IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 07.15

YOUR NEXT DRONE

Professional-quality aerial videography is now within reach for amateurs P. 32

IEEE

Smarter Embedded Designs,

Faster Deployment

The combination of NI LabVIEW system design software and reconfigurable I/O (RIO) hardware helps small design teams with varied expertise develop demanding embedded applications in less time. Using this graphical system design approach, you can take advantage of the same integrated platform to program embedded processors and FPGAs for faster application development in industries ranging from energy to transportation, manufacturing, and life sciences.

LabVIEW system design software offers ultimate flexibility through FPGA programming, simplifies code reuse, and helps you program the way you think–graphically.

>> Accelerate your productivity at ni.com/embedded-platform

800 453-6202

NATIONAL INSTRUMENTS

IEEE SPECTRUM

On the Cover **Photograph by XFly Systems**

DJI

DEPARTMENTS_07.15

IEEE SPECTRUM

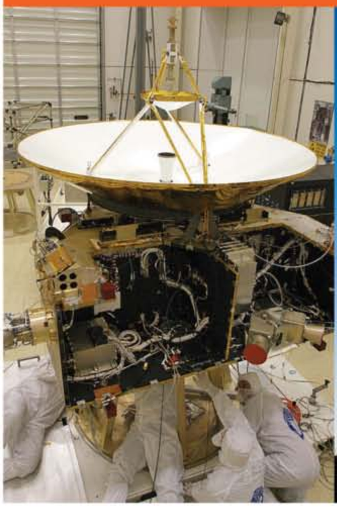

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- 3-D Printing and Entrepreneurship—9 July
- Requirements Management for Project Life Cycle—Available on Demand
- Predictive Analytics UX: A Case Study—Available on Demand
- Engineering for the Internet of Things: From Physical Product to Digital Experience—Available on Demand

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

The Institute

Available on 8 July 2015

GAMES GET SERIOUS Researchers are working on developing a game to train physicians and medical students on how to better communicate with patients.

THE FUTURE OF GAMING Computational intelligence methods are expected to have a big impact on game technology and development, assisting developers and enabling new types of computer games. This topic and more will be explored next month at the IEEE Conference on Computational Intelligence and Games, in Tainan, Taiwan.

CHANGING CAREERS Breaking into a new industry doesn't come without its challenges. Read about several ways to reassure an organization that you're a risk-free hire, even if you have no direct experience with the position you're applying for.

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2015 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 52, issue No. 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_

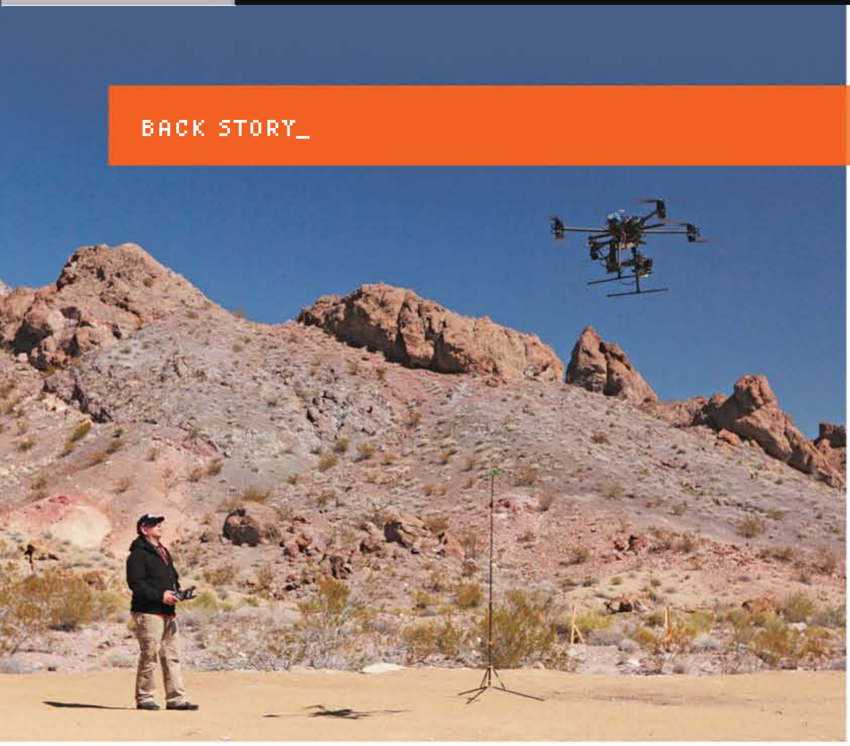

XFLY SYSTEMS

Boom Times for the New Camera Boom

THERE'S NOTHING LIKE BEING THERE when a new technology takes off—a metaphor that's particularly appropriate when you're talking about small drones. The U.S. Federal Aviation Administration was slow to embrace their commercial use, but this past February it proposed rules that will soon allow them to become an integral part of many industries, including construction, agriculture, and—especially—film.

So how might you become a drone cinematographer, either for fun or profit? This month's feature "Lights, Drone...Action" gives you the basic outline. Its author, T.J. Diaz, had long been consulting in IT software for a living and flying radio-controlled helicopters on the side as a hobby. Then in 2008, Canon introduced digital single-lens reflex cameras that made it possible to obtain cinema-quality video on a modest budget.

"I just connected the dots," says Diaz, who immediately went to work mounting such cameras on his helicopters and soon added gyrostabilized gimbals to dampen their motions in the air. "Once I saw the video, I knew we were onto something big."

Diaz now works entirely with drones. As an instructor with the Phoenix-based Unmanned Vehicle University, he trains new drone pilots. He also designs equipment for drone cinematography through his own company, Denver-based Xfly Systems. In the photo above, Diaz is demonstrating a video drone to about 100 camera operators, production assistants, and other professionals attending the 2014 convention of the National Association of Broadcasters in Las Vegas. Such people realize that small drones are becoming as essential as traditional camera cranes and want to be sure they're on top of the new technology. ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, And Now This... is in *IEEE Spectrum*, Vol. 52, no. 7 (INT), July 2015, p. 52, or in *IEEE Spectrum*, Vol. 52, no. 7 (NA), July 2015, p. 64.

07.15

IEEE SPECTRUM

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

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

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

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

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

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

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

SENIOR ASSOCIATE EDITORS
Rachel Courtland, r.courtland@ieee.org
Eliza Strickland, e.strickland@ieee.org
ASSOCIATE EDITOR
Celia Gorman (Multimedia), celia.gorman@ieee.org
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

EXECUTIVE PRODUCER, SPECTRUM RADIO Sharon Basco

ADMINISTRATIVE ASSISTANTS
Ramona L. Foster, r.foster@ieee.org
Nancy T. Hantman, n.hantman@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

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

EDITORIAL ADVISORY BOARD
Susan Hassler, *Chair*; Gerard A. Alphonse, Jason Cong*, Sudhir Dixit, Limor Fried, Robert Hebner, Joseph J. Helble, Chenming Hu*, Grant Jacoby, Norberto Lerendegui, John P. Lewis, Steve Mann, Jacob Østergaard, Umit Ozguner, Thrasos Pappas, John Rogers, Jonathan Rothberg, Umar Saif, Takao Someya, Jeffrey M. Voas, Kazuo Yano, Larry Zhang*, 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

CONTRIBUTORS_

Robert Colburn

Colburn is the research coordinator for the IEEE History Center. In this issue, he recounts the paper punch card's surprising longevity [p. 52]. Colburn learned Fortran programming using punch cards as a student at St. Andrew's School, in Middletown, Del. "I never dropped a compiler or subroutine deck"—which could take up several hundred cards—"but I saw a fellow student do it," he recalls. "The cards went everywhere, and he sat on the floor with this awful expression on his face."

Pedram Khalili

Khalili is an adjunct assistant professor at UCLA. In this issue, he and coauthor Kang L. Wang outline progress in spintronics, which uses the magnetism of electrons to make ultra-energy-efficient memory and logic [p. 24]. Khalili's earlier focus was on fairly large microwave magnetic devices for RF applications. Moving to atomic scales was a bit of an adjustment, he says. Now he's excited to see what big things these tiny devices can accomplish.

Diane Proudfoot

Proudfoot heads the department of philosophy at the University of Canterbury, in New Zealand. Her interest in artificial intelligence led her naturally to Alan Turing. For a talk several years ago, she wanted to replay excerpts of Turing's BBC radio broadcasts, but no recordings existed, so she recorded actors reading from an original script. "This was my short-lived career as a theatrical producer," says Proudfoot. Her study of Turing's radio scripts became the basis for "What Turing Himself Said About the Imitation Game" [p. 38].

Eric E. Sabelman

Sabelman and coauthor Roger Lam work at Kaiser Permanente. In "The Real-Life Dangers of Augmented Reality" [p. 44], they discuss how research on visual and physical impairments is being applied to Google Glass and other AR devices. Sabelman is nearsighted and slightly walleyed (a condition known as strabismus), which theoretically makes it easier to use Google Glass. "But my experiment looking at Glass with my right eye while the left eye looks straight ahead didn't work, because I can't focus on the display," Sabelman says.

Paul Wallich

In Wallich's previous Hands On columns for *IEEE Spectrum,* he described how to build a kid-tracking drone and how to automate your home without rewiring it. For this issue, he created a laser system for his bicycle that warns of cars coming from behind [p. 18]. "The biggest thing I learned was the importance of packaging. Getting the mounting brackets right, fitting the available space in a solid enclosure, even a practical on-off switch turned out to be much more work than the device itself," says Wallich.

07.15

PHOTO-ILLUSTRATIONS BY **Gluekit**

FROM MODEL

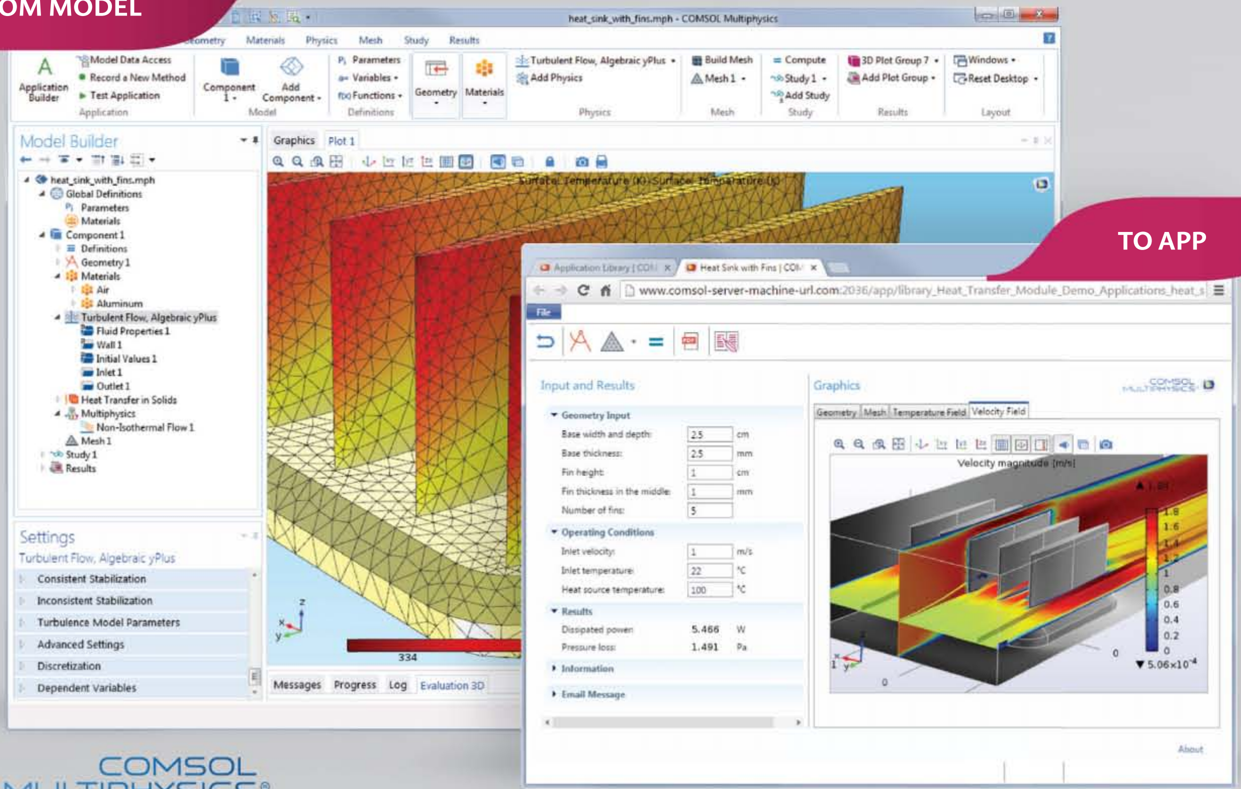

TO APP

COMSOL MULTIPHYSICS®

COMSOL SERVER™

How do you create the best design and share your simulation expertise?

THROUGH POWERFUL COMPUTATIONAL TOOLS.
WITH SIMULATION APPS THAT CAN BE EASILY SHARED.

comsol.com/5.1

PRODUCT SUITE

› COMSOL Multiphysics®
› COMSOL Server™

ELECTRICAL
› AC/DC Module
› RF Module
› Wave Optics Module
› Ray Optics Module
› MEMS Module
› Plasma Module
› Semiconductor Module

MECHANICAL
› Heat Transfer Module
› Structural Mechanics Module
› Nonlinear Structural Materials Module
› Geomechanics Module
› Fatigue Module
› Multibody Dynamics Module
› Acoustics Module

FLUID
› CFD Module
› Mixer Module
› Microfluidics Module
› Subsurface Flow Module
› Pipe Flow Module
› Molecular Flow Module

CHEMICAL
› Chemical Reaction Engineering Module
› Batteries & Fuel Cells Module
› Electrodeposition Module
› Corrosion Module
› Electrochemistry Module

MULTIPURPOSE
› Optimization Module
› Material Library
› Particle Tracing Module

INTERFACING
› LiveLink™ *for* MATLAB®
› LiveLink™ *for* Excel®
› CAD Import Module
› Design Module
› ECAD Import Module
› LiveLink™ *for* SOLIDWORKS®
› LiveLink™ *for* Inventor®
› LiveLink™ *for* AutoCAD®
› LiveLink™ *for* Revit®
› LiveLink™ *for* PTC® Creo® Parametric™
› LiveLink™ *for* PTC® Pro/ENGINEER®
› LiveLink™ *for* Solid Edge®
› File Import *for* CATIA® V5

COMSOL

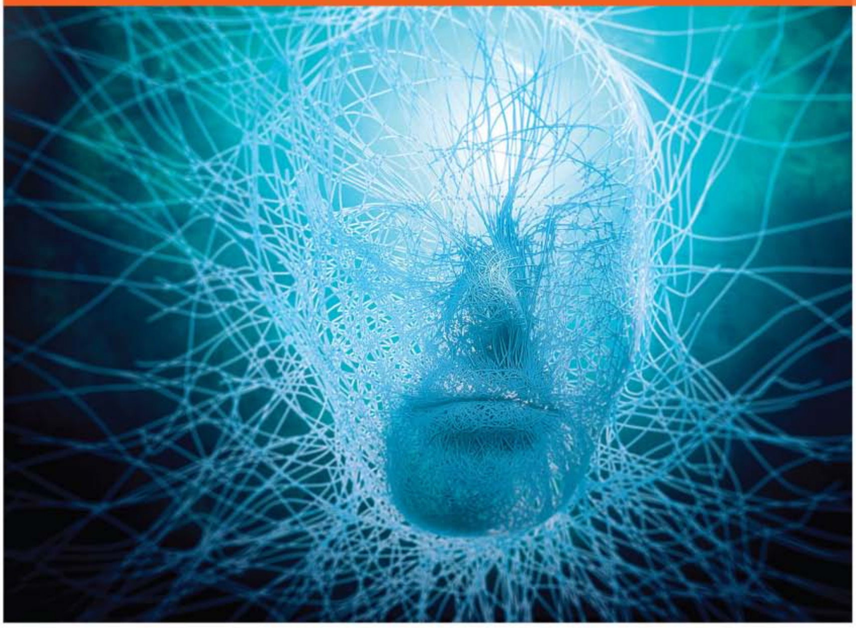

Let's Shape AI Before AI Shapes Us

It's time for a global conversation

ARTIFICIAL INTELLIGENCE IS LIKE A BEAUTIFUL suitor who repeatedly brings his admirer to the edge of consummation only to vanish, dashing hopes and leaving an unrequited lover to wonder what might have been.

Once again, big shots are hearing the siren song of AI and warn of hazards ahead. Visionary entrepreneur Elon Musk thinks that AI could be more dangerous than nuclear weapons. Physicist Stephen Hawking warns that AI "could spell the end of the human race." Even Bill Gates, who usually obsesses on such prosaic tasks as eliminating malaria, advises careful management of digital forms of "super intelligence."

Will today's outsized fears of AI become fodder for tomorrow's computer comedy?

Past AI scares now seem silly in hindsight. In the 1950s, building on excitement over the advent of digital computers, scientists foresaw machines that would instantly translate Russian to English. Rather than 5 years, machine translation took more than 50. "Expert systems" similarly have experienced a long gestation, and even now these programs, built around knowledge gained from human experts, deliver little. Meanwhile, HAL, the Terminator, Ava, and other computer-generated rivals remain the stuff of Hollywood.

In recent years, some claims for AI seem to have been realized. Computers now can literally pick faces out of a crowd and unerringly provide customer service over the phone by simulating a real conversation. Driverless cars and package-delivering drones promise to revolutionize the movement of things and people. Bombs that select their own targets and robots that kill are within reach. Daily life already seems impossible without digital devices that record, alert, and advise their owners on actions and plans.

The beautiful suitor is back, more fetching than ever. Now the human embrace of robots is closer, and probably inevitable.

Because betrayal is central to romance, humans, jealous of AI, worry about the loss of their own supremacy in the world. Digital minds may emulate, then resent, and finally attack humans.

From this dark space, thoughts arise of existential doom. An evil genius might conquer the world with a malevolent AI army. Software agents could knock out the essential systems from which a stricken society could not recover. To Daniel H. Wilson, author of *How to Survive a Robot Uprising,* humans need not wait for the first AI catastrophe in order to install a "steel-reinforced panic room" to which they can escape from disobedient digital servants.

Dark fantasies, however, distract attention from more urgent questions. How will AI affect employment, especially higher-paying work? When will robot writers and artists alter the way humans consume creative content? Who will be held accountable for accidents when humans are no longer in the decision or action loop?

Instead of "wolf" criers of the Musk sort, humans need a serious discussion about new norms and practices that will shape and govern AI. Here are a few suggestions on how to direct the global conversation in ways fruitful rather than fearful:

- **Embrace the precautionary principle:** Bans rarely work; careful testing and technical revisions do. Civil society needs safe spaces to conduct experiments with synthetic intelligences. Governments should encourage controlled tests by private actors on the condition that data and analyses be widely shared.
- **Engineer in equity and diversity:** The disturbing truth is that the digital frontier is dominated by men living in North America, Europe, Japan, and China. Adventures in AI ought to reflect the aspirations of women as much as men and also reflect culture and values from the global South.
- **Help the losers:** The spread of AI will hurt some, mostly by reducing the demand for human labor. Responses should include helping people to use their time in useful and appealing ways. Governments might also consider giving every human, as a legal right, dominion over a set number of robots. Such policies would at least create a measure of equity, because surely wealthy folks will invest in assembling their own armies of bots. —G. PASCAL ZACHARY

G. Pascal Zachary is a professor of practice at Arizona State University and author of *Endless Frontier: Vannevar Bush, Engineer of the American Century* (MIT Press, 1999).

ANDRZEJ WOJCICKI/CORBIS

NEWS

0 of 60 GB – 1.5 years remaining

1.5 YEARS: HOW LONG IT WILL TAKE NEW HORIZONS TO TRANSMIT 60 GIGABITS OF PLUTO DATA TO EARTH

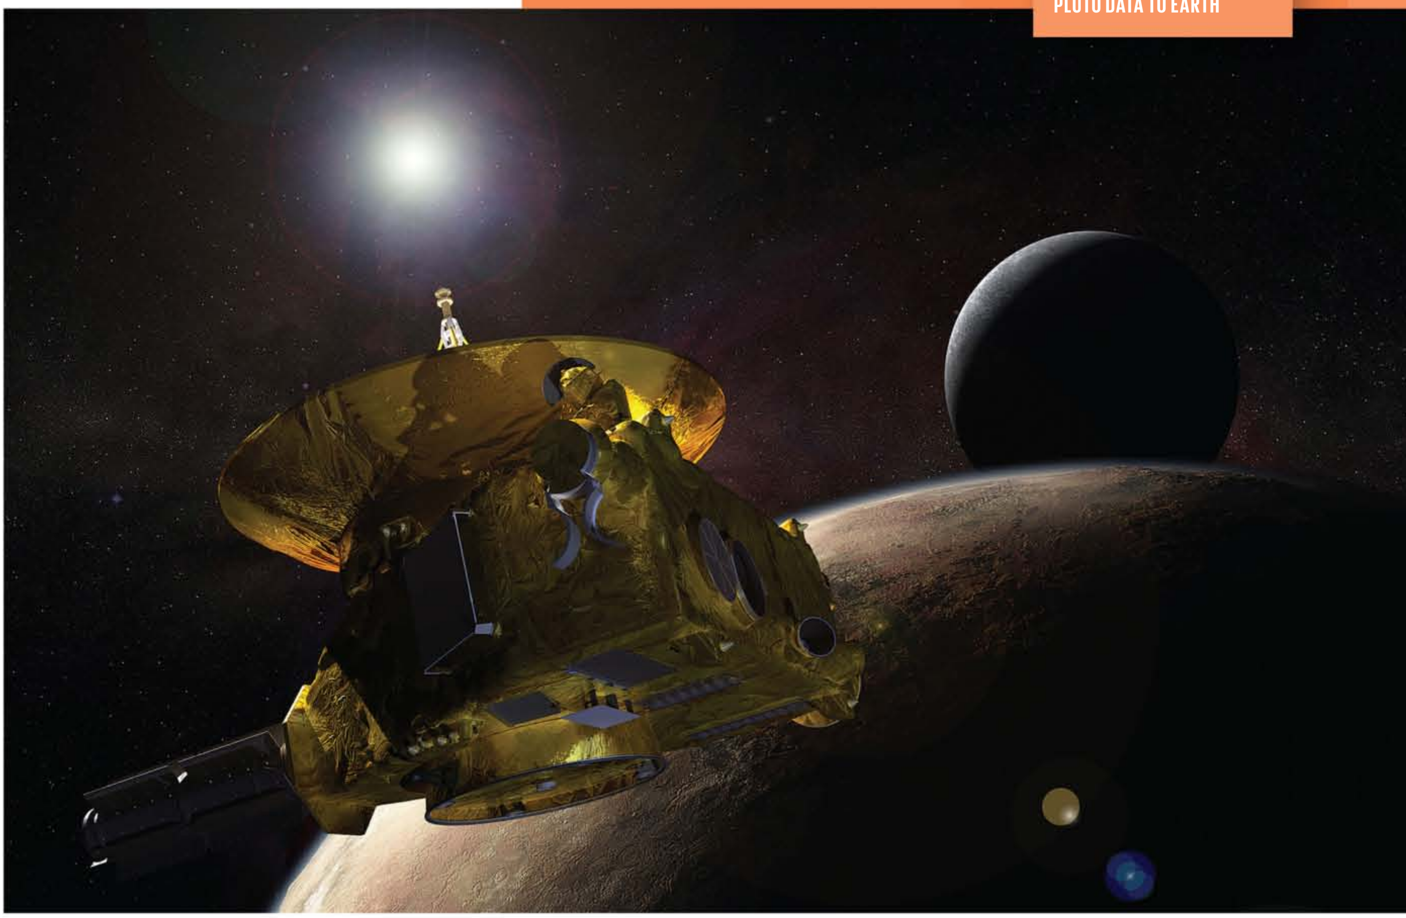

JOHNS HOPKINS UNIVERSITY APPLIED PHYSICS LABORATORY/SOUTHWEST RESEARCH INSTITUTE

LONG WAY FROM HOME: New Horizons will use radio waves from Earth to analyze Pluto's atmosphere.

