IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 10.13

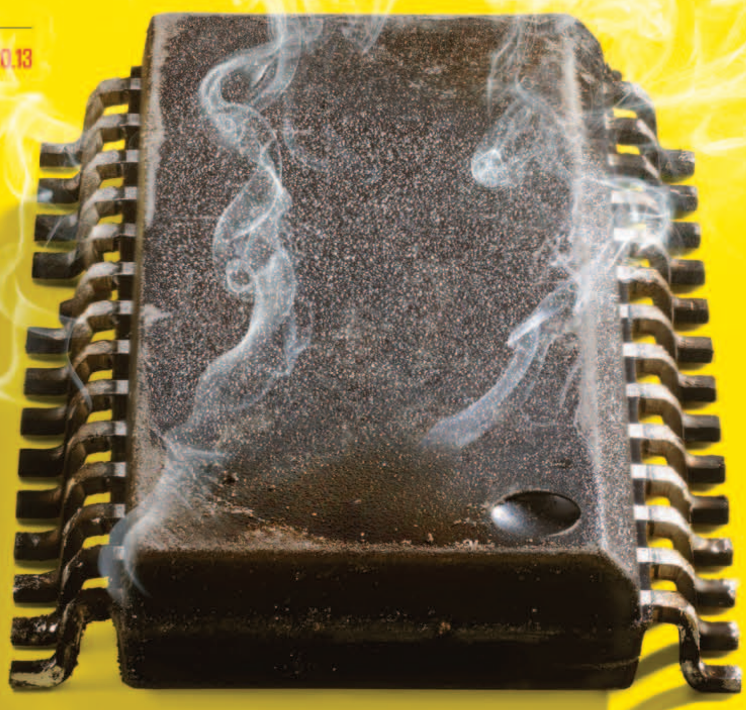

FAKE

Scavenged chips passed off as new are cropping up in U.S. military equipment. Are they in passenger jets and cars, too? P. 36

IEEE

FEATURES_10.13

IEEE SPECTRUM

SEBASTIAN SCHERER

On the Cover **Photograph for *IEEE Spectrum* by Adam Voorhes**

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

Infinite Designs, One Platform

with the only complete system design environment

NI LabVIEW is the only comprehensive development environment with the unprecedented hardware integration and wide-ranging compatibility you need to meet any measurement and control application challenge. And LabVIEW is at the heart of the graphical system design approach, which uses an open platform of productive software and reconfigurable hardware to accelerate the development of your system.

LabVIEW system design software offers unrivaled hardware integration and helps you program the way you think–graphically.

>> Accelerate your system design productivity at ni.com/labview-platform

800 453 6202

NATIONAL INSTRUMENTS™

DEPARTMENTS_10.13

IEEE SPECTRUM

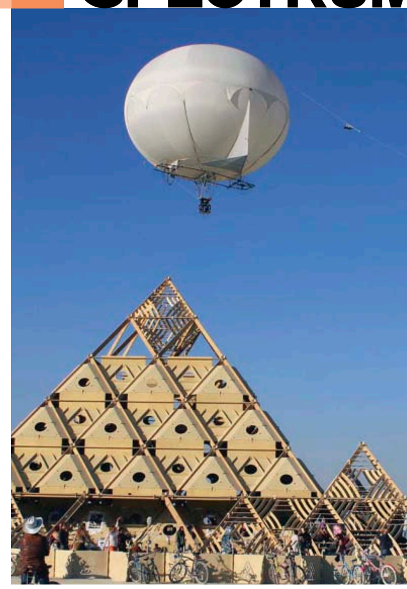

FROM LEFT: GOOGLE; SWINGTIP; SUSAN KARLIN

ADDITIONAL RESOURCES

Webinars & Events

Available at spectrum.ieee.org/webinar

- High-Speed Serial Link: Full-Wave EM Modeling Methodology and Measurement Correlation—1 October
- Acoustics Simulation in Microphones, Mobile Devices, and Hearing Aids—3 October
- Electromagnetic Compatibility (EMC) Design and Verification for Embedded Systems—9 October
- Democratization of Embedded System Design—17 October
- Electromagnetic Simulation of Composite Materials and Cable Harnesses in Aircraft—31 October
- **MASTER BOND WHITE PAPER LIBRARY** http://spectrum.ieee.org/static/masterbond-whitepaper-library
- **IBM SYSTEMS ENGINEERING RESOURCE LIBRARY** http://spectrum.ieee.org/static/ibm-systems-engineering-resource-library

The Institute

Available 7 October at theinstitute.ieee.org

IEEE INVESTIGATES THE WEB 3.0 The next iteration of the Internet is here. Known as the Internet of Things, it's expected to link approximately 50 billion machines and devices by 2020. IEEE is tackling the topic in 2014 with a new journal and conference.

NEXT-GENERATION BROADBAND The IEEE 802.3 Ethernet Working Group is pursuing a new physical layer standard that will enable cable operators to offer an end-to-end Ethernet network capable of efficiently supporting next-generation services.

HISTORY OF RAILROAD SIGNALS The invention of the train presented a new challenge for engineers: How could they contact conductors when trains traveled faster than the existing communication technology? Read about how the telegraph and other innovations revolutionized train travel.

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2013 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 50, issue No. 10, 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 American Business Media, the Magazine Publishers of America, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

BACK STORY_

Fore!

SOMETIMES unusual roadblocks crop up when you test consumer electronics, as Senior Editor Joshua J. Romero found out when *IEEE Spectrum* got the chance to review the SwingTip golf training aid. The task naturally fell to Romero, who's not an avid golfer but enjoys a round when he can. And so we packed him off with his clubs and the new gadget to the busy driving range at New York City's Chelsea Piers athletics complex [see "Gadgets for Golfers," in this issue].

The SwingTip is a small black box that clips onto the shaft of a golf club. The box transmits information about a golfer's swing to an app running on a nearby iOS device. Setting himself up with an iPad on the second of the range's four levels, Romero ran through a quick series of 10 swings, aided by the range's automatic teeing mechanism. But while stepping back through the data on the iPad, he found that only four swings had been detected. The last of these showed an animation in which the club seemed to go through a bizarre set of gyrations, with a peak recorded speed of 204 kilometers per hour.

Checking his club, Romero realized that the SwingTip had detached midswing. However, it had continued to faithfully transmit information during its subsequent trajectory, while the software did its best to map the ballistic flight data to a golfer's motions. Sure enough, when Romero peered down at the ground far below, he saw "a little black box among hundreds of golf balls."

It might have been the work of a moment to retrieve a detached device on a golf course or on a less busy driving range. But the Chelsea Piers range operates from early morning to midnight, and not even employees dare venture out amid the nonstop hail of balls. The SwingTip was never recovered. Romero believes it probably got chewed up by the range's armored ball-collection truck. We subsequently contacted the SwingTip's manufacturer, which kindly supplied another unit so Romero could complete the testing for his article. ■

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, Dataflow is in *IEEE Spectrum*, Vol. 50, no. 10 (INT), October 2013, p. 68, or in *IEEE Spectrum*, Vol. 50, no. 10 (NA), October 2013, p. 80.

10.13

PHOTOGRAPH BY Randi Silberman Klett

CONTRIBUTORS_

Lyle Chamberlain

Since graduating from Caltech in 2005, Chamberlain has been developing aircraft sensors that would do away with pilots, though he himself flies sailplanes. Together with Sebastian Scherer, a researcher at Carnegie Mellon University, he founded Near Earth Autonomy, which is helping the U.S. Navy develop an autonomous flight-control package for helicopters. In "Robocopters to the Rescue" [p. 24], they describe the testing stages of one such system.

Benjamin Gross

Gross, a research fellow at the Chemical Heritage Foundation, in Philadelphia, wrote about RCA's attempts to commercialize the liquid-crystal display in *IEEE Spectrum*'s November 2012 issue. He returns this month to describe artifacts from the Sarnoff Collection [p. 42], objects that "illustrate the evolution of telecommunications through the 20th century," he says. "The Sarnoff Collection is one of the great treasure troves of American electronics history."

Suzanne Kantak

Kantak is a photographer based in Brooklyn, N.Y., who enjoys working on still lifes, "where the objects can also tell a story," she says. Shooting the Sarnoff Collection [p. 42] provided many such opportunities. She gave a Russian Constructivist treatment to a plastic phonograph intended to play U.S. Cold War propaganda records. Using a style suggestive of the opposing ideology, Kantak says, "shows off the record player's quirkiness."

Adrian Sampson

Lead author of "Good-Enough Computing" [p. 50], Sampson studies computer science and engineering at the University of Washington, in Seattle, where his two coauthors, Luis Ceze and Dan Grossman, are his Ph.D. advisors. One specializes in programming languages, the other in computer architecture. Sampson likens their often disparate guidance to "having an angel perched on one shoulder and the devil on the other," though he won't say which is which.

Alan Seabaugh

In "The Tunneling Transistor" [p. 30], Seabaugh describes a new kind of transistor that uses quantum tunneling for switching. An IEEE Fellow and a professor of electrical engineering at the University of Notre Dame, Seabaugh entered the field in the late 1980s, when he worked at Texas Instruments. "The aim was to leapfrog Moore's Law, but we kept getting run over by silicon technology," he says. "Today, the rules are changing."

John Villasenor

Over the past 20 years, Villasenor, a professor at UCLA and senior fellow at the Brookings Institution, has seen globalization benefit the electronics industry. He's also witnessed an unwelcome by-product: the rise in counterfeit chips, an issue he and coauthor Mohammad Tehranipoor discuss in "Chop-Shop Electronics" [p. 36]. Nowadays, he says, the hardware supply chain is "even more opaque than I and many other people suspected."

10.13

IEEE

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

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

IEEE STAFF EXECUTIVE, PUBLICATIONS Anthony Durniak

IEEE BOARD OF DIRECTORS
PRESIDENT Peter W. Staecker, president@ieee.org
+1 732 562 3928 FAX: +1 732 465 6444

PRESIDENT-ELECT Roberto de Marca
TREASURER John T. Barr **SECRETARY** Marko Delimar
PAST PRESIDENT Gordon W. Day

VICE PRESIDENTS
Michael R. Lightner, Educational Activities; Gianluca Setti, Publication Services & Products; Ralph M. Ford, Member & Geographic Activities; Karen Bartleson, President, Standards Association; Robert E. Hebner, Technical Activities; Marc T. Apter, President, IEEE-USA

DIVISION DIRECTORS
Cor L. Claeys (I); Jerry L. Hudgins (II); Douglas N. Zuckerman (III); Jozef Modelski (IV); James W. Moore (V); Bogdan M. Wilamowski (VI); Cheryl "Cheri" A. Warren (VII); Roger U. Fujii (VIII); Jose M. Moura (IX); Stephen Yurkovich (X)

REGION DIRECTORS
Peter Alan Eckstein (1); Parviz Famouri (2); David G. Green (3); Karen S. Pedersen (4); James A. Jefferies (5); Michael R. Andrews (6); Keith B. Brown (7); Martin J. Bastiaans (8); Gustavo A. Giannattasio (9); Toshio Fukuda (10)

DIRECTORS EMERITUS Eric Herz, Theodore W. Hissey

IEEE STAFF
EXECUTIVE DIRECTOR & COO James Prendergast +1 732 502 5400, james.prendergast@ieee.org
HUMAN RESOURCES Betsy Davis, SPHR +1 732 465 6434, e.davis@ieee.org
PUBLICATIONS Anthony Durniak +1 732 562 3998, a.durniak@ieee.org
EDUCATIONAL ACTIVITIES Douglas Gorham +1 732 562 5483, d.g.gorham@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski +1 732 562 5504, c.jankowski@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios +1 732 562 3820, constantin@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER Eileen Lach, +1 212 705 8990, e.m.lach@ieee.org
CORPORATE STRATEGY Matthew Loeb, CAE +1 732 562 5320, m.loeb@ieee.org
CHIEF MARKETING OFFICER Patrick D. Mahoney +1 732 562 5596, p.mahoney@ieee.org
CHIEF INFORMATION OFFICER Alexander J. Pasik, Ph.D. +1 732 562 6017, a.pasik@ieee.org
CHIEF FINANCIAL OFFICER Thomas R. Siegert +1 732 562 6843, t.siegert@ieee.org
TECHNICAL ACTIVITIES Mary Ward-Callan +1 732 562 3850, m.ward-callan@ieee.org
MANAGING DIRECTOR, IEEE-USA Chris Brantley +1 202 530 8349, c.brantley@ieee.org

IEEE PUBLICATION SERVICES & PRODUCTS BOARD
Gianluca Setti, *Chair;* John B. Anderson, Robert L. Anderson, John Baillieul, Silvio E. Barbin, Herbert S. Bennett, Don C. Bramlett, Stuart Bottom, Thomas M. Conte, Samir M. El-Ghazaly, Sheila S. Hemami, Lawrence O. Hall, David A. Hodges, Donna L. Hudson, Elizabeth T. Johnston, Hulya Kirkici, Khaled Letaief, Carmen S. Menoni, William W. Moses, Michael Pecht, Vincenzo Piuri, Sorel Reisman, Jon G. Rokne, Curtis A. Siller, Ravi M. Todi, H. Joel Trussell, Leung Tsang, Timothy T. Wong

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

PHOTO-ILLUSTRATIONS BY **Gluekit**

COLIN ANDERSON/GETTY IMAGES

Rules for the Digital Panopticon

The technologies of persistent surveillance can protect us only if certain boundaries are respected

For centuries, we humans have lacked the all-knowing, all-seeing mechanisms to credibly predict and prevent bad actions by others. Now these very powers of preemption are perhaps within our grasp, thanks to a confluence of technologies.

In the foreseeable future, governments, and perhaps some for-profit corporations and civil-society groups, will design, construct, and deploy surveillance systems that aim to predict and prevent bad actions—and to identify, track, and neutralize people who commit them.

And when contemplating these systems, let's broadly agree that we should prevent the slaughter of children at school and the abduction, rape, and imprisonment of women. And let's also agree that we should thwart lethal attacks against lawful government.

Of late, the U. S. government gets most of the attention in this arena, and for good reason. The National Security Agency, through its vast capacity to track virtually every phone call, e-mail, and text message, promises new forms of preemption through a system security experts call *persistent surveillance*.

The Boston Marathon bombing, in April, reinforced the impression that guaranteed prevention against unwanted harm is elusive, if not impossible. Yet the mere chance of stopping the next mass shooting or terror attack persuades many people of the benefits of creating a high-tech version of the omniscient surveillance construct that, in 1787, the British philosopher Jeremy Bentham conceived as a *panopticon*: a prison with a central viewing station for watching all the inmates at once.

Some activists complain about the potential of such a system to violate basic freedoms, including the right to privacy. But others will be seduced by the lure of techno fixes. For example, how could anyone object to a digital net that protects a school from abusive predators?

Ad hoc surveillance will inevitably proliferate. Dropcam and other cheap surveillance programs, already popular among the tech-savvy, will spread widely. DIY and vigilante panopticons will complicate matters. Imagine someone like George Zimmerman, the Florida neighborhood watchman, equipped not with a gun but with a digital surveillance net, allowing him to track pretty much anything—on his smartphone.

With data multiplying exponentially and technology inexorably advancing, the question is not whether an all-encompassing surveillance systems will be deployed. The question is how, when, and how many.

In the absence of settled laws and norms, the role of engineers looms large. They will shoulder much of the burden of designing systems in ways that limit the damage to innocents while maximizing the pressures brought to bear on bad guys.

But where do the responsibilities of engineers begin and end?

It is too early to answer conclusively, but engineers would do well to keep a few fundamental principles in mind:

1. Keep humans in the loop, but insist they follow the "rules of the road." Compiling and analyzing data can be done by machines. But it would be best to design these surveillance systems so that a human reviews and ponders the data before any irreversible actions are taken. If citizens want to spy on one another, as they inevitably will, impose binding rules on how they do so.

2. Design self-correcting systems that eject tainted or wrong information fast and inexpensively. Create a professional ethos and explicit standards of behavior for engineers, code writers, and designers who contribute significantly to the creation of panopticon-like systems.

3. Delete the old stuff routinely. Systems should mainly contain real-time data. They should not become archives tracing the lives of innocents.

Engineers acting responsibly are no guarantee that panopticons will not come to control us. But they can be part of getting this brave new world right.
—G. Pascal Zachary

G. Pascal Zachary is the author of *Endless Frontier: Vannevar Bush, Engineer of the American Century* (Free Press, 1997). He teaches at Arizona State University.

NEWS

6.1 PERCENT: THE PENETRATION RATE IN 2013 OF FIXED BROADBAND SERVICE IN DEVELOPING COUNTRIES

GOOGLE

FIVE WAYS TO BRING BROADBAND TO THE BACKWOODS

These hopeful schemes aim to deliver high-speed data service to remote places

BROADBAND BY BALLOON: It may seem loony, but in June Google engineers tested the delivery of wireless broadband by high-altitude balloon in New Zealand.

Bringing broadband to rural communities and remote outposts has always been a problem. Telecom companies simply can't afford to string cables kilometers across badlands and through jungles to sign up just a few hundred more customers. Traditional cellular and satellite systems are similarly constrained by the cost of equipment as well as data capacity.

Without high-speed Internet access, billions of people worldwide continue to miss out on opportunities for a higher standard of living that includes online commerce, distance learning, and telemedicine. But if broadband services are slow to reach the hinterlands, it isn't for lack of ideas.

Here are five of the most innovative new efforts to make broadband available and affordable everywhere—including locations as diverse as ships at sea and mountain villages in Myanmar. Some of them, such as Microsoft's experiments with "white space" spectrum in Africa, stem from familiar solutions that are now starting to gain momentum. Others, such as Google's vision of globe-trotting balloons or Titan Aerospace Corp.'s aspirations to fly satellite-like drones, are more ambitious. But they just might prove crazy enough to work. —ARIEL BLEICHER

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

LOW-ORBITING SATELLITES

In June, a rocket launched from French Guiana shot four new satellites into orbit. Owned and operated by O3b Networks, based on the island of Jersey, the spacecraft are the first in a planned constellation that the company says will provide cheap, fast broadband to the planet's most poorly connected communities.

The O3b orbiters are unlike other broadband satellites. Today's high-bandwidth satellites occupy geostationary orbits about 36 000 kilometers above Earth, and while the availability of wider spectrum bands now enables some providers to offer data rates and prices comparable to those of cable service, coverage is limited and latency is high. To solve this problem, O3b–short for "the other 3 billion"–will operate its satellites in mid-Earth orbit, about one-fourth the altitude of the geostationary orbiters. This proximity means that more satellites are needed to guarantee coverage of a region, but the company says it will reduce round-trip data delays from an average of 638 milliseconds to less than 150 ms.

When this article went to press in September, O3b had plans to launch another four satellites by the end of the month. Each orbiter will support a total capacity of 12 gigabits per second (6 Gb/s in each direction). The company plans to start selling broadband wholesale to local Internet service providers in Africa, Asia, Latin America, and the Middle East.

METAMATERIAL ANTENNAS

Even as prices for fixed-satellite broadband drop, customers on mobile devices out of the range of cellular towers and Wi-Fi hot spots–such as on ships or in disaster areas–still face steep fees for satellite service. A large part of the cost covers the phased arrays or other equipment needed to keep a moving antenna locked on to a satellite.

But engineers at Kymeta Corp., based in Redmond, Wash., think they can make such beam steering a lot cheaper using antenna arrays made of metamaterial elements. Metamaterials are synthetic substances that can bend electromagnetic waves in ways that natural materials can't, making them prime candidates for such sci-fi-esque applications as cloaking devices. Kymeta claims to have come up with a proprietary technique for dynamically tuning elements in an array to emit radiation in different directions. The resulting interference produces a beam that can track a satellite, creating an unbroken broadband link.

Because metamaterial elements can be printed using standard photolithography, Kymeta's antennas could be made on the cheap. At the moment, engineers are still perfecting a prototype, but in May the company reported the first successful download from a broadcast satellite. And last month, Kymeta and O3b Networks began developing antennas together for the latter company's satellite broadband system.

BROADBAND BALLOONS

Acknowledging that the scheme sounded a bit bonkers, Google engineers named it Project Loon. In their vision, miners, ranchers, and seamen would stream videos and surf the Net via thousands of high-pressure balloons circling the globe. Chief technical lead Rich DeVaul pitched the idea two years ago as a member of Google's secretive X lab, the origin of self-driving cars and Project Glass. But Google didn't unveil the project until June, when it tested 30 balloons in a small farming community in New Zealand.

Filled with helium, the transparent balloons surf the winds of the stratosphere, at twice the elevation of passenger planes. Each one carries a solar panel and batteries, antennas, and computers. Together, they form a mesh network, relaying signals until the data reaches a base station on the ground with a fiber link to the Internet. Google engineers say one balloon could serve hundreds of people within a 20-km radius. Data speeds, they add, would be at least as fast as those of 3G cellular systems.

Google may have the technical chops to pull off DeVaul's plan, but the political hurdles will be steep. Simply getting permission to fly balloons in the airspaces of many countries will be no small feat.

SOLAR-POWERED DRONES

Another company betting on airborne broadband is Titan Aerospace Corp. The New Mexico-based start-up is building unmanned solar-powered planes that it says could soar through the stratosphere for up to five years, serving as inexpensive, upgradable alternatives to conventional satellites. The company unveiled the first prototypes in August and is now building a commercial product: a lightweight 104-km/h drone called the Solara 50.

Boasting a wingspan as long as an Olympic swimming pool, the plane is enveloped in more than 3000 photovoltaic cells. As a broadband link, it could provide coverage over about 17 000 square kilometers, an area equivalent to the reach of more than 100 cellular towers. Titan says it aims to sell the Solara for around US $1 million and already has customers lined up to buy the first three in early 2014.

TELEVISION "WHITE SPACES"

Internet firms and entrepreneurs have long touted unused television frequencies, or white space, as a solution to rural broadband. Signals in this spectrum easily penetrate trees, buildings, and other obstacles that block traditional Wi-Fi signals, and remote communities tend to have a lot of frequencies freely available.

Only recently, however, have regulators started opening up white space for commercial use. As a result, pilot broadband networks using these bands are now beginning to pop up in the United States and elsewhere. In April, California-based Carlson Wireless Technologies rolled out base stations among the mountain towns near Lake Tahoe. Also this year, Microsoft set up trials in Kenya, South Africa, and Tanzania.

Because white space spectrum is unlicensed, companies can deploy networking technologies quickly and cheaply. Microsoft, for instance, estimates that it could deliver white space broadband in Africa to each subscriber for $5 or less per month–about one-tenth the price for fixed broadband.


BROADBAND BALLOONS


LOW-ORBITING SATELLITES


SOLAR-POWERED DRONES

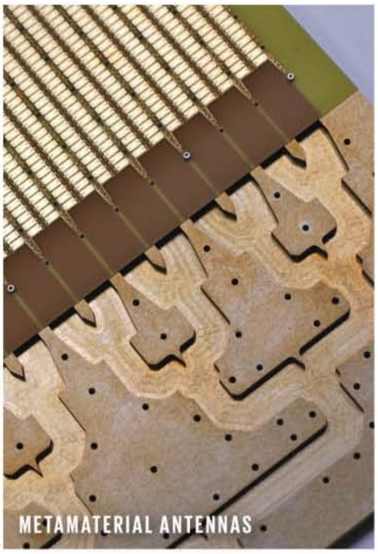

NEWS

SECRET SOFTWARE

A new form of encryption could make practically unhackable code

Software can hold invaluable secrets within orderly lines of code. Algorithms now predict which Amazon.com products you're likely to buy next, whether an early movie script will be a box-office hit, and even whether or not a legal case will go to trial. Naturally, coders of such software don't want outsiders to be able to reverse engineer the programs and learn their secret formulas. Now, computer scientists from the University of California, Los Angeles, IBM Research, and the University of Texas at Austin have begun to pave the way toward eliminating that threat.

The researchers say they've developed a "mathematical obfuscation" scheme to encrypt such valuable software. They hope that one day this scheme will allow users to run programs normally while transforming the underlying code into math puzzles that would take hundreds of years to solve.

To see how software obfuscation works, you need to understand how programs ordinarily operate. People first write programs in languages that humans can understand, and then a compiler translates that script into machine code—the instructions the CPU can execute. The program can then receive inputs and produce the proper outputs.

In the researchers' obfuscation scheme, the compiler first translates code into an intermediate form; then an obfuscating compiler translates this form into pieces of what UCLA computer science professor Amit Sahai calls a "mathematical jigsaw puzzle." A special jigsaw verifier program, written in machine code, takes these puzzle pieces—essentially sequences of random numbers—and all the inputs and tries to assemble them. If the pieces "fit," then the completed "puzzle" will tell the CPU how to produce the correct output.

If the pieces don't fit—for instance, because a hacker has altered the code in an attempt to figure out how the software works—the resulting output is useless. "The modified software would not give

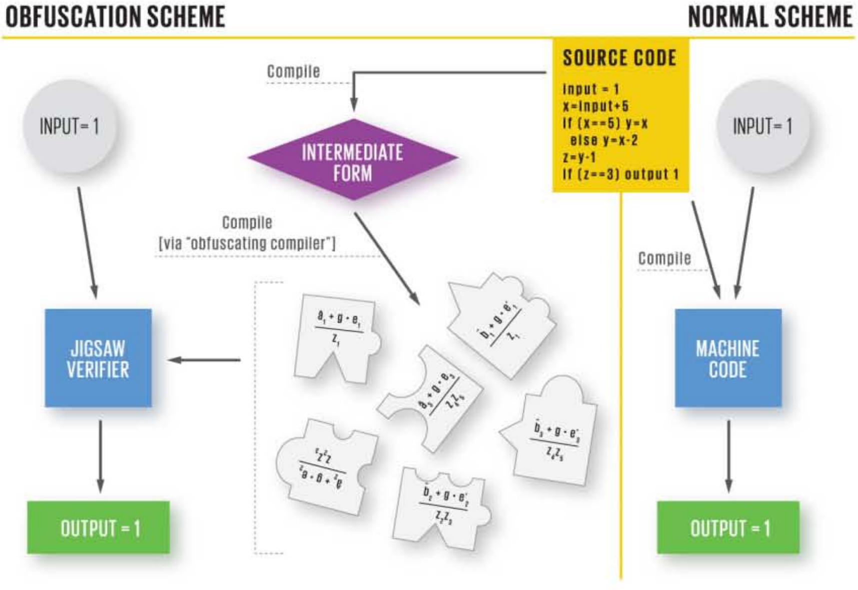

PUZZLE PROCESSING: Indistinguishability obfuscation hides the inner workings of software by turning it into a "mathematical jigsaw puzzle" that's nearly impossible for an attacker to solve.

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

you any insight into how the original software works," Sahai says.

In the scheme's current version, obfuscation requires too much computation to make it feasible, but its developers think a practical version is possible.

For decades, true software encryption eluded computer scientists. In 2001, seminal research by Boaz Barak at the Weizmann Institute in Rehovot, Israel, and his colleagues proved that there are some programs for which an ideal version of obfuscation is impossible. This notion of obfuscation, called virtual black box, demands that an encrypted program's inner workings be completely hidden. While disappointing, this research still held out the possibility that almost all programs could be encrypted using another type of scheme, which is called indistinguishability obfuscation, or IO. This type would also obscure the program; however, the source code could still, technically, be deciphered if you were willing to put in an impractical amount of time and resources.

"We didn't know if it was possible or not to achieve indistinguishability obfuscation" in 2001, says Barak, now a senior researcher at Microsoft Research New England.

What Sahai and his team—Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova of IBM Research, and Brent Waters, an assistant professor of computer science at the University of Texas at Austin—say they've proved with their mathematical jigsaw puzzle trick is that indistinguishability obfuscation is possible. Their research will be presented this month at the 54th annual IEEE Symposium on Foundations of Computer Science, in Berkeley, Calif.

Such obfuscation would be particularly useful for software whose inner workings could reveal security flaws to hackers. When app makers issue security patches, hackers often grab them and quickly dig into the code to figure out a system's Achilles' heel, uncovering the vulnerability that the patch was designed to fix. "The crucial technical goal for us was to prevent this adversary from being able to do this," Sahai says.

According to Sahai, his team's research throws the doors open to new possibilities in the field of cryptography. In one possibility, called functional encryption, entire computable functions may be designed to run differently for individual users depending on their identity or purpose. For instance, a medical researcher might have a unique key that allows access to sensitive hospital data while keeping patients' personal information private.

Sahai's team says the next step in their research is to reduce the amount of computation needed to perform obfuscation. Based on previous trajectories in cryptographic research, the system could be practical in about a decade, Sahai says.

Microsoft Research's Barak, for one, extols the group's accomplishment. "This new work...gives a proof of concept that it could be achievable," he says. "Even if it's not yet practical...they showed how to use IO to achieve not all but many of the applications of full-fledged obfuscation," he says.

Shafi Goldwasser, a professor of electrical engineering and computer science at MIT, says that while the team used a looser definition of obfuscation than earlier research did, it's laudable that their system works on all programs under this new definition. "It's clear that this is a hammer that works on a lot of nails; it can be used to address many other open problems of cryptography," she says. "And from that point of view, it's a big deal." —DAVEY ALBA

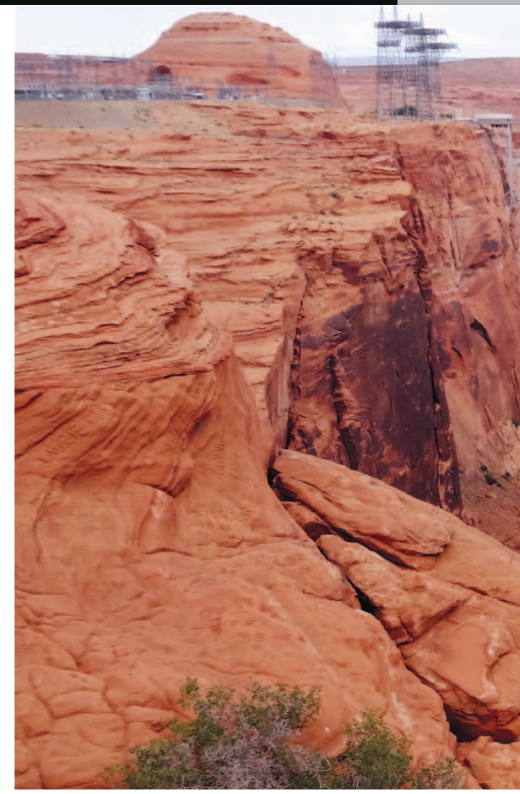

Last year, the Hoover Dam hydroelectric plant installed the first of five wide-head turbines. These are designed to work efficiently even as the Colorado River shrinks under a record-long drought. The dry spell affecting the dam's power source has outlasted any other in the 77 years that the structure has generated electricity. By the time the fifth turbine is installed in 2016, Hoover Dam will likely need them all.

Lake Mead, which sits on the border between Nevada and Colorado behind Hoover Dam, is expected to drop 2.4 meters in 2014, as less and less water flows downstream from Lake Powell, which straddles Utah and Arizona. The sharp decline comes about because the U.S. Bureau of Reclamation needs to cut Lake Powell's water release by nearly 1 billion cubic meters to 9.2 billion m^3 for the 2014 water season, the smallest release since the lake was filled in the 1960s. The flow of water to Lake Powell from key tributaries has been decreasing for more than a decade, and the bureau's forecasters expect that the reservoir could hit an all-time low this season.

"This is the most extreme drought since measurements began in the early 1900s," says Jack Schmidt, a professor of watershed

PREVIOUS PAGE: CLOCKWISE FROM TOP LEFT: GOOGLE; UNIVERSITY OF STRATHCLYDE; TITAN AEROSPACE CORP.; KYMETA CORP.; O3B NETWORKS

CHARLES PLATIAU/REUTERS

LOW ENERGY: There's less water than ever behind Glen Canyon Dam and other hydropower generators on the Colorado River.

COLORADO RIVER HYDROPOWER FACES A DRY FUTURE

Drought is hindering output from the river's iconic dams

sciences at Utah State University and current chief of the U.S. Geological Survey's Grand Canyon Monitoring and Research Center. A heavy snowfall this winter could change everything, but "no one knows when this will end," he adds.

The five new wide-head turbines being installed at the Hoover Dam are meant to keep the power plant working with less water in the lake. "We're trying to increase the power we can get from decreasing levels of water," says Rob Skordas, area manager of the Lower Colorado Dams office of the Bureau of Reclamation. Skordas says the new turbines should function well even if the water elevation falls to 305 meters above sea level, far below the historical average of 358.

Lake Mead, however, was already down to 337 meters in late August, when power capacity at Hoover Dam was at 1735 megawatts, down from a full capacity of 2074 MW. As water levels continue to decline, power output could fall even farther.

Upstream at the Glen Canyon Dam, power production is expected to be 8 percent lower than in 2013 as a result of the lower water, according to Jane Blair, manager of the bureau's Upper Colorado power office. The bureau estimates that the Western Area Power Administration will have to spend about US $10 million to meet its electricity supply obligations.

If the situation doesn't improve, Glen Canyon Dam could have even bigger problems. When the water level drops below 1063 meters, just about 30 meters below its August levels, vortex action would draw air into the turbines and damage them. Power generation would then likely cease at Glen Canyon, says Blair. Currently, engineers at Glen Canyon aren't looking to install any wide-head turbines like those at Hoover Dam.

If the drought cycles become longer and more severe, hydropower and other power needs will continue to take a backseat to water supply for the southwest region of the United States and California. On the Colorado River, there is a total hydropower generating capacity of 4178 MW, but many of the plants are already operating below their measured capacities because of the drought. Nearly 30 million people depend on the river for drinking water and irrigation. "Producing hydropower is clearly essential to maintaining a secure energy system," says Schmidt. "But in the grand scheme of things, water only comes from one place, and electricity comes from lots of places."

If electricity has to come from somewhere else, delivering drinking water to some of the largest cities in the western United States could be particularly problematic. Nearly 30 percent of the energy from Hoover Dam goes to the Metropolitan Water District of Southern California, which provides drinking water to nearly 19 million people across 26 cities and water districts. Less power also means less money for various water quality and environmental studies that inform how the water from the Colorado River should be allocated.

Many experts would like this year to mark the end of the drought cycle, but all they can do now is hope for snow while planning for withering water resources. "I think it's fair to say everybody involved with the river is hoping this coming winter is the snowiest winter on record," says Schmidt.
—KATHERINE TWEED

NEWS

NEWS

KEPLER'S CONTINUING MISSION

NASA's exoplanet hunter may be permanently disabled, but researchers say the best results are yet to come

In early August, the moment that Bill Borucki had been dreading finally arrived. As the principal investigator of NASA's Kepler space telescope, Borucki had been working with his colleagues since May to restore the spacecraft's ability to precisely point itself. The planet-hunting telescope has four reaction wheels–essentially, electrically driven flywheels–and at least three must be functional to maintain positioning. But by mid-May 2013 two had failed, effectively ending science operations. After a few months of recovery efforts, the telescope team was finally forced to call it quits, six months after the mission was originally scheduled to finish but years before they hoped it would.

The failures mark the end of an era for Kepler. With only two reaction wheels, the telescope can't steady itself well enough to ensure that light from each star hits the same fraction of a pixel on its charge-coupled devices for months on end without deviation. That's what Kepler needs in order to detect, with high precision, the transit of a planet: the slight dip in the brightness of a star that occurs when an orbiting planet crosses in front of it.

But the Kepler spacecraft might still have its uses, and the data it has already gathered almost certainly will. The telescope's managers are currently evaluating proposals for what might be done with a two-wheeled spacecraft. And the telescope's analysis team is gearing up for the rest of the science mission: a two- to three-year effort to systematically crawl through the four years of data that Kepler has collected since its launch in 2009.

That analysis effort, which will incorporate new machine-learning techniques and a bit of human experimentation, could yield as many as a thousand new potential planets on top of the 3500 that Kepler has found so far. If all goes well, the revised hunt might even uncover the first handful of terrestrial twins–or at the very least, near cousins: roughly Earth-size planets on nearly year-long orbits around sunlike stars.

Uncovering those Earth analogues won't be easy. The orbits are slow and the planets themselves are small. "You're looking for a percent of a percent" dip in the brightness of a star, says Jon Jenkins, the telescope's analysis lead. "That's a very demanding and challenging measurement to make."

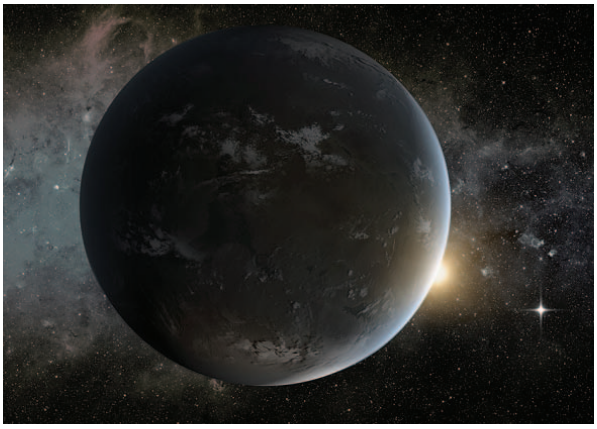

NEW WORLD: An artist's rendition shows Kepler-62f, a "super-Earth" in the habitable zone of a star 1200 light-years from Earth.

AMES/JPL-CALTECH/NASA

The task will be made even more difficult by an unexpected complication: Stars vary in brightness due to sunspots and flares, and Kepler's observations reveal that these variations are greater than scientists had previously estimated. Those fluctuations can hide the presence of a planet, reducing the telescope's sensitivity to terrestrial transits by 50 percent.

In April 2012, NASA granted Kepler a four-year extension that would have compensated for the extra noise. But with the failure of the reaction wheels, Jenkins and his colleagues now must find a different way to uncover planetary signals.

Earlier this year, they moved the data processing from a set of computer clusters containing 700 microprocessors to the Pleiades supercomputer at the NASA Ames Research Center, in Moffett Field, Calif., where they have the use of up to 15 000 of the machine's more than 160 000 cores. The team is also working on implementing a machine-learning process using an algorithm called the random forest, which will be trained with data already categorized by Kepler scientists. Once it's up and running, the software should be able to speedily differentiate false positives and data artifacts from promising candidates. Eventually, Jenkins says, the analysis team will insert fake data into the pipeline to test the performance of both the humans that ordinarily do the processing and the automated algorithms. "We need to know for every planet we detect how many we missed," Jenkins says.

No one can predict exactly how many planets Kepler will find. The telescope's main goal was to determine how common planets are in and around the habitable zones of stars–the areas around stars with the right temperature range for liquid water to be present. Such statistics could help astrophysicists decide how practical it would be to build a space telescope capable of directly detecting light from Earth-like planets, which is necessary to determine whether they have atmospheres that could support life.

For Earth-size planets in settings similar to our own, developing a good statistical estimate will be difficult. With small numbers, the uncertainty in the size of the overall population will be large, says Sara Seager, a professor of planetary science and physics at MIT and a participating scientist on the team.

Even if Kepler finds no Earth analogues, Seager says, the mission is a success. "Kepler revolutionized exoplanet science and, arguably, big-data astronomy," she says. "We'll see the data being mined for years to come."
–RACHEL COURTLAND

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

RESOURCES

$5.6 BILLION: WHAT AMERICAN GOLFERS SPENT ON GOLFING EQUIPMENT IN 2011, ACCORDING TO SRI INTERNATIONAL.

GADGETS FOR GOLFERS

COMPUTERIZED SWING ANALYSIS FOR THE REST OF US

RESOURCES_TOOLS

The Nintendo Wii game controller proved that you needed just a couple of acceler-ometers to capture the basic motion of a golf swing. Now, inexpensive products like the SwingTip put those sensors on real clubs, promising to improve your real-life game by offering detailed feedback. • The SwingTip is a small plastic box that clips around the shaft of any golf club. When paired with a smartphone or tablet, it tracks the position, speed, and angle of the club as it moves through a swing. The device is part of a new class of golf training aids; it's now one of about half a dozen swing trackers available for less than US $200. • In many ways these devices are just sport-specific versions of the activity monitors, such as the FitBit, that are popular with self-quantifiers . They provide more precise data, however, because capturing the kinematics of a club is easier than measuring the complex motions of the human body. • The SwingTip has a mounting plate that simply snaps to the shaft of a club, with two double-throw latches to hold it tight. (But do make sure that it's fully secured—it's not fun to see your new gadget fly 30 meters out onto the driving range. See Back Story, in this issue.) The SwingTip works with a free app available for Android and iOS devices, and you pair it with your phone or tablet as you would any other Bluetooth accessory. • The SwingTip marks each swing by detecting contact with the ball. That means it's not suitable for practice swings in your living room, as some of its competitors are. After you watch your ball sail away, it takes several seconds for the app to crunch the data. (If you use the optional video-capture mode to create a recording of your swing, it takes substantially longer.) When the phone or tablet finishes processing, you can watch an animated avatar re-create your swing. • Taking out my iPad at the driving range ▸

PHOTOGRAPH BY Randi Silberman Klett

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

RESOURCES_HANDS ON

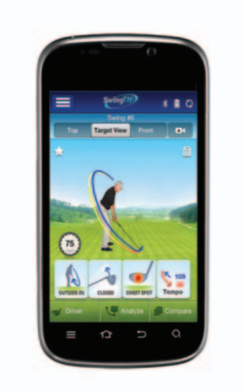

AVATAR IN ACTION: Golf-swing data is mapped to the motions of an animated golfer, allowing you to examine your swing from the back, the side, and above.

made me feel a bit self-conscious. But once I got used to checking it, the driving range felt more like a video game. It was fun to see the path of my swing—even with all its flaws—from three different angles. In addition to the animation, the app provides data such as the speed and angle of the club at impact. It also has an "analysis" feature, which tries to identify errors at nine points in the swing. I now had specific numbers and targets that I could improve on, or try to repeat.

After numerous swings, however, I found myself questioning the validity of the SwingTip's results. I noticed that the starting point of the swing in the animations (where the club head was right behind the ball) and the ending point (where the club met the ball) were often several inches apart. As the ball doesn't move until it's struck, this discrepancy is likely due to sensor drift.

The onscreen avatar is also a bit misleading, because its movements are just a stock animation that has nothing to do with what my body actually does. In this sense, the video capture was more informative because it showed how my body and the club moved together. But there are numerous golf apps available that do video capture alone for a few dollars.

Still, I found that the system was valuable for identifying what went wrong in a less-than-ideal swing. Alas, such knowledge is only the first step to improvement. Although the app provides a library of tips and a few video lessons, it can't tell you why you tend to take too long of a backswing, or how to break the habit.

Golf is perhaps uniquely suited to data-based training aids: It's a quantitative sport, with success measured in strokes, and there's always room for improvement. But after a long session of constantly replaying, analyzing, and reliving each swing, it feels good to turn off the screen and just hit some balls.

—JOSHUA J. ROMERO

POINT-AND-SHOOT WEATHER DATA

AUTOMATE DATA COLLECTION WITH A CAMERA

LEFT: SWINGTIP; RIGHT: PAUL WALLICH

A WHILE BACK, I GOT A CHEAP WIRELESS LACROSSE brand weather station, with sensors for temperature, humidity, wind speed and direction, and rainfall. The stand-alone display shows weather conditions using seven-segment LCD numerals. The box the station came in promised that it "Connects to your PC!," which appealed to me because I'd be able to automatically log the data and pass it around my home network. Well, it turns out that this promise could be honored only if I had a PC of the precise operating system and vintage that the station's USB-to-wireless dongle called for. I'd also need just the right version of some proprietary weather software. Unfortunately, my Macintosh setup met none of these requirements. ● Sure, if I had access to an already working installation (and a lot of time), I might have tried reverse engineering the embedded hardware and software. In theory, I could have figured out what the weather-station components were telling the dongle, what the dongle was telling the PC, and what commands might be flowing in the other direction. Then, with some more work and time, I could have reproduced that conversation on a Linux or OS X box. But I didn't have the access or the time. ● So I decided to make do with the station's stand-alone display. I've been learning how to use OpenCV, the

WEATHER CAM

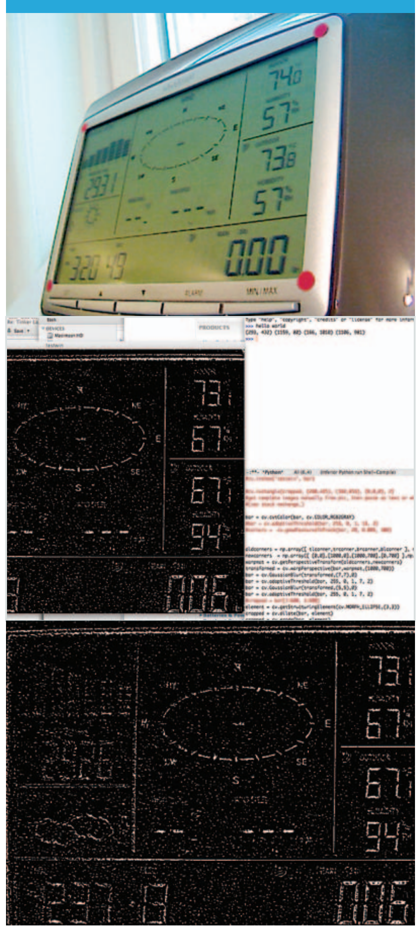

I placed red dots around the edges of the weather station's display [top] to help my modified software [center] identify the limits of the display and correct for perspective. The software then passes the image to OCR software for data extraction [bottom].

PAUL WALLICH (3)

comprehensive image-processing library initially created at Intel and now supported mostly by Itseez and the open-source roboticists at Willow Garage. The library is designed to make it easy to extract information about a scene from the raw images coming from a camera. It's available for many platforms, including Mac.

With such software at hand, and with webcams being cheap to the point of being disposable, I thought it should be a simple matter to take a picture of the screen, extract the text displaying the weather conditions, run the text through an optical character recognition (OCR) system to get data in numeric form, and log the result. Because I needed to make a log entry only every few minutes, I didn't need a lot of computational horsepower, so I could do the processing using spare background cycles on my Mac.

An early issue was figuring out which of the many possible ways I could use OpenCV to do this most effectively. At first I thought that it would mostly be a matter of correcting the image for distortions caused by my camera's perspective (my webcam is off to one side to let me still read the display by eye) and uneven lighting. Once I had something to pass to OCR software, the hard work would be done. However, it turns out that almost all free OCR software is designed for printed text, where each letter forms a continuous contour. It's not at all good for numbers and letters made up of disconnected segments, as in my station's display.

So then I was entranced by the idea of template matching: comparing small "model" images of digits with those from the webcam feed, seeing where and if they matched, and collating that with the positions of the temperature, wind speed, and other indicators on the display. But that would have meant waiting until all the digits from 0 to 9 appeared on the local weather display at least once so I could save their images to a file. Or I could have drawn sample digits by hand, but OpenCV's standard template-matching function is unforgiving of mismatches in size or orientation.

Then I found a program called SSOCR, or Seven-Segment Optical Character Recognition, which was developed by Erik Auerswald of the Technical University of Kaiserslautern, in Germany. SSOCR makes it possible to paste one-time-password codes from a security fob into a Web page log-in screen. Optimized for the fob's single fixed line of six digits, SSOCR turned out to be a bit too specialized. It requires a close-up image under unchanging lighting. However, the light on my weather station's screen varies depending on the time of day, and the camera has to be far enough back to capture the station's wider, multiline screen. So I stole some ideas about how to slice up seven-segment images from SSOCR and wrote my own recognizer, based simply on which segments had enough black pixels to be considered "on."

I still had to correct for the camera's perspective, which I thought would be easy, as software for solving this exact problem is already available. However, the weather-station display is made of gray plastic, and the LCD is dark gray on light gray, so there aren't any distinctive points for the corner-finding routines normally used for this kind of correction to lock onto. Finally, I just cut out four little circles of red construction paper and glued them to the display; the correction routines have no trouble finding those. Once an image has been acquired, my OCR routine does its work, and the results are saved into a text file for further processing.

What am I going to do with all my weather data? Over the long term, I'm going to match it to the data from the nearest official weather station, so that I can figure out how the weather there correlates with the weather here. But the first thing I need to do is to create a weather page on my Mac's personal Web server. Then I can use my tablet or phone to see just how stormy it is outside while I'm still lying in bed at the other end of the house from the station. The easiest way is to just fire up a browser, but with MIT (formerly Google) App Inventor, which allows drag-and-drop assembly of Android apps, I should need just a couple of hours to write a program—once I get yet another development environment set up. **—PAUL WALLICH**

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

GONÇALO ABECASIS
HIS SOFTWARE MINES GENOMES FOR BIOMEDICAL GOLD

Gonçalo Abecasis's sole intention was to become a geneticist while using his computer programming prowess to pay for university. But years later, those skills have become an integral part of his research.

Today, as the Felix Moore Collegiate Professor of Biostatistics at the University of Michigan's School of Public Health, in Ann Arbor, Abecasis develops statistical tools and computational methods that help determine the genetic fingerprints for such diseases as psoriasis, macular degeneration, diabetes, and heart disease.

"I want to figure out why people get sick, find better ways to treat and manage disease, and learn how best to use computational modeling to get that information," he says.

Computer modeling is now an unavoidable component of Abecasis's genetic research, as the volume of available data has increased at a faster rate than computer processing power.

"The same amount of money can generate four times as much data as the year before," says Abecasis, who estimates that his lab has amassed some 5 petabytes of data. "Computer processing power only doubles in that time, so the cost of generating data has gone down much faster than the cost of computing it. So you constantly need more efficient ways of compressing, storing, assessing, and prioritizing the data." [For more about this problem, read "The DNA Data Deluge," *IEEE Spectrum*, July 2013.]

"We're working on more efficient mathematical models for data—such as leveraging repeated patterns in data—and compressing data to store it," he adds. "We're also thinking differently about what data is important and what we can discard."

Over the years, Abecasis and his staff have developed their own software, written mostly in C++, to handle the voluminous data and visualize it in colorful 2-D pictures and graphs instead of tables of numbers. The tools are distributed free and widely used in the disease genetics community.

Once Abecasis and his team have used their tools to isolate genetic variants associated with a disease, they look for a cascade of small steps leading to disease, starting with the variants and then including inputs such as diet and environment. "We're trying to determine the steps that start with a small genetic defect and eventually manifest into the disease," he says.

Abecasis grew up with 11 siblings in Portugal and Macao, a former Portuguese colony in China near Hong Kong. He taught himself programming as a teenager with some assistance from a software engineer who ran a youth program. Later, computer programming jobs helped pay his way through the University of Leeds, in the United Kingdom, where he earned a bachelor's in genetics in 1997, and the University of Oxford, where he graduated with a Ph.D. in human genetics in 2001.

That might have been the end of his programming had Abecasis not experienced a computing bottleneck firsthand: A project he was working on at Oxford generated more data than the school's computer could handle. "At some point, I coded something to be more efficient, and then I became a supervisor for the computational side of the research," he says. "The tools we ended up developing were more popular than the established methodology, so I began focusing on that side of things."

After Oxford, Abecasis headed to the University of Michigan as a research faculty member, working his way up to full professor in 2009. Along the way he garnered several awards, including this year's Overton Award from the International Society for Computational Biology for exceptional contributions made by a researcher in the early to middle stages of his or her career.

Abecasis's work often has him bridging the distinct cultures of engineering and biology. "Both are attracted to big data, but they think differently," he says. "Biologists like to figure out cause and effect that they can replicate in experiments, while computer scientists are happy to create models that make predictions but they don't necessarily need to understand why. Over time, differences can be a good thing. Things I thought weren't important turned out to be useful, and vice versa." —**SUSAN KARLIN**

PETER SMITH


Enhancing entertainment
in everyday life
life.augmented
Experience by STMicroelectronics. Simulation by ANSYS.
ANSYS®

RESOURCES_TOOLS

GOOGLE GLASS
THIS WEARABLE COMPUTER AUGMENTS THE SELF, NOT REALITY

Google Glass is a polarizing device. Positive opinions have often focused on the possibilities for augmented reality (AR), while the negative ones tend to focus on the privacy ramifications. But I'm going to focus on what makes the current Glass hardware—dubbed the XE or Explorer Edition—groundbreaking, and what we can expect from Glass-like devices in the near future.

Glass XE is not a consumer product. It is an experimental, voice-activated wearable computer equipped with a camera and a head-up display. In its current iteration, despite the early hype of Glass as an augmented-reality device, the small screen can overlay information on only a portion of your field of view. In addition, although the onboard 1-gigahertz dual-core processor is reasonably powerful, it's held back by a relatively small battery. These issues make true AR applications impractical.

However, despite its small size, Glass's head-up display is oddly immersive. Because Glass spends most of the time in standby with the screen off—and therefore transparent—you can easily forget you're wearing it. When Glass wants your attention, it uses an integrated bone conduction speaker to notify you. If you choose to engage, tilting your head 30 degrees upward activates the screen, a frictionless process that feels natural. Most of the time you'll be interacting with Glass either by using its touch pad or by voice command.

The features currently available cover a significant portion of the tasks you would normally use your smartphone for. You can receive notifications (from Gmail or SMS, for instance) and reply, search using Google, get directions, and make calls (including video calls). Its point-of-view camera allows you to take pictures or shoot videos with ease.

ALWAYS ON: Google Glass makes it easy to capture the details of daily life, like your favorite restaurant's menu. Such images could someday be automatically parsed to create searchable databases, so you could recall a particular dish.

Glass's brilliance lies in its form factor and multimodal natural interface, which translates into a user experience fundamentally different from anything I've experienced before. Wearing and interacting with Glass doesn't feel anything like using a conventional mobile device. Rather, it feels more like having a computer inside my head or having artificial senses spliced into my existing ones.

As I mentioned in my previous article about wearable computers [see "Build Your Own Google Glass," *IEEE Spectrum*, January 2013], our brains are eager to incorporate new streams of information into our mental models of the world (one fascinating example involves wearing an ankle bracelet that vibrates to indicate north). After a period of adaptation, these streams of information fade into the background as conscious attention is replaced with mostly automatic behavior.

For me, it took my brain roughly two weeks to fully incorporate Glass into my model of the world. Now, I often find myself tilting my head to activate Glass even when I'm not wearing it.

What's the near-future potential of wearable point-of-view computers? Future versions of Glass will enable a wide range of augmented cognition applications—combining the natural strengths of the human brain, the computational power of the cloud, cheap storage, and machine learning.

For example, once we deal with the (admittedly nontrivial) privacy constraints around continuously recording video with Glass, hardware iterations with improved battery life could record everything you see and hear and upload it to the cloud, where the data would be sifted, salient features extracted, and transcripts generated, thus making your audiovisual memory searchable. Imagine being able to search or summarize every conversation you ever had, or extract statistics about your life from aggregated visual, aural, and location data.

Ultimately, given enough time, those digital memory constructs will evolve into what can be loosely described as external brains in the cloud—picture a semiautonomous process that knows enough about you to act on your behalf in a limited fashion.

A couple of days ago I was stopped by a stranger who asked me, "What can you see through Google Glass?" To which I replied, only partly tongue in cheek, "I can see the future."

—ROD FURLAN

An extended version of this article appeared online in August.

TOP: KIM FURLAN; BOTTOM: ROD FURLAN

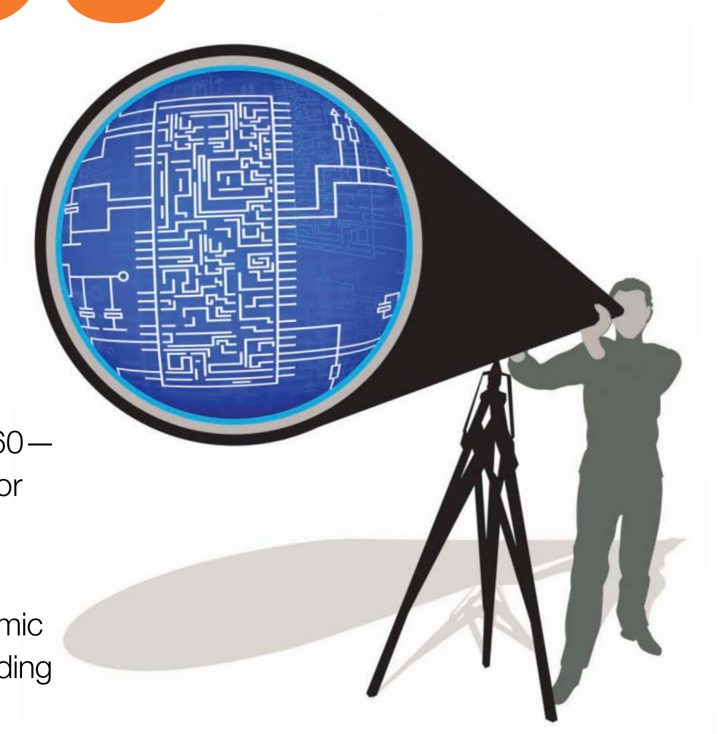

MEET THE DATASEXUAL

The datasexual looks a lot like you and me, but what's different is their preoccupation with personal data. They are relentlessly digital, they obsessively record everything about their personal lives, and they think that data is sexy....Their lives—from a data perspective, at least—are perfectly groomed.
—Dominic Basulto, Big Think (a Web "knowledge forum")

YOU MIGHT REMEMBER (or more likely have been actively trying to forget) the *metrosexual*, that much-discussed urban male with the legendarily strong aesthetic sense and the equally legendary amount of time and money to spend on his appearance and lifestyle. As maligned as the stereotype came to be, it proved to be a powerful generator of new terms, including its opposite, the **retrosexual**, the **ecosexual** (an environmentally conscious person with a strong aesthetic sense), and the **technosexual** (a male who combines aesthetic flair with a love of technology). The latest incarnation is the **datasexual**, a person who's an obsessive *self-tracker* [see "Tracking the Quantified Self," *IEEE Spectrum*, August 2013], not just to enhance self-knowledge but also to embellish self-presentation, especially on social networks. • The datasexual spends a good part of the day sending out chunks of digital flotsam that fall under the rubric of the **narb**, which refers to any item of personal information posted online, particularly as it contributes, often unwittingly, to a personal narrative that the individual is creating online. (The word is a blend of *narrative* and *bit*.) The difference between your garden-variety quantified-selfer and a datasexual is the latter's emphasis on public self-embellishment. While a **QSer** might use a pedometer to track the number of steps she takes each day, a datasexual will wear a Nike+ FuelBand on his wrist to *display* the number of steps he takes each day, and he'll post that number to his online friends. The datasexual transforms self-obsession into conspicuous **oversharing**.

A typical self-monitoring data junkie will take the good with the bad (and use the bad to improve), but our friend the datasexual is almost always into **success theater**, the posting of images and stats designed to make others believe he is more successful than he really is. The classic elements of such a strategy include the flattering **selfie**, which is a photographic self-portrait, particularly one taken with the intent of posting it to a social network (the extreme here is the **Facebook face-lift**, cosmetic surgery designed to improve how a person looks in photos posted to social networking sites). Bonus points are awarded if the posting includes a **humblebrag** (an ostensibly humble comment that also demonstrates one's wealth, fame, or importance) or a **vanity metric**, a numeric value or data point that serves to highlight some positive aspect of one's life, such as one's health (for example "My resting heart rate is 55!") or one's standing ("Just passed the 10 000-follower mark on Twitter!") or if the photo includes someone famous.

Management guru Tom Peters calls such things one's **braggables**, and they represent the datasexual's stock-in-trade. They're what motivate datasexuals to engage in their daily regimen of **data hygiene** and **data grooming**.

Off-line, the datasexual is sure to engage in **stage-phoning**, which is the attempt to impress nearby people by casually including envy-inducing personal stats while talking on a cellphone in a theatrical manner. With phones omnipresent in the social landscape, and would-be thespians appearing at every airport waiting lounge, coffee shop, and street corner, we see that indeed, 400 years after Shakespeare declared it, all the world really is a stage. After a few minutes of their overly loud and overly proud boasting, we suppress an urge to offer them the traditional actor's send-off—"break a leg"—because, well, this time we might mean it. ■

ILLUSTRATION BY Greg Mably

THE SENSOR TECHNOLOGY under the nose of Boeing's Unmanned Little Bird (ULB) helicopter maps terrain, plans routes, discerns safe landing sites, and avoids obstacles—all by itself.

GUTTER CREDIT GOES HERE

ROBOCOPTERS *to the* RESCUE

THE NEXT MEDEVAC HELICOPTER WON'T NEED A PILOT

By LYLE CHAMBERLAIN & SEBASTIAN SCHERER

WE'RE STANDING ON THE EDGE OF THE HOT Arizona tarmac, radio in hand, holding our breath as the helicopter passes 50 meters overhead. We watch as the precious sensor on its blunt nose scans every detail of the area, the test pilot and engineer looking down with coolly professional curiosity as they wait for the helicopter to decide where to land. They're just onboard observers. The helicopter itself is in charge here. ● Traveling at 40 knots, it banks to the right. We smile: The aircraft has made its decision, probably setting up to do a U-turn and land on a nearby clear area. Suddenly, the pilot's voice crackles over the radio: "I have it!" That means he's pushing the button that disables the automatic controls, switching back to manual flight. Our smiles fade. "The aircraft turned right," the pilot explains, "but the test card said it would turn left." ● The machine would have landed safely all on its own. But the pilot could be excused for questioning its, uh, judgment. For unlike

the autopilot that handles the airliner for a good portion of most commercial flights, the robotic autonomy package we've installed on Boeing's Unmanned Little Bird (ULB) helicopter makes decisions that are usually reserved for the pilot alone. The ULB's standard autopilot typically flies a fixed route or trajectory, but now, for the first time on a full-size helicopter, a robotic system is sensing its environment and deciding where to go and how to react to chance occurrences.

It all comes out of a program sponsored by the Telemedicine & Advanced Technology Research Center, which paired our skills, as roboticists from Carnegie Mellon University, with those of aerospace experts from Piasecki Aircraft and Boeing. The point is to bridge the gap between the mature procedures of aircraft design and the burgeoning world of autonomous vehicles. Aerospace, meet robotics.

The need is great, because what we want to save aren't the salaries of pilots but their lives and the lives of those they serve. Helicopters are extraordinarily versatile, used by soldiers and civilians alike to work in tight spots and unprepared areas. We rely on them to rescue people from fires, battlefields, and other hazardous locales. The job of medevac pilot, which originated six decades ago to save soldiers' lives, is now one of the most dangerous jobs in America, with 113 deaths for every 100 000 employees. Statistically, only working on a fishing boat is riskier.

These facts raise the question: Why are helicopters such a small part of the boom in unmanned aircraft? Even in the U.S. military, out of roughly 840 large unmanned aircraft, fewer than 30 are helicopters. In Afghanistan, the U.S. Marines have used two unmanned Lockheed Martin K-Max helicopters to deliver thousands of tons of cargo, and the Navy has used some of its 20-odd shipborne Northrop Grumman unmanned Fire Scout helicopters to patrol for pirates off the coast of Africa.

SO WHAT'S HOLDING BACK UNMANNED helicopters? What do unmanned airplanes have that helicopters don't?

It's fine for an unmanned plane to fly blind or by remote control; it takes off, climbs, does its work at altitude, and then lands, typically at an airport, under close human supervision. A helicopter, however, must often go to areas where there are either no people at all or no specially trained people—for example, to drop off cargo at an unprepared area, pick up casualties on rough terrain, or land on a ship. These are the scenarios in which current technology is most prone to fail, because unmanned aircraft have no common sense: They do exactly as they are told.

If you absentmindedly tell one to fly through a tree, it will attempt to do it.

PREVIOUS PAGES: SEBASTIAN SCHERER; THIS PAGE: LEFT: NORTHROP GRUMMAN CORP.; RIGHT: LOCKHEED MARTIN

THE U.S. NAVY'S carrier-based copter, the Northrop Grumman Fire Scout [far left], can fly itself onto and off of a moving deck. The U.S. Marine Corps's K-Max ferries cargo to soldiers, dangling a "sling load" to work around its weakness in landing on rough ground. To rescue people, however, a robocopter needs better eyes and judgment.

One experimental unmanned helicopter nearly landed on a boulder and had to be saved by the backup pilot. Another recently crashed during the landing phase. To avoid such embarrassments, the K-Max dangles cargo from a rope as a "sling load" so that the helicopter doesn't have to land when making a delivery. Such work-arounds throw away much of the helicopter's inherent advantage. If we want these machines to save lives, we must give them eyes, ears, and a modicum of judgment.

In other words, an autonomous system needs perception, planning, and control. It must sense its surroundings and interpret them in a useful way. Next, it must decide which actions to perform in order to achieve its objectives safely. Finally, it must control itself so as to implement those decisions.

A cursory search on YouTube will uncover videos of computer-controlled miniature quadcopters doing flips, slipping vertically through slots in a wall, and assembling structures. What these craft are missing, though, is perception: They perform inside the same kind of motion-capture lab that Hollywood uses to record actors' movements for computer graphics animations. The position of each object is precisely known. The trajectories have all been computed ahead of time, then checked for errors by software engineers.

If you give such a quadcopter onboard sensors and put it outdoors, away from the rehearsed dance moves of the lab, it becomes much more hesitant. Not only will it sense its environment rather poorly, but its planning algorithms will barely react in time when confronted with an unusual development.

True, improvements in hardware are helping small quadcopters approach full autonomy, and somewhat larger model helicopters are already quite far along in that quest. For example, several groups have shown capabilities such as automated landing, obstacle avoidance, and mission planning on the Yamaha RMax, a 4-meter machine originally sold for remote-control crop dusting in Japan's hilly farmlands. But this technology doesn't scale up well, mainly because the sensors can't see far enough ahead to manage the higher speeds of full-size helicopters. Furthermore, existing software can't account for the aerodynamic limitations of larger craft.

Another problem with the larger helicopters is that they don't actually like to hover. A helicopter typically lands more like an airplane than most people realize, making a long, descending approach at a shallow angle at speeds of 40 knots (75 kilometers per hour) or more and then flaring to a hover and vertical descent. This airplane-like profile is necessary because hovering sometimes requires more power than the engines can deliver.

THE NEED FOR FAST FLYING HAS A LOT TO do with the challenges of perception and planning. We knew that making large, autonomous helicopters practical would require sensors with longer ranges and faster measurement rates than had ever been used on an autonomous rotary aircraft, as well as software optimized to make quick decisions. To solve the first problem, we began with ladar–laser detection and ranging–a steadily improving technology and one that's already widely used in robotic vehicles.

Ladar measures the distance to objects by first emitting a tightly focused laser pulse and then measuring how long it takes for any reflections to return. It creates a 3-D map of the surroundings by pulsing 100 000 times per second, steering the beam to many different points with mirrors, and combining the results computationally.

The ladar system we constructed for the ULB uses a "nodding" scanner. A "fast-axis" mirror scans the beam in a horizontal line up to 100 times per second while another mirror nods up and down much more slowly. To search for a landing zone, the autonomous system points the ladar down and uses the fast-axis line as a "push broom," surveying the terrain as the helicopter flies over it. When descending nearer to a possible landing site, the system points the ladar forward and nods up and down, thus scanning for utility wires or other low-lying obstacles.

Because the helicopter is moving, the ladar measures every single point from a slightly different position and angle. Normally, vibration would blur these measurements, but we compensate for that problem by matching the scanned information with the findings of an inertial navigation system, which uses GPS, accelerometers, and gyros to measure position within centimeters and angles within thousandths of a degree. That

TOPOGRAPHY *for* ROBOTS

THE UNMANNED HELICOPTER can land in an unprepared site to evacuate a patient. Using a ladar system, the craft can construct maps, select landing sites, and generate safe approaches to them, giving itself a fallback option in case the situation changes.

1 The helicopter flies over the landing area.

2 The optical sensors under its nose scan the terrain to create a cloud of data points.

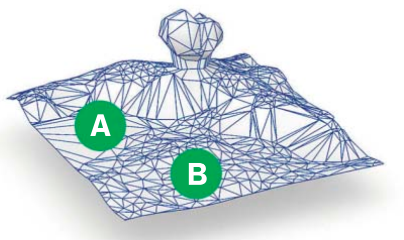

3 The scanner maps the cloud in a high-definition 2-D form, then ranks a large set of potential landing sites.

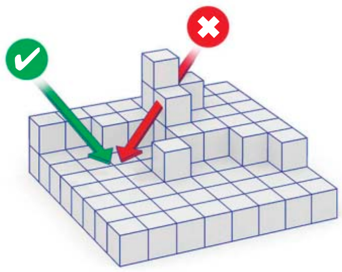

4 Meanwhile, the scanner also constructs a lower-definition 3-D map, to which it transposes the best landing areas. The system then searches the map for obstacle-free approaches to those areas.

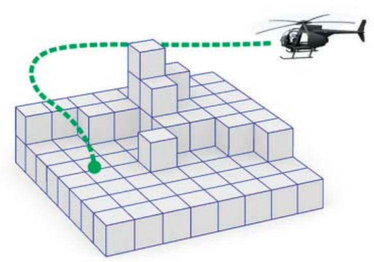

5 The system chooses the best route to the primary landing area.

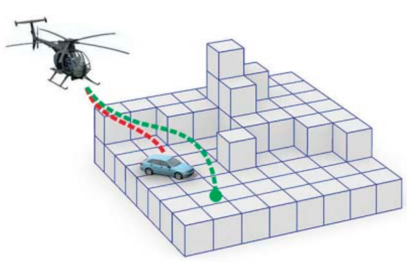

6 If some new obstacle should suddenly block the way, it will opt for a different route, leading to a secondary landing area.

way, we can properly place each ladar-measured reflection on an internal map.

To put this stream into a form the planning software can use, the system constantly updates two low-level interpretations. One is a high-resolution, two-dimensional mesh that encodes the shape of the terrain for landing; the other is a medium-resolution, three-dimensional representation of all the things the robot wants to avoid hitting during its descent. Off-the-shelf surveying software can create such maps, but it may take hours back in the lab to process the data. Our software creates and updates these maps essentially as fast as the data arrive.

The system evaluates the mesh map by continually updating a list of numerically scored potential landing places. The higher the score, the more promising the landing site. Each site has a set of preferred final descent paths as well as clear escape routes should the helicopter need to abort the attempt (for example, if something gets in the way). The landing zone evaluator makes multiple passes on the data, refining the search criteria as it finds the best locations. The first pass quickly eliminates areas that are too steep or rough. The second pass places a virtual 3-D model of the helicopter in multiple orientations on the mesh map of the ground to check for rotor and tail clearance, good landing-skid contact, and the predicted tilt of the body on landing.

In the moments before the landing, the autonomous system uses these maps to generate and evaluate hundreds of potential trajectories that could bring the helicopter from its current location down to a safe landing. The trajectory

ILLUSTRATION BY Emily Cooper

includes a descending final approach, a flare—the final pitch up that brings the helicopter into a hover—and the touchdown. Each path is evaluated for how close it comes to objects, how long it would take to fly, and the demands it would place on the aircraft's engine and physical structure. The planning system picks the best combination of landing site and trajectory, and the path is sent to the control software, which actually flies the helicopter. Once a landing site is chosen, the system continuously checks its plan against new data coming from the ladar and makes adjustments if necessary.

That's how it worked in simulations. The time had come to take our robocopter out for a spin.

SO IT WAS THAT WE FOUND OURSELVES on a sunny spring afternoon in Mesa, Ariz. Even after our system had safely landed itself more than 10 times, our crew chief was skeptical. He had spent decades as a flight-test engineer at Boeing and had seen many gadgets and schemes come and go in the world of rotorcraft. So far, the helicopter had landed itself only in wide-open spaces, and he wasn't convinced that our system was doing anything that required intelligence. But today was different: Today he would match wits with the robot pilot.

Our plan was to send the ULB in as a mock casualty evacuation helicopter. We'd tell it to land in a cleared area and then have it do so again after we'd cluttered up the area. The first pass went without a hitch: The ULB flew west to east as it surveyed the landing area, descended in a U-turn, completed a picture-perfect approach, and landed in an open area close to where the "casualty" was waiting to be evacuated. Then our crew chief littered the landing area with forklift pallets, plastic boxes, and a 20-meter-high crane.

This time, after the flyover the helicopter headed north instead of turning around. The test pilot shook his head in disappointment and prepared to push the button on his stick to take direct control. But the engineer seated next to him held her hand up. After days of briefings on the simulator, she had begun to get a feel for the way the system "thought," and she realized that it might be trying to use an alternative route that would give the crane a wider berth. And indeed, as the helicopter descended from the north, it switched the ladar scanner from downward to forward view, checking for any obstacles such as power lines that it wouldn't have seen in the east-west mapping pass. It did what it needed to do to land near the casualty, just as it had been commanded.

SEBASTIAN SCHERER

THE BOEING ULB has both the sensors and the judgment needed to find its way to an unprepared site, avoid collisions—even with obstacles that might arrive at the last minute—and land near a casualty. This allows people on the ground to evacuate a wounded person (who might not be mobile) without putting a pilot's life at risk.

This landing was perfect, except for one thing: The cameras had been set up ahead of time to record an approach from the east rather than the north. We'd missed it! So our ground crew went out and added more clutter to try to force the helicopter to come in from the east but land further away from the casualty. Again the helicopter approached from the north and managed to squeeze into a tighter space nearby, keeping itself close to the casualty. Finally, the ground crew drove out onto the landing area, intent on blocking all available spaces and forcing the machine to land from the east. Once again the wily robot made the approach from the north and managed to squeeze into the one small (but safe) parking spot the crew hadn't been able to block. The ULB had come up with perfectly reasonable solutions—solutions we had deliberately tried to stymie. As our crew chief commented, "You could actually tell it was making decisions."

THAT DEMONSTRATION PROGRAM ENDED three years ago. Since then we've launched a spin-off company, Near Earth Autonomy, which is developing sensors and algorithms for perception for two U.S. Navy programs. One of these programs, the Autonomous Aerial Cargo/Utility System (AACUS), aims to enable many types of autonomous rotorcraft to deliver cargo and pick up casualties at unprepared landing sites; it must be capable of making "hot" landings, that is, high-speed approaches without precautionary overflight of the landing zone. The other program will develop technology to launch and recover unmanned helicopters from ships.

It took quite a while for our technology to win the trust of our own professional test team. We must clear even higher hurdles before we can get nonspecialists to agree to work with autonomous aircraft in their day-to-day routines. With that goal in view, the AACUS program calls for simple and intuitive interfaces to allow nonaviator U.S. Marines to call in for supplies and work with the robotic aircraft.

In the future, intelligent aircraft will take over the most dangerous missions for air supply and casualty extraction, saving lives and resources. Besides replacing human pilots in the most dangerous jobs, intelligent systems will guide human pilots through the final portions of difficult landings, for instance by sensing and avoiding low-hanging wires or tracking a helipad on the pitching deck of a ship. We are also working on rear-looking sensors that will let a pilot keep constant tabs on the dangerous rotor at the end of a craft's unwieldy tail.

Even before fully autonomous flight is ready for commercial aviation, many of its elements will be at work behind the scenes, making life easier and safer, just as they are doing now in fixed-wing planes and even passenger cars. Robotic aviation will not come in one fell swoop—it will creep up on us. ■

POST YOUR COMMENTS at **http://spectrum.ieee.org/robocopter1013**

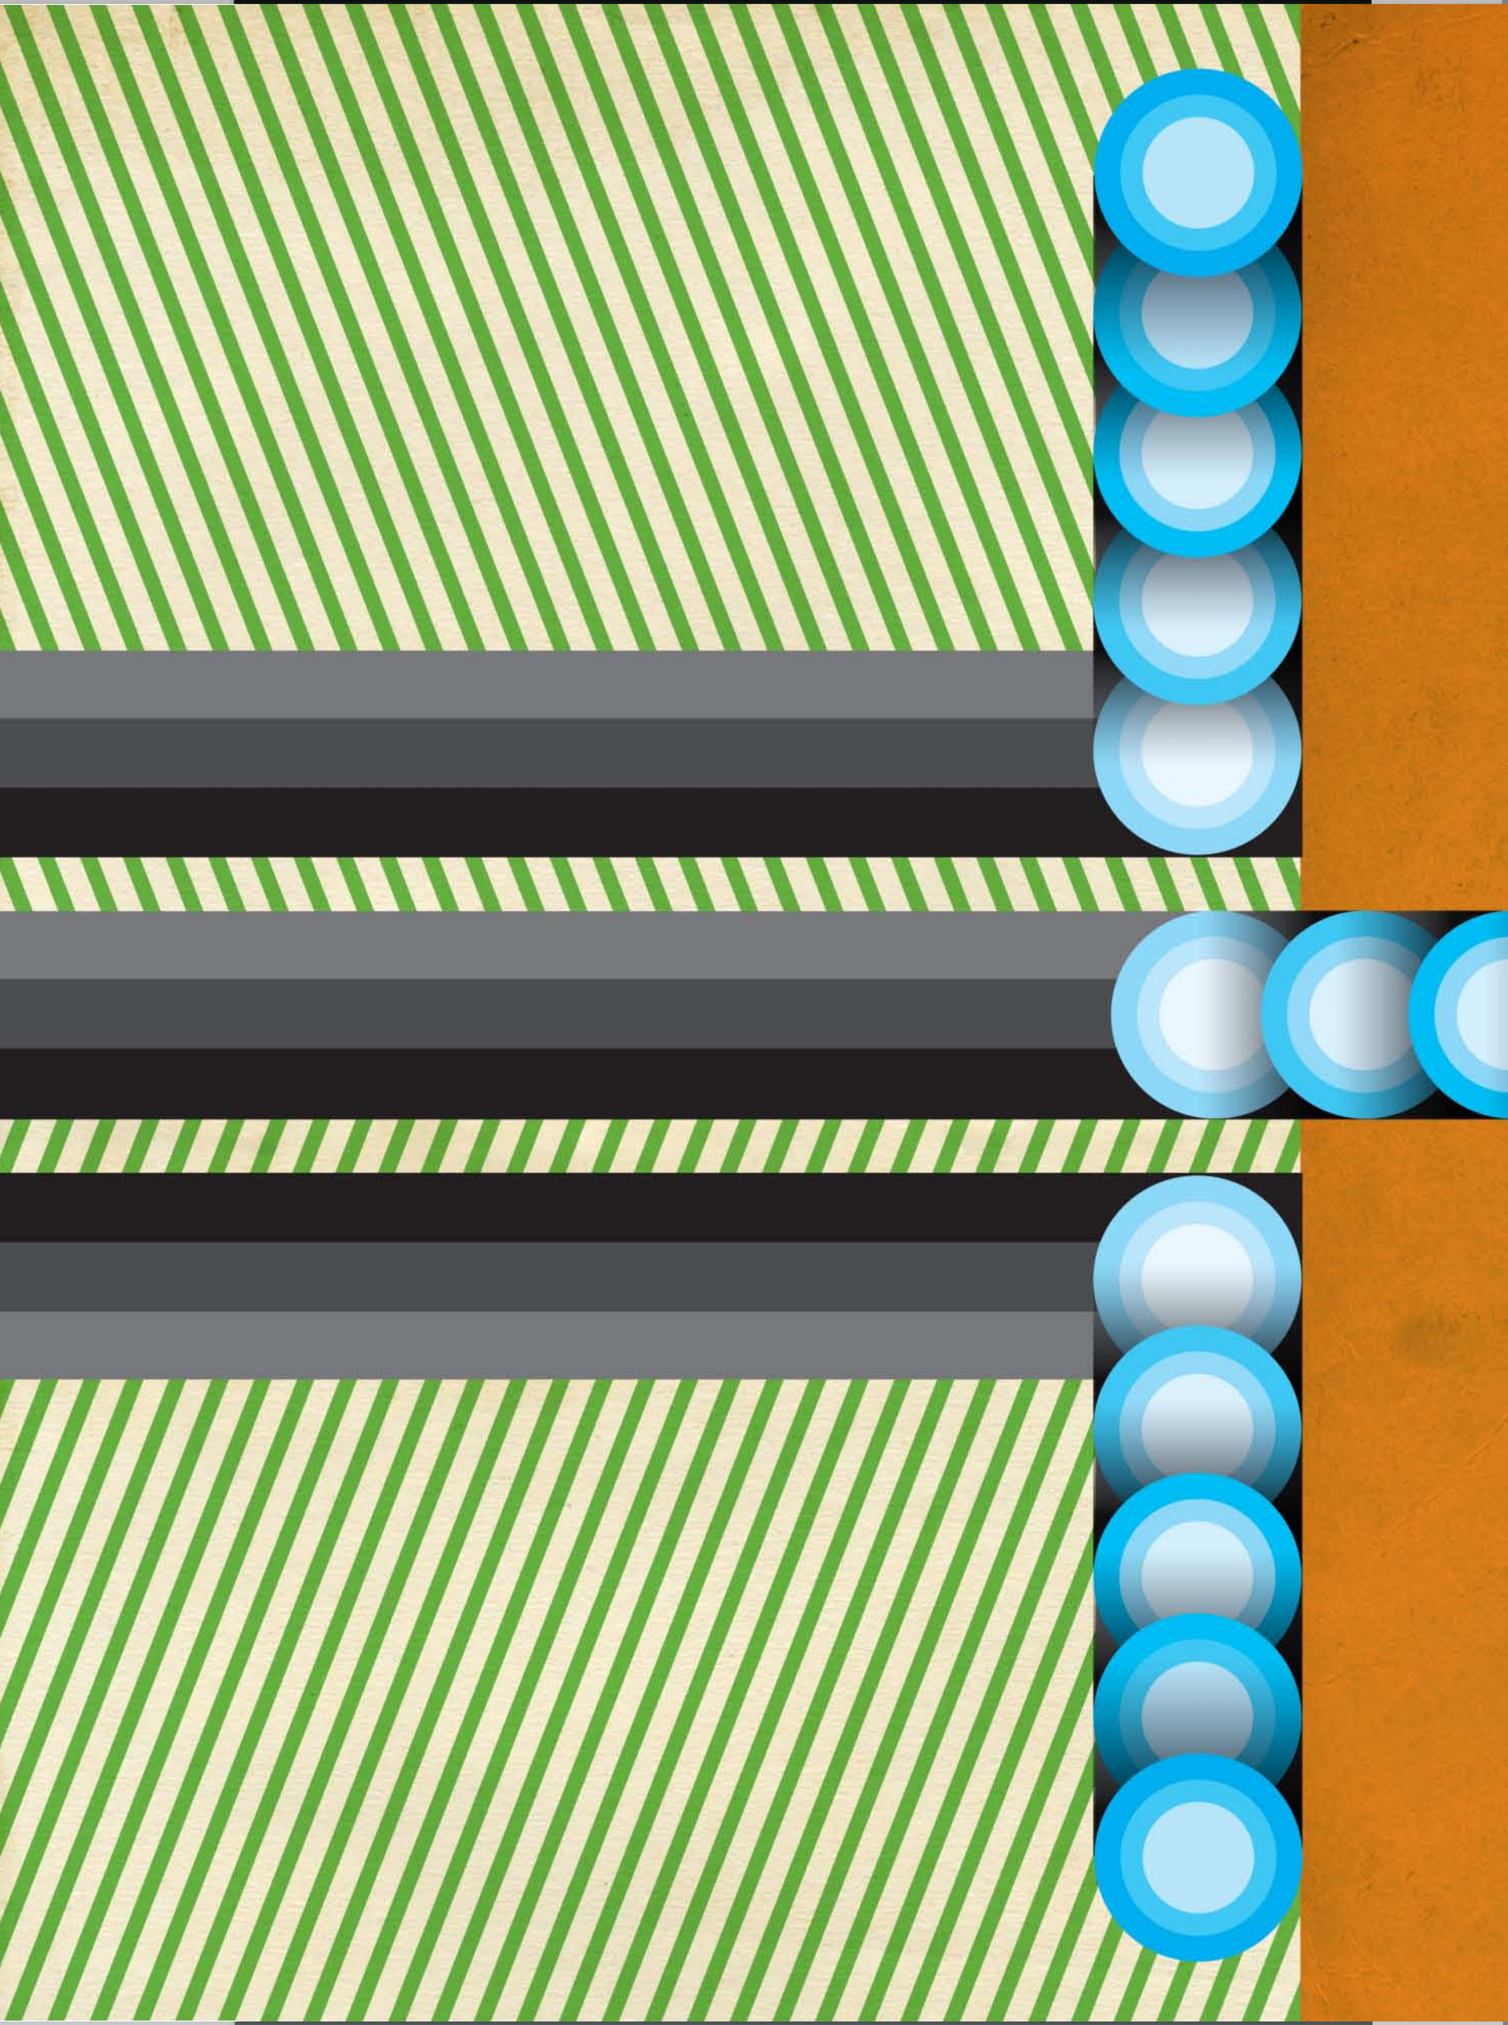

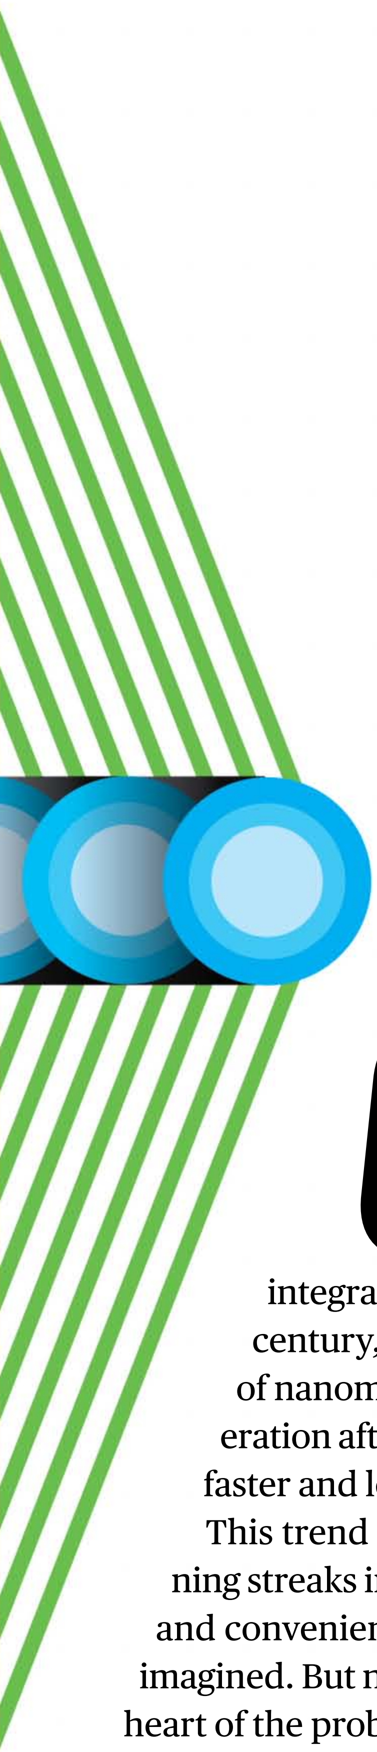

The Tunneling Transistor

QUANTUM TUNNELING IS A LIMITATION IN TODAY'S TRANSISTORS, BUT IT COULD ALSO MARK THE PATH FORWARD ••• *BY ALAN SEABAUGH*

Our always-on world of PCs, tablets, and smartphones has come about because of one remarkable trend: the relentless miniaturization of the metal-oxide-semiconductor field-effect transistor, or MOSFET. This device, which is the building block of most integrated circuits, has shrunk a thousandfold over the past half century, from the tens-of-micrometers scale in the 1960s to tens of nanometers today. And as the MOSFET has become tinier, generation after generation, the chips based on it have become much faster and less power hungry than their predecessors.

This trend has given rise to one of the longest and greatest winning streaks in industrial history, bringing us gadgets, capabilities, and conveniences that previous generations could scarcely have imagined. But now this steady progress is under threat. And at the heart of the problem lies quantum mechanics.

ILLUSTRATION BY Mark Montgomery

BACK OR THROUGH: In classical electrodynamics, an electron [blue] would bounce back from an energy barrier [orange] if its energy did not exceed the barrier height. In fact, electrons have a finite probability of passing through the energy barrier. The thinner the barrier, the higher the probability that such a tunneling event might occur.

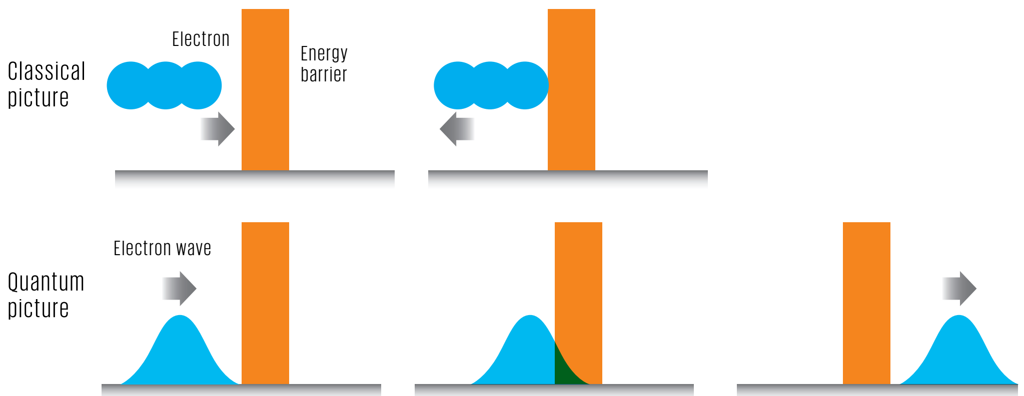

The electron has a pesky ability to penetrate barriers—a phenomenon known as quantum tunneling. As chipmakers have squeezed ever more transistors onto a chip, transistors have gotten smaller, and the distances between different transistor regions have decreased. So today, electronic barriers that were once thick enough to block current are now so thin that electrons can barrel right through them.

Chipmakers have already stopped thinning one key transistor component—the gate oxide. This layer electrically separates the gate, which turns a transistor on and off, from the current-carrying channel. Make this oxide thinner and you can induce more charge in the channel, boost the current, and make the transistor faster. But you can't reduce the oxide thickness to much less than roughly a nanometer, which is about where it is today. Beyond that, too much current will flow across the channel when the transistor is "off," when ideally no current should flow at all. And that's just one of several leakage points.

It has long been hard to pin down the precise year when size reductions will end. Industry road maps now project the miniaturization of the MOSFET out to 2026, when gates will be just 5.9 nanometers long—about a quarter the length they are today. This timeline assumes that we'll be able to find better materials to stanch leaks. But even if we do, we'll need to find a replacement for the MOSFET soon if we want to continue getting the performance enhancements we're used to.

We can't stop electrons from tunneling through thin barriers. But we can turn this phenomenon to our advantage. In the last few years, a new transistor design—the tunnel FET, or TFET—has been gaining momentum. Unlike the MOSFET, which works by raising or lowering an energy barrier to control the flow of current, the TFET keeps this energy barrier high. The device switches on and off by altering the likelihood that electrons on one side of that barrier will materialize on the other side.

That's a huge departure from the way traditional transistors work. But it might be just the thing to pick up where the MOSFET leaves off, paving the way for faster, denser, and more energy-efficient circuits that will extend Moore's Law well into the next decade.

IT WOULDN'T BE THE FIRST TIME the transistor has changed form. Initially, semiconductor-based computers used circuits made from bipolar transistors. But only a few years after the silicon MOSFET was demonstrated in 1960, engineers realized they could make two complementary switches. These could be combined to make complementary metal-oxide-semiconductor (CMOS) circuits that, unlike bipolar transistor logic, consumed energy only while switching. Ever since the first integrated

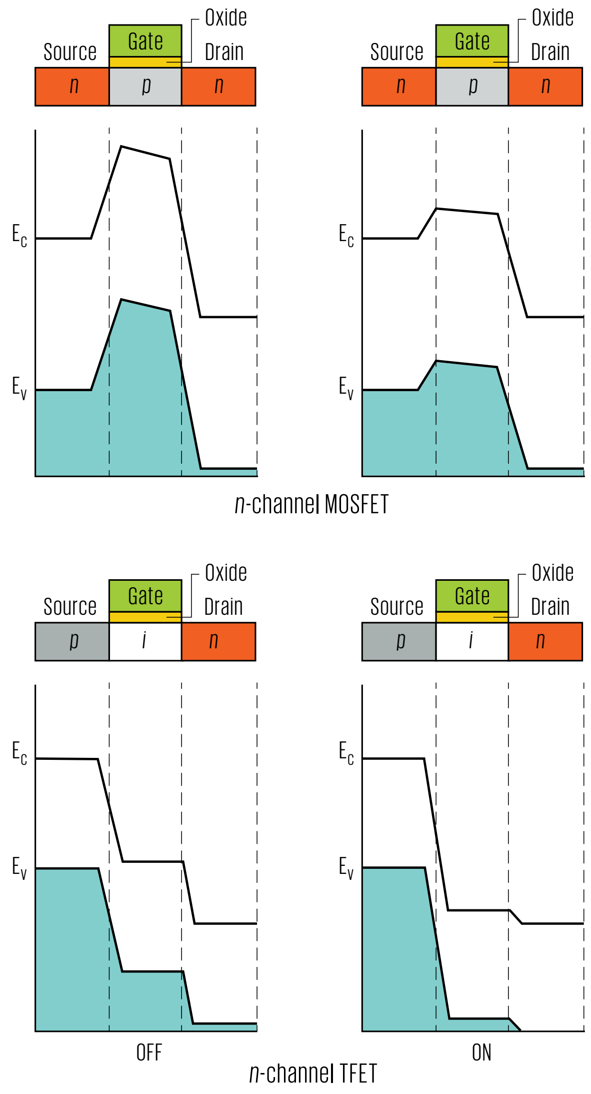

OFF AND ON: In an *n*-channel MOSFET, electrons move in the conduction band (E_c) from source to drain. The device's state can be switched from off [top left] to on [top right] if enough voltage is applied to draw down the energy barrier between the two regions. In an *n*-channel TFET, electrons originate in the source's valence band (E_v). A small voltage applied to the gate lowers the conduction band of the channel so it overlaps in energy with the source valence band, allowing electrons to tunnel into the channel.

circuits based on CMOS emerged in the early 1970s, the MOSFET has dominated the marketplace.

In many ways, the MOSFET wasn't a big departure from the bipolar transistor. Both control the current flow by raising and lowering energy barriers–a bit like raising and lowering a floodgate in a river. The "water" in this case consists of two kinds of current carriers: the electron and the hole, a positively charged entity that's essentially the absence of an electron in the outer energy shell of an atom in the material.

There are two allowable energy ranges, or bands, for these charge carriers. Electrons with enough energy to flow freely through the material are in the conduction band. Holes flow in a lower-energy band, called the valence band, and they move from atom to atom, much as an empty parking space might migrate around a nearly full parking lot as neighboring cars pull in and out.

These bands are fixed, but we can shift the energies associated with them up or down by adding impurities, or dopant atoms, to alter the conductivity of the semiconductor. *N*-type semiconductors, which are doped to contain an excess of electrons, conduct negatively charged electrons; *p*-type semiconductors, which are doped to produce a deficit of electrons, conduct positively charged holes.

If we put these two semiconductor types together, we get a junction with bands that are misaligned, thus creating an energy barrier between them. To make a MOSFET, we insert one type of material between two of the complementary type, in either an *n-p-n* or a *p-n-p* configuration. This creates three regions in the transistor: the source, where charges enter the device; the channel; and the drain, where they exit.

The two *p-n* junctions in each transistor provide electronic barriers to the flow of charges, and the transistor can be switched on by applying a voltage to the gate on top of the channel. A positive voltage applied to an *n*-channel MOSFET makes the channel more attractive to electrons, because it decreases the amount of energy an electron needs to have in order to move into the channel. A negative voltage applied to a *p*-channel MOSFET will perform the same task for holes.

This simple barrier-lowering strategy is the most widely used current-control mechanism in semiconductor electronics. Diodes, lasers, bipolar transistors, thyristors, and most field-effect transistors all take advantage of it. But it has a physical limit: Transistors need a certain amount of voltage to be switched on or off. This arises from the fact that electrons and holes are in constant motion due to their thermal energy, and the most energetic among them spill over the barrier. At room temperature, the current flowing over the barrier increases by a factor of 10 when the energy barrier is lowered by 60 millivolts; every "decade" of current change requires a change of 60 mV.

All this current leakage occurs below the device's threshold voltage, which is the voltage needed for the transistor to turn on. Device physicists call this barrier-lowering region the subthreshold region, and 60 mV per decade is known as the minimum subthreshold swing. To keep power consumption down, subthreshold swing should be as low as possible. The device will then need less voltage to be switched on, and it will leak less current when it's off.

Subthreshold swing wasn't much of an issue in the past, in the days when chips ran with higher voltages. But now it's starting to interfere with our ability to drive down power. That's partly due to the fact that circuit designers want to make sure their logic components have a big difference between the current that is used to define a "0" and the current that defines a "1." Transistors are typically designed so that when they're on they carry about 10 000 times as much current as they leak when they're off. That means at least 240 mV must be applied to the transistor to turn it on: four decades of current, each decade requiring at least 60 mV.

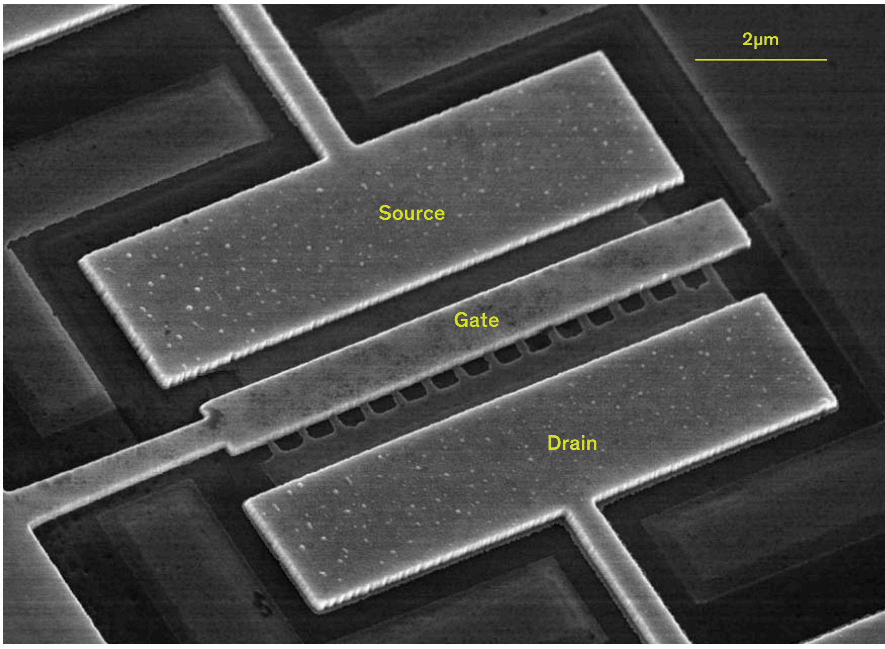

BRIDGE AND TUNNEL: Pairing semiconductors made of different compounds can dramatically boost current. This TFET uses aluminum gallium antimonide to make the source and drain regions of the device. A comblike air bridge, made of indium arsenide, is used to connect the channel to the drain for better electrical isolation. Metal air bridges are also used to wire the source, gate, and drain.

R. LI/ UNIVERSITY OF NOTRE DAME

In practice, the operating voltages used in CMOS circuits are typically much higher, closer to 1 volt. That's because the fundamental logic circuit in CMOS, the inverter, uses two transistors in series. The NAND gate takes three transistors in series, which means it requires even more voltage than the inverter. If you adjust for process variability–which means you need to set wider voltage margins to account for variations from device to device–you arrive at the voltages needed today to guarantee operation.

These voltage requirements, coupled with the leakage problems, mean we're in the waning days of MOSFET miniaturization. There's no way around it. If we want to lower voltage further to cut down on power consumption, we have two equally unattractive options. We can lower the current we drive through the device, which lowers the switching speed and thus sacrifices performance, or we can keep the current high and allow more current to leak through the device when it's supposed to be off.

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

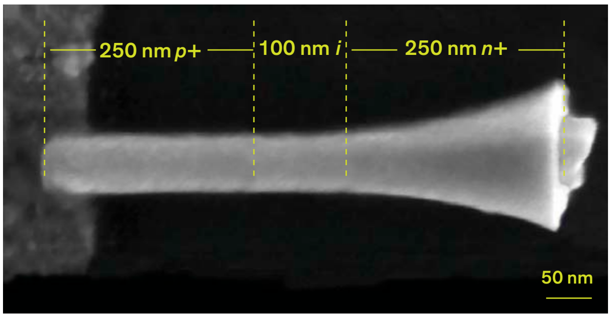

ALL-AROUND DEVICE: Today's cutting-edge transistors are three-dimensional, with gates that drape around three sides of a finlike channel. The TFET above employs a gate that wraps entirely around the channel. In this device, charge carriers move from left to right through source, channel, and drain regions made of indium gallium arsenide.

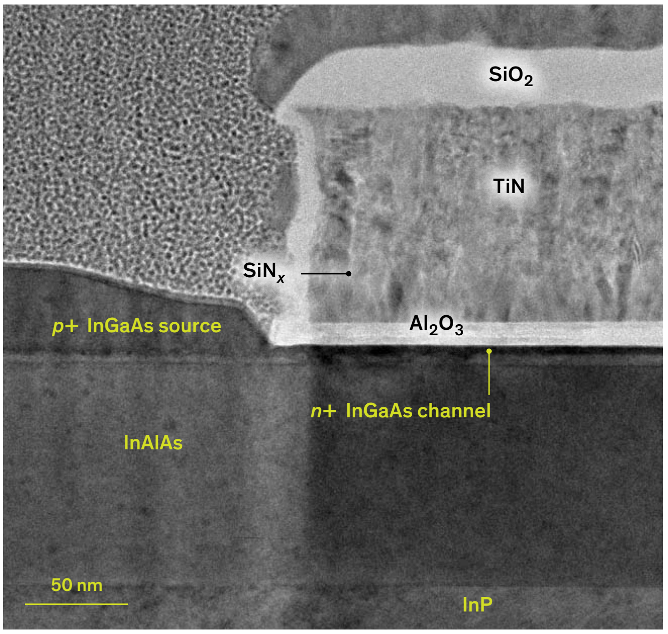

LAYER ON LAYER: The same etch-and-deposition processes used to make metal oxide gate stacks in today's silicon chips can be used to make TFET transistor regions [left]. This close-up view shows the region containing the source and channel of a TFET (the drain is at the right, out of view). The source and channel are made of oppositely doped indium gallium arsenide. The device is controlled by a gate made of titanium nitride, which is isolated from the channel by a layer of aluminum oxide.

That's where the tunnel FET comes in. Instead of raising or lowering the physical barrier between the source and drain as you would in a MOSFET, we use the gate to control the effective, electrical thickness of the barrier and thus the probability that electrons can slip through it.

The trick to doing this is again the *p-n* junction—with a bit of a twist. In a TFET, we arrange semiconducting material in *p-i-n* and *n-i-p* configurations. The i stands for "intrinsic," and it means that the channel has as many electrons as there are holes. The intrinsic state corresponds to the maximum resistivity that a semiconductor can have. It also pushes up the energies associated with the bands in the channel, introducing a thick energy barrier that charge carriers in the source are unlikely to traverse.

Electrons and holes obey the laws of quantum mechanics, which means they have a fuzzy, uncertain size. When an energy barrier has a thickness below about 10 nanometers, there is a small but nonzero probability that an electron that starts on one side of the barrier will appear on the other.

In the TFET, we boost this probability by applying a voltage to the transistor gate. This causes the conduction band in the source and the valence band in the channel to overlap, opening up a tunneling window. Note that in a TFET, the electrons tunnel between conduction and valence bands as they move into the channel. This is in contrast to what happens in a MOSFET, in which electrons or holes travel primarily in either one band or the other all the way from source to channel to drain.

Because the tunneling mechanism isn't controlled by the flow of carriers over a barrier, TFETs should be able to switch with a much smaller voltage swing than that required in a MOSFET. You have to apply only enough voltage to create or remove an overlap, crossing and uncrossing the bands [see "Off and On," bottom half of illustration].

AS A DEVICE MECHANISM, tunneling is not a new idea. The flash memory inside our USB sticks, cellphones, and other gadgets uses tunneling to inject electrons across oxide barriers into charge-trapping regions. Tunnel junctions like the one used in the TFET are also widely used to connect multijunction solar cells and to trigger semiconductor-based quantum cascade lasers. And tunneling governs the way current flows across metal-semiconductor contacts, an essential part of every semiconductor device.

The *p-n* tunnel junction has also been around a while. It was first demonstrated and explained by Nobel Prize winner Leo Esaki in 1957. But it took a fundamental impediment to get the industry to think seriously about how tunneling might be applied to logic.

The first TFET papers were written only about nine years ago, when chipmakers started to see computer clock speeds stall and struggled with the problem of removing heat from denser, leakier chips.

Joerg Appenzeller and his colleagues at IBM were the first to demonstrate that current swings below the MOSFET's 60-mV-per-decade limit were possible. In 2004, they reported they had created a tunnel transistor with a carbon nanotube channel and a subthreshold swing of just 40 mV per decade. Within a few years, groups at UC Berkeley, CEA-Leti, Imec, and Stanford had followed suit. They showed that switches that consume less than 60 mV per decade could be made using semiconducting materials that are staples of the chip industry: silicon and germanium.

That got the community excited, because although the current-control mechanism in the TFET is new to the semiconductor industry, the device bears a strong resemblance to the MOSFET. It has the same basic configuration of source, drain, and gate and similar electrical behavior when wired into circuits. The semiconductor design infrastructure does not need to change.

BUT SOME CHANGES ARE REQUIRED. It turns out that silicon and germanium aren't great for tunneling. It's for the same reason that these materials don't make good light emitters and lasers. Silicon and | CONTINUED ON PAGE 56

LEFT: N. GOEL/C. PARK/SEMATECH/UNIVERSITY OF NOTRE DAME; RIGHT: T. MAYER/PENN STATE

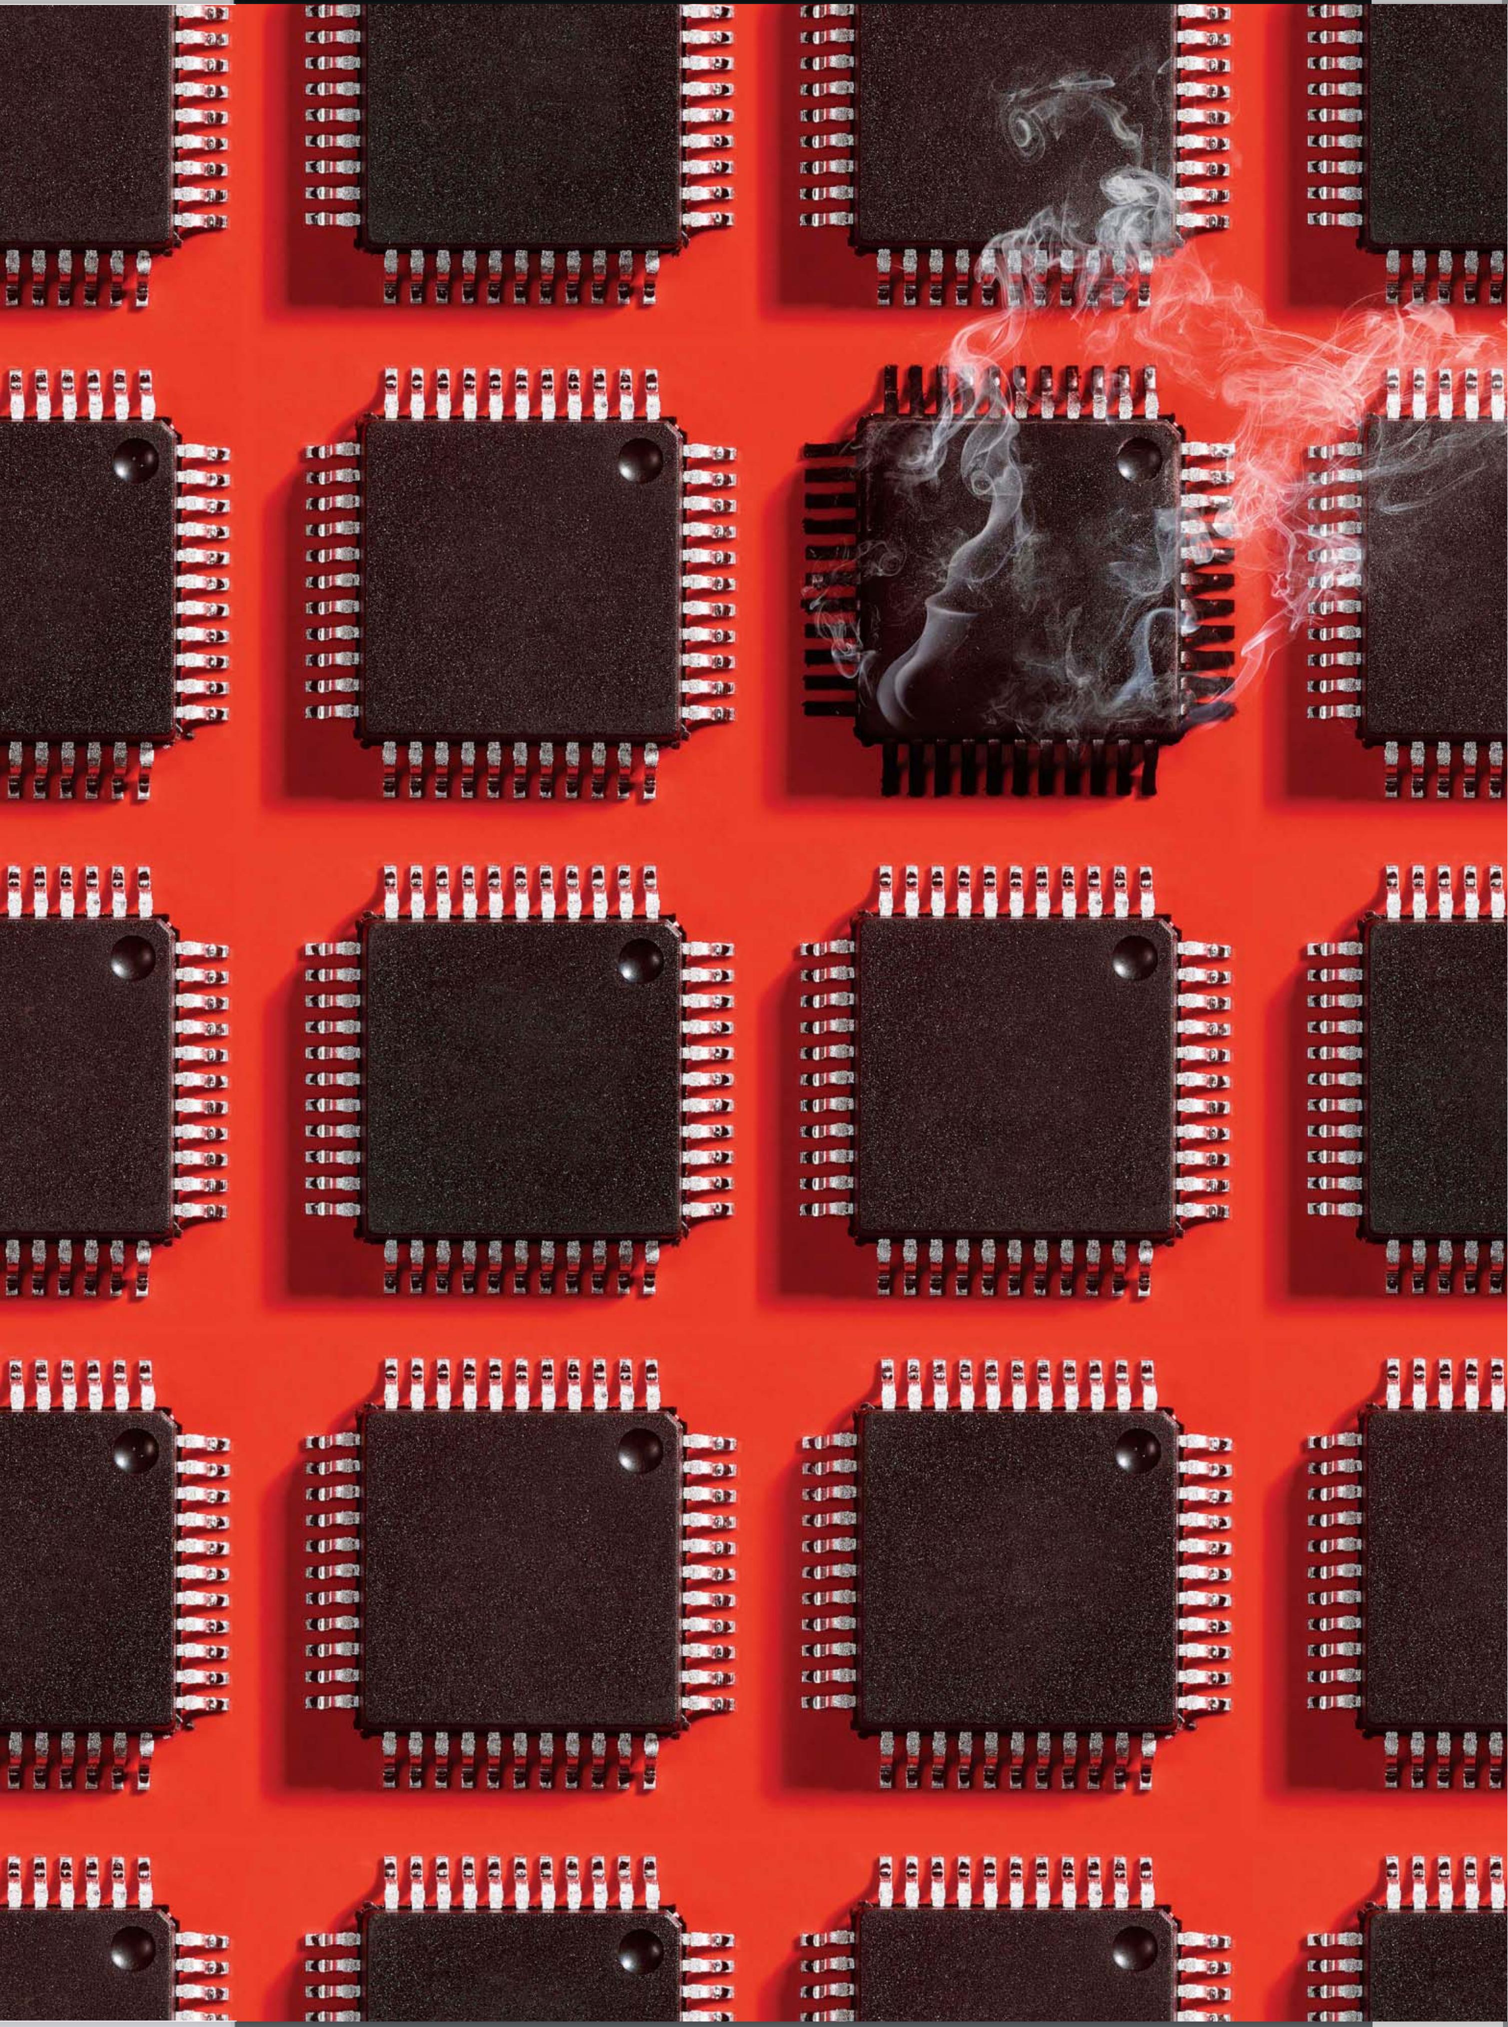

CHOP-SHOP ELECTRONICS

CLEVER COUNTERFEITERS SELL OLD COMPONENTS AS NEW, THREATENING BOTH MILITARY AND COMMERCIAL SYSTEMS

BY JOHN VILLASENOR & MOHAMMAD TEHRANIPOOR

ON 17 AUGUST 2011, BOEING WARNED the U.S. Navy that an ice-detection module in the P-8A Poseidon, a new reconnaissance aircraft, contained a "reworked part that should not have been put on the airplane originally and should be replaced immediately." In a message marked "Priority: Critical," the company blamed the part, a Xilinx field-programmable gate array (FPGA), for the failure of the ice-detection module during a test flight.

How could this have happened? Xilinx, based in San Jose, Calif., is a highly respected manufacturer of FPGAs, and Boeing bought the ice-detection module containing the suspect

PHOTOGRAPH BY Adam Voorhes

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

part from BAE Systems, a reputable British defense company. The trouble occurred somewhere in the supply chain upstream from BAE, which wound through companies in California, Florida, Japan, and China. However, retracing that FPGA's path led not to Xilinx but to a Chinese company called A Access Electronics. It apparently had turned a quick profit by selling used Xilinx parts as new. BAE ended up purchasing about 300 suspect FPGAs, many of them untested. Fortunately, most had not yet been installed on planes.

To a suspicious buyer, a number of clues might have suggested that the parts were not brand new. For example, parts stamped with the same manufacturing lot code had different ceramic package shapes and four different date codes. Some of the pins were shorter than the length specified in the manufacturer's data sheet. Some packages themselves were chipped. But somewhere along the supply chain, someone accepted used parts as new, and they ended up on U.S. Navy airplanes.

This incident, described in a hearing held by the U.S. Senate Committee on Armed Services in November 2011, is only the tip of the iceberg. The global trade in recycled electronics parts is enormous and growing rapidly, driven by a confluence of cost pressures, increasingly complex supply chains, and the huge growth in the amount of electronic waste sent for disposal