulators leave some frequency bands relatively open. Users do have to comply with technical requirements, including power limits, but they do not have to apply for specific permission. There are several of these bands out there, but home Wi-Fi networks primarily operate in the 2.4-gigahertz and 5-GHz bands, because these are the only available parts of the radio spectrum that have the range and bandwidth needed. The 2.4-GHz band works the best here: It easily penetrates walls and furniture, and signals generally travel farther at the same power level as they do in the 5-GHz band.

In the United States, in the 2.4-GHz band, the FCC has authorized roughly 80 megahertz for Wi-Fi use. The channels operating there under the IEEE 802.11 standard are 20 or 22 MHz wide, so you can fit in only three nonoverlapping channels: 1, 6, and 11. The situation is only slightly different in Europe, where 13 channels allow for still just three nonoverlapping channels at a time, and in Japan, where 14 channels allow for four nonoverlapping channels.

So when you scan for available networks on your phone or computer in the United States, if you can see more than three 2.4-GHz routers (a likely scenario for anyone not living in a rural environment), or if you see only three but any one of them is not on channel 1, 6, or 11, there is interference.

Signals in the 5-GHz band have a shorter range in the home, mostly because of the walls and furniture, but that band, which extends from 5.180 to 5.825 GHz, has 24 nonoverlapping 20-MHz-wide channels in North America, with five fewer in Europe and Japan. That's a huge number of additional lanes on our crowded wireless highway. But roughly half of these channels–more in North America–are allocated for primary use by weather and military radar. Inserting Wi-Fi into this radar-priority spectrum requires special technology, so to date most consumer routers ignore these lanes. But they are important, and we'll come back to them later.

In either band, we have a set number of channels that don't interfere. As more and more routers come online, and more and more devices connect with them, interference becomes the norm. In the Wi-Fi world, when two conversations collide, all the devices go quiet and then try to talk again a little while later. The amount of time they wait is determined

Wi-Fi's World of Colliding Channels

In the 2.4-gigahertz region of radio spectrum allocated to Wi-Fi, each Wi-Fi channel can span 20 or 22 megahertz. So although it may appear that your Wi-Fi router has 11 or more channels to choose from, only a few of these channels can be used at the same time without potentially interfering with one another.

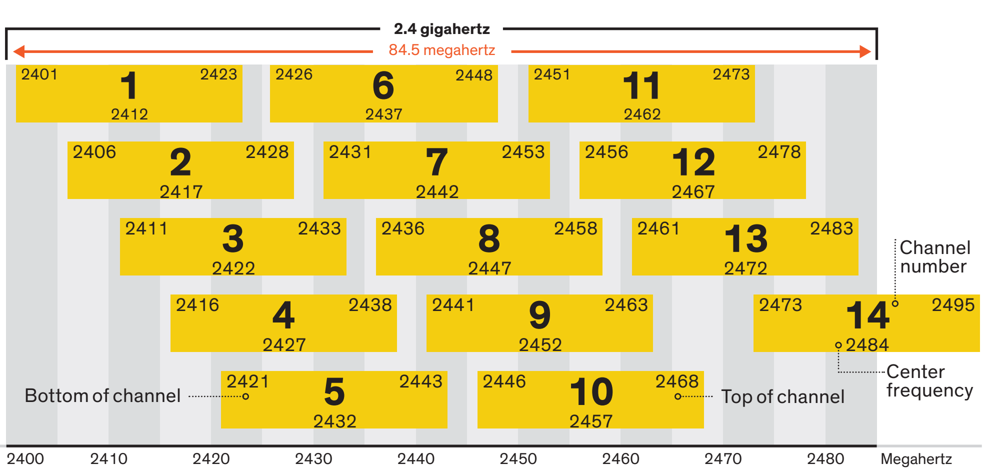

IEEE SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

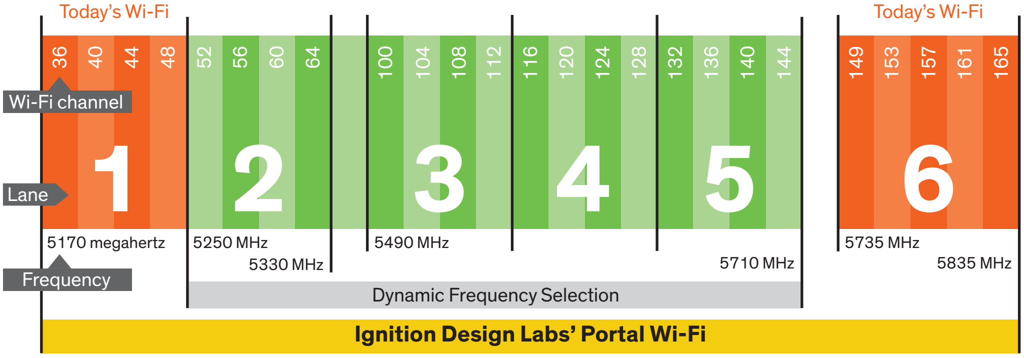

Wi-Fi's Express Lanes

Most of the traffic in the 5-gigahertz band crowds into just one or two lanes. In the remaining lanes, Wi-Fi shares the road with radar signals, and access is restricted to routers like Portal that use Dynamic Frequency Selection to detect the signals and get out of their way.

by an exponentially increasing time delay, known as a backoff. With more collisions, the backoff increases, and the Wi-Fi becomes slower and less reliable.

Today, congestion has gotten so bad in many regions that it has pretty much made the 2.4-GHz band unusable for transferring data at high rates. Several broadband-service providers (including AT&T, British Telecom, and Comcast) no longer use 2.4 GHz for video or voice, and almost all smartphone makers, including Apple, no longer recommend using their smartphones at 2.4 GHz. The latest, and fastest, variant of Wi-Fi, IEEE 802.11ac, provides for operations in only the 5-GHz band, although most Wi-Fi equipment includes both bands to accommodate older mobile devices.

So we have a limited number of lanes on the Wi-Fi highway in which your signals can travel without slowing down traffic overall. But it's not just devices that are contributing to congestion. It's the networks themselves.

Moving Wi-Fi communications from 2.4 GHz to 5 GHz initially helped with congestion but sacrificed coverage, so many consumers turned to simple boosting solutions—like network range extenders or mesh networks—to get Wi-Fi to every room in their homes. These extenders or mesh networks are placed at the range limits of a router where the transmission signals fade away. These devices listen to all transmissions, then rebroadcast them at a higher power level, sometimes on a different channel. Now you have even more Wi-Fi signals overlapping in the same frequency ranges.

Making matters worse has been the introduction of public hotspots—that is, places where Wi-Fi is available to the general public or some subset of users (like subscribers to a particular Internet service). In 2005, Spain-based Wi-Fi provider Fon Wireless pioneered the concept of community hotspots—that is, hotspots that piggyback on a private router—and they have become more and more common around the world. Today, Internet providers like AT&T, Comcast, and Verizon in the United States are quickly rolling out such hotspots, accessible to any of their subscribers, by piggybacking this functionality on the wireless gateways installed in their customers' homes. Juniper Research, based in England, estimates that by 2017 one in three home gateways around the world will allow community access by incorporating a second network identifier and allowing some of the Wi-Fi spectrum available to that gateway to be shared, typically without the awareness of the people in whose homes the gateways are installed.

Wi-Fi congestion will soon get even worse, thanks to the mobile-phone carriers, which have exhausted much of their exclusive spectrum. These wireless carriers are planning to off-load mobile-data transmissions, as much as 60 percent within the next three years, onto the unlicensed spectrum used by Wi-Fi.

The technology for that is called either LTE-Unlicensed (LTE-U) or Licensed Assisted Access (LAA). It uses 4G LTE radios and routers to send and receive data via the same 5-GHz frequencies as Wi-Fi. While carriers downplay interference to Wi-Fi users, some organizations like Cable Television Laboratories, Google, and Microsoft claim that LTE-U and LAA will absolutely increase congestion on Wi-Fi channels and degrade Wi-Fi service. In the United States, Verizon and T-Mobile have begun trial deployments for LTE-U to determine its impact on Wi-Fi. Carriers in Europe and Asia are planning similar trials.

On top of all of this, the latest variant of the Wi-Fi standard—IEEE 802.11ac—is actually reducing the number of lanes on the radio highway.

IEEE 802.11ac satisfies a growing need for speed—speed to stream high-definition videos and to allow mobile devices to conserve battery life by transmitting at high rates for only a limited time. This form of Wi-Fi delivers 1.3 gigabits per second compared with 450 megabits for 802.11n, the previous generation.

To allow data to move that quickly, 802.11ac has to merge channels. In its highest-performance configuration, IEEE 802.11ac Wave 3, it combines the entire available Wi-Fi spectrum into two 160-MHz-wide channels. This merger means that only two pairs of devices can communicate on the widest of channels simultaneously without interfering. So if one of your neighbors, say, is using one of those two channels to watch a movie, and your other neighbor uses the other one, there may be nothing left for you. Suddenly, all those extra noninterfering channels that made 5 GHz an improvement over 2.4 GHz are gone.

With developments of this kind on so many fronts, Wi-Fi connectivity is likely to soon go from often annoying to completely broken. In 2013, Ofcom, the Brit-

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

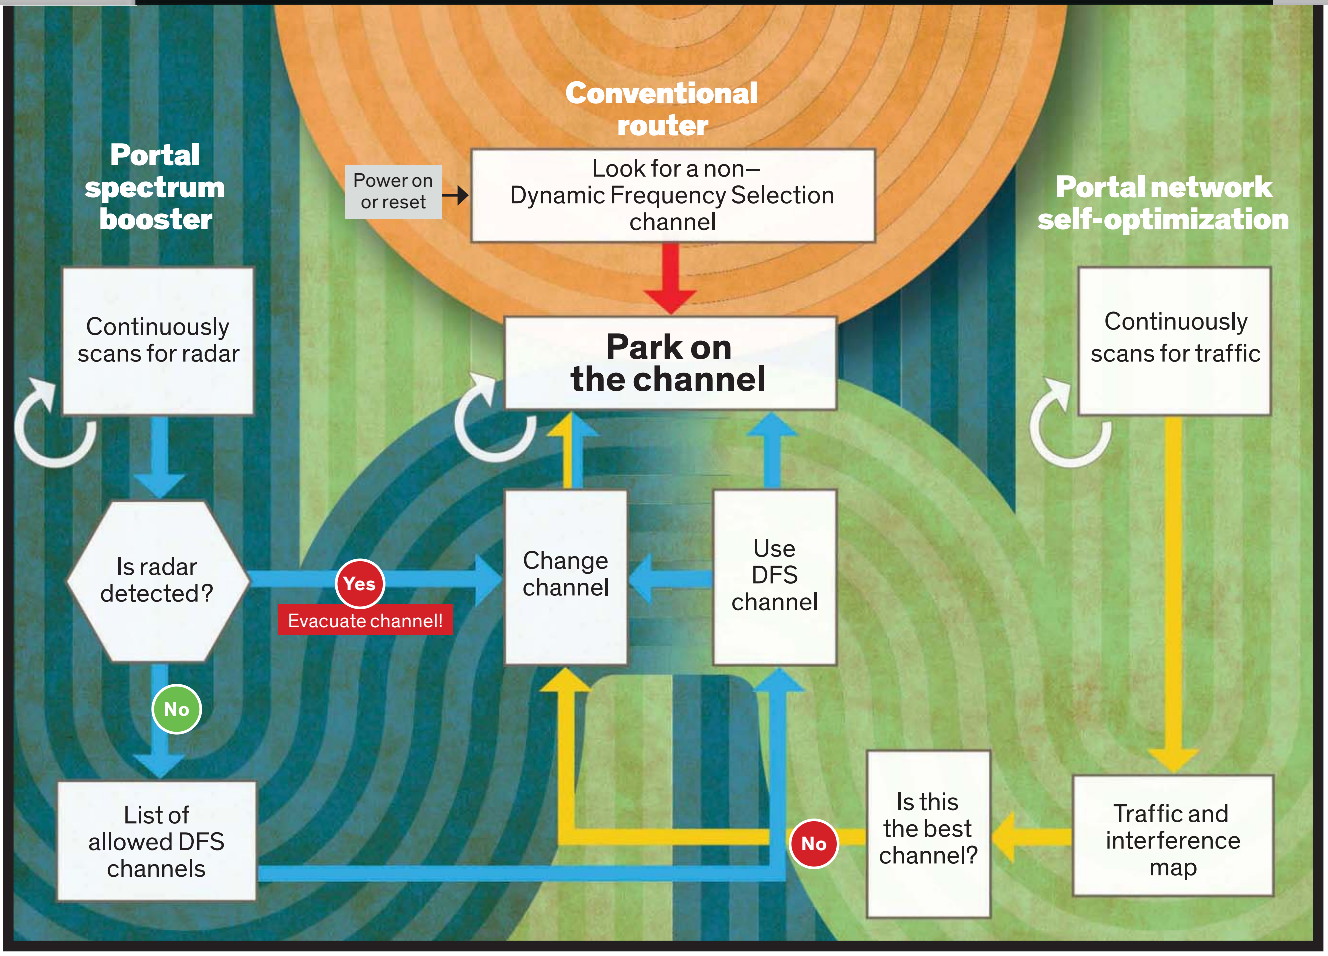

ish national telecommunications agency, published a study, "The Future Role of Spectrum Sharing for Mobile and Wireless Data Services," which predicted that Wi-Fi and mobile Internet airwaves could become critically congested by 2020–now just four years away.

To date, technical-standards developers and router manufacturers, who have worked hard to improve speed over the past 15 years, have all but ignored these issues. In particular, they haven't addressed the fact that a widespread rollout of 802.11ac, with its ability to offer wider but fewer channels, will make the congestion problem vastly worse.

There is, however, a near-term fix. Remember that chunk of 5-GHz spectrum I mentioned, where radar has first dibs and which requires special technology to use for Wi-Fi? Today, the makers of consumer Wi-Fi routers are ignoring these channels. Opening up these other frequencies to consumers would make a huge difference.

This additional spectrum was made available for Wi-Fi traffic by the FCC and other regulators around the world in 2007. The regulators realized that radar–for example, the Terminal Doppler Weather Radar system, which warns of low-altitude wind shear at airports–is not located everywhere and does not run 24/7. So the Wi-Fi industry could move Wi-Fi communications into these frequencies, as long as the devices that use these channels implement a mechanism called Dynamic Frequency Selection (DFS) to stay out of the way of radar signals.

DFS acts as a high-speed traffic cop–when it spots a radar signal in one of these protected channels, it quickly shifts all Wi-Fi traffic to another lane. There are a few rules about how this traffic cop works: It must listen for radar for at least 60 seconds before declaring a channel free to use, then continue to listen while Wi-Fi traffic is on the channel. If the mechanism detects even a 1-microsecond radar pulse, the Wi-Fi transmitter must clear the channel within 10 seconds and stay off it for half an hour.

The vast majority of mobile devices introduced in the last three or four years have radios that can operate in these bands and have the required software to respond to instructions from what is called a DFS master. But they need that DFS master built into the router to tell them when it's all right to use a radar-priority channel and when they need to move aside.

DFS-master technology isn't trivial to implement. Radar pulses are hard to detect because they are very, very fast (each pulse lasts just about one half of 1 microsecond) and can be present at very low power levels (as low as -62 to -64 decibel-milliwatts). Incorporating radar-detection tools eats up bandwidth–as much as 17 percent–because a router must listen on a channel for a minimum

A Wi-Fi Router That Dodges Radar and Interference

To optimize its use of the Wi-Fi spectrum while protecting radar systems sharing the same frequencies, Ignition's Portal router continually looks for radar traffic, vacating a channel immediately if it detects radar, and making that channel available for use again when appropriate. It also looks for general interference to select the least-crowded channel.

of 60 seconds before concluding that it is clear for transmitting, and then continue to listen during and in between normal transmissions.

Currently, DFS-master technology is available in expensive routers, the kind that are typically installed only by large businesses. It is migrating to some lower-cost consumer-level routers in Europe and Japan. But both the expensive business versions and the cheaper consumer versions are not all that clever: When they detect radar, they quickly shift traffic back to a set default channel in the non-DFS part of the 5-GHz band—a crowded spot. And they don't go back to using radar-priority channels until their users reboot their routers. In a business environment, that is often scheduled to happen daily, but in a residential setting, it might take weeks or months before the user realizes that the router is not performing well and needs to be reset. So even routers with DFS capabilities are staying out of those express lanes, at least most of the time.

Still, it's in these DFS channels where the solution to Wi-Fi congestion lies. The trick is creating a less costly and more effective radar-detection technology. My colleagues and I at Ignition Design Labs, located in San Jose, Calif., think we've figured that out.

We have designed an enhanced router, called Portal, which incorporates a full-spectrum radio scanner and a CPU dedicated to radar detection and channel management alongside standard router hardware. The scanner continuously sweeps the entire 5-GHz band for radar as well as for Wi-Fi traffic and general interference. Making this detection system completely separate from the Wi-Fi sending and receiving radio solves a lot of the problems of current radar-detection technology, which shares the main processor and Wi-Fi radio for radar detection as well as communications.

In such standard hardware, the radio sees what is happening only within a single channel in the DFS band at a time, so a DFS master can monitor only one specific DFS channel at a time. And when

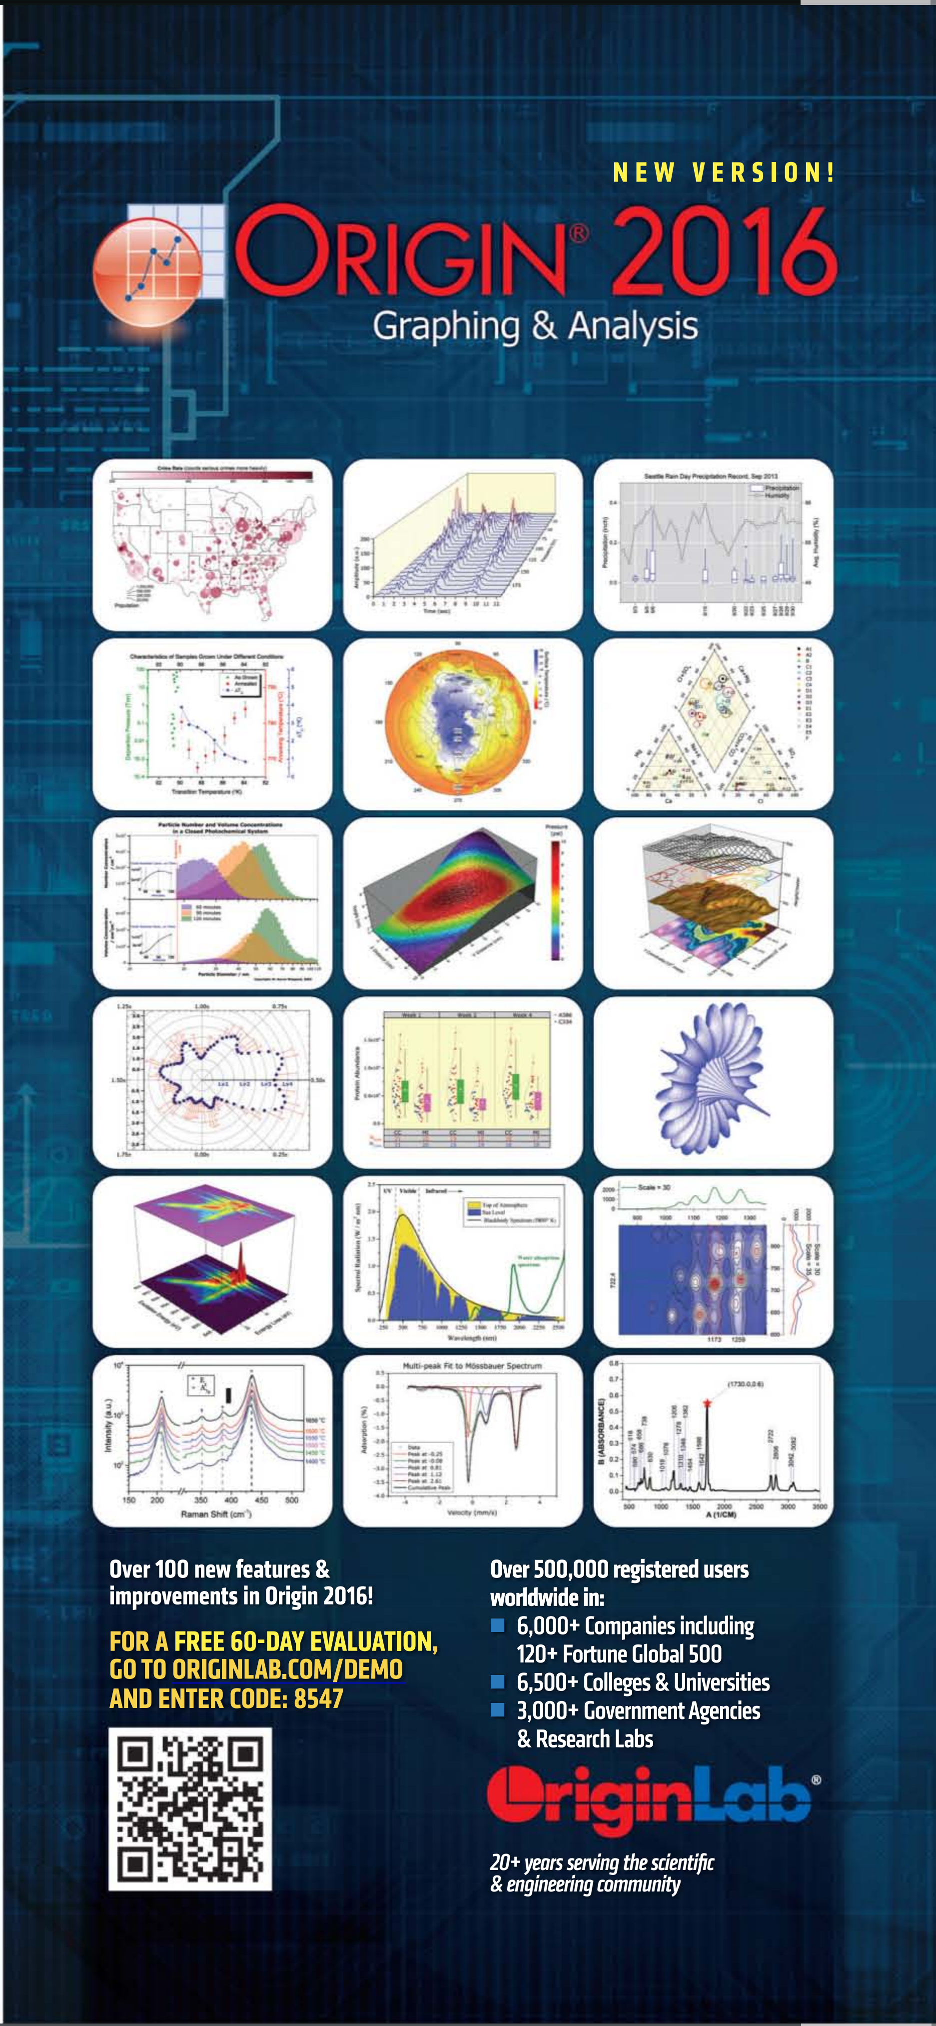

a DFS master's radio first monitors a DFS channel, the FCC requires that the radio refrain from transmitting on any channel for at least 60 seconds to make certain it does not interfere with the receiver looking for radar. To avoid this kind of shutdown, most of these radio designs look for an open DFS channel only when the router is reset.

A second radio dedicated to the task of sensing radar signals removes that barrier. It can also regularly scan through all the channels. So when a radar pulse is detected in a channel being used for a transmission, the system will know whether or not another DFS channel is clear of radar, in which case it can move data traffic over there instead of to a preset default channel. And when the DFS master does require the router to abandon a channel because it detects radar traffic, it can automatically go back and recheck that channel, after a 30-minute required waiting period, without having to shut down ongoing transmissions.

Meanwhile, the dedicated CPU can minimize the number of false radar alerts, reducing the times that Wi-Fi traffic has to vacate a channel. Today, when the processor in the router is handling a large number of Wi-Fi streams–you're watching Netflix, your child is playing a game, other family members are listening to music or browsing Facebook–it doesn't have the processing power left to analyze radio energy it detects in the protected channels to determine whether it fits the pattern of a radar signal. So it errs on the side of caution–if it detects any interference in its channel (which might just be Wi-Fi traffic from a neighbor's router), it vacates the channel until the router is reset.

We're trying to make the process of assigning channels even more intelligent by gathering information not only on radar but also on all sorts of interference, and sending that information, along with data on general Wi-Fi and radar-traffic patterns, to a cloud server; there, our software analyzes the data and makes adjustments to how our Portals behave: We call it network self-optimization.

With this system, we can determine the best channels for Wi-Fi devices to use in different places. For example, let's say we know that at 8 p.m. in Europe, where DFS channels are already being used by consumers in a limited way, the default DFS channel, channel 100, gets very busy. We can then shift one user's traffic to channel 132 and his neighbor's traffic to channel 154. This kind of coordination can have a huge impact on the quality of Wi-Fi communications.

We have received approval for our technology from the FCC and will ship our first products later this summer in North America and late autumn in Europe. We are also working with some manufacturers of Wi-Fi equipment and Internet providers in anticipation that they will eventually incorporate our hardware into their Wi-Fi routers and gateways.

It is essential to get this kind of comprehensive, intelligent system for managing Wi-Fi resources out into the world before Wi-Fi becomes so unreliable as to be unusable.

Those of us who work in the communications industry didn't manage the way Wi-Fi devices use the 2.4-GHz spectrum, and we essentially exhausted it. But we were able to move to the 5-GHz spectrum, so consumers didn't really notice when 2.4 GHZ ran out. But when we exhaust the 5-GHz spectrum, we will be out of real estate, at least for the immediate future.

In the longer term, technology will be developed to move some of the traffic to other types of communications networks that are not compatible with current Wi-Fi. The FCC has several frequencies under consideration for possible spectrum reallocation, including small amounts at 5.9 GHz, 4.9 GHz, and 3.5 GHz. But this process of reallocating spectrum can take years, even decades. And all these frequency ranges encompass radar and other primary uses (like public-safety communications). So if these ranges are approved, they will also require intelligent use of DFS technology.

That's why finding affordable, robust technology to allow us to use all the lanes on today's wireless freeways is the only way out of the data traffic jam. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/wi-fi0716

Computing With Carbon Nanotubes

CONTINUED FROM PAGE 29

groups, including our own, have used these approaches to create both *n*-type and *p*-type CNFETs. Moreover, we have demonstrated that these two devices can perform at comparable levels, an important consideration when designing circuits.

Another source of concern has been the electrical connections between carbon nanotubes and the metal source and drain contacts. When these connections aren't perfect, they can have high electrical resistance, which consumes a lot of energy and slows computation. But this problem plagues any transistor as its dimensions get smaller and smaller. Pairing materials at the atomic scale is difficult and unforgiving. Fortunately for carbon-nanotube transistors, models suggest that we don't need to have ideal contacts in order to make high-performance CNFET circuits. What's more, recent work at IBM shows that big improvements can be made to the connection between carbon nanotubes and their metal contacts by essentially fusing the two materials, creating a near-seamless transition from carbon to metal.

We'll have to bring together all these pieces of carbon-nanotube device fabrication, including doping and contact construction, if we hope to make these chips competitive with silicon. Further progress in decreasing variability will also be a priority, since a chip's performance isn't limited by its best transistors but by the worst ones. These are big challenges, but there will be a considerable reward for overcoming them. Models calibrated to experimental data suggest that, even with variations and other imperfections, microprocessors made from carbon nanotubes could outperform future silicon circuitry by a factor of 10 when it comes to the energy-delay product, or EDP, a metric of circuit speed and energy efficiency. The benefits will be even greater as transistors continue to shrink, because silicon devices will reach their size limit well before carbon-nanotube devices do.

Carbon nanotubes have other advantages, too. An emerging area of chip design is 3D monolithic integration, which creates a high-rise of circuits built directly on top of one another, connected by a dense forest of ultrafine vertical wires.

There is a strong motivation for this sort of structure. Today's computers spend a very large amount of energy and time shuttling data inside a processor and between the processor and off-chip memory, across millimeters or centimeters of space. Today we can create 3D chips that mix logic and memory by making individual 2D silicon chips and then stacking and bonding them together. The layers are connected through

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

extremely large and sparse vertical wires. Creating 3D structures with far denser connections would open up bottlenecks that exist today between logic and memory, which would greatly benefit applications that depend on processing vast amounts of data.

Such monolithic 3D integration is very difficult in silicon; the high temperatures commonly used to manufacture the transistors can melt metal wiring and damage the transistors beneath it. But the carbon-nanotube transfer technique we described earlier allows us to build circuits at temperatures that won't damage underlying circuits. That's because the approach essentially separates the high-temperature growth phase of the material—which happens at around 900 °C—from the rest of the circuit fabrication process. This isn't speculation; we have already demonstrated several working CNFET-based monolithic 3D integrated circuits. The process is compatible with silicon and can even be used to add carbon-nanotube-device layers on top of a base of silicon transistors.

Projections suggest that combining carbon-nanotube transistors and emerging memory devices in a 3D system could yield a 100- to 1,000-fold benefit in energy efficiency over a traditional silicon chip with off-chip memory (this estimate assumes that both chips are manufactured with similar feature sizes).

Of course, as carbon-nanotube research continues, so do other efforts to make better transistors. A relatively new approach, the negative capacitance FET, uses material between the gate and the channel to boost the effect of the gate voltage, reducing the energy needed to switch. This design is compatible with carbon nanotubes and could be used to provide even greater benefits. Other alternative technologies, such as transistors that exploit electron tunneling and logic based on manipulating the spins of electrons, are also being pursued. But CNFETs remain the only emerging nanotechnology that has achieved such a high level of performance, complexity of system demonstration, and manufacturing maturity.

So when will we see this carbon-nanotube circuitry in our smartphones and data centers? Like any new technology, the carbon nanotube will merge onto the computing technology highway only if chipmakers decide to take it up in earnest. Thanks to these recent improvements in carbon-nanotube fabrication and design—and the fact that the material can actually complement silicon—that goal is becoming more and more realistic. This is one upstart that could overtake the silicon steamroller, if we'd only give it some gas. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/carbonnanotube0716

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

ONE CHIP, TWICE THE BANDWIDTH

CONTINUED FROM PAGE 40

can be programmed in the field, it can reconfigure its cancellation so that it meets the demands of a changing electromagnetic environment.

Our prototype achieved an isolation of 50 decibels—that's five orders of magnitude, or 100,000-fold—over a 300-MHz bandwidth. This level of isolation represents a 1,000-fold improvement over the isolation that is achieved without the polarization-based cancellation. Even when we put a strongly reflecting metallic plate near the antenna to heighten the self-interference, we were still able to reconfigure the system and fully restore the cancellation.

Our method of coordinating antenna design with cancellation techniques is easy to adapt to higher frequencies—30 GHz and above. In that part of the spectrum, the wavelengths are measured in mere millimeters, so the antennas that receive and transmit them are small as well.

These high frequencies are particularly promising for next-generation communication networks because they offer significantly wider bandwidths than radio frequency. We have implemented a 60-GHz full-duplex transceiver IC that employs both our reconfigurable polarization-based antenna cancellation and RF and digital cancellation. The device achieves self-interference suppression of nearly 80 dB (or 100 million-fold) over 1-GHz bandwidth. It thus enabled the world's first millimeter-wave full-duplex link over a distance of almost a meter. That's a respectable distance for millimeter-wave links, which are being considered for various short-range applications, such as wireless USB connections.

Our system can work in both Wi-Fi and cellular networks. The cellular application is the harder one because every link must be carefully coordinated by the base station, which might assign one person to one frequency and another person to another, timing everything perfectly. But Wi-Fi is a free-for-all, with none of these constraints. Also, cellphones transmit over a kilometer or more, but Wi-Fi reaches only a few score meters, producing signals that are fewer in number and much lower in power. All this makes it easier to perform echo cancellation.

One of the authors (Zussman) and his student Jelena Marašević were able to analyze the benefits from full duplex running on real chips (not merely in an idealized scenario, as others had done). They found that you can't always assume that the cancellation is perfect: The received signal needs to be a lot stronger than the weak echo that remains even after you've done all your cancellations.

In the meantime, several problems will have to be solved before we can declare wireless full duplex fully achieved. First, to take full advantage of the compactness of our duplex-on-a-chip, we have to construct what's called

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

a circulator—a device that shares a single antenna between the transmitter and the receiver. That's not easy because such a circulator must be nonreciprocal—that is, it must process a signal coming in differently from the way it processes a signal going out. Only in that way could a single antenna serve as a transmitter and a receiver simultaneously.

But treating a signal differently depending on which way it's moving violates the basic physical precept of Lorentz reciprocity. Fortunately, this precept applies only to *most* materials and systems. Ferrite materials are one exception, and indeed researchers have used them to build nonreciprocal circulators for years.

The ferrite materials work by twisting an electromagnetic wave clockwise with respect to the axis of motion when it's going one way, counterclockwise when it's going the other way. But because ferrite circulators don't fit on a chip, we came up with another way: We twist the wave with switches—transistors, in other words.

This past April, Krishnaswamy and his Ph.D. student Negar Reiskarimian did manage to build a nonreciprocal circulator that uses transistors to mimic the behavior of ferrites. This was the first such unit ever built on a chip, and we integrated it with a single-chip, full-duplex, echo-canceling receiver. The result: single-antenna full duplex.

A second challenge is to extend our method of canceling self-interference to transceivers that have multiple inputs and multiple outputs (MIMO). Such transceivers are common in base stations, and they can transmit over multiple parallel streams, greatly increasing the data rate. Unfortunately, self-interference plagues every possible transmitter-receiver pair, and if you try to implement the filters to handle every possible pair, the complexity will scale quadratically with the number of MIMO elements. This is a real problem, and though we have some ideas on how to proceed, there is a lot of work ahead.

But even now, we expect to be able to use full duplex on a chip to improve key aspects of existing wireless systems. One of the first applications, probably within two years, should come in short-range wireless links and in systems, such as Wi-Fi, in which the received signal-to-noise ratio is generally high. Getting the technology into cellphones, with their more powerful self-interference signals, will probably take about five years. Another application that will benefit is fixed, point-to-point microwave and millimeter-wave backhaul and relays, which form the backbone of our telecommunications network.

Many major players in the wireless industry have expressed interest in our work, and some, such as Qualcomm, have even funded us. We have also received funding from the U.S. Defense Advanced Research Projects Agency and the U.S. National Science Foundation.

Five years down the line, we expect to be using full duplex for small cellular networks. We are confident that the day will finally come when all our wireless devices will speak and listen at one time, on one channel, through one antenna, within one chip. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/fullduplex0716

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Dean of the School of Engineering
at Ecole polytechnique fédérale de Lausanne (EPFL)

EPFL – the Swiss Federal Institute of Technology Lausanne – invites applications for the position of Dean of the School of Engineering.

With an international reputation for excellence, EPFL is one of the foremost European institutions of science and technology. Known for the highest standards of academic distinction and community contribution, EPFL fosters transdisciplinary research collaborations and technology transfer with a first class infrastructure. Deeply engaged regionally, nationally and internationally, its campus provides a unique and exciting learning environment located in the French-speaking area of Switzerland, next to Lake Geneva, at the foot of the Alps. For more information about EPFL, please visit www.epfl.ch.

The exceptional quality of teaching and research of the School of Engineering is internationally recognized. The School has a strong track record of major technological advances and strategic ties with industry, academia and government. With more than 120 faculty members, 2'200 students in bachelor and master programs, and 700 PhD candidates, the School's extensive research program is funded through public and private sources including the Swiss Confederation, the European Union, private foundations and industrial partners. Additional information about the School is available at sti.epfl.ch.

Reporting to the President as part of EPFL's senior management, the Dean provides vision and leadership across the School's teaching and research activities, and oversees its administration. The successful candidate will have an exemplary academic record, recognized achievements in research, teaching, scholarly activities and service, demonstrated leadership in a collegial university setting, and the ability to inspire. Possessing outstanding management and interpersonal skills, s/he will be committed to fostering interdisciplinary teaching and research and to promoting the endeavors of the School at the campus, regional, national, and international level. Experience in building diverse and collaborative teams, relating to a range of internal and external partners, and experience in fundraising will be important assets in this role. A prior working knowledge of French is not required.

The position offers competitive personal compensation, tenure at the full professor level, and financial support for the candidate's research program. The candidate should be willing to act as Dean for at least one term of 4 years and to start as early as possible in 2017.

Please submit a curriculum vitae, a vision statement and the names of up to five professional references by **August 31st, 2016** using the following website:

https://academicjobsonline.org/ajo/jobs/7314

Inquiries, nominations, and expressions of interest can be addressed to:

Prof. Michael Unser
Chair of the Search Committee
Email: michael.unser@epfl.ch

EPFL is committed to expanding the ranks of women on its faculty, qualified women are enthusiastically encouraged to apply.

IEEE Membership Can Help You **Reach Your Personal and Professional Goals**

Gain access to the latest IEEE news, publications and digital library. Give and receive personal mentoring. Network locally and globally with IEEE members. And that's only the beginning. Discover how IEEE can help jumpstart your career.

Visit **www.ieee.org/join** today.

"Participating in IEEE has developed me as a well-rounded engineer and helped me shine during networking events."

-Likhitha Patha
Electrical Engineering Student, IEEE Brand President, Virginia Polytechnic Institute and State University

IEEE

HOT! HOT! HOT!

Constructed in 1949 in the citadel of Mont-Louis in the French Pyrenees, this solar-powered death ray used concentrated photonic energy to melt just about anything that wandered into its 3,000 °C beam. It was never used in combat due to an easily exploitable 12-hour safety lockout (otherwise known as nighttime) and was eventually repurposed to fire decorative pottery. ■

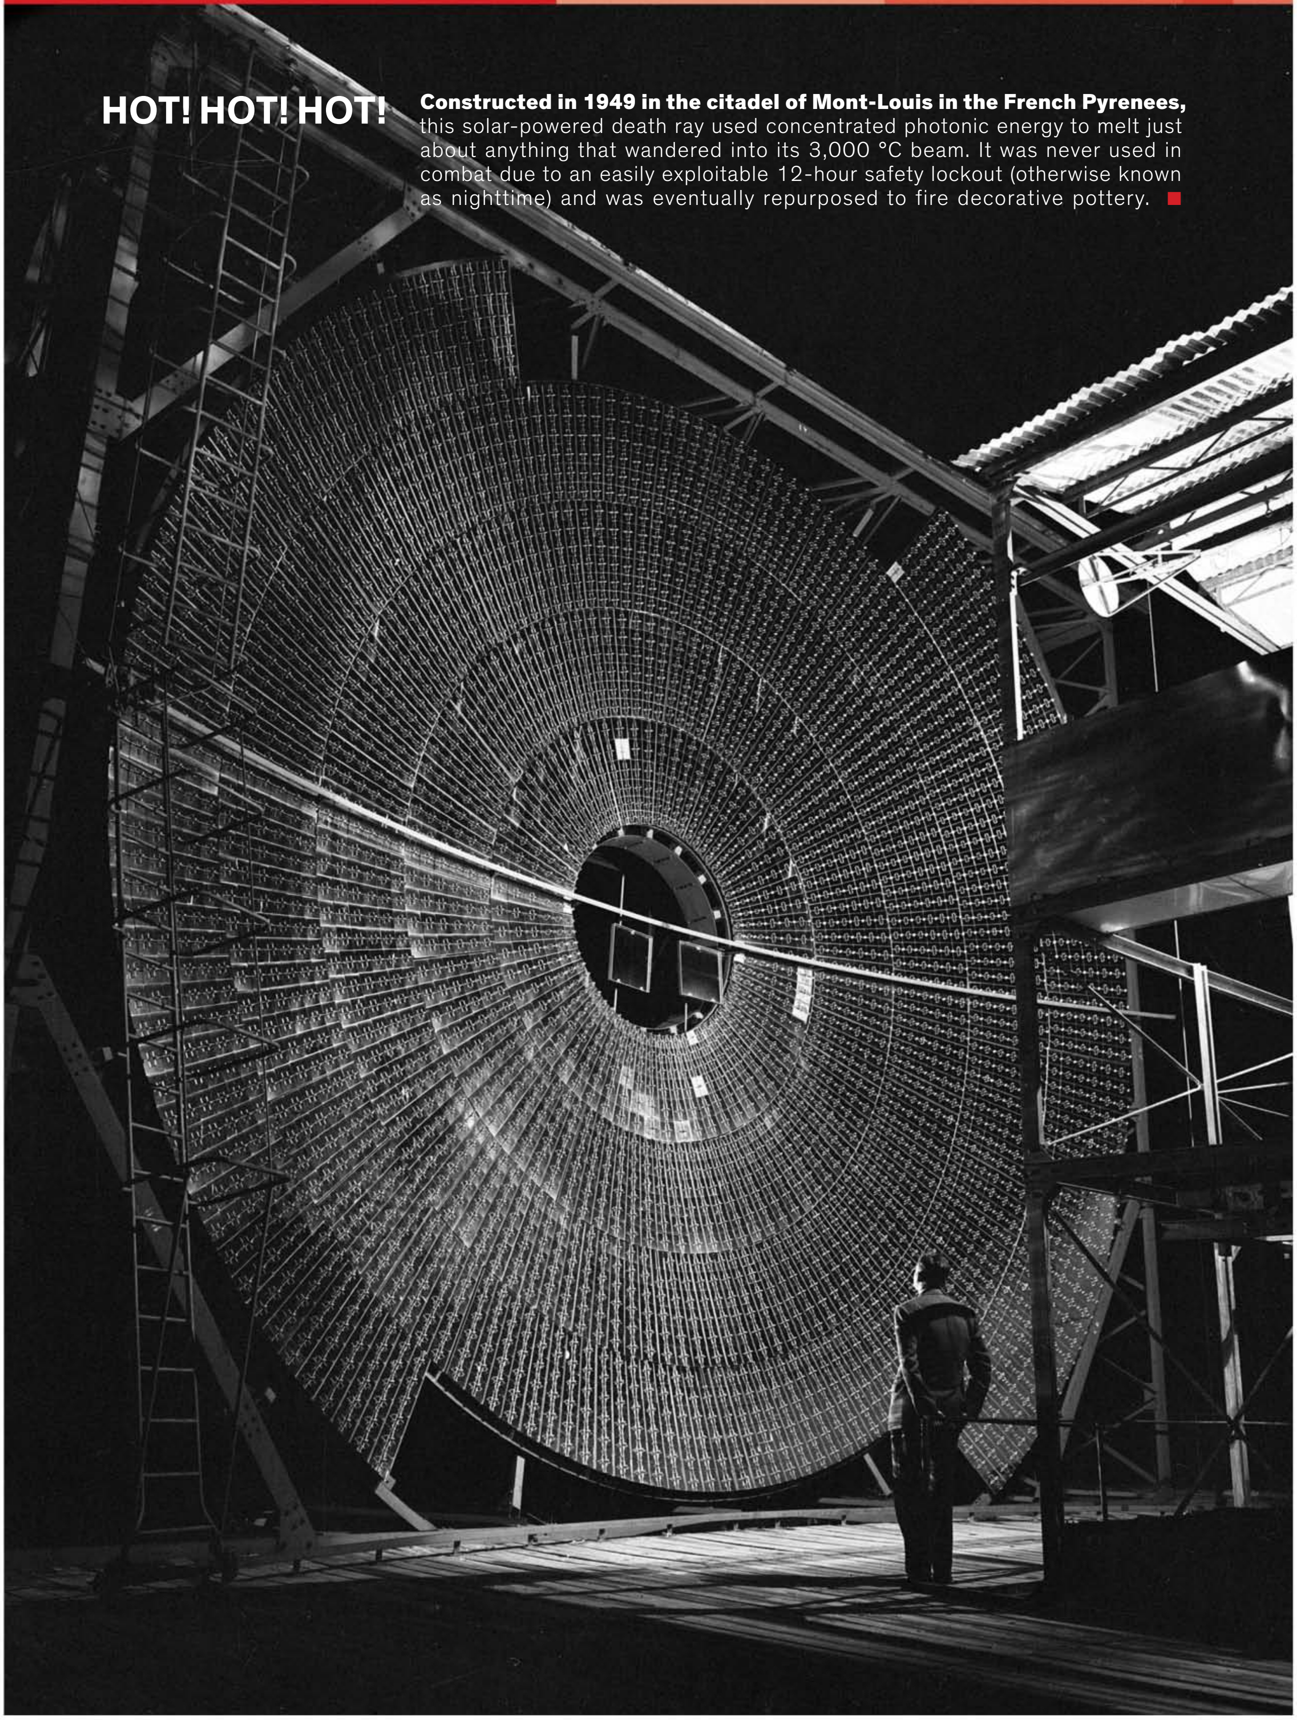

MANUEL LITRAN/PARIS MATCH/GETTY IMAGES