CLOSE ENCOUNTER OF THE NINTH KIND

After almost a decade in transit, New Horizons will have little time to measure Pluto's atmosphere

In a few short weeks, the NASA New Horizons spacecraft will finish a long, lonely trek through the outer solar system and zip past Pluto. It will gather the first close-up views of the erstwhile ninth planet and examine, among other things, Pluto's thin atmosphere by measuring sunlight and radio waves that pass through it.

The 14 July encounter has been a long time in the making. It has taken New Horizons more than nine years to make its way from Earth to Pluto, which was demoted and reclassified by the International Astronomical Union as a "dwarf planet" while the spacecraft was en route.

The US $700 million mission has fared well so far, hibernating for long stretches along the way. But much of New Horizons' success will hinge on a few short hours around its closest approach, when the probe will thread between »

Pluto and the orbit of its innermost moon, Charon, pass some 12,500 kilometers above Pluto's surface, and cross into the shadow cast by the dwarf planet.

"It's getting busy–exhausting as well," says mission system engineer Chris Hersman of the Johns Hopkins University Applied Physics Laboratory (APL), which built the craft and manages the mission. Key personnel will have cots in their offices to stay close to the mission as it reaches Pluto, he says.

But with a 4.5-hour communications delay in each direction, New Horizons will be on its own during the most crucial hours. Seven days out, Hersman says, the spacecraft will enter a special "encounter" mode designed by APL to prioritize data collection. If there is a problem, instead of switching into safe mode the spacecraft will shift over to redundant systems to keep data gathering going. Once the spacecraft enters this state, Hersman says, its actions will be set, although the team will be able to shift the spacecraft's data-taking sequence to earlier or later times, as the exact distance between New Horizons and Pluto becomes clearer.

Some of the spacecraft's most time-sensitive observations occur during a series of occultations: As it passes behind Pluto, its view of the sun and then the Earth are blocked. Two of the spacecraft's seven science instruments–an ultraviolet imaging spectrograph called Alice and the radio signal-processing Radio Science Experiment (REX)–will observe the limb, or outer edge, of Pluto as the spacecraft enters and leaves its shadow.

The measurements are expected to yield deeper insight into Pluto's tenuous atmosphere, whose pressure at the surface is less than one-thousandth of Earth's. Alice will pick apart the sunlight that shines through Pluto's upper atmosphere, looking for indicators of absorption that could reveal more about its composition.

REX will be the only instrument used for lower altitudes, Hersman says, where the atmosphere, although still just a wisp of molecules, will be too thick to permit ultraviolet radiation. The instrument, which consists of two circuit boards integrated into the New Horizons' telecommunications system, will look for radio signals transmitted hours before by antennas at two NASA Deep Space Network tracking stations: the Goldstone complex in California and another outside Canberra, Australia.

Pluto's atmosphere will refract the waves, slowing them down and altering their phase. If all goes well, the REX team will be able to use this information to measure the temperature and pressure all the way to the surface of the dwarf planet. When New Horizons is in the shadow of Pluto, where it will spend about 10 minutes, REX should also be able to pick up the radio waves given off by the heat of the dwarf planet's surface, allowing scientists to measure surface temperature. As the probe emerges on the other side, signals from Earth will help better pin down Pluto's size. REX and Alice will perform the same sorts of occultation measurements on Charon about an hour later, as New Horizons passes into the shadow cast by that moon.

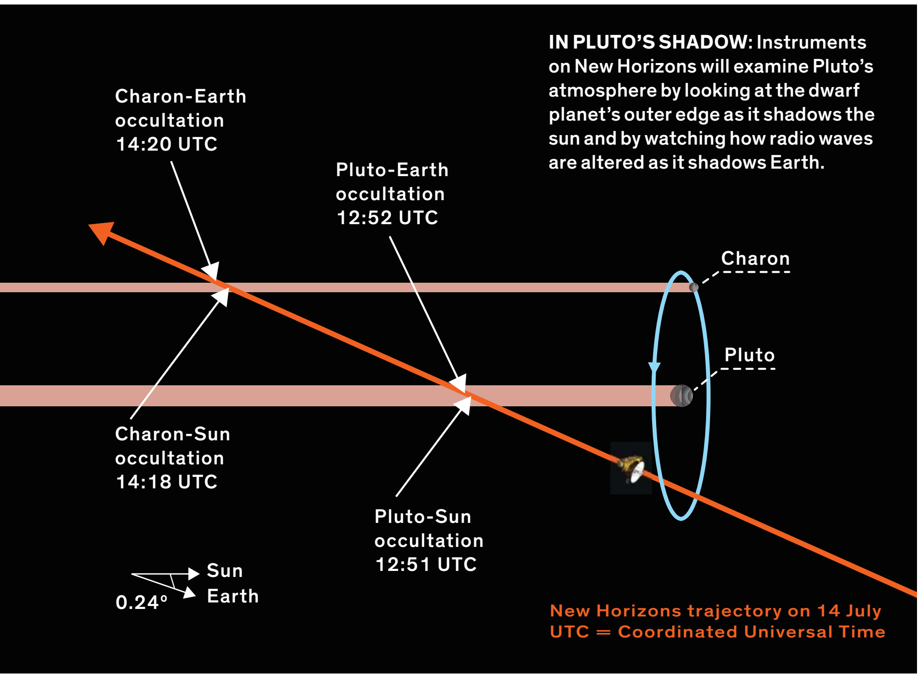

IN PLUTO'S SHADOW: Instruments on New Horizons will examine Pluto's atmosphere by looking at the dwarf planet's outer edge as it shadows the sun and by watching how radio waves are altered as it shadows Earth.

All told, REX's experiment will be over in a little more than 2 hours, says Ivan Linscott of Stanford University, one of the instrument's principal investigators. "It's a one-time measurement," Linscott says. "I think of it like a Thanksgiving dinner. You [do] all this preparation, you invite all the guests, you have the meal, and then suddenly, it seems, it's over."

But New Horizons' job will continue even after it leaves Pluto far behind. The spacecraft can transmit on average only about 2,000 bits per second of data when it is so far away, Hersman says. With the available radio telescope time on Earth, he says, it will take about 1.5 years for the spacecraft to successfully transmit the 60 gigabits of unique data it collected during the encounter.

During this time, New Horizons may be well on its way to a new target. Later this year, the team hopes to fire the spacecraft's engines in order to direct it to another target in the Kuiper Belt, the ring of icy objects beyond Neptune that Pluto calls home. –RACHEL COURTLAND

INFOGRAPHIC BY Erik Vrielink

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

CHEERS! Erik Sorto uses a brain-controlled robotic arm to take a drink. "It gives me great pleasure to be part of the solution for improving paralyzed patients' lives," he says.

SPENCER KELLIS & CHRISTIAN KLAES/CALTECH

TELEKINESIS MADE SIMPLE

A brain implant reads a paraplegic man's intentions to move a robotic arm

Erik Sorto had been paralyzed for 10 years when he volunteered for a bold neural engineering experiment: He would receive a brain implant and try to use the signals it recorded to control a robotic arm. Erik had no qualms about signing up for brain surgery, but his mother wasn't happy about it. "She was just being a mom," Sorto says with a smile. "She was like, 'Your brain is the only part of your body that works just fine. Why would you mess with that?' "

Instead of reducing his capacities, the surgery gave Sorto superhuman abilities. In the experiments, Sorto simply imagines reaching out to grab an object and the robotic arm carries out his commands. While a handful of paralyzed people have previously used brain-computer interfaces (BCIs) to control robotic limbs, those subjects' implants recorded signals from the primary motor cortex, which is linked directly to the spinal cord and muscles. Sorto's implant was the first to record instead from the posterior parietal cortex, a brain region involved in planning movements.

Lead researcher Richard Andersen, a neuroscience professor at Caltech, says that tapping into a subject's intentions may provide more intuitive control of the robotic limb, and with less lag time. When the scientists asked Sorto to focus on the object he wished to grab, the system was able to identify that goal very quickly, says Andersen, usually within 200 milliseconds. The BCI then combined that high-level signal with smart robotics functions such as object recognition and simultaneous localization and mapping, or SLAM. "Erik worked with the robotic intelligence to manage the limb's movement, rather than try to control all the kinematic details," Andersen says. He and his colleagues revealed their novel approach recently in an article in *Science*.

To prepare for Sorto's surgery in April 2013, the researchers first used functional magnetic resonance imaging to identify two precise regions of his parietal cortex that activated when he imagined reaching and grasping motions. The surgeons implanted two tiny microelectrode arrays, each with 96 electrodes that could record the electrical activity of single neurons. The grids linked to two metal "pedestals" that jutted out of Sorto's skull. Within one month of surgery, Sorto was ready to get to work. The researchers connected cables to the pedestals, bringing the neural signals to a computer that analyzed them and sent commands to the robot arm. In the first experiment, a scientist made a simple twisting motion with his own wrist, and Sorto imagined himself imitating the gesture. And that's just what the robot arm did. "It was pretty much effortless," Sorto says.

From the population of neurons tapped by those electrodes, the researchers could distinguish cells whose activity coded for the location that Sorto wanted to reach, movement trajectories, and

NEWS

particular types of movement. Each day's experiments began with system calibration, because the electrodes shifted slightly within Sorto's head and picked up different neurons. "Our decoding algorithms took that into account," Andersen says. If a given electrode was no longer contributing useful information to the decoding of a goal location, for example, the algorithm would ignore its signal and substitute other inputs. Andersen thinks such adaptive algorithms may enable a wide range of BCIs to record reliably over time.

Most of the prior studies in which paralyzed people used implanted BCIs were conducted by John Donoghue, director of Brown University's Institute for Brain Science and a pioneer in the use of implants in the motor cortex. Donoghue says that the new research advances our understanding of the parietal cortex's role in generating movements and proves that it can provide a useful control signal. However, he's not convinced it's inherently a better signal than that provided by the motor cortex. "You get pretty good control from all these places, but it's not clear how to get really good control," he says.

While Andersen has suggested that combining the signals from the parietal and motor cortices might yield clearer commands for a robotic limb, Donoghue says that combination wouldn't necessarily add value. "When you move your hand, you're probably using 80 percent of your brain," Donoghue says. The parietal and motor cortices provide similar signals, Donoghue says, so some other brain region may hold the key information that would result in truly natural movement. "It's one of the great mysteries," Donoghue says. "There's something we still haven't detected about the way the brain does movement."

Sorto, meanwhile, has continued to work with what he has: More than two years after his surgery, his electrodes are still functioning and his enthusiasm is undimmed. In his second year of experiments, he mastered precise reaching and grasping movements by dint of an unusual exercise. "I played over 6,700 rounds of rock-paper-scissors," Sorto says with agonized emphasis. "I want everybody to know that I worked hard." It paid off, though, at the end of that second year, when he reached a long-held goal: He used the robot arm to lift a bottle of Modelo beer to his mouth and take a long swig.

—ELIZA STRICKLAND

Sorto played over 6,700 rounds of rock-paper-scissors while training to use a new kind of brain-machine interface

INNOVATION AMID A RAUCOUS ROOFTOP SOLAR SQUABBLE

Utility grid fees in Arizona have installers redesigning systems to keep residential solar profitable

The explosive expansion of U.S. residential rooftop solar installations—which increased by 51 percent last year—threatens utilities' traditional role as electricity suppliers. The resulting backlash has been most intense in sun-bleached Arizona, which is second only to California in installed solar capacity. The battle has the potential to blunt the state's solar growth, but there are signs that it could also inspire grid-friendly technologies that expand solar's role.

The utility backlash got real in recent months as Arizona utilities levied or proposed new fees for customers installing rooftop solar systems. Tempe-based Salt River Project (SRP), which serves much of greater Phoenix, has seen applications to connect solar systems drop 96 percent since it announced a new rate structure in December 2014.

SRP now exacts a monthly "demand charge" based on the maximum level of grid power that solar customers consume during its 1 to 8 p.m. peak demand period. Altogether, SRP's fee changes add about US $50 per month to solar users' bills, wiping out the economic gains of producing rooftop solar power.

Southern Arizona will see the slowdown next if state regulators approve a request by investor-owned Tucson Electric Power to pay wholesale rates for surplus rooftop generation rather than retail rates, which will add an estimated $22 to bills. Meanwhile, investor-owned Arizona Public Service Co. (APS), the state's largest utility, is seeking permission for a solar user fee of about $21 per month.

The rooftop solar boom began in Arizona six years ago and is still growing in many states, thanks

SANDY HUFFAKER/CORBIS

SUNSHINE AND STRIFE: It will take smarter rooftop solar installations to get around Arizona utility rules that have halted growth in the state.

to "net metering," whereby surplus power exported by solar panels earns their owners' retail credit against the grid power they consume at night and in the winter. Arizona's utilities, however, say net metering is not a sustainable plan because it means solar customers rely heavily on the grid but pay little or nothing to maintain it.

The utilities' new fees have spurred high-profile negative media campaigns by solar advocates, backed by research showing that rooftop solar is a net plus for the grid. The solar leasing firm SolarCity, based in San Mateo, Calif., is suing SRP, alleging that its solar fees are discriminatory and anticompetitive.

Behind the scenes, however, the battle is also advancing technological solutions that seek to align rooftop solar with new rate structures and the grid's needs. "It's making innovation happen a lot faster," says Russ Patzer, CEO for Chandler, Ariz.-based solar provider Sun Valley Solar Solutions.

Many solar providers are designing energy storage into rooftop systems. Joy Seitz, president and CEO for Scottsdale-based solar supplier American Solar and Roofing, says batteries make solar viable again under SRP's demand charges by minimizing the use of grid power during peak hours. Battery power spans a 6 to 8 p.m. gap, when solar output tends to wane but many Arizonans' air conditioners are still running hard.

Solar output will, however, decline as customers forego some panels to pay for pricey batteries. With one-third fewer panels, Seitz estimates, such systems will cost one-third more at $30,000 and take 14 years instead of 9 to pay for themselves.

Patzer plans to offer solar-plus-storage systems but is also preparing off-grid options that evade the utilities' solar fees altogether. For example, Patzer predicts growing demand for stand-alone air conditioners running on direct current from their own dedicated solar panels. Isolating the solar panels from the grid should protect buyers from the utilities' solar fees.

For their part, Arizona's investor-owned utilities are innovating by jumping into the residential solar market. By early this month, APS expects to install the first of 1,500 planned systems on customer rooftops to serve as a unique test bed for remote control of distributed solar. "This is state of the industry. There isn't another utility doing this at this scale," according to Marc Romito, APS's manager for renewable energy.

The arrays will connect to six distribution feeders in central Phoenix via advanced inverters akin to those recently specified for use in California.

Romito foresees operators using a secure communications link to connect the inverters to APS's grid management system and using them for a wide range of actions, including suppressing solar output on lines that are oversupplied and releasing extra power from associated batteries where the grid is straining to meet demand. Ultimately, he anticipates that remote control of distributed resources will enable APS to defer upgrades to congested power lines and substations.

Kris Mayes, acting director of Arizona State University's Utilities of the Future project, says Arizona's utilities and regulators need more information to strike the right balance. What is missing, she says, is an independent study documenting what distributed solar is costing the utilities, and what benefits it offers their grids. As Mayes said on a state radio talk show: "Utilities haven't proven the costs, and we haven't analyzed the benefits. It's like doing a pros and cons list where you only do the cons." –PETER FAIRLEY

A CHEAPER WAY FOR ROBOCARS TO AVOID PEDESTRIANS

Artificial-intelligence video analysis could soon be fast enough to replace expensive lidar

Google's self-driving cars roam the sunny streets of Mountain View, Calif., but much of the technology that powers them has never seen the light of day. In May, attendees at the IEEE International Conference on Robotics and Automation in Seattle got a rare glimpse into what the tech giant is working on.

Anelia Angelova, a research scientist at Google, presented a new pedestrian detection system that works on video images alone. Recognizing, tracking, and avoiding human beings is the critical function of any driverless car, so Google's vehicles are duly festooned with lidar (laser ranging), radar, sonar, and video sensors to ensure that they identify all pedestrians within hundreds of meters.

But that battery of sensors is expensive, in particular the spinning lidar unit on the roof, which can cost US $8,000 (or more for multiple units). If autonomous vehicles could reliably locate humans with cheap video cameras alone, it would lower their cost and, observers hope, usher in an era of robotic crash-free motoring all the sooner. But video cameras have their issues. "Visual information gives you a wider view [than radars] but is slower to process," Angelova says.

Or at least it used to be. Many leading video analysis systems use deep neural networks, machine learning algorithms that can be trained to classify images with extreme accuracy. Neural networks are made up of software that's inspired by the human brain and composed of virtual neurons. In real brains, neurons are arranged in layers. A deep neural network copies this arrangement. For image recognition, the bottom layer of artificial neurons learns features of the pixels of an image. The next layer learns combinations of those features, and so on up the stack, with more sophisticated correlations gradually emerging. The top layer makes a guess about what the system is looking at.

WELL SPOTTED! A new kind of neural network can identify pedestrians using an ordinary camera instead of expensive lidar.

Modern deep networks can outperform humans at recognizing faces, with accuracy rates of 99.5 percent. But traditional deep networks are also very slow, according to Angelova, dividing each street image into 100,000 or more tiny patches and then analyzing each in turn. This can take seconds or even minutes per frame, making them useless for navigating city streets. A car using such a system might run over a pedestrian long before it identifies him.

Angelova's new, high-speed pedestrian detector has three separate stages. The first is a deep network, but one that slices up the image into a grid of just a few dozen patches rather than tens of thousands. This network is trained to do multiple detections simultaneously at multiple locations, picking out what it thinks are pedestrians. The second stage is another network that refines that result, and the third is a traditional deep network to deliver the final word on whether the car is seeing a person or, say, a mailbox.

However, because that slow, accurate network analyzes only a small portion of the image where pedestrians are likely to be, the whole process runs much faster–between 60 and 100 times as fast as the best previous networks, says Angelova. Running on graphics processors similar to those in Google's self-driving cars and fed street images from a standard database, the system was trained in about a day. It could then accurately identify pedestrians in around 0.25 second.

"That's still not the 0.07 second needed for real-time use," admits Angelova. Self-driving cars need to know almost instantly whether they are facing pedestrians or not in order to safely take evasive action. "But it means they could be complementary in case other sensors fail," she says.

Deep networks face some stiff competition. Methods that use cascades of simple visual detectors are also improving quickly, with errors dropping nearly 50 percent over the last year. "The fastest 'cross talk' cascades run faster than the method presented here, with some slight disadvantage in accuracy," says Dan Levi, a senior researcher in the Smart Sensing and Vision Group at the General Motors Research Advanced Technical Center, in Israel. "But the expectation in the vision community is that clever implementations of deep networks, or the use of dedicated hardware, will eventually lead to deep networks replacing traditional methods," he says.

By the time self-driving cars are available to the general public, their distinctive spinning lidars may have disappeared altogether. –MARK HARRIS

ANELIA ANGELOVA/GOOGLE

Industrial Ethernet Switches

Switch to the future: SCALANCE XM-400

Secure and reliable data transmission is indispensable for industrial communication networks – not only today, but in the future as well. That makes it all the more important to be equipped with technology capable of keeping up with increasing requirements – as it is then much simpler to adapt or expand a system.

Industrial Ethernet Switches SCALANCE XM-400 are our innovative answer to the challenges of tomorrow – their modular design, fast mobile diagnostics for smartphones/ tablets via NFC and existing WLAN, and upgradeable Layer 3 functionality using KEY-PLUG will ensure that your automation network is powerfully and flexibly equipped for the future.

siemens.com/x-400

PHOTOGRAPH BY **China Foto Press/Getty Images**

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

HOME ON THE HIGH SEAS

THE SEARCH FOR OIL and gas deposits has moved ever farther from the world's coasts, making it impractical to shuttle workers to and from offshore platforms each day. This has created a demand for long-term accommodation vessels such as the *Arendal Spirit*, shown here in March before making the trip from Zhoushan, China, where it was built, to the waters off Brazil, where it will be leased by Petrobras on a three-year charter contract. The "floatel" has 500 beds in 248 cabins with en suite bathrooms. Other onboard amenities include office areas with TV, telephone, and Internet connections, plus lounge areas, a coffee shop, and a movie theater.

THE BIG PICTURE NEWS

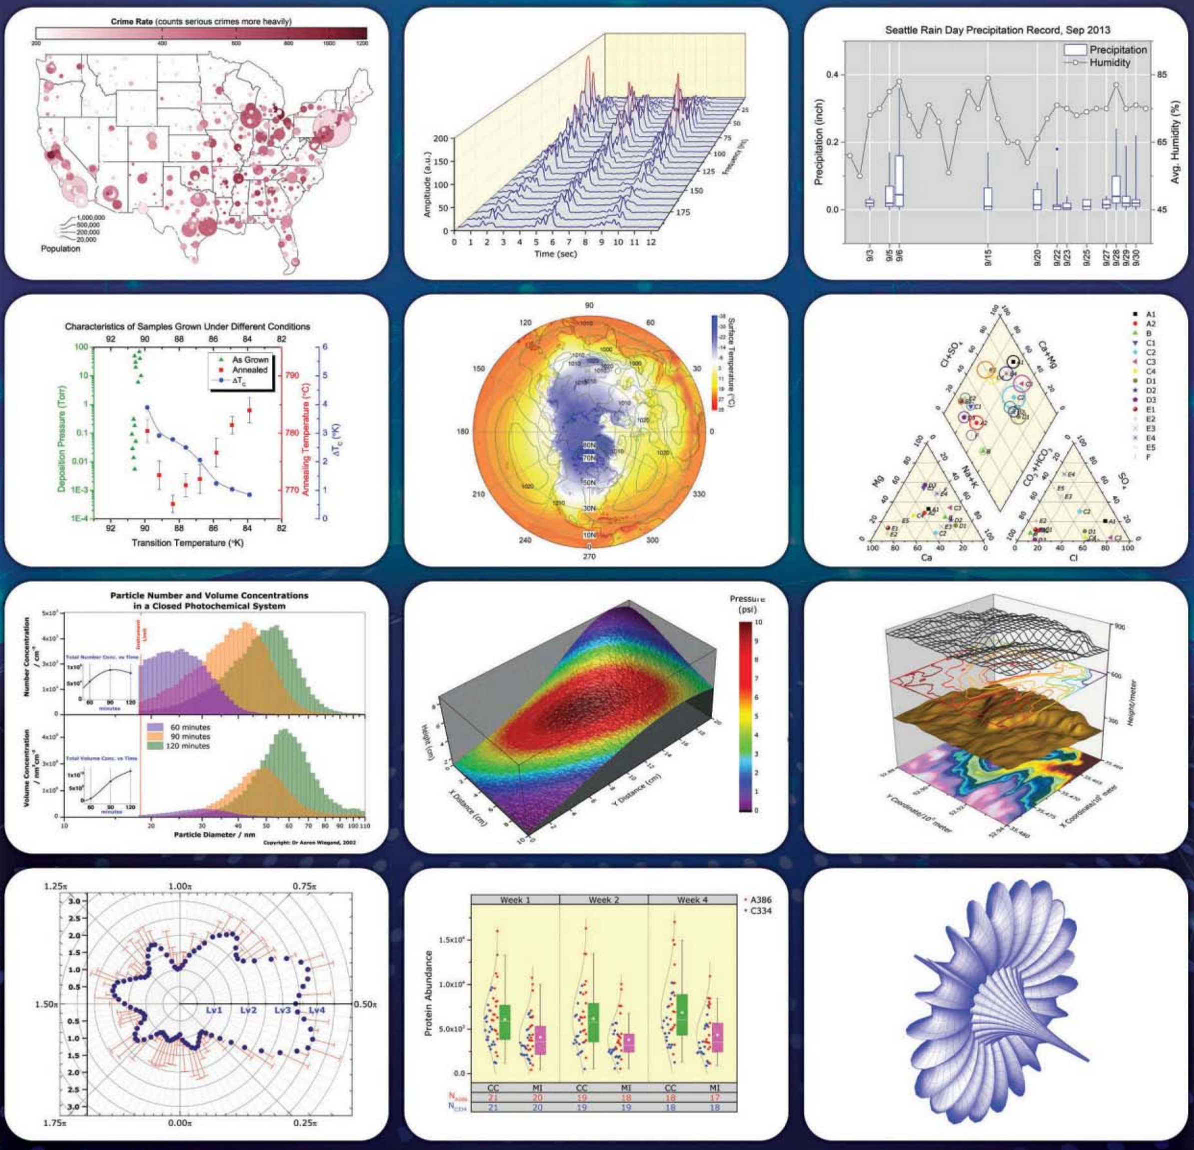
Crime Rate (counts serious crimes more heavily)
Population
Seattle Rain Day Precipitation Record, Sep 2013
Precipitation
Humidity
Characteristics of Samples Grown Under Different Conditions
Transition Temperature (°K)
Particle Number and Volume Concentrations in a Closed Photochemical System
Particle Diameter / nm
Protein Abundance
Week 1
Week 2
Week 4

RESOURCES

306,100

THE NUMBER OF EEs EMPLOYED IN THE UNITED STATES IN 2012, ACCORDING TO THE MOST RECENT DATA FROM THE U.S. BUREAU OF LABOR STATISTICS

ROBERT CHURCHILL/GETTY IMAGES

RESOURCES_CAREERS

WHERE THE JOBS ARE: 2015

IEEE SPECTRUM'S ANNUAL ROUNDUP OF THE HOT SECTORS IN ENGINEERING

This should be a good year for those entering the workforce from college, or midlevel engineers looking to change jobs in the United States. As the U.S. economy continues to grow, American employers plan to hire 9.6 percent more college graduates this year than they did last year, according to the National Association of Colleges and Employers (NACE) in Bethlehem, Pa. And according to tech recruiting firm Randstad Engineering, in April there were over 130,000 engineering job openings to fill across the country, an average of 17 openings per candidate. ● The top starting salaries for new electrical engineers are in pharmaceuticals manufacturing and computer/electronics manufacturing, averaging US $69,958 and $67,227 respectively, according to NACE. For computer engineers, the utilities and computer/electronics manufacturing industries offer top starting salaries of $69,250 and $66,938. The average starting salary for EEs this year is $64,081. ● There are increasingly more and better jobs for graduating electrical engineers in the areas of software and IT. Most of the on-campus hiring for EEs has been in those areas, says Holly Evarts, director of communications and media relations at Columbia University's engineering school. ● For those already in the workforce, many companies need engineers with 5 to 10 years' experience, says David Findley, senior vice president for Randstad. Because of the recession five years ago, fewer engineers entered the workplace and began gaining relevant experience. As a result, there is extreme competition for midlevel engineers, he says. ● The most recent data from salary analysis firm PayScale shows that the highest-paying jobs for EEs in 2015 are in IT and computer networking. Staffing firm Robert Half identifies mobile applications development, big data, and wireless networking as providing some of the highest-paying careers in the United States and Canada with the largest starting salary increases this ▸

year over last year. The firm predicts that the supply of highly skilled tech professionals in these areas will remain below demand for the foreseeable future.

Electrical engineers are also in high demand in aerospace and defense, engineering services, and consumer goods companies, says Greg Margin, Randstad's vice president of embedded engineering. U.S. companies looking to hire the most electrical engineers include Black & Veatch, Burns & McDonnell, L-3 Communications, Raytheon, and Rockwell Collins. EEs should also be able to use their talent in the areas of renewables and energy efficiency, while important areas of expertise in manufacturing include automation and 3-D printing, according to employment company Kelly Services.

Outside the United States, in India, IT giant Tata Consultancy Services says it plans to hire 55,000 people in 2015, while Infosys plans to hire 25,000 to 30,000 people in the next three years. But entry-level salaries at these top IT firms have been stagnant for the past seven years due to a glut of fresh graduates.

In the United Kingdom, the news is mixed. The number of engineering jobs there is growing at its fastest rate since the recession, but salaries seem to have leveled out, according to a study by accounting firm Nixon Williams. The study found that the number of engineering jobs increased last year by 17.7 percent to 159,000, which is 15,000 shy of 2009's number. However, the median annual pay this year is only 2 percent higher than it was in 2009.

The nonprofit EngineeringUK released a report in January that claimed there was an annual shortfall of 55,000 engineers in the U.K. Germany, with its numerous high-tech companies, has a large number of vacancies for EEs: over 12,000 at the end of 2014, according to the Association of German Engineers.

To fill the skill gap, Germany, the U.K., and other European countries are starting to welcome more foreign engineers as immigrants. —**PRACHI PATEL**

AN EARLY-WARNING SYSTEM FOR YOUR BIKE

LOW-COST LIDAR CAN DETECT APPROACHING CARS

PAUL WALLICH

IN THE SUMMER MONTHS, I REALLY ENJOY RIDING a bike along some of the flatter roads and trails around my small Vermont town. But sharing some of those roads with cars can be a dicey proposition. There's one dirt road in particular, where no matter how often I scan my mirror or sneak a look over my shoulder, it seems that sooner or later I get surprised by a car zipping by. ● I've often imagined fitting an ultrasonic sensor to the bike to warn me of incoming traffic, but most of those have a range of less than 10 meters, which is worse than even my mediocre human senses. A lidar—similar to a radar system, but with laser pulses instead of radio waves—would have the longer range I need, but lidars have traditionally been too expensive for my biking needs. ● This changed last year when startup PulsedLight, of Bend, Ore., announced Lidar-lite, a crowdfunded lidar detector less than half the size of a deck of cards, with a range of up to 40 meters and an accuracy of plus or minus 2.5 centimeters, all for US $89. ● When the package arrived, my first step was to make a set of labels for the lidar's four signal wires, as for some reason PulsedLight ended up with a color-coding scheme consisting of black, black, black, and black. I plugged the little thing into an Arduino microcontroller, ran the example software provided online, and darned if I didn't see a stream of distance measurements running back to my PC over the Arduino USB connection. Pretty quickly I moved on to attaching a board with four alphanumeric LED displays to the Arduino. Once I'd done the programming, I had myself a handy little pocket range finder. (FYI, the cavernous basement of the nearest maker space to me is about 45 meters long.)

The next step was to build something that could fit on the back of a bicycle and wouldn't be shaken apart. That was a little more difficult, but luckily, in addition to putting manuals, data sheets, and sample Arduino code on the Web, the folks at PulsedLight also make available a 3-D model of the Lidar-lite's case. An hour or so with Openscad, my 3-D design tool of choice, plus a few hours at the 3-D printer, yielded an outer shell. Into the shell I was able to fit the Lidar-lite, the microcontroller, and a little strip of Neopixel RGB LEDs that act as a warning display, with room for a uniform 6-millimeter layer of padding. (Of course, my first shell design didn't fit on the back of the bike, but a few more iterations fixed that.) The bike mount built into the shell is tilted to point a few degrees upward and to the left, so that as the lidar beam diverges with distance it will be more likely to impinge upon approaching cars rather than just road surface and shrubbery.

One thing I hadn't really considered when I first started designing the shell was power consumption issues. A USB port was fine for powering the breadboard system, but of course the real version needed to run off a battery. I have a bunch of nifty DC-DC converters that will make 5 volts from any power source that can provide 1 to 4 V. I also found a really nice Openscad library that let me incorporate a holder for two AA batteries into my shell design. But, in addition to the 27 milliamperes my microcontroller draws while running, the Lidar-lite module draws an average of 100 mA while in use (and this can spike to 200 mA). Plus, the warning LEDs can draw up to 500 mA. Unfortunately my DC-DC converters have a hard current limit of 200 mA.

So, after sketching out a few designs with two or three DC converters in parallel, I ditched the elegant, power-efficient AA battery approach for one using a two-cell rechargeable lithium battery that provides 7.4 V. A 7805-style linear voltage regulator gets the voltage down to 5 V, by throwing away roughly a third of the energy that goes through it. Conveniently enough though, the new lithium battery is almost exactly the same volume as a couple of AAs.

Does it work? I haven't been run over yet. But time will tell. The first time I saw the LED blips that indicated a sedan approaching from 35 meters behind me, I was very nearly too fascinated by the display to remember to get out of the way. I also need to fabricate a new case so that I can adjust the lidar angle (the current aim point misses cars that pull slightly to the left to pass me, which is probably not a safety problem but still unnerving). And while I'm at it, something that I don't have to disassemble completely when I want to recharge the battery would be nice.

The range is a little shorter than I would like, especially around blind curves, and my code has a tendency toward false alarms that I could reduce at the cost of slightly slower response. But those shortcomings are probably a good thing, since I should be looking behind me on a regular basis anyway, rather than relying on a hacked-together gadget as a first line of safety. **—PAUL WALLICH**

FRICKIN' LASERS!

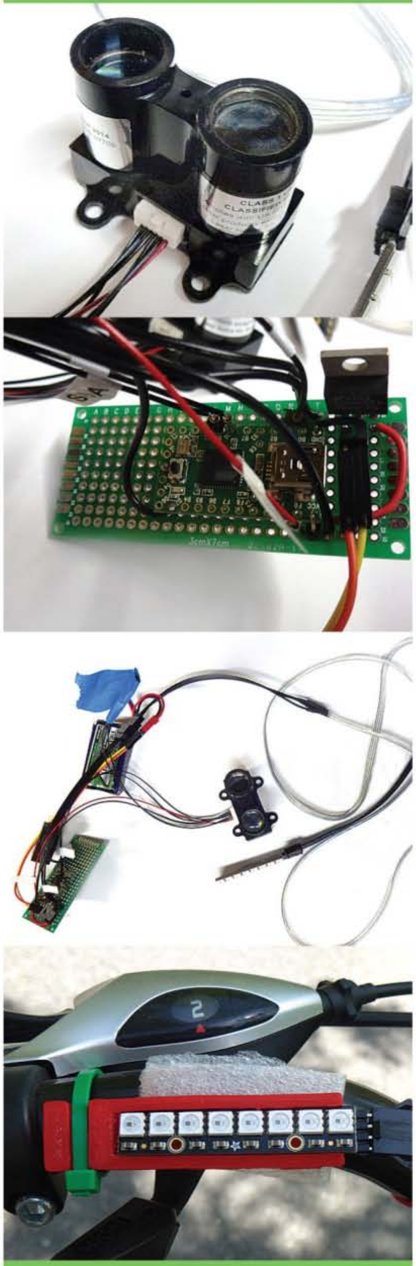

PULSEDLIGHT'S lidar [top] easily interfaces with an Arduino-compatible controller [second from top]. The controller is powered from a battery pack and wired to a row of LEDs [second from bottom]. The LEDs are attached to the bike's handlebars [bottom] and display how close cars are to its rear.

PAUL WALLICH (4)

HIGH TECH DECO

THE WYNDHAM NEW YORKER HOTEL'S MINI-MUSEUM WILL DRAW DIESELPUNKS AND TESLA FANS

What began as a trickle has become routine for Joseph Kinney, the chief building engineer and unofficial historian for the Wyndham New Yorker Hotel, whom we profiled in November 2014. Several times a week, Kinney escorts visitors from around the world on a tour of the two hotel rooms and art deco surroundings where the father of alternating current power, Nikola Tesla, lived for the last 10 years of his life.

But Kinney's tour will soon include a few new stops. Construction is almost complete on a history exhibit in the hotel's lower lobby. This will celebrate the New Yorker's connection to Tesla, as well as its former role as the world's most technically advanced building when it opened on 2 January 1930. (The other new stop will be a peek inside an underground bar that will incorporate an old bank vault—the vault was used as a location for such films as Denzel Washington's *Inside Man* and Will Smith's *I Am Legend*.)

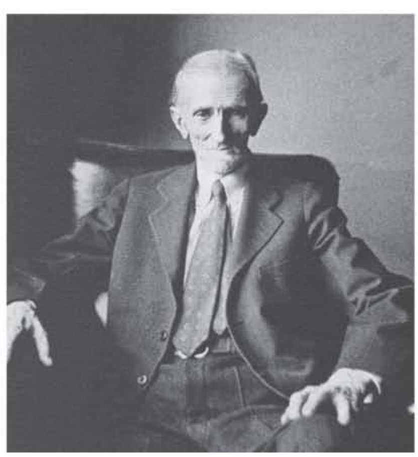

TEMPLE OF TECH: The original 1929 control board [top] of the Wyndham New Yorker's DC generation plant. Nikola Tesla [bottom], circa 1940, sitting in his room at the New Yorker Hotel, where he lived the final years of his life.

In 2008, IEEE recognized the hotel's technological past by declaring it a Milestone in Electrical Engineering and Computing for the 1929 installation of what was then the largest private power plant in the United States. Steam-driven, the plant was capable of supplying electric power for 35,000 people, with excess steam used for heating. Somewhat ironically, given Telsa's contributions to AC, the plant produced direct current. After the inventor took up residence at the New Yorker in 1933, "one can only imagine Tesla down in the boiler room asking to see it and giving advice on ways to improve the systems," says Kinney, a 63-year-old bespectacled man with a professorial air.

Kinney believes that the history exhibit will reflect the hotel and New York City's past, both in good times and bad. "Tesla is the best-known historical connection, but the building hosted concerts by all the great big bands, was a last unforgettable experience for many GIs who were heading to Europe to fight in World War II, and is where notable happenings occurred," he says. (Among the photos are those of Rosa Parks and Coretta Scott King sitting side by side at a luncheon for the National Association of Negro Business and Professional Women's Club in 1956 and Muhammad Ali recovering in suite 2549 after his defeat by Joe Frazier at nearby Madison Square Garden in 1971.)

For Tesla aficionados, the exhibit will include Serbian sculptor Bojan Mikulic's recently commissioned bronze bust of the inventor. It will also have several rarely seen photos of Tesla at the hotel and images of personal items such as hats, pens, and cigars left in the inventor's rooms following his death. The miniature museum will also reflect some of the other once-cutting-edge technologies used in the New Yorker's early years, which should make it something of a dieselpunk mecca: The artifacts on display will speak to a telephone system that once required 95 operators, elevator cars that rose at almost 245 meters per minute, a 42-chair barbershop, and the world's largest electric laundry and dishwashing systems. Interactive kiosks are planned that will allow visitors to listen to big band music and other historical broadcasts made from the hotel. —**BILL GLOVIN**

WYNDHAM NEW YORKER HOTEL

RESOURCES_REVIEW

CODE: DEBUGGING THE GENDER GAP

THE SEARCH FOR SOLUTIONS TO A TECH DIVERSITY PROBLEM

CRAFTING *CODE*: Director and coproducer Robin Hauser Reynolds [left] recording an interview about the work lives of female coders.

Code: Debugging The Gender Gap is a documentary with an ambitious mission. It asks why there are relatively few female programmers, explores the consequences of this state of affairs, and examines some attempts to increase their numbers. The film, which premiered in April at the Tribeca Film Festival in New York City, is intended to be the centerpiece of educational programs in schools and corporations. It features interviews with many notable female (and some male) coders, educators, and policymakers.

Code opens by pointing out that in the early days of programming, women dominated the field. This was in no small part because male engineers thought this "soft" work was relatively trivial compared to building actual machines such as ENIAC. And in the decades that followed, more and more women entered computing: by the mid-1980s, women made up more than 35 percent of computer science majors in U.S. colleges. Then began a falloff that persists to this day, with U.S. computer science majors being about 15 percent female.

Why the decline? *Code* alludes to the introduction of the trope of the male whiz-kid hacker in 1980s popular culture, as exemplified in movies such as *WarGames* and *Weird Science*, and the rise of video games targeted at boys. But the documentary is far more interested in how current popular, educational, and corporate cultures—especially Silicon Valley's—discourage women from being programmers and fail to recognize the female talent that is available. "It's human nature to be able to identify with someone else who is like us. That's why startups tend to hire other white 25-year-old men—'Oh, you went to Stanford? I went to Stanford!' " explains Robin Hauser Reynolds, the director and coproducer of *Code*.

Code paints a grim picture of the results, with so-called brogrammers creating unwelcoming workplaces; any woman who complains can expect backlash. This leads to the fundamental tension in the film: It's trying to portray the real problems women face, while still urging girls and women to enter the field. "This was our biggest challenge," says Reynolds. "There are a lot of women in tech who would not be very happy with us if we skirted those issues.... So how then do you do that, and yet encourage women and young girls to want to go into coding?" A key part of Reynolds's strategy was to show female coders who are happy in their work, including women working in places other than the stereotypical startup. "That's why it was so important for us to show Danielle Feinburg, [a director of photography] at Pixar," says Reynolds, "and to really talk about how collaborative and creative the space can also be."

I was surprised that *Code* made no reference to Gamergate, the ongoing online farrago that began last August, ostensibly about "ethics in gaming journalism." It has in fact been dominated by a relentless campaign of attacks on female game developers, including Zoe Quinn, Brianna Wu, and others. Reynolds says that Gamergate occurred too late in the making of the movie for her team to really delve into it: "It's better to take an issue and drill deep rather than become too broad and sort of be a grab-bag film.... We really felt as though we wouldn't be able to do [Gamergate] justice, to do it quickly."

When it comes to solutions for reducing gender disparity, *Code* argues that computer science courses in primary and secondary schools should be universally available, so that girls don't fall behind the informal education that boys are more likely to receive. At the other end of the pipeline, companies should think beyond their conventional hiring practices. Etsy, the (sometimes controversial) online craft marketplace, is prominently featured for significantly increasing the share of women on its development staff by sponsoring a local educational program that has acted as a feeder for recruitment.

Code doesn't have all the answers, of course. But ultimately, it does make a good case that everyone should think deliberately about diversity in their hiring.

—STEPHEN CASS

CODE

THE MIRACULOUS 1880s

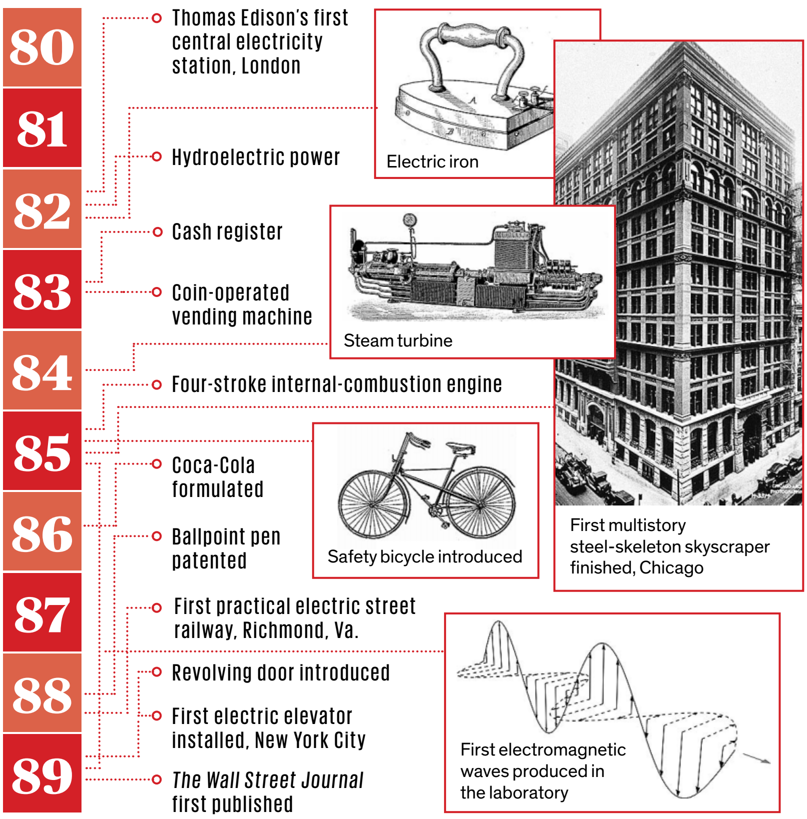

ACCORDING TO THE WORSHIPPERS OF THE E-WORLD, the late 20th century brought us an unprecedented number of profound inventions. But that is a categorical misunderstanding, as most recent advances have been variations on the microprocessor theme and on the parsing of the electromagnetic spectrum. ● Perhaps the most inventive time was the 1880s. Have any two sets of primary inventions and epochal discoveries shaped the modern world more than electricity and internal combustion engines? ● Electricity alone, without microchips, is enough to make a sophisticated and affluent world (we had one in the 1960s). Yet a microchip-governed e-world is utterly dependent on an electricity supply whose fundamental design remains beholden to thermal- and hydropowered-generation systems, both reaching the commercial market in 1882, which still provide more than 80 percent of the world's electricity. And we aspire to make it available at least 99.9999 percent of the time, so that it can serve as the cornerstone of everything electronic. ● Add to that the feats of Benz, Maybach, and Daimler, whose success with gasoline-fueled engines inspired Rudolf Diesel to come up with a more efficient alternative just a decade later. By the end of the 19th century we also had conceptual designs of the most efficient of all internal-combustion engines, the gas turbine. And it was in the 1880s that Heinrich Hertz proved the existence of electromagnetic waves (which had been predicted by James Clerk Maxwell decades earlier). Hertz thus paved the way to our wireless world.

But the 1880s are also embedded in our lives in many smaller ways. A decade ago, in *Creating the 20th Century*, I traced many daily American experiences to mundane artifacts and actions that stem from that miraculous decade. A woman wakes up today in a U.S. city and makes a cup of Maxwell House coffee (launched in 1886). She considers eating her favorite Aunt Jemima pancakes (sold since 1889) but goes for packaged Quaker Oats (available since 1884). She touches up her blouse with an electric iron (patented in 1882), applies antiperspirant (available since 1888) but cannot pack her lunch because she has run out of brown paper bags (the process to make strong kraft paper was commercialized in the 1880s).

She commutes on the light rail system (descended directly from the electric streetcars that began serving U.S. cities in the 1880s), is nearly run over by a bicycle (the modern version of which, with equal wheels and chain drive, was another creation of the 1880s), then goes through a revolving door (introduced in a Philadelphia building in 1888) into a multistory steel-skeleton skyscraper (the first one was finished in Chicago in 1885). She stops at a newsstand on the first floor, buys a copy of *The Wall Street Journal* (published since 1889) from a man who rings it up on his cash register (patented in 1883). Then she goes up to the 10th floor in an elevator (the first electric one was installed in a New York City building in 1889), stops at a vending machine (introduced, in its modern form, in 1883) and buys a can of Coca-Cola (formulated in 1886). Before she starts her work she jots down some reminders with her ballpoint pen (patented in 1888).

The 1880s were miraculous: They gave us such disparate contributions as antiperspirants, inexpensive lights, reliable elevators, and the theory of electromagnetism—although most people lost in their ephemeral tweets and in Facebook gossip are not even remotely aware of the true scope of this quotidian debt. ■

CLOCKWISE FROM TOP LEFT: UNITED STATES PATENT AND TRADEMARK OFFICE; CHICAGO ARCHITECTURAL PHOTOGRAPHING COMPANY/WIKIPEDIA; LENNYWIKIDATA/WIKIPEDIA; WIKIPEDIA; PROJECT GUTENBERG/WIKIPEDIA

INFOGRAPHIC BY Erik Vrielink

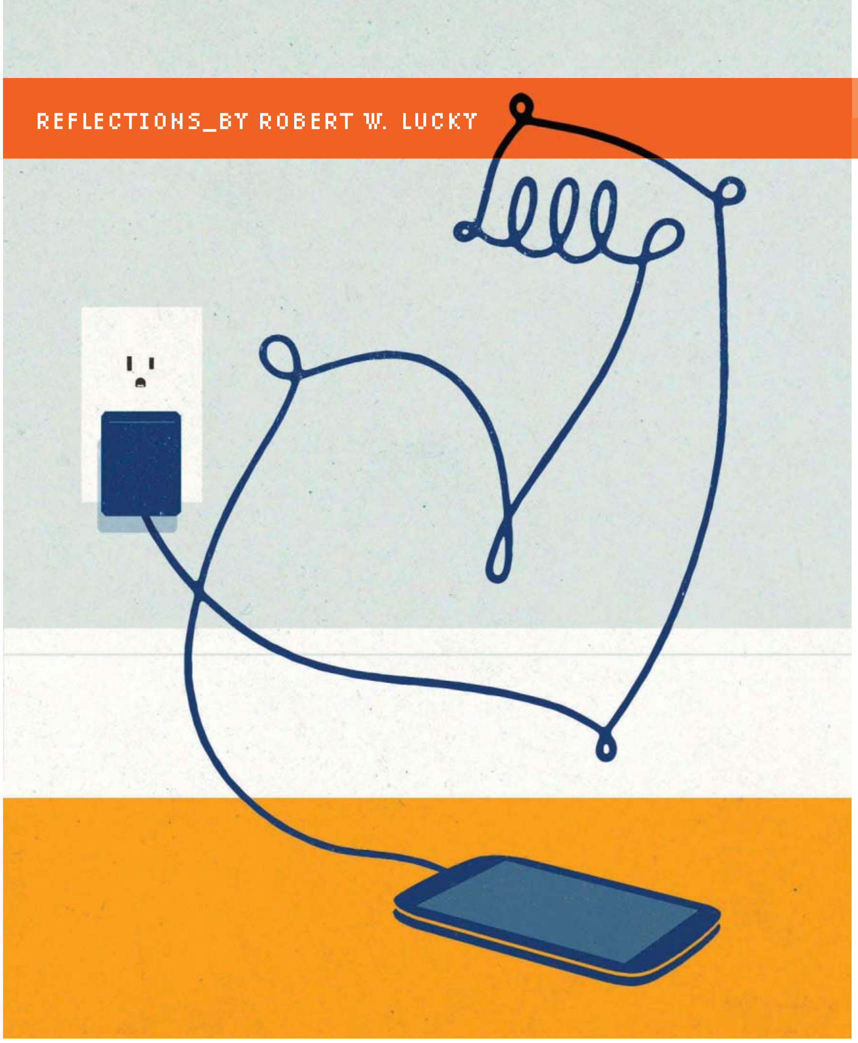

LITTLE THINGS MEAN A LOT

A humble device reminds me of electrical engineering's foundations

THE MOST COMPLEX ELECTRONIC DEVICE COMMONLY FOUND in households is undoubtedly the smartphone, which epitomizes the brilliance and wonder of current electrical engineering technology. The depth of complexity in a smartphone's circuitry, software, and algorithms is such that any single engineer can understand only a fraction of its entirety. • However, today I write in praise of the simplest piece of electronics in households. A little gadget that is both ubiquitous and unnoticed. It is something that is unappreciated until you forget to take it with you when you're traveling. • It's the little AC adapter that converts household power to the 5 volts that keeps your electronic gadgets happy. • There are at least a half dozen of these adapters with their connecting cords tangled at my feet right now. More are scattered around the house, and upstairs there is a box of older adapters that I'm afraid to throw out. I have no idea what they connect to, but maybe I'll find that I need one and that it is irreplaceable. All different kinds of connecting plugs are represented, and few of them are labeled with any information that indicates what device they are supposed to support. Other AC adapters are embedded in electronics such as televisions and audio equipment. And there are a lot of other AC adapters around me that are hidden in the bases of the LED lightbulbs. • The older adapters are mostly heavy wall plugs of the kind that use up your sockets while blocking adjacent ones. They use sizeable transformers and are probably linear power converters. Newer adapters got smaller and smaller, and the newest ones–those that are associated with the present generation of cellphones–are no bigger than standard AC plugs. Obviously, they use switching supply technology.

The appeal of the AC adapter to me as an electrical engineer is that it is a return to the roots of engineering. So many of us today work at much higher levels than are represented by the simple circuitry of AC power conversion. No one could design a cellphone from scratch. A hierarchy of complexity allows an engineer to interconnect existing modules of software and hardware to create some new functionality without having to understand all the detail at the lower levels.

In contrast, the circuitry of an AC adapter looks like a diagram from an EE 101 textbook. There is a full-wave rectifier whose output is chopped by an oscillator to produce a higher-frequency wave, which is then regulated and converted through a split-core transformer that serves the purpose of both energy storage and isolation. A feedback circuit adjusts the chopper to produce the correct output voltage irrespective of the input AC voltage. There are obvious trade-offs between size, cost, and capacity, but also design alternatives–chopper frequency, transformer size, output ripple, efficiency with and without load, risk of core saturation, electrical noise, and so forth. This all makes the adapter actually complicated enough that a detailed analysis involves many issues covered in an engineering education (especially analog circuitry) that may seldom crop up in modern design.

I'm really tempted to cut open one of the very latest adapters, but I'm afraid of what I'd find. Maybe there wouldn't be anything inside at all. ■

ILLUSTRATION BY Dan Page

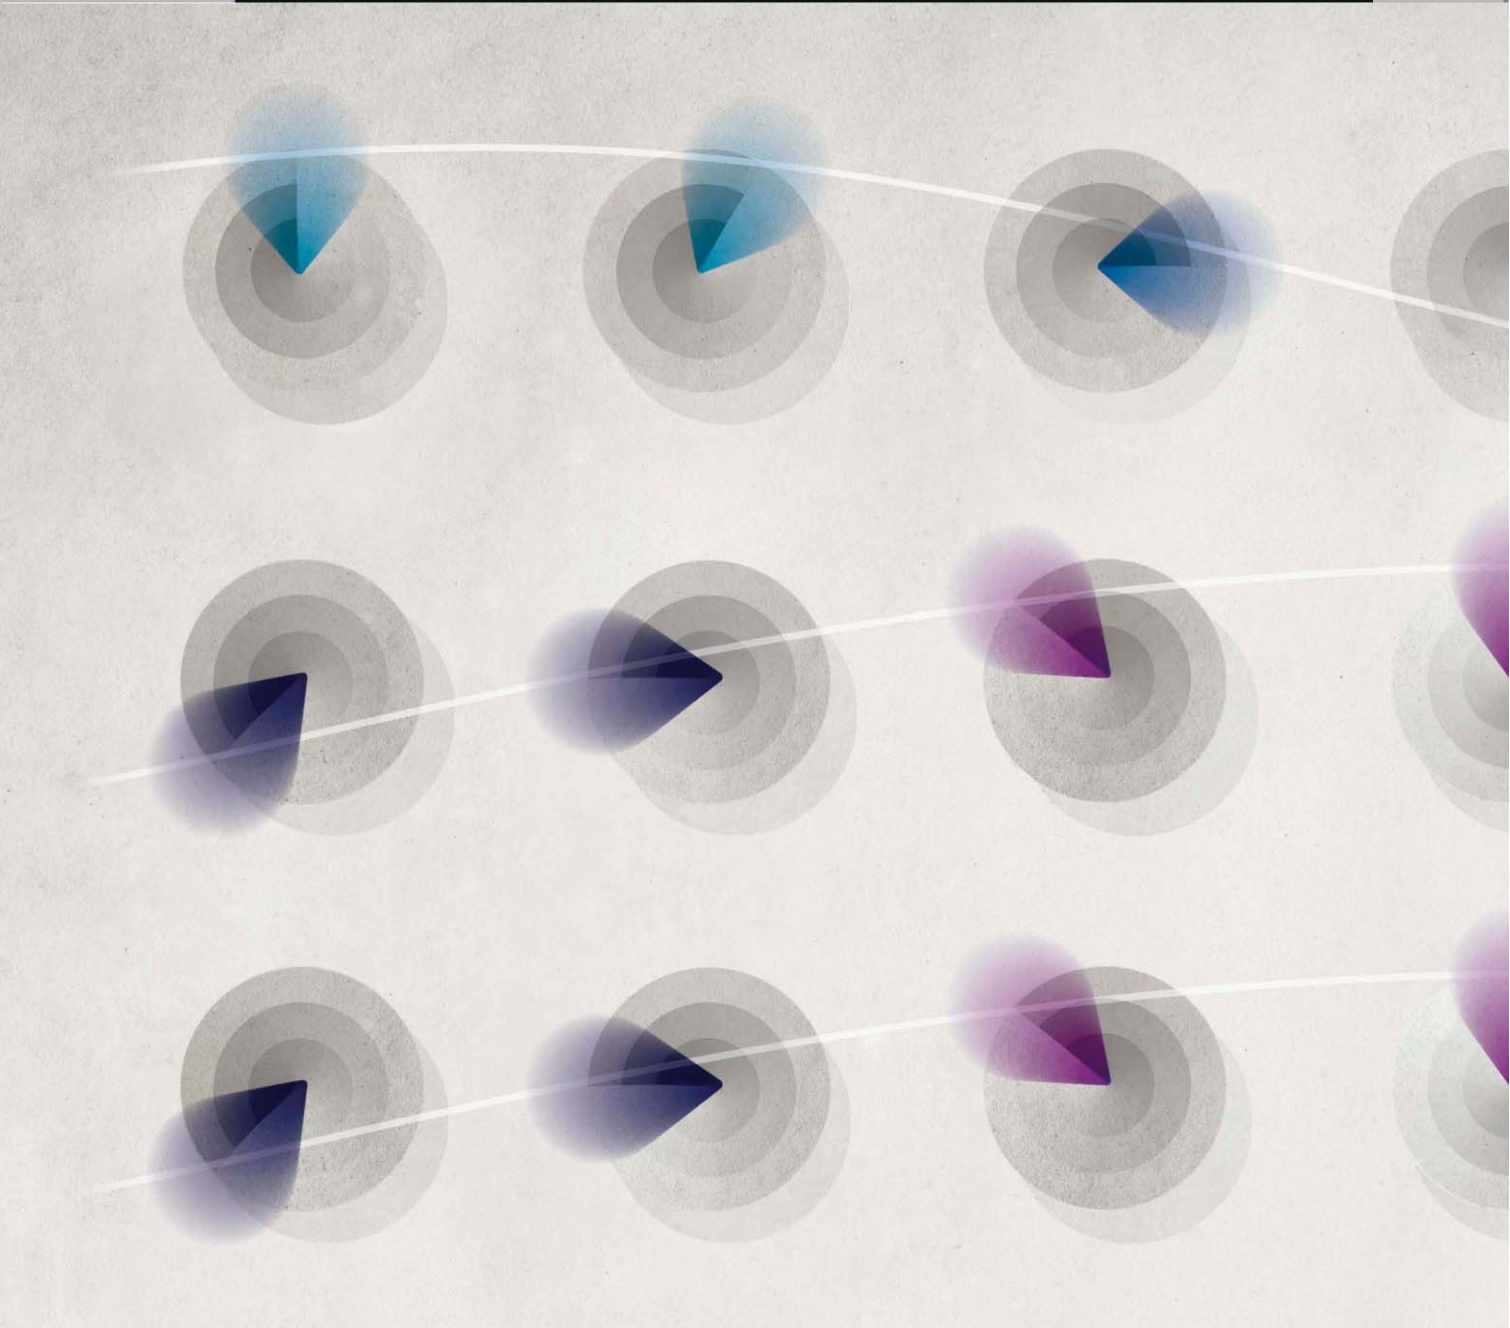

The Computer Chip That Never Forgets

Melding **spin-based** logic and memory could lead to low-power, instant-on electronics

IN 1945, MATHEMATICIAN JOHN VON Neumann wrote down a very simple recipe for a computer. It would contain two key components: a central processing unit to perform calculations and logical operations, and a memory bank to store instructions and data.

Our computers and microprocessor-equipped gadgets still follow this basic recipe. But under the hood, of course, they are far more complex. No existing form of memory is good at everything. So to move instructions and data as fast as possible, engineers have had to compromise. Today's computers use

ILLUSTRATION BY Chad Hagen

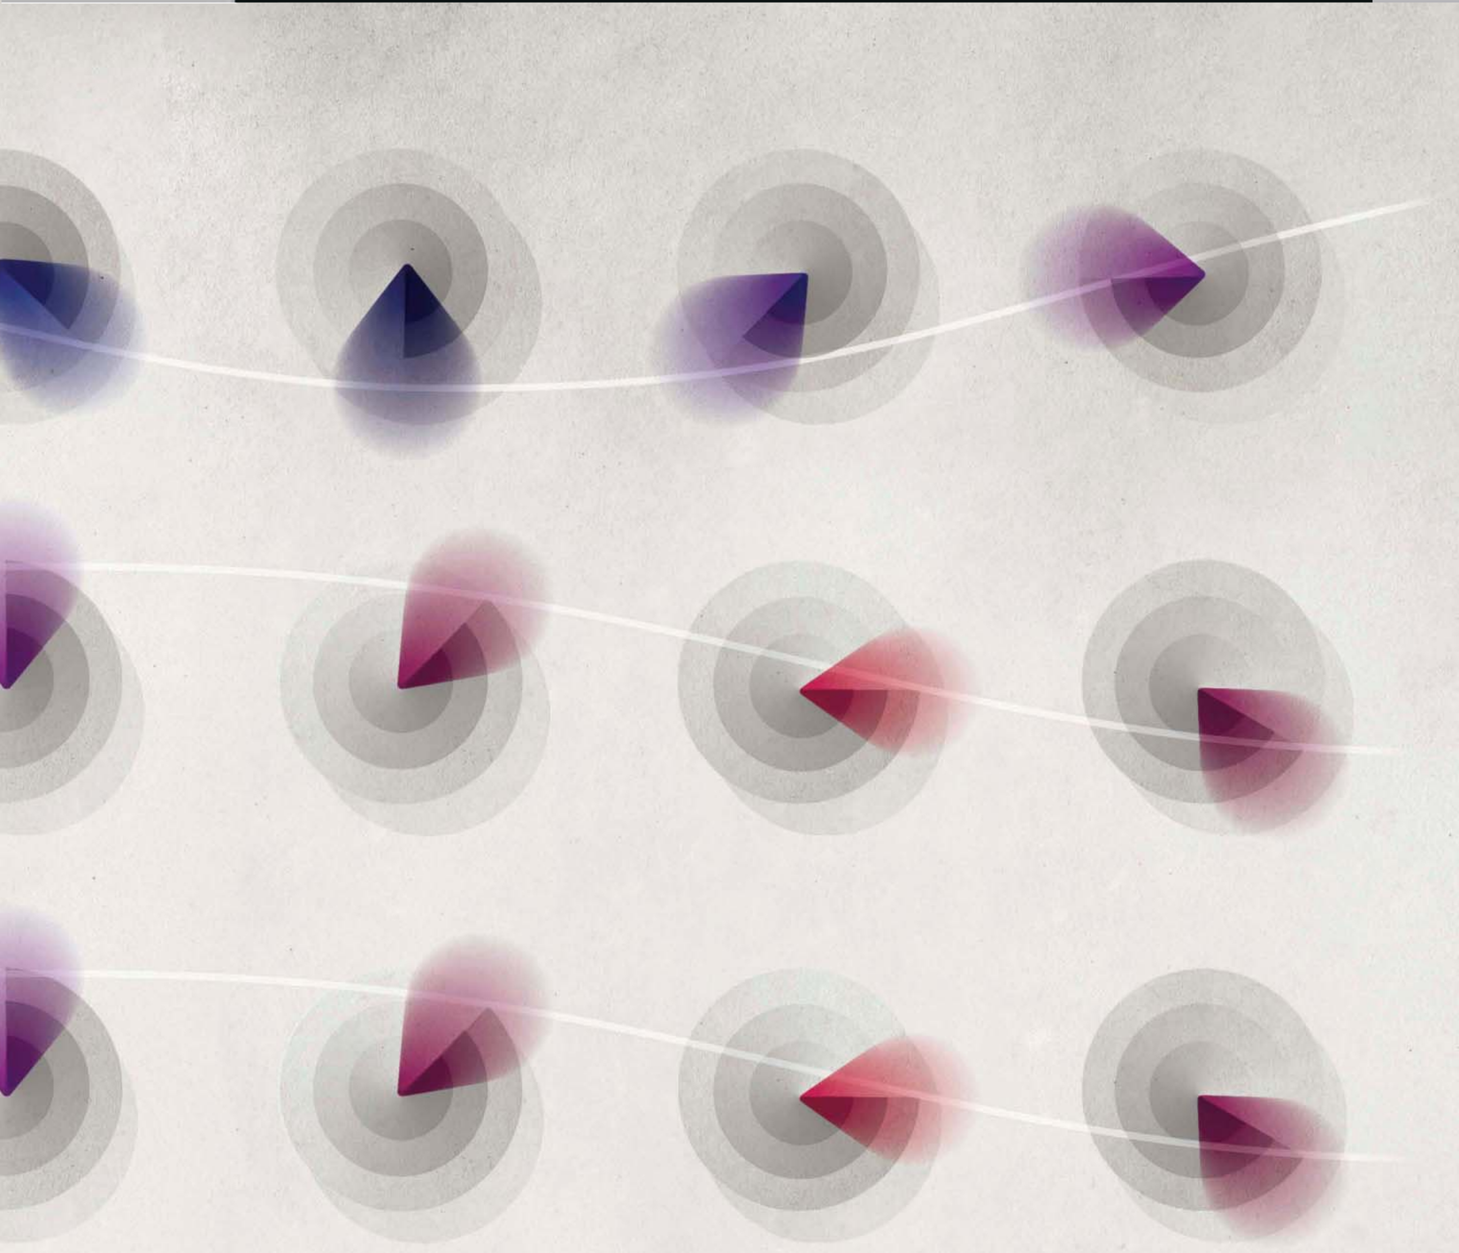

By Pedram Khalili & Kang L. Wang

a smorgasbord of different memory technologies, exploiting the best parts of each.

This hodgepodge has worked quite well for decades, but it's far from ideal. One of the biggest energy drains comes from shuttling data around the CPU and all the levels of memory that surround it. What's more, our fastest memories lose their data if power isn't continuously supplied to them. And the memories that are most compact–those that can store large numbers of bits in a small area–are slow, which is the main reason it takes so long for our gadgets and computers to wake up from sleep.

Those of us who work on alternative memories have long sought a way past these limitations. We've dreamed of creating a single, "universal" memory that could do everything well and could therefore replace the many kinds of memory we have now. It must be fast, to minimize the delays associated with reading and writing data, and it must consume little power every time it is used. We should be able to manufacture it on the same chip as a CPU, allowing us to put it close to the computational action, and be able to make it dense enough to compete with existing memories on cost. At the same time, we would like it to be nonvolatile–able to retain data without having to continuously draw power–so it can be shut down when it isn't needed.

Such a memory is years away. But this vision has been enough to inspire a range of ideas for alternative memories that might one day fill the bill. Our team at the University of California, Los Angeles, has been working on what we believe is one of the most promising candidates: a form of magnetic memory called magnetoelectric random access memory (MeRAM). We have already created small arrays of this memory, and it has recently started to attract the interest of chip manufacturers.

In parallel, we and others have also been exploring how to use the same physical process that lets us write data to MeRAM to do something much more radical: Overhaul the von Neumann blueprint of the computer, by eliminating the long-standing distinction between logic and memory. With MeRAM technology, it's possible to create logic that is its own form of memory–a switch that can both perform computations and remember. This switch would retain its state even when it's powered down.

Such "nonvolatile logic" may one day spur a full overhaul of the microprocessor, allowing us to build chips that can very quickly shut off parts that aren't being used in order to save energy, freeze their state if they lose power, and remember exactly what they were last doing the instant you turn them back on. In our lab we're fond of calling this new form of instant-on electronics "Instantonics," and we think it could dramatically enhance the speed and battery life of computers, tablets, and smartphones. It could also provide a massive boost to some of the most memory-intensive computing tasks, such as video and multimedia signal processing, pattern recognition, virtual reality, and machine learning.

THE HIERARCHY OF MEMORY in a computer starts with static RAM, which is the fastest and is usually found on a microprocessor chip next to a computing core. Next, often on a separate chip or chips, is dynamic RAM, which is somewhat slower than SRAM, but also considerably cheaper. Like transistor-based logic, both of these memories require power to retain their data. In fact, DRAM bits need to be continuously refreshed, or rewritten, to prevent them from losing data. Because of this limitation, long-term storage in computers is accomplished by two memories that do not need power to retain information: the traditional magnetic hard drive and NAND flash.

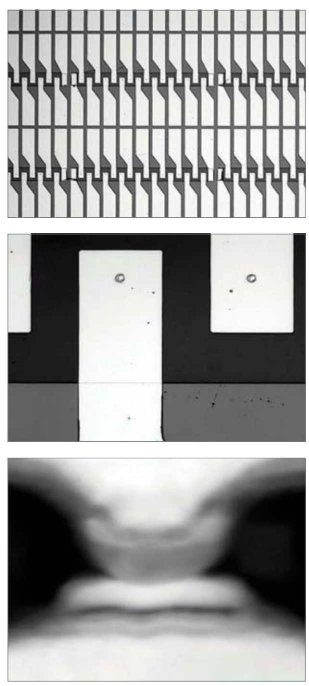

Making **Magnetic Bits**

A portion of a 1-kilobit array of magnetoelectric RAM (MeRAM) as viewed from above [top]. The top electrodes of the bits are made from tantalum and gold [middle]. The small circular dot on each electrode denotes the place where the electrode makes contact with the underlying magnetic tunnel junction, which contains the thin magnetic layer where information is stored. The magnetic tunnel junction [bottom] is shown from the side.

These nonvolatile data storage devices have their own drawbacks. Flash, which stores information as charge that is added to or removed from a transistor structure, can be made at high density and thus low cost. But the process of writing information to each bit is very slow–easily slower by a factor of hundreds than it is in DRAM. It also requires large voltages and can be done only perhaps 100,000 times. Hard disk drives, which store data in the form of the magnetic orientation of different patches, or domains, on a ferromagnetic disk, have better endurance. But they are even slower than flash because they rely on mechanically moving parts to read and write bits.

These shortcomings have led memory researchers to consider alternative, nonvolatile memories that use spin. Spin is a basic quantum-mechanical property of subatomic particles, such as electrons. In magnetic materials such as iron, cobalt, and nickel, it is the spins of the electrons that give those metals their overall magnetic properties–that is, their north and south poles. A particle's spin is closely related to intrinsic angular momentum, which is the property that causes the particle to interact in particular ways with magnetic fields. But despite its name, spin does not have anything to do with actual physical rotation. Yes, the concept of spin is pretty abstruse, but for our purposes here you really need to know just two things. One is that spin has a direction; in other words, if a particle has spin, then that spin points somewhere. The other important concept is that in a material that has been magnetized, most of the individual electrons have spins that point in the same direction. That is what gives the material its magnetization.

Ordinary hard disk drives are, in fact, a form of spin-based memory. On the disk, each bit of data is stored as a microscopic patch of magnetic material, where all of the electrons have the same spin direction. To select a bit to read or write, the disk physically rotates underneath a read/write head, which also moves. To change a bit from a 1 to a 0, the head reverses the polarity of the magnetic material in the tiny patch, so that the spins in it point in the opposite direction.

FROM TOP: UCLA/INSTON (2); UCLA

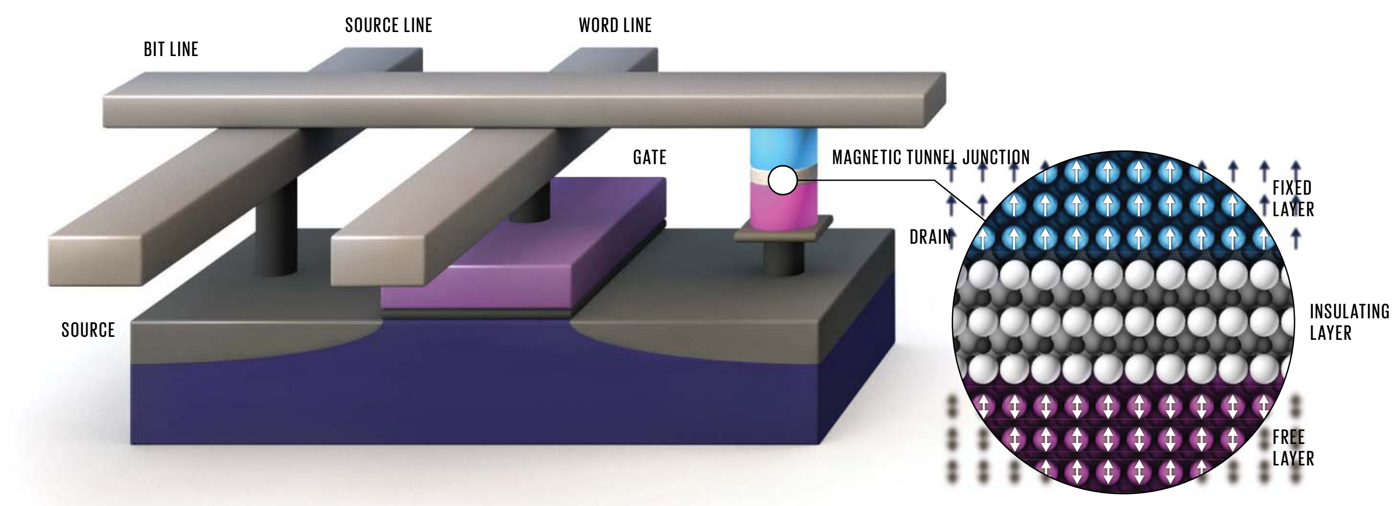

Storing the Information

Both spin transfer torque MRAM (STT-MRAM) and magnetoelectric RAM (MeRAM) can use the same basic architecture to store data in the orientation of electron spin. Each bit in an array can be accessed at the intersection of two lines of electrodes—a source line and a bit line. A third electrode—the word line—is used to control voltage supplied to the bit. A single bit of information can be stored in the free layer of each magnetic tunnel junction. In STT-MRAM, current flows directly through the junction in order to write the bit. MeRAM, which boasts a thicker insulating layer, does not permit current to flow as readily and instead uses voltage-associated effects to change the state of the bit. In each case, the magnetization of the free layer can be flipped [inset, right]. When the magnetization of the free layer and that of the fixed layer (which serves as a reference) point in the same direction, resistance is relatively low. It is higher when the two magnetizations point in opposite directions. Note that a MeRAM transistor is much smaller than an STT-MRAM transistor, since it does not have to provide as large a current. As a result MeRAM cells are smaller overall, and arrays of them can be made more dense.

The term "spintronic memory" refers to a specific kind of memory, also called magnetic RAM, of which there are several variants. These memories have no moving parts. The basic memory element in MRAM is a magnetic tunnel junction, a nanoscale sandwich of two magnetic layers separated by a thin insulating dielectric barrier. (In practice, several other layers are added to boost performance and provide contacts to the outside world.)

In the junction, one of the magnetic layers is pinned, which means the direction of its magnetization is fixed to serve as a reference. The other magnetic layer, which is referred to as the free layer, is where information is stored. The free layer's magnetization can be switched so that it's either oriented in the same direction as the pinned layer or 180 degrees in the other direction. The orientation of this free layer affects how readily current can quantum mechanically "tunnel" across the device, through the insulating barrier. So the value of the resistance of the device indicates the orientation of the free-layer magnetization, and thus whether the bit is 0 or 1.

Overall, the space needed for the wiring and other components makes spintronic memories less dense than hard disk drives. But they're also much faster, less energy hungry, and more reliable because there is no mechanical motion involved.

The process that's used to write to the memory bit, by altering the orientation of the free layer, turns out to be where most MRAM technologies have diverged. One of the earliest MRAM variants, which emerged commercially about 10 years ago, used the magnetic field generated by a nearby current-carrying wire to write data. This type of MRAM was the first to go into commercial production, and it can deliver fairly fast writing speeds on the order of a few tens of nanoseconds, hundreds of times as fast as flash and on par with DRAM. The drawback is that large currents need to be driven through the conducting wires to generate the magnetic fields needed for switching. This requirement prevents the memory cells from being scaled down to the small dimensions you need for modern chips, as it becomes harder and harder to drive sufficient currents through smaller wires and, as cells get closer and closer together, to ensure that magnetic fields affect only the target bit.

A second variant of MRAM, which is being pursued by many of the biggest memory manufacturers and chipmakers, uses a different physical phenomenon, called spin transfer torque (STT), to write information. In STT-MRAM, information is written to the magnetic bit by driving current directly through the device, as opposed to using the magnetic field created by a nearby wire. If the cur-

ILLUSTRATION BY Emily Cooper

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

rent running through an STT-MRAM bit is large enough, the electrons in it can be used to tug on the spins in the free layer and wrench them into the desired direction—either parallel or antiparallel to the spins in the fixed, reference layer.

Over the past few years, a number of research groups including our team at UCLA have been able to show that STT-MRAM can be written in just 100 picoseconds, and with as little as a couple of hundred femtojoules of energy. That's better than the first form of MRAM and on par with SRAM, but it is still not enough to rival the logic switches at the heart of the CPU. A typical CMOS transistor nowadays expends only about 1 femtojoule of energy per switch. Moreover, STT-MRAM does not allow much room for improvement in energy efficiency. The reason is a fundamental one: The device is essentially a wire. As current is driven through it, energy is lost to heat. And as with a wire, the narrower the device gets, the higher its resistance. Also, each memory cell in STT-MRAM needs a transistor to drive the write current through the device. Because the transistor needs to provide relatively large currents, it can't easily be shrunk. So although STT-MRAM is picking up steam, its memory cells will have to stay fairly large—perhaps three to five times the size of DRAM cells.

It turns out, though, that many of the limitations of MRAM and STT-MRAM can be avoided by designing a device that uses voltage instead of current to switch magnetization. This approach to writing information is being actively pursued by several teams, including ours. It promises switching energies of about a femtojoule or even lower—less than 1 percent of what's required with STT-MRAM.

USING AN APPLIED VOLTAGE is standard practice in semiconductors; that's how modern MOS transistors work. It's a voltage, after all, that opens or shuts the gate to the flow of charge carriers in the semiconducting channel beneath it.

Even so, it wasn't immediately obvious that you could use a similar, voltage-based approach to manipulate spins. In order to store information, spintronic devices need materials that have a permanent magnetization, or magnetic moment. The best candidates are usually metals such as iron and cobalt. Unfortunately, conductive materials like these are also pretty good at blocking electric fields, preventing them from penetrating deep into the metal. The reason is that the natural motion of the electrons in such conductive materials easily scrambles electric fields that enter the material, "screening" the effects of an applied voltage much as a fast-moving river might overwhelm the flow of an incoming stream.

As a result, some of the earliest efforts to use voltage in spintronics aimed to make the devices with as little metal as possible. One way to do this is to take a semiconducting material like silicon or germanium and scatter magnetic atoms through it, creating what's called a dilute ferromagnetic semiconductor. If a voltage is applied to this material, the resulting electric field won't be screened as readily as it would in a metal. Instead, the field will pull or push on charges in the vicinity of the voltage gate, which in turn changes the magnetic properties of the bit.

The difficulty with this approach is that it is hard to get ferromagnetic semiconductors to work at room temperature, because their magnetic properties are too weak for their spins to withstand thermal fluctuations without losing alignment.

Fortunately, there's another option. About 10 years ago, theorists began exploring what would happen if you tried to use a metallic magnetic material to make a very thin electric-field-controlled structure—about a nanometer, or less than 10 atoms, thick. This thickness is still greater than the distance an electric field can penetrate into a metal before being largely scrambled. But it's thin enough that small changes at the surface of the material can have a big impact on the film's overall properties, including the natural inclination of the spins to align in a particular direction. This idea was met with great interest and motivated researchers at Osaka and Tohoku universities, in Japan; at Johns Hopkins University; in our group at UCLA; and elsewhere to make memory devices based on this scheme.

The magnetoelectric RAM, or MeRAM, device we've been working on is one of these thin-film spintronic memories. In many ways it's similar to STT-MRAM. For one thing, it can be made with the same cobalt-iron alloys. The main change is in the device structure and engineering of its interfaces. The layer that acts as the bit is very thin. And the insulating layer that in STT-MRAM would pass a current is made thicker, so that very little current can flow. This change in structure essentially turns the magnetic tunnel junction into a capacitive—instead of a resistive—device. When a voltage is applied across this capacitor, the resulting electric field alters the magnetic properties of the device's free layer. This will change how strongly the bit's spin is locked into an up or down orientation. The spin direction will begin to swing, and if the voltage pulse is cut off at the right time, the bit's spin will end up pointing 180 degrees opposite its initial orientation. A small magnetic field or current can also be used to nudge the spin into the desired final state.

The benefits of this arrangement are impressive. Because MeRAM does not require a large current to switch, its transistor can be made much smaller than in STT-MRAM, resulting in much denser arrays. And we have already demonstrated that data can be written in less than a nanosecond, using just a tenth of the energy needed to write to STT-MRAM. As

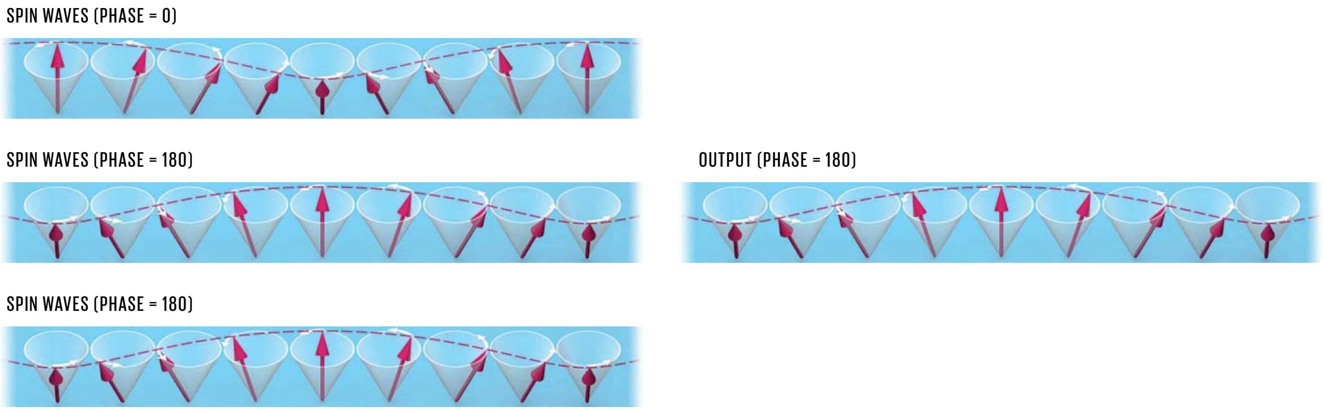

Riding the Waves

A voltage pulse can be used to knock an electron's spin out of alignment, causing it to precess. This effect can be used to transmit information along a strip of magnetic material. The first perturbed spin will affect its neighbor and cause a wave of precessing spins to propagate down the line of electrons [top]. These waves can be made to take on two different phases, separated by 180 degrees, to represent a binary 0 and 1. Two waves with opposite phases can be made to interfere with one another, canceling out the spin perturbation. Logic gates, such as the majority gate, can take advantage of this interaction. The majority gate "votes" for whatever state constitutes the majority of inputs. This can be accomplished as shown [right], by causing three lines of spin-wave inputs to meet and interact. It can also be implemented along a single line, by introducing the voltage pulses that trigger spin-wave perturbations at three different points. In either case, MeRAM memory cells can be built on top of the logic gate to retain input data, trigger spin waves, and read out the result.

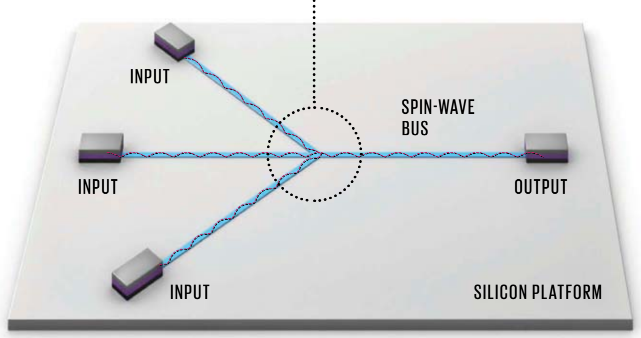

MAJORITY LOGIC GATE

the devices get smaller and materials improve, we expect to be able to further cut down on the energy by as much as a factor of 100.

Last year, our team built a 1-kilobit array of this memory. That's a small step compared with conventional memories, which now extend into the gigabits or even terabits. But this array demonstration, made possible in part with the support of the U.S. National Science Foundation, is an important proof of concept, and we're starting to see some indication of commercial interest. Makers of semiconductor fabrication equipment, including Applied Materials, Canon Anelva, and Singulus Technologies, are working on special sputter-deposition and etching machines that can lay down and pattern the nanometer-thick layers of magnetic compounds needed to make high-density arrays of these devices.

The next step is to show that these memory arrays can be integrated onto existing CMOS chips. This integration will be done at relatively low temperature at the back end of the chip manufacturing process, when layers of metal wiring are used to connect the components of the chip to one another and to the outside world. This basic "back end of line" approach is already used to make MRAM and STT-MRAM and could be adapted to MeRAM as well. (There is, incidentally, a very different memory, resistive RAM, which also exploits these metal layers; it's likely to be more useful for cheap, higher-density storage.)

If MeRAM can be integrated in conventional processors, it could potentially supplant all but the fastest SRAM on a chip. And because it should be easy to make compact, this memory will fit in the metal layers, where plenty of room is available. Thus this technique could bring memory functions that are now done off-chip in DRAM onto the processor. Memory cells would then be separated from the CPU by millimeters or even micrometers instead of centimeters, drastically cutting down on delay and heat-dissipation problems.

ILLUSTRATION BY Emily Cooper

WE CAN ALSO APPLY THE SAME voltage-switching approach used in MeRAM to do something even more radical: Make logic that is nonvolatile and so acts as its own form of memory. This idea got a boost in 2010, when our team received support from the U.S. Defense Advanced Research Projects Agency (DARPA) to investigate the possibility.

The spintronic logic devices we've designed take advantage of the way spins in a material tend to interact with one another so that they all line up in the same direction. Knocking one or more of these spins out of kilter with an electric or magnetic field will cause a wave of spin disturbances–or "spin wave"–to move through the material, like ripples moving out from a splash in a pond.

These disturbances can be created using voltages like those used to control MeRAM. But in this case, the voltage pulses are weaker, perturbing the spins instead of reorienting them by 180 degrees. Our experimental logic devices exploit these waves. In these devices, 0 and 1 do not take the form of the presence or absence of current. Instead they are represented by two different phases of a spin wave–a peak at a given position and time would be a 0, for example, and a valley would be a 1 (or vice versa).

To give you a sense of how spin waves can be used to perform calculations, we can take the example of the majority gate. A majority gate works by taking on the value of the majority of its inputs; typically there are three of them. Including the output, such a gate has four terminals.

In the device, the three voltage-controlled input terminals are used to create spin waves with a phase of either 0 or 180 degrees. The device is designed so that these waves meet and either interfere with or reinforce one another. At the output terminal, the detected signal will be the sum of the spin waves from the different inputs. If all three waves have the same phase, the output will share that phase. If there are two waves with opposite phases, they will cancel each other out, and the phase of the remaining spin wave will determine the phase of the output. Thus two out of three will win. In order to add nonvolatility to this logic operation, MeRAM magnetic memory bits can be integrated on each input and output terminal; they can be used to both input and store data.

Together with our collaborators at the University of California, Riverside, and the University of Massachusetts Amherst, we have simulated these majority gate devices and found that it's possible to use them to create general-purpose computers with just one additional type of logic gate: an inverter.

Spin-wave logic's ability to exploit phase doesn't really have a counterpart in conventional logic. But the simulations suggest that logic built with this technology could have speeds in the gigahertz range, more or less comparable to those of today's chips. At the same time, the spin logic would consume much less energy per operation, on the order of a hundredth of the amount consumed by CMOS.

Other approaches to nonvolatile logic are also being explored. One, which was also supported by the DARPA nonvolatile logic program, transmits information through magnetic perturbations between closely set, discrete nanomagnets. Others are investigating the movement of the walls between magnetic patches. In both of these approaches, it's looking like the most energy-efficient operation may come from using an applied voltage.

Should one of these approaches succeed, we could start thinking about doing dramatic things to the microprocessor. We could overhaul the traditional von Neumann architecture and create circuits that don't need a separate memory to perform calculations and don't have to shuttle data to some external chip to make sure it's not lost.

The implications would go well beyond improving the battery life of our personal gadgets. Chips that can run at extraordinarily low power could be very useful in devices that are physically difficult to access and therefore can't be easily recharged. Examples include medical implant chips and sensors in hard-to-reach places, such as at high elevations, in space, below ground, underwater, or in environments otherwise dangerous to humans. With vast amounts of low-power, on-chip memory, these devices could also be substantially better at particular memory-intensive computing tasks, such as pattern recognition and machine learning. Defense and space applications will benefit from instant-on systems, which will mean power failure, radiation, and other disruptions will be far less likely to result in a loss of critical information.

There are even more exotic circuits still on the horizon. Topological insulators, materials that transmit current only along their surfaces, could potentially be used to make memories that switch with just a thousandth the energy needed by metal-based spintronic devices. But researchers are still exploring the fundamental properties of these materials. In the meantime, these recent spintronic results make it clear that there are things we can do now to greatly improve computing.

Taking full advantage of these technologies will require a willingness to move beyond the strange and somewhat inelegant mix of logic and memory we've relied on for decades. But once we get over the shock of pushing aside the status quo, we will find we can accomplish great things. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/meram0715

Lights, Drone...

IT'S A TAKE: A DJI Spreading Wings S900 drone swoops in for the shot.

.Action

AN EXPERT GUIDE TO DRONE VIDEOGRAPHY + By T.J. Diaz

The camera swoops in on the face of a cliff rising more than 100 meters from the surrounding scrublands. Its view reveals a young woman in a fluorescent orange tank top perched halfway up a wall of ochre quartzite. The camera then cuts to just a few meters away from the rock, close enough to make out the woman's gaze as she looks for the next handhold.

How on earth did they get those amazing shots? You might indeed be wondering that if the title of this YouTube video hadn't given the answer away: "Rock Climbing with the DJI Phantom 2." The Phantom 2 is a tiny electric helicopter you can buy for about US $700.

The ability to capture footage of this sort is not exactly new: For many years the film industry has taken advantage of radio-controlled (RC) helicopters to carry movie cameras to places that cranes and dollies couldn't go. But now small multiple-motor helicopters, or multicopters, have largely supplanted the more mechanically complex single-engine model helicopters used in the past for this task.

These camera-equipped electric drones have been getting cheaper and easier to use, in large part because of the proliferation of smartphones, which led to the development of inexpensive gyroscopes and accelerometers and made it possible to mass-produce very capable yet inexpensive autopilots. Indeed, the cost of camera-equipped multicopters is now well within the means of the countless people who want to take aerial videos for a variety of purposes: journalism, wildlife observation, search-and-rescue operations, real-estate photography—or perhaps just plain fun.

But before you rush out and buy one of these systems for your own pet project, you should educate yourself about the capabilities and limitations of current technology. At least learn enough to know which features are important to your application and which are not. And above all, find out what is safe and legal to do. Here's a guide to get you started.

Although drones have historically been associated with military operations, small remotely piloted aircraft are increasingly used for peaceful purposes. Already in Europe, green-power

PREVIOUS PAGES: DJI

DRONE DAZZLER SHOTS: These aerial platforms provide a unique perspective for taking in the majesty of natural wonders, such as Pakistan's Baltoro Glacier [opposite page]. And what better than a flying camera platform to capture flying sports action, such as a daring bicycle jump [left].

producers fly drones to inspect the blades of wind turbines, farmers use them to survey crops, and oil companies inspect their installations with them. One U.K. company has even used a drone of this kind for close inspection of a tower used to burn off combustible gases—even as the flames rose just meters away.

Multicopters are a type of drone with at least three independent motors and propellers, although the number of motors generally ranges between four and eight, depending on the load to be carried and how much redundancy is sought. A small onboard computer controls the distribution of power to the different motors, which is how most multicopters are steered through the air.

To advance, you send more power to the motors in back and less to those in front, which makes the vehicle nose down slightly, thus propelling it forward. Conversely, providing more power to the motors in front will make the vehicle pitch and move backward. And varying the power laterally causes the drone to roll to the left or right.

You can also rotate the drone about a vertical axis (yaw) by boosting power to half the props—the ones that rotate one direction—and diminishing power to those rotating in the opposite direction, always adjusting the power so that the redistribution doesn't cause unwanted pitch or roll. The resultant yaw motion is a reaction to the change in overall angular momentum of the spinning props.

Shifting power to four or more motors in this way would be impossible to do manually, but the onboard electronics readily handle the task. Indeed, the electronics found in most systems include an inertial-measurement unit (IMU) that can stabilize the drone even in the face of changing winds, which are what makes flying a traditional RC helicopter so challenging.

To keep a traditional model helicopter hovering in the same spot requires a skilled RC pilot who must continually make adjustments. With most multicopters, though, those moment-by-moment corrections are automated. And with the help of onboard GPS receivers, barometric-pressure sensors, and flight-control systems, a multicopter can hover in the same area of the sky with ease, a capability that makes it straightforward to capture video using a camera mounted on the airframe.

What kind of camera and airframe you use will, of course, depend on your exact aims. For serious cinematography, you'll need a multicopter that can carry the considerable weight of a pro-level camera, which probably means a multicopter with long frame arms, up to 12 motors and props, and larger batteries. But you can also get quite good results with less sophisticated cameras and a correspondingly smaller drone.

In general, to carry a Micro Four Thirds–format camera (such as the Panasonic Lumix GH4) will require a multicopter with arms that measure roughly 30 centimeters long. To lift a digital single-lens reflex (DSLR) camera like the Canon EOS 5D or 7D requires something like 45-cm arms. And if you intend to use a true cinema camera, such at the Red Epic, Arri Alexa, or Canon EOS C500, look for a really hefty multicopter with arms that measure from 80 to 110 cm.

You'll also need to think about how you are going to stabilize your camera, which will likely require an electronically controlled gimbal. There are many such gimbals on the market, most nowadays incorporating an IMU and brushless motors, which can rotate the camera smoothly, either to point it at the subject or to compensate for motions of the multicopter.

Pay attention to this component: The gimbal is usually the weakest link in the system. But if you get this right, you'll obtain video that's so steady it's almost mind-bending.

Now that I've sketched the general outlines of what's involved, let me offer some concrete recommendations. Many other choices are available and reasonable, of course, but you won't go too far wrong following these leads.

THESE PAGES: LEFT, DAVID KASZLIKOWSKI/DJI; RIGHT: DJI

Although it might appear that the first thing to decide is which camera you want to use, I'd advise starting with the gimbal instead. That's because some of the best gimbals are designed for specific camera models.

In selecting a gimbal, you'll first need to consider whether a two-axis system is sufficient for your needs. For beginners or hobbyists, this may be perfectly adequate. But if you intend to do professional filming, a three-axis brushless gimbal is de rigueur. That's because a two-axis system will not stabilize the pan axis during flight, and the resulting motions will show up in your videos as a disturbing "yaw wag."

The most popular "prosumer" systems used today include the DJI Zenmuse and Freefly MoVi gimbal systems. The Zenmuse line, which includes two-axis and three-axis variants, is the best choice for stabilizing smaller cameras. A great thing about these gimbals is that they come prebalanced and tuned to handle specific models, including GoPro cameras and DSLRs from Blackmagic, Canon, Panasonic, and Sony. Pick the right camera and this gimbal will work perfectly out of the box. The MoVi gimbals, which would be needed to stabilize a heavier, cinema-quality camera, require somewhat more setup.

HIGH-END EQUIPMENT: This octocopter (an xFold Cinema X8, made by the author's company) is suitable for professional videographers, who are increasingly using such equipment in their work.

XFLY SYSTEMS

Drone-video hobbyists generally use small cameras, like one of the GoPro models or the Sony RX100. This allows them to use smaller, cheaper gimbals (such as the Zenmuse H4-3D line) and airframes (like DJI's extraordinarily popular Phantom 2 quadcopter). The disadvantage is that these cameras don't offer complete manual control of aperture, ISO, and shutter speed. Also, the GoPro's fixed wide-angle lens produces an unpleasant amount of fish-eye distortion. So more serious drone videographers prefer to use DSLRs or even larger cinema-quality cameras.

Most people use fixed focal-length lenses and leave the focus set at infinity with as large a depth of field as possible. But it is possible to put together a system that allows you to adjust zoom and even focus remotely. Of course, the more magnification you use, the more prone your video system will be to the effects of vibration.

Having chosen a camera and gimbal combination, you'll be set to go shopping for a multicopter. For the GoPro and Zenmuse H4-3D, the DJI Phantom 2 is a natural–and much nicer than the earlier-designed Phantom 1. The newer version has a better battery, providing up to 25 minutes of flight time (almost twice that of the original Phantom), and it's also a lot easier to swap in and out. The Phantom 2's props are larger and provide greater lift and efficiency compared with those of its predecessor. The flight controller has been updated and now interfaces with a Mac. The Phantom 2 also comes prewired for a Zenmuse gimbal. These improvements in flight time, user friendliness, and Mac support all make it worth getting the second version of this very capable little multicopter.

Another good option that should give similar if not better performance is the newly released Phantom 3, which comes with a gimbal and video camera and, given DJI's reputation, ought to fly pretty well straight away.

For those who want to put together a system that can carry a DSLR, I'd suggest another of DJI's products: the Spreading Wings S900, which is suitable for advanced amateurs or professionals. This hexacopter ($3,400 for the high-end configuration) is straightforward to assemble and is very portable when broken down.

If you're looking to have all possible shooting options and want a top-of-the-line product, you could purchase one of the xFold line from my own company, Denver-based Xfly Systems. This high-end multicopter is superb for shooting video for HDTV or even footage for feature films. It'll run you more than $10,000, though, for a suitable model. And once you add an RC transmitter, gimbal, video and data telemetry,

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

and a battery charger, your bill will approach double that. Still, it will be considerably less expensive than the kind of cinema-quality camera you're likely be flying with it.

Having the necessary gear is a start, but you'll also need to develop the necessary skills. Getting good video with a multicopter demands attention to lighting conditions, framing, composition, and using the correct camera settings–just as any videography would. But there are some special considerations when you use a drone to gather footage.

First and foremost, the craft itself needs to be stable and properly set up so it flies well. In particular, the feedback-gain settings of the copter's flight-control system must be tuned so that the craft reacts promptly enough to counter gusts of wind and other outside influences without becoming too twitchy. It's also important that the copter not respond too quickly to the operator's control inputs. A little sluggishness will produce smoother flights, which translates into smoother footage.

The best way to set the proper gains is to have them remotely adjustable using a spare channel on the radio used to control the copter. Every time you fly, then, you can adjust the gains to match environmental conditions. You might also need to retune the flight controller's gain settings if you switch cameras or change prop size. Do this diligently. Failing to get your gains set correctly could well result in shaky footage no matter how well your gimbal works.

Another requirement for obtaining good results is carefully balancing the props, which otherwise could cause vibrations in the video. Drone videographers call this all-too-common effect "jello."

Of course, getting the perfect shot will also hinge on your ability as a drone pilot to fly smoothly, accurately, and safely around your subject. Ideally, a second operator will use video telemetry sent from the camera to keep the subject properly framed, panning and tilting the camera as needed by remote control while you concentrate on flying.

Having two people working together in this way is the best strategy by far. Indeed, many experienced drone videographers believe having a separate camera operator is of paramount importance to flying safely and is really the only way to get consistently good results. Others, though, manage to take impressive footage on their own. They generally leave the camera aimed toward the nose of the craft, panning the view by turning the whole copter.

After reading this guide and doing some research online, you might be eager to take the plunge and purchase a multicopter and associated gear. But if you live in the United States, don't hang out your aerial-videographer shingle just yet. Commercial drones, unlike those flown only as a hobby, are still not allowed in the United States, except in limited circumstances. The U.S. Federal Aviation Administration has, however, recently proposed rules for commercial drones weighing less than 25 kilograms (55 pounds). Best guesses are that these rules, or something like them, will go into effect within the next two or three years.

In the meantime, you could go ahead and familiarize yourself with the basics of piloting multicopters by flying them as a hobby. Many forward-thinking aerial photographers are learning to do that now. You could do it on your own or seek more formal tutoring. The Unmanned Vehicle University in Phoenix, where I am a senior instructor, provides one source of expert help. Companies such as Stampede Presentation Products, in Amherst, N.Y., and my own employer, Xfly Systems, offer similar services as well. The Drone Media Group, which I helped found, also offers workshops and training for budding multicopter videographers.

Drone-pilot training, wherever you get it, will no doubt include ample instruction about how to operate multicopters safely. If you're teaching yourself, pay considerable attention to the dangers involved in lofting something high into the air that can weigh several kilograms and has a bunch of rather sharp blades whirling around. A little common sense goes a long way.

Whether or not you end up buying and flying one of these new camera platforms, expect to see more and more stunning aerial sequences in the movies and on television as professional videographers embrace all the things that multicopters can do. ■

DIAZ'S PICKS FOR GOOD DRONE PICS

Here are three models that the author knows work well. These examples span a range of budgets and required capabilities.

FROM LEFT: DJI (2); XFLY SYSTEMS

	DJI Phantom 2	DJI Spreading Wings S900	xFold Cinema X8 Basic
Classification	Consumer	Prosumer	Professional
Number of motors	4	6	8
Propeller diameter (centimeters)	23	38	41
Frame-arm length (centimeters)	18	36	24/28
Battery capacity (ampere-hours)	5.2	10–15	20–44
Payload (kilograms)	up to 0.3	up to 5	up to 21
Flight time (minutes)	25	18	30
Base cost	US $699	$1,400	$10,335

POST YOUR COMMENTS at http://spectrum.ieee.org/drones0715

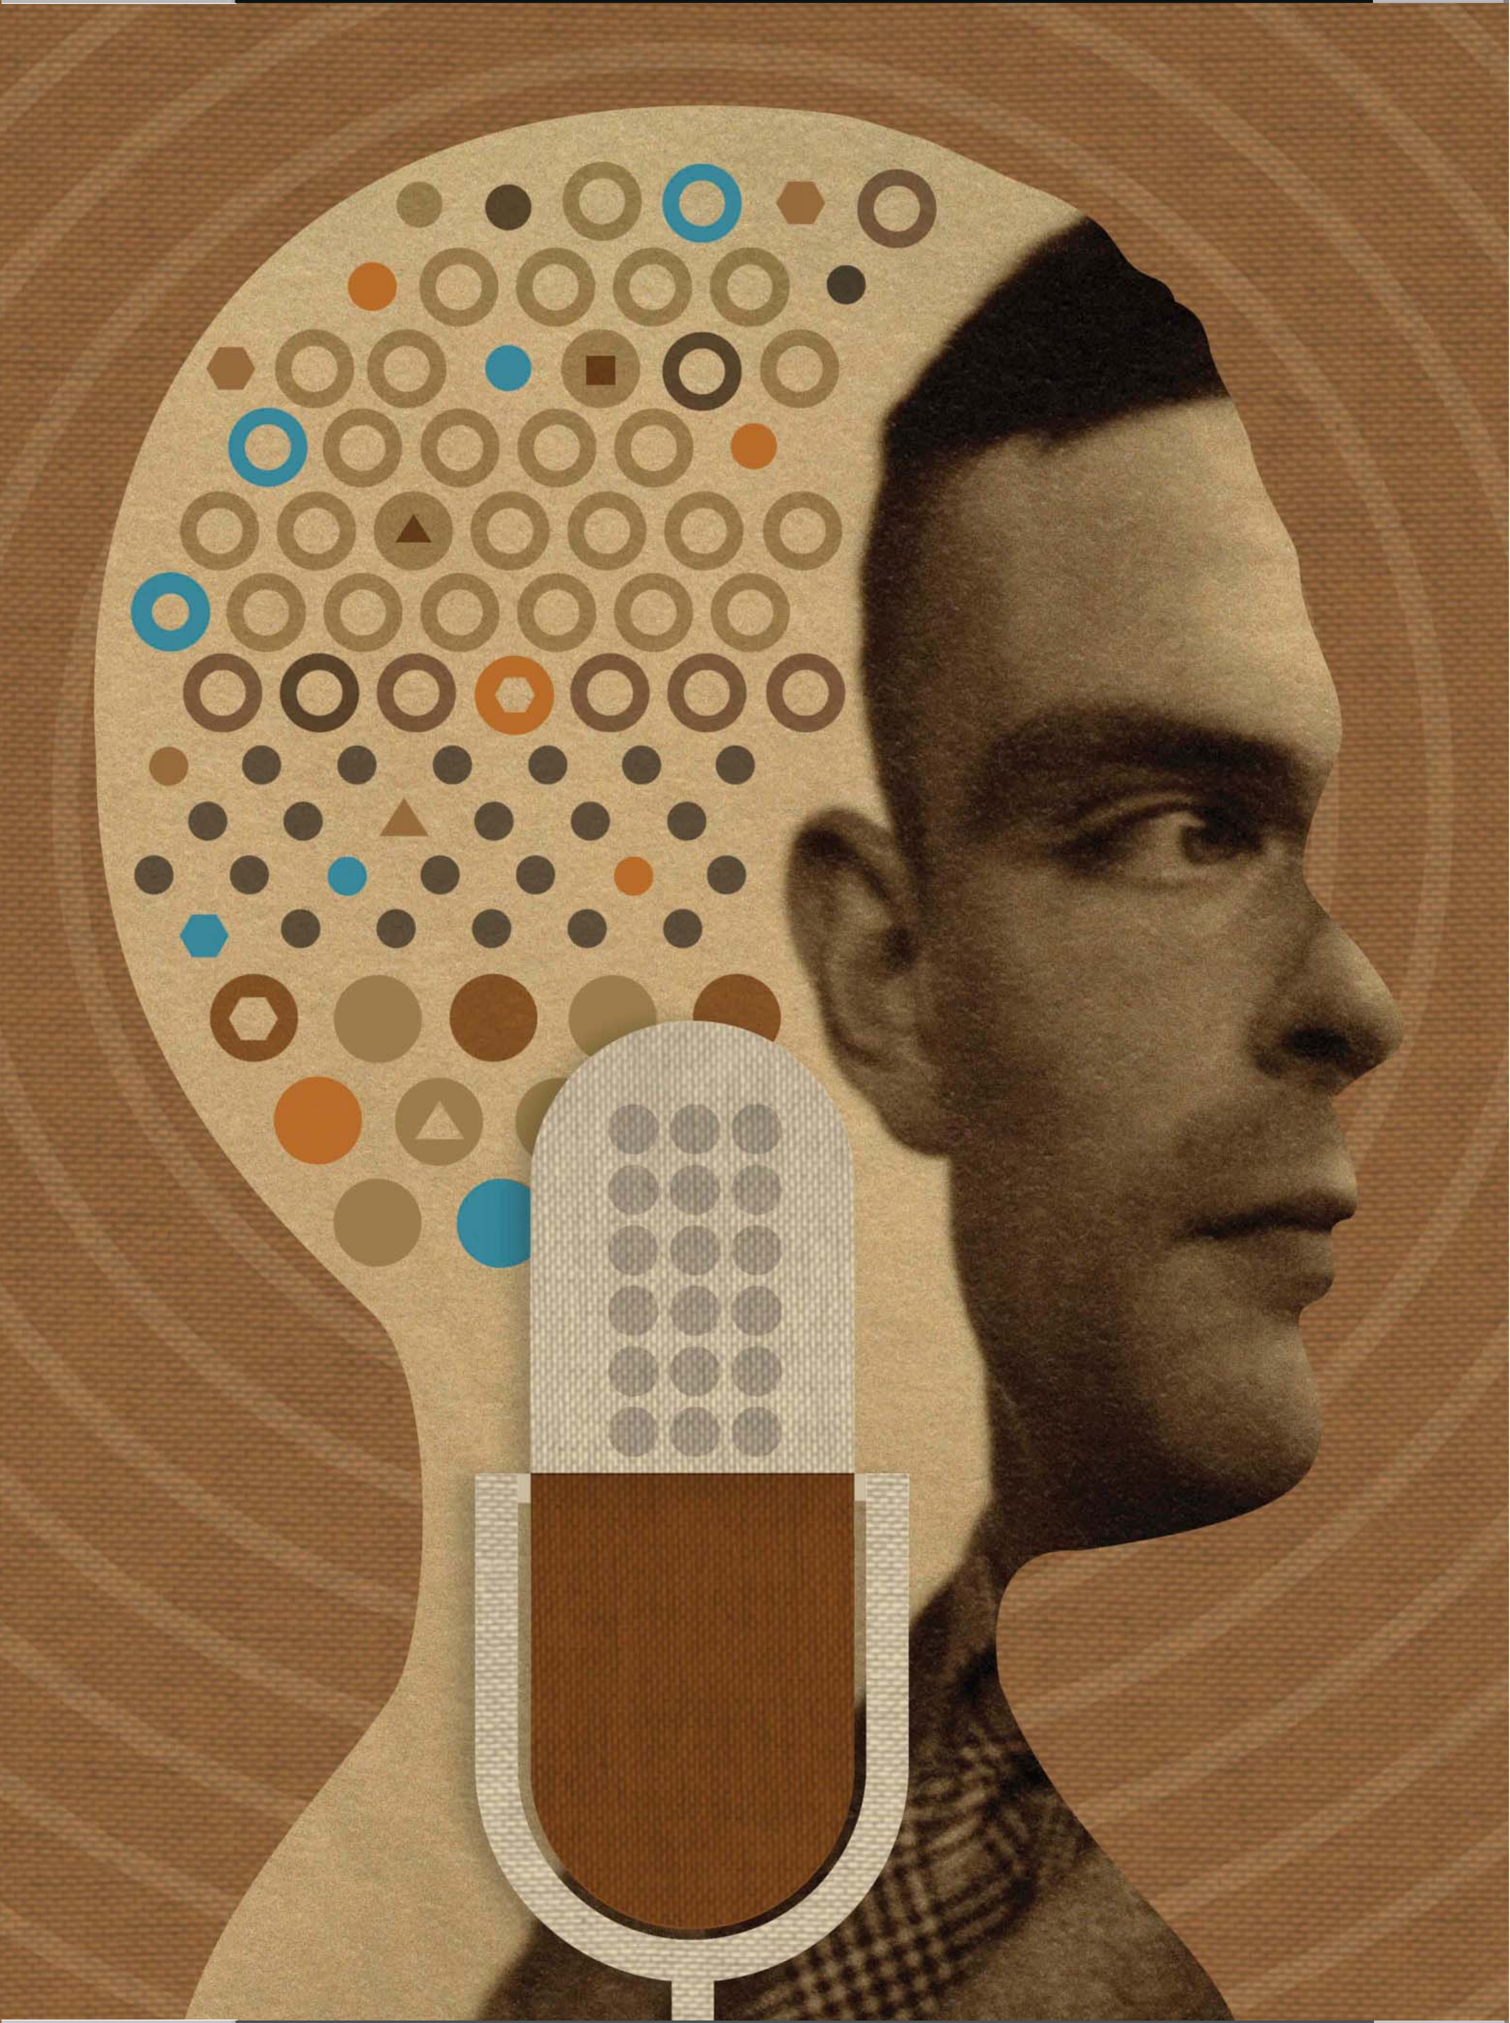

BY DIANE PROUDFOOT

WHAT TURING HIMSELF SAID ABOUT THE IMITATION GAME

The mathematician and cryptanalyst explained his famous test of computer intelligence during two BBC radio broadcasts in the early 1950s

THE IMITATION GAME, the recent biopic about Alan Turing's efforts to decipher Nazi naval codes, was showered with award nominations. It even won the 2015 Academy Award for Best Adapted Screenplay. One thing it won't win any awards for, though, is its portrayal of the "imitation game" itself–Turing's proposed test of machine thinking, which hinges on whether a computer can convincingly imitate a person. The Turing test, as it is now called, doesn't really feature in the film. (Given that the movie gets so much of the history wrong, perhaps that's a good thing.)

Turing described his now-famous test in a scholarly article, "Computing Machinery and Intelligence," which was published in the quarterly journal *Mind* in 1950. But his last recorded statements on the topic were, in fact, aired on the radio in 1952, in a broadcast that very nearly didn't happen. What he had to say to the British public, though little known today, reveals quite a bit about his thoughts on the topic.

The story goes like this: In 1950, a BBC producer named Archibald Clow, on the lookout for new talent, met Turing in Manchester. Clow was unimpressed. Writing to Christopher Holme, a top executive at the BBC, he stated, "He undoubtedly has a lively mind, but I am very doubtful about him as a speaker." Clow reported that Turing seemed to have "a definite hesitation in his speech" and thus would not recommend him.

What's more, Clow grumbled, Turing "doesn't seem keen to talk along the lines of his paper in *Mind*.... He would rather stick to the mechanico-factual side and I have a suspicion that the reaction to the *Mind* article may have been unfavourable and this at present is making him excessively cautious."

Fortunately, Holme knew that Clow was not the best judge of broadcasters. The previous year, Clow had advised against using the astronomer Fred Hoyle on account of his strong Yorkshire accent. Hoyle turned out to be one of the BBC's most successful

TURING PHOTO: ARCHIVIO GBB/CONTRASTO/REDUX

PHOTO-ILLUSTRATION BY Chad Hagen

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

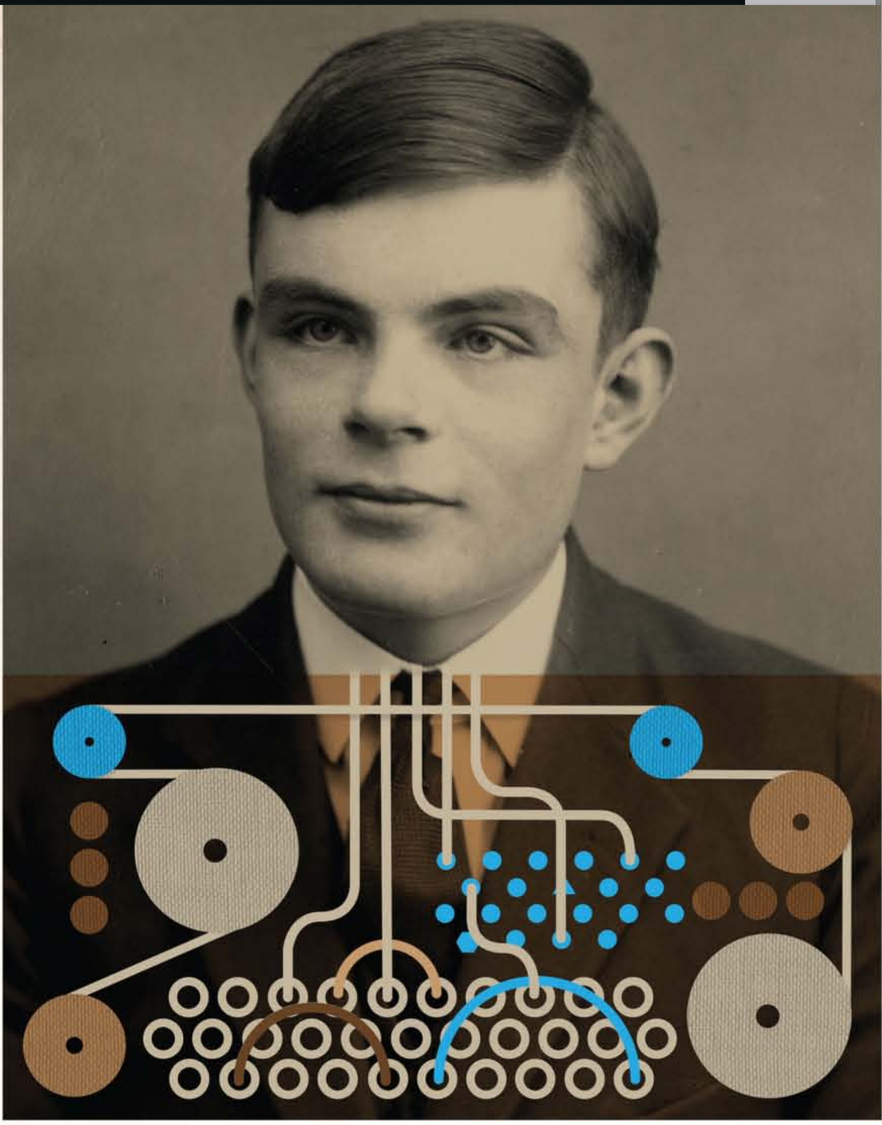

broadcasters—it was actually in a BBC talk on the origin of the universe that Hoyle coined the term "big bang."

Holme took just as much notice of Clow's reservations about Turing: "Turing, who is by training a mathematical logician rather than a scientist, should be given the fullest opportunity to develop his argument in the Programme even if it requires more than one talk," he wrote back to Clow. While Holme took a more limited view of Turing's aims in the *Mind* article than people do today—"Basically Turing does no more than declare his faith in what is in reach of practical achievement in the immediate future"—he nevertheless thought that the ideas in the article "could be made intelligible to a fairly wide audience."

Holme was right, and in the end Turing recorded two talks, both broadcast on a BBC radio network called the Third Programme, one in May 1951, the other in January 1952.

The Third was a network that couldn't exist in today's ratings-conscious climate. Born on the heels of World War II of what lofty BBC types called a "virtually insatiable demand for serious literature and drama, for good music and intelligent discussion," it was (they said) a chance to "experiment boldly…without regard for…mass-appeal." According to its first head, George Barnes, the Third's goal was to provide the intellectual stimulation that had been so lacking during Britain's many years of war.

The Third was demanding. It expected a listener to "meet the performer half-way by giving his whole attention to what is being broadcast." This was exactly the kind of outlet Turing needed for his radio debut. And we should be grateful to the BBC executives who engineered the opportunity, because Turing's radio talks—archived only in the form of transcripts, alas—provide a wonderful window into his thinking.

TURING'S FIRST BROADCAST, titled "Can Digital Computers Think?," was one of a series of five radio talks by early pioneers of computing. The other speakers were Douglas Hartree and Max Newman, both mathematicians, and Maurice Wilkes and Freddie Williams, who were electronics experts.

Like Turing, these men were scarcely household names. All had spent the recent war doing ultrasecret work that, if only they could have spoken about it, would have awed their listeners. The average person would have seen these scientists' names only in newspaper articles or letters to London's *Times* about the new "electronic brains." Yet it's no exaggeration to say that Turing and Newman (also one of the leaders in decrypting Nazi communications) contributed at least as much to the defeat of Hitler as any of Churchill's celebrated generals.

So it was only fitting that these men had a chance to speak to the British public. And fortunately, Holme's confidence in Turing as a broadcaster was justified. The producer of his first talk said that Turing spoke "quite naturally and unaffectedly." The BBC then asked him to make a second broadcast, this time as a participant in a symposium titled "Can Automatic Calculating Machines Be Said to Think?" There were three other participants: Newman, philosopher Richard Braithwaite, and neurosurgeon Geoffrey Jefferson, whose views Turing had attacked in his *Mind* article.

The BBC recorded this symposium in London on 10 January 1952. The company paid Turing's train fare and hotel expenses, plus a fee of 20 guineas, which is about US $800 in today's money—probably quite a bit less than the movie's lead actor, Benedict Cumberbatch, would receive for appearing on a British talk show.

TURING PHOTO: HERITAGE IMAGES/GETTY IMAGES

PHOTO-ILLUSTRATION BY Chad Hagen

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

The symposium first aired at 9:35 p.m. on Monday, 14 January, a prime evening slot. In line with the Third's goal of "presenting the great classical repertoire," the talk was preceded by French chamber music and followed by a Schubert piano sonata. The actual symposium lasted nearly 45 minutes–four scholarly voices against a background of crackle and hiss. (The Third wasn't always easy to tune in to: Even the BBC acknowledged that there were "too many places where to hear it is physically impossible or where reception is so bad that listening to it is a penance.")

Although Turing's audience was not huge by today's standards, it was nevertheless large–approximately 100,000 people, according to BBC estimates. The folks who tuned in were likely a more diverse group than you might expect. About a third of them were "working class," according to a report on the Third's market. Some probably lived in continental Europe, where the Third was popular. Indeed, Turing would have been audible as far away as Switzerland, even though 30 percent of Britain wouldn't have been able to hear him clearly or even at all.

HOW DID ALL THESE PEOPLE react to Turing's radio talks? Mathematicians were as prone to caricature in the 1950s as they are today. "Mathematicians are queer folk," said the president of the London Mathematical Society on BBC radio a year before Turing's first broadcast. "Mathematicians live in intellectual isolation," he told his listeners. "They are usually shy and reserved about their work, and in the company of other learned men they tend to feel isolated and to look upon their studies as an incommunicable and delightful secret." In some respects, he could have been talking about Turing, who was undoubtedly shy.

The late Sarah Baring, Lord Mountbatten's goddaughter and one of the female "slaves," as Turing called the women who assisted the almost entirely male code breakers at the Government Code and Cypher School at Bletchley Park, said that "on the rare occasions when he was spotted like a protected species, he would be shambling down to the canteen in a curious sideways step, his eyes fixed to the ground." Nevertheless, Turing seemed self-assured in his radio broadcasts, and he addressed his audience directly and in an accessible way.

The public still didn't have a clue about Bletchley Park's computing machines. But the postwar electronic brains being built by Turing, Wilkes, Williams, and others were well known, thanks to a blaze of publicity about these superfast machines in the national dailies and even the local papers.

The computer is "permitted all sorts of tricks so as to appear more man-like"

In the growing debate about the possibility of computer intelligence, Turing was almost a lone voice on the side of the machines. His BBC broadcasts gave him a chance to get his viewpoint across to "the man in the street" (Turing's words), and his scripts reveal him to be a good communicator–a man dedicated to explaining difficult ideas to the public and who did this rather well. *The Imitation Game* portrays Turing as unable to understand jokes, but his broadcasts display a nice, if dry, sense of humor.

TURING'S 1952 BROADCAST shreds three modern myths about the Turing test, as my colleague Jack Copeland at the University of Canterbury, in Christchurch, New Zealand, has argued. First, according to many accounts–including Andrew Hodges's biography of Turing, which inspired the screenplay of *The Imitation Game* –Turing's aim was to provide an operational definition of thinking. In fact, he expressed the opposite intention, saying on the radio, "I don't want to give a definition of thinking," and adding, "I don't really see that we need to agree on a definition at all." Indeed, several listeners polled by the BBC complained that Turing and his fellow speakers had avoided giving a definition.

The second myth is that Turing predicted a machine would pass his test around the beginning of this century. What he actually said on the radio in 1952 was that it would be "at least 100 years" before a machine would stand any chance with (as Newman put it) "no questions barred."

The third myth is that Turing's test is flawed, because it can be passed by machines that obviously don't think–for example, a computer that merely searches through a huge database of ready-made conversations, looking for matching responses to the judge's questions. But Newman and Turing made it clear in the broadcast that brute-force searches like this might take the computer "thousands of millions of years."

During the broadcast, Braithwaite, Jefferson, and Newman obligingly drew Turing out on key points of his proposed test of a machine's ability to think, including such questions as *Is the computer allowed to lie?* Turing said: The computer is "permitted all sorts of tricks so as to appear more man-like." This even gives the machine the option of trying to evade snarly questions by pretending to be a foreigner, with a limited grasp of local culture and vocabulary. (The winner of a recent computer-imitates-human game, which feigned being "Eugene Goostman," supposedly a 13-year-old Ukrainian boy, did just this–

but even then the program did not pass the Turing test, despite much media hype to the contrary.)

Other issues emerged during the discussion, such as *What sort of thing is the judge permitted to ask?* "Anything," Turing said and added, "And the questions don't really have to be questions, any more than questions in a law court are really questions."

Won't a computer be easy to detect, since it is infallible? No, because, as Turing said, the computer can use tricks to avoid being spotted. In any case, his view was that "computing machines aren't really infallible at all." Turing knew very well that computers make mistakes.

What kind of computer program could do well at this test? Turing's idea, brilliantly innovative at this time, was that a computer set up in the right way could learn from its teacher and from its experiences, as a human child does. He called these "child-machines" and mentioned his own experiments along these lines.

Won't a computer stand out as particularly unemotional? Turing suggested that child-machines might even develop humanlike emotions, which he said are "likely to occur as a sort of by-product of genuine teaching." If he was right about this, computers need not pretend in order to appear emotional.

Isn't an electronic brain very different from a human brain? The speakers agreed that there are indeed large differences, such as in the size of the memory store. But Turing challenged the idea that the electronic brain would have to be similar to the human brain in all respects to count as something that could think: "To take an extreme case, we are not interested in the fact that the brain has the consistency of cold porridge." Indeed, his proposed test, where the machine is interviewed by teleprinter, filters out the many differences that are irrelevant to the issue of whether computers can think.

Did listeners buy Turing's arguments? Audience reaction was mixed, with comments ranging from "This was a pure gem" to "I think it sad to hear men of such learning wasting 45 minutes in such a fashion." Some listeners raised exactly

Mocking AI Panic

Turing poked gentle fun at those who worried about intelligent machines taking over

Many people today are concerned by the prospect of out-of-control artificial intelligence. Some call it "killer AI," "evil AI," or "malevolent AI." Billionaires throw money at the "existential risks" posed by ultraintelligent machines: In January, Elon Musk, creator of PayPal and CEO of SpaceX and Tesla Motors, donated US $10 million to the Future of Life Institute, in Cambridge, Mass., which is "focusing on potential risks from the development of human-level artificial intelligence." Other new research institutes with apocalyptic names explore the dangers of "singularity" scenarios, and Google has recently formed a hush-hush AI ethics board.

Actually, history is repeating itself. In the mid-1940s, public reaction to reports of the new "electronic brains" was fearful. Newspapers announced that "the controlled monster" (a room-size vacuum-tube computer) could rapidly become "the monster in control," reducing people to "degenerate serfs." Humans would "perish, victims of their own brain products."

Alan Turing, the father of modern computing, added to the consternation. In his first BBC radio talk he noted, "If a machine can think, it might think more intelligently than we do, and then where should we be?" In another talk, he said that "it seems probable that once the machine thinking method had started, it would not take long to outstrip our feeble powers." But 21st-century pundits have usually overstated the progress AI has made since Turing's day, suggesting that ultraintelligent computers are just around the corner.

ILLUSTRATION BY Chad Hagen

Turing thought the prospect of such machines "remote" but "not astronomically remote." He noted, "If it comes at all it will almost certainly be within the next millennium." And then, he said, "we should have to expect the machines to take control."

Stephen Hawking united with Nobel laureate physicist Frank Wilczek, Wilczek's MIT colleague Max Tegmark, and Berkeley computer scientist Stuart Russell to declare last year in *The Huffington Post*: "Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last." And this past December, Hawking warned in a BBC interview that humans would be "superseded" by AIs.

In his 1951 radio broadcast, Turing used that very word—as did Samuel Butler, who wrote in his 1872 novel *Erewhon* (a satire Turing mentioned approvingly) that we must put "an immediate stop to all further mechanical progress." He was, of course, ridiculing opponents of the machine age.

Turing (following Butler) poked fun at the fear of out-of-control AI. When he predicted in the London *Times* that machines could "enter any one of the fields normally covered by the human intellect, and eventually compete on equal terms," the media protested at the "horrific" implication of these ideas—namely, "machines rising against their creator." But Turing said drily, "A similar danger and humiliation threatens us from the possibility that we might be superseded by the pig or the rat." He joked—but with a good pinch of his usual common sense—that we might be able to "keep the machines in a subservient position, for instance by turning off the power at strategic moments."

Turing knew that developments in AI worried some scientists, as well as other folk. Pioneering cyberneticist William Grey Walter had said as early as 1948 that there was "something sinister" about the new "mechanical monsters."

Alarm today spreads further and more quickly. Last year Musk told his 1.7 million Twitter followers that AI is potentially "more dangerous than nukes." Even escaping to Mars—presumably on one of his own rockets—wouldn't help, he said: "The AI will chase us there pretty quickly." In January, Bill Gates told Reddit's millions of visitors, "I agree with Elon Musk and some others on this and don't understand why some people are not concerned."

Turing's response to AI panic was gentle mockery. All the same, there was a serious edge to his humor. If runaway AI comes, he said, "we should, as a species, feel greatly humbled." He seemed almost to welcome the possibility of this humiliating lesson for the human race.

In Turing's view, humans are not exceptional: We, too, are machines. He said matter-of-factly on the BBC, "It is customary, in a talk or article on this subject, to offer a grain of comfort, in the form of a statement that some particularly human characteristic could never be imitated by a machine." But, Turing ended chillingly, "I cannot offer any such comfort." *—D.P.*

the same objections Turing anticipated and countered in his 1950 article in *Mind*. Thinking, some said, is "the God-given prerogative of man" and computers "will never be anything other than a cleverly constructed mass of metal." Yet overall 54 percent of listeners polled by the BBC gave the broadcast an A or A+.

It seems that these people didn't detect (or object) that these experts were reading previously prepared comments, as was standard practice in those days. Prior to airtime, all four departed from the original script, with Turing making the most changes. The final version–the one actually used–is, fortunately, preserved at the BBC Written Archives Centre, which holds a vast and only partly cataloged collection of old scripts and other historical papers.

POST YOUR COMMENTS at http://spectrum.ieee.org/turing0715

Turing's biggest single change to the script is to his discussion of learning by analogy, such as when cosmologists explain the idea of the expanding universe using the analogy of an inflating balloon. Turing was interested in how an analogy can, as he said, "ring the bell" in the brain. Today's skeptics of what is now called strong artificial intelligence say that a computer is incapable of these lightbulb moments. Turing, by contrast, thought that if only scientists could discover a mechanical explanation of how analogical thinking works in the human brain, they could program a computer to do the same.

Turing's fundamental message about thinking in his 1952 radio broadcast was that we shouldn't set the bar any higher, or any differently, for computers than we do for people. We don't decide that our fellow human beings think by putting their brains under a microscope–ordinary everyday interaction is enough. That was Turing's astute observation. If you want to tell whether a machine thinks, try communicating with it. ■

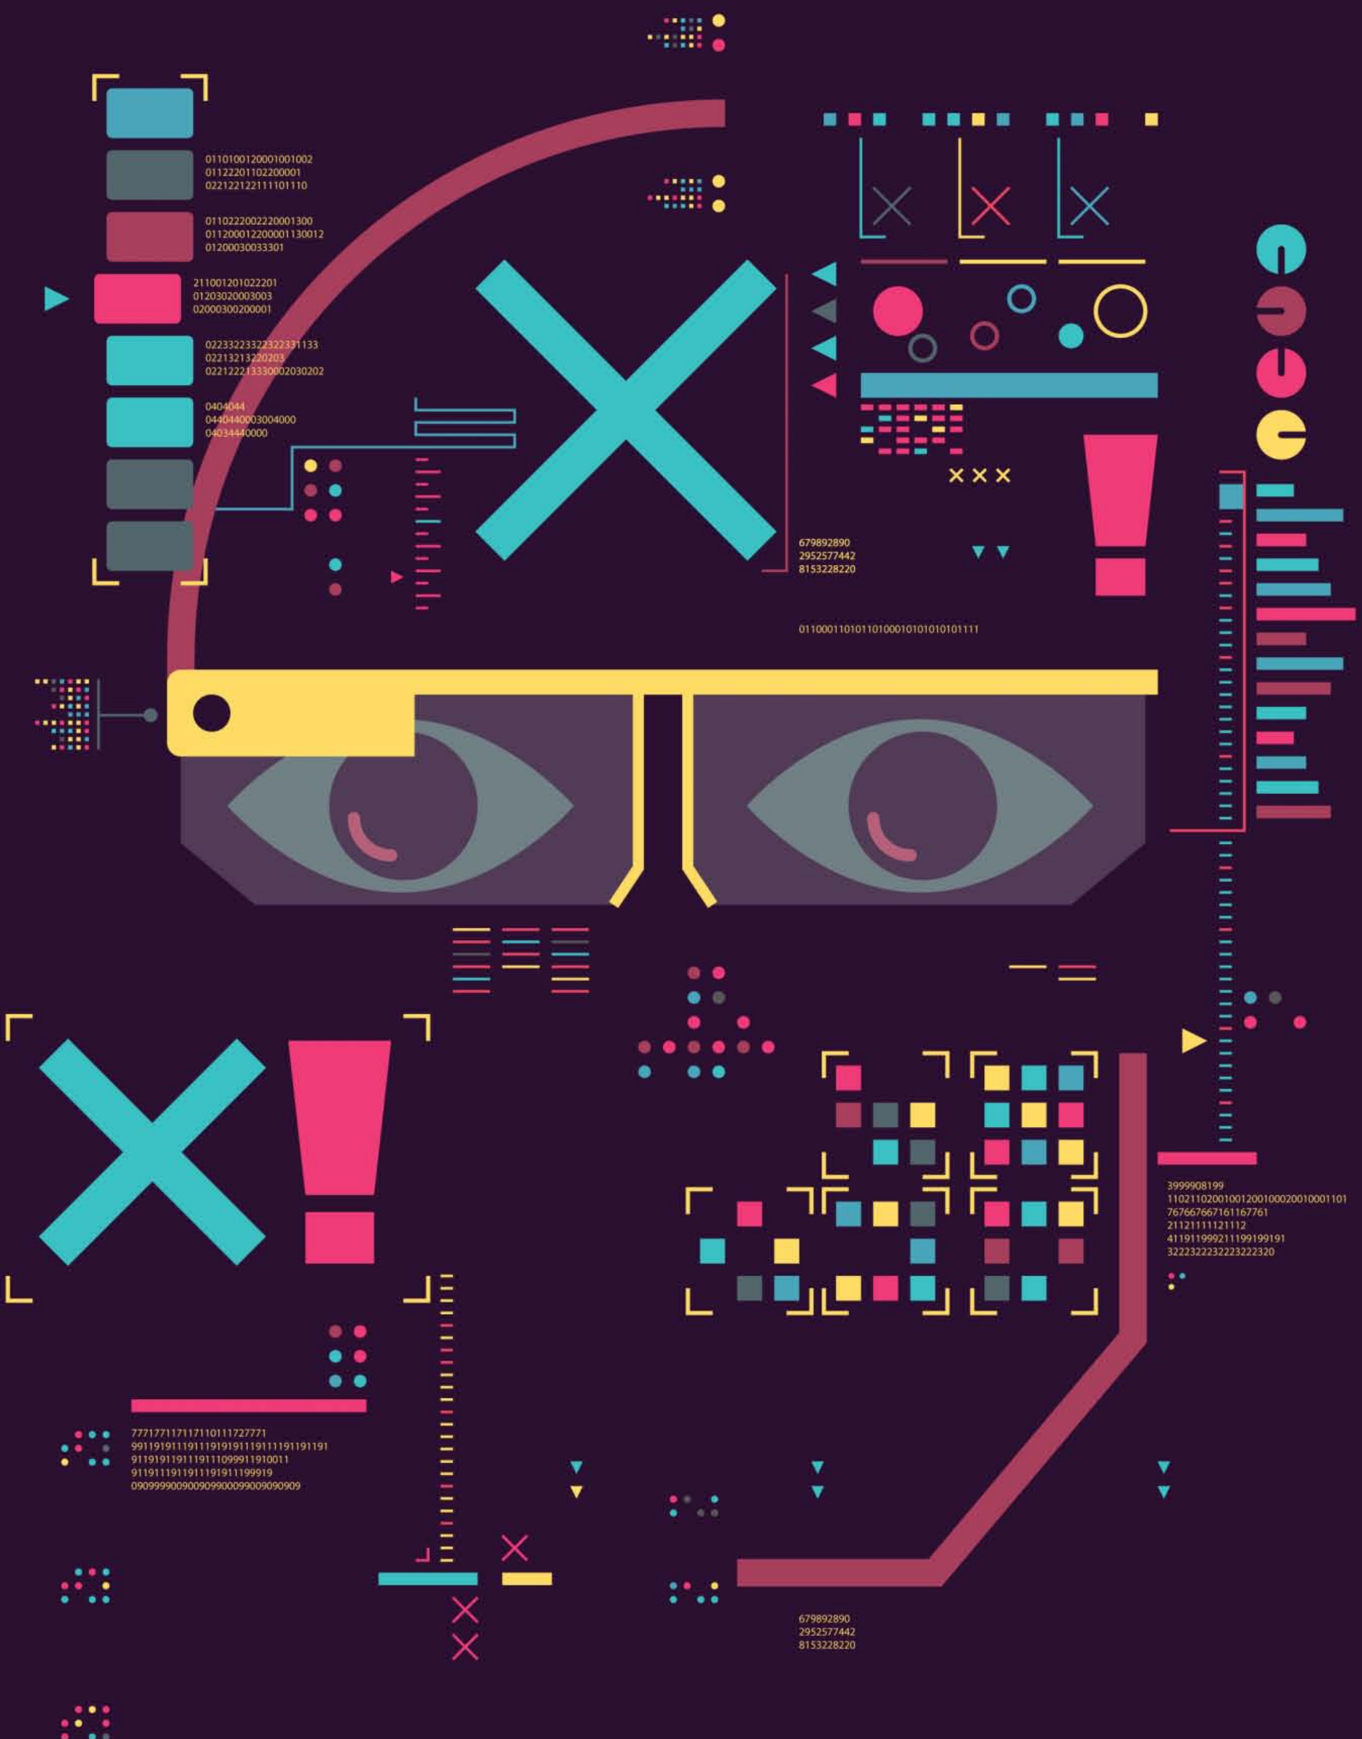

ILLUSTRATION BY Eric Frommelt

679892890
2952577442
8153228220

THE REAL-LIFE DANGERS OF AUGMENTED REALITY

Augmented reality can impair our perception, but good design can minimize the hazards

By Eric E. Sabelman & Roger Lam

YOU KNOW YOUR CELLPHONE can distract you and that you shouldn't be texting or surfing the Web while walking down a crowded street or driving a car. Augmented reality—in the form of Google Glass, Sony's SmartEyeglass, or Microsoft HoloLens—may appear to solve that problem. These devices present contextual information transparently or in a way that obscures little, seemingly letting you navigate the world safely, in the same way head-up displays enable fighter pilots to maintain situational awareness. ● But can augmented reality really deliver on that promise? We ask this question because, as researchers at Kaiser Permanente concerned with diseases that impair mobility (Sabelman) and with using technology to improve patient care (Lam), we see dangers looming. ● With augmented-reality gear barely on the market, rigorous studies of its effects on vision and mobility have yet to be done. But in reviewing the existing research on the way people perceive and interact with the world around them, we found a number of reasons to be concerned. Augmented reality can cause you to misjudge the speed of oncoming cars, underestimate your reaction time, and unintentionally ignore the hazards of navigating in the real world. And the worst thing about it: Until something bad happens, you won't know you're at greater risk of harm.

There's a simple way to fix this, of course. The GPS receivers built into wearables already detect the speed of motion (at least outdoors); designers could use them to stop notifications when the user is moving. And many AR wearables have cameras, so image analysis could likewise trigger a safety mode indoors in situations likely to cause trouble. So technically, there are straightforward solutions. But they aren't likely to be used: The last thing the people buying wearables want is to stop the flow of information—ever. The whole point of the devices is to stay connected no matter what you are doing.

So there are—and will continue to be—hazards engendered by AR wearables—at least when the user is on the move. We turned to studies conducted with visually impaired people and others using early versions of wearable AR to find out exactly what those hazards are. We also found that some of these studies suggest that augmented reality has the potential to help some people with disabilities to overcome their impairments.

Why would augmented reality be bad for you but good for a fighter pilot? After all, a head-up display, like augmented-reality glasses, overlays information on a person's view, obscuring it to some extent and potentially causing distraction. The difference is that an aircraft head-up display typically shows information in a highly symbolized and minimalistic way, with little text and no images of people (we'll talk later about why that is important). And pilots go through extensive training to be able to interpret this information quickly.

To understand how AR wearables affect the way a typical person perceives the world, we considered various natural impairments to vision. Presbyopia, farsightedness, and nearsightedness all affect your ability to focus. Diabetes, glaucoma, and retinitis pigmentosa can create tunnel vision, masking objects in the peripheral visual field. Age-related macular degeneration causes the reverse, leaving only items in peripheral vision clearly defined. A poorly designed AR interface could interfere with vision to the same degree as these diseases.


Eyeglasses with thick frames


Tunnel vision


Central vision loss


Diabetic retinopathy


Leakage into aqueous humor


Cataracts

RIGHT IN FRONT OF YOUR EYES: Augmented reality is new, but vision problems (or corrective glasses) that obscure a user's view of the world are not. AR can mimic some of these conditions.

First consider the general ability to focus. People with an impaired ability to focus either wear corrective lenses or face varying degrees of difficulty in getting around. Have you ever struggled to read a street sign in the distance, or do you find yourself driving more cautiously at night because it's harder to focus your eyes in dimmer light? Wearables like Google Glass require the ability to quickly shift focus from the real world in the distance to the images presented by the device, which are projected on the retina as if they were about 2.5 meters away. Objects at this distance should be easy for most people with decent distance vision or corrective lenses to handle without strain—but learning how to comfortably adjust focus to clearly see an AR display has a learning curve, similar to that of adapting to bifocals. And as people age, presbyopia makes it harder to rapidly change focus. In our tests of Google Glass on several hundred people, we discovered that roughly 5 to 10 percent experience so much eye strain that they gave up. They'd struggle to focus for about 20 seconds and then look away—it was just that uncomfortable. Another 25 percent of the people we tested had difficulty focusing but were able to stick with it, and eventually most were able to handle Glass.

Based on these tests—which were fairly informal and should be repeated in a broader and more rigorous study—we believe that a significant minority of Glass wearers, at least initially, will have difficulty adjusting, and during that period they will be much more distracted, with much longer reaction times, than other users.

But changing focus is not the only issue. AR wearables also obscure your vision. The loss of central vision is so obviously bad that designers of augmented-reality wearables carefully avoid blocking it—at least when the user is looking straight ahead. That's why AR displays tuck notifications off to the side. This doesn't avoid the distraction problem—the same one that makes people look at their phones while driving—because you're meant to shift your gaze, at least briefly, to check out these alerts.

PHOTO-ILLUSTRATION BY Erik Vrielink

DISTRACTION OR ENHANCEMENT?
The simple lines and numbers on the pilot's display [top] are designed to improve the ability to focus on the task at hand, not distract from it. A map displayed on Google Glass may provide helpful information, but it obscures part of the user's view of the real world.

TOP: TODD LAPPIN/FLICKR; BOTTOM: ROGER LAM

Shifting your view to the side for too long can create problems, of course. But even if you avoid the temptation to glance at a notification appearing at the edge of your vision—waiting, perhaps, until you finish crossing a street—these intrusions still present a danger.

This can be seen from a study conducted as part of the Salisbury Eye Evaluation and published in 2004. Researchers rated 1,344 adults between the ages of 72 and 92 on their ability to keep track of central and peripheral objects on a display screen. Then they had the subjects walk through an obstacle course and scored them on the number of things they bumped into, taking into account the size of their visual fields, balance, and other factors. All else being equal, a 10 percent decrease in the score on the vision test predicted a 4 percent increase in the number of collisions with obstacles. This suggests that those who couldn't keep proper track of things in their peripheral vision would be more prone to falls.

Peripheral vision is more important than you might think, because it provides a wealth of information about speed and distance from objects. Central vision, despite the great detail it offers, gives you only a rough estimate of movement toward or away from you, based on changes in size or in the parallax angle between your eyes. But objects moving within your peripheral vision stimulate photoreceptors from the center of the retina to the edge, providing much better information about the speed of motion. Your brain detects objects in your peripheral field and evaluates if and how they (or you) are moving. Interfering with this process can cause you to misjudge relative motion and could cause you to stumble; it might even get you hit by a car one day.

It's ironic, really. You buy an AR device to make you more able, yet you're likely to experience some of the same problems faced by visually impaired individuals: reduced depth of focus, distance and speed perception, and reaction time. Indeed, AR users may be at more risk than someone with a permanent vision problem, because they have developed no compensatory strategies for lost vision.

A 2008 study at Johns Hopkins University, in Baltimore, would suggest just this. Researchers used two sets of subjects—one group with long-standing retinitis pigmentosa and another with normal vision who had their peripheral vision temporarily blocked. Each person walked through a room-size virtual space with statues placed throughout it. Then they each navigated the virtual room one more time, without the objects. The second time the subjects were asked to go to the locations where the statues had been placed.

It turned out that the people with the narrowest natural fields of view overestimated the distance between themselves and the remembered objects compared with normal or less-impaired subjects. This tendency to give yourself more room, in effect, is a way of compensating for poor vision. If you know you might run into an object ahead and you're concerned about your ability to see it, you'd probably slow your steps down sooner and approach more cautiously than people confident in their ability to dodge it. Subjects who had their vision artificially narrowed, however, didn't automatically add the extra distance—they trusted their abilities, even though this trust was misplaced.

Let's talk about those fighter pilots again. They need to pay full attention to what is in front of them and to be able to judge the speed and distance of objects

Clearing the FOG

Freezing of gait (FOG), a sudden inability to take that next step, puts Parkinson's disease patients at risk of falls. Visual cues can, however, prevent FOG in some patients. At the Department of Veterans Affairs Rehabilitation R&D Center in Palo Alto, Calif., researchers used motion sensors and LEDs attached to eyeglasses and connected to a wearable computer to determine if augmented reality could help prevent FOG.

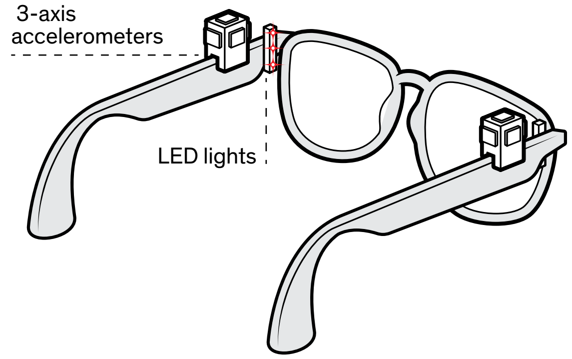

1. A SUBJECT WITH PARKINSON'S DISEASE walks down a 10-meter test track. Sensors on eyeglasses measure head movement, and sensors on a belt measure body movement.

2. THE SYSTEM DETECTS an abortive forward movement at the waist combined with the head-down tilt that is an early sign of FOG.

3. IT THEN TRIGGERS LEDS on the eyeglass frames to alternately blink at the rate the subject had previously been walking. This visual clue helps get the subject out of FOG faster.

4. TURNING CAN INCREASE the risk of FOG. If the system detects FOG during a turn, it blinks the LEDs on the outside of the turn faster than the ones on the inside, augmenting the motion of the visual field.

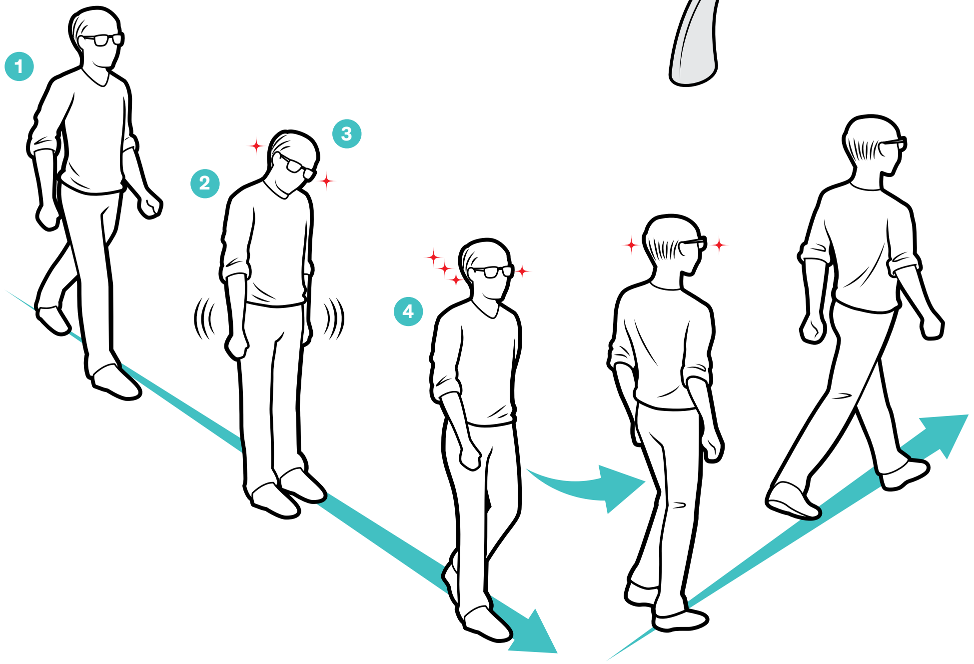

around them accurately. Head-up displays don't interfere with the ability of peripheral vision to help in this task because they present information in the central field. This works only because that information is stripped down to essential lines and symbols.

Newer AR wearables—like Microsoft HoloLens and, reportedly, the system being developed by stealth startup Magic Leap, in Dania Beach, Fla.—appear to be using the central visual field to display objects integrated into the real world for the most part, not relegating them to the peripheral vision. It appears, however, that these will be detailed, full-color, realistic objects, not simple graphics. And that brings up another concern.

AR is really engaging, and designers are likely to make it more and more so, with sophisticated graphics that go far beyond the simple lines and numbers that appear in a fighter pilot's display. You will likely still be able to see the real world through the projected image. But our innate neural wiring prefers images of people to objects—and that is the case even if the people are virtual and the objects are real. So if you're looking at something not human in the real world, but the AR image includes people (or even simple shapes resembling the human form) AR will win the struggle for attention.

We know about this phenomenon through experiments conducted since the 1970s with dot patterns that resemble human stick figures. Researchers call these point-light walkers. These patterns of 11 to 15 points illustrate how our perception is biased to recognize human forms with minimal prompting. With just a dozen or so dots, we can identify a figure as a walking human within 200 milliseconds—even though it takes a person between 1 and 2 seconds to take two steps.

The point-light walker studies demonstrate just how easy it is for an image to grab someone's attention. And you can bet that content providers are going to try to make attention-grabbing apps. The more apps include humanlike shapes, the more users will focus on them, disregarding the real world and increasing the hazards for the users.

We understand the effects of humanlike images presented through augmented reality. We know less about other, nonhuman images that may also be problematic. At Kaiser Permanente,

ILLUSTRATION BY MCKIBILLO

we have been testing virtual images for use in controlled-exposure therapy, intended to help people overcome certain fears. Some of those fears—like a fear of spiders or other living things—can be triggered by very-low-resolution images. Designers of aviation head-up displays don't use biological forms to present information, but the people creating AR apps might well do that. So for a small percentage of the population, these apps could trigger surprisingly negative responses.

The news is not all bad, though. Used properly, augmented reality can help those who already have difficulty navigating the world. This is possible because an AR device such as Google Glass carries sensors—at the most basic, a camera and an accelerometer—that can allow the system to track the user's surroundings and movement in real time and generate cues that improve the person's safety.

To aid blind individuals, researchers at RMIT University, in Melbourne, Australia, are developing a speeded-up robust features (SURF) algorithm to allow AR devices to recognize traffic lights and pedestrian signals and to predict potential collisions with moving people and objects from video data. To date, the algorithm can recognize 90 percent of the people in view and 80 percent of the objects. If installed in an AR device, along with appropriate audible output—that doesn't mask the real-world sounds essential to navigation without vision—this technology could be of great benefit to blind users. Even though these people would not need the visual display AR glasses provide, such wearables are still a fine choice for this application because the form factor is suitable, the cost of the mass-produced hardware is relatively low, and the camera is perfectly positioned to "see" what's ahead.

AR wearables can also help patients with Parkinson's disease, which causes tremors and stiffness. These people often have a perplexing characteristic known as motor block or freezing of gait (FOG)—the sudden inability to initiate a step or stride—putting them at risk for serious falls. FOG occurs more frequently when someone is coming to a turn, changing direction, maneuvering in a closed space, or passing through a doorway.

Some Parkinson's disease patients with FOG can walk with a nearly normal stride when presented with the proper visual cues. In the 1990s, podiatrist Thomas Riess, of San Anselmo, Calif., who himself has Parkinson's, developed several early augmented-reality devices for superimposing such cues. He patented a head-mounted array of LEDs that projects a scrolling ladder of bars on a transparent screen in front of the patient. In tests, Riess's device was able to improve his own and other Parkinson's disease patients' ability to walk without freezing.

In 2002, at the Department of Veterans Affairs Rehabilitation R&D Center in Palo Alto, Calif., with Riess's cooperation, one of us (Sabelman) also conducted a test to determine whether computer-generated cues from wearables could reduce the time Parkinson's patients spent in FOG. In this experiment, the virtual cues were presented only when needed. To do so, Sabelman placed LEDs and 3-axis accelerometers on the corners of eyeglass frames to measure head motion and to flash lights when the computer detected the onset of FOG. This method takes advantage of the fact that early in a FOG state patients tilt their heads forward to visually confirm the position of their feet.

The computer used the head-tilt angle and other sensor inputs to determine the intended motion, at which point it flashed the lights to simulate continuing that motion. If the system determined that a patient froze while trying to turn left, for example, the lights on the left side would flash more slowly than those on the right, because when you are turning you see more motion on the outside of your turn than on the inside. If a patient froze while moving forward, the LEDs would alternately flash in what had been the rhythm of the patient's steps before the freeze. In a 10-meter walking task, this system reduced FOG by almost 30 percent. We think that Google Glass, or newer symmetric AR devices with displays for the left and right eyes, paired with significantly faster processors, could do much better.

POST YOUR COMMENTS at http://spectrum.ieee.org/ar-dangers0715

Elsewhere, researchers at Brunel University in London are developing systems that project lines in front of patients. In their current state of development, these are clumsy and difficult to use, but they show the power of AR to solve this particular problem.

The message of all this research is clear. Because AR hardware and applications can impair their users' vision, their designers owe it to the world to be careful. What's more, they need to test their products on people of all ages and physical abilities. They need to evaluate users' reaction times while their apps are running. They need to put wearers through real-world obstacle courses with opportunities to stumble. And they need to determine just how much content you can present via an AR app before it becomes dangerous.

The manufacturers also need to educate the people who buy and use these devices of the hazards, perhaps using mandatory training games so they will learn the hazards before stepping out into the real world. After all, when we get our hands on new gadgets, our tendency is to start using them right away and to read the instructions later—if at all. But new users of wearable AR devices should start slowly—perhaps practicing at home, or in a park or other safe environment. That way users would know there are risks and accept them—or not—consciously.

Software for augmented-reality and virtual-reality wearables is emerging that includes cognitive assessment tools, so any app developer could, if he or she desired, build in diagnostics to measure things like reaction time and balance. With these tools, apps could guide users safely through the initial training and make them understand the dangers. The future of this technology will hinge on whether such improved tools are used to rigorously test new products. ■

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

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

Joint Institute of Engineering

Faculty Positions available in Electrical and Computer Engineering

Sun Yat-sen University and **Carnegie Mellon University** have established the **SYSU-CMU Joint Institute of Engineering (JIE)** as a conduit for innovative engineering education and research. Our mission is to nurture a passionate and collaborative global community and network of students, faculty and professionals advancing the field of engineering through education and research.

The JIE enrolled its first cohort of dual-degree M.S. and Ph.D. students in Electrical and Computer Engineering in fall 2014. All current JIE faculty members have been recruited worldwide and we continue to seek **full-time tenure-track faculty** in all areas of electrical and computer engineering. Candidates should have a doctoral degree in electrical and computer engineering, computer science or related areas, with a demonstrated record of or potential for research, teaching and leadership. The position includes an initial year at Carnegie Mellon University in Pittsburgh to establish educational and research collaborations before relocating to Guangzhou, China.

This is a worldwide search open to qualified candidates of all nationalities. We offer an internationally competitive compensation package.

Please visit jie.cmu.edu for details and to apply online.

SHUNDE INTERNATIONAL

Joint Research Institute

Research positions available in Electrical and Computer Engineering

SYSU-CMU Shunde International Joint Research Institute (JRI) is located in Shunde, China. Supported by the provincial government and industry, JRI aims to form high-level teams of innovation, research and development, transfer research outcomes into products, develop advanced technology, promote industrial development and facilitate China's transition from labor-intensive industries to technology-intensive and creative industries.

JRI is seeking **full-time research faculty and research staff** who have an interest in the industrialization of science research, targeting electrical and computer engineering or related areas.

Candidates with industrial experiences are preferred. Application review will continue until the position is filled.

Applicants should include a full CV, three to five professional references, a statement of research and teaching interests and copies of up to five research papers.

Email applications or questions to sdjri@mail.sysu.edu.cn.

SUN YAT-SEN UNIVERSITY Carnegie Mellon University

AND NOW THIS...

OUR PERFORATED PAST

The punch card's centuries-long reign

Paper punch cards, each representing a line of code or data, were how programmers got information into computers from the 1950s on. The concept, though, is a lot older. In 1884 Herman Hollerith borrowed the idea from the Jacquard loom and applied it to his U.S. census machines. Hollerith later founded a company that became IBM, and punch cards became known as IBM cards. Compilers, subroutines, and programs were coded into thick decks of cards, which were then fed through card readers like this ITT Cannon machine, featured in a March 1964 *IEEE Spectrum* ad. Pity the poor programmer who dropped her deck.

Decades after other forms of data storage came along, punch cards held on in highway toll collection, warehouse inventory, and payroll systems. Cardamation Co., the last U.S. company to make punch card equipment, quietly closed in 2012. The technology's last moment in the public eye came during the 2000 U.S. presidential election, when incompletely punched ballot cards brought the "hanging chad" into the zeitgeist.
—ROBERT COLBURN, IEEE HISTORY CENTER

PHOTOGRAPH BY Randi Klett


MATLAB speaks Arduino
so you don't have to
You can design, build, test and run a system–
on Arduino, Raspberry Pi, LEGO and more–
without writing traditional code.
Download free MATLAB and Simulink
hardware support packages at
hardware.mathworks.com
MATLAB® & SIMULINK®
©2015 The MathWorks, Inc
MathWorks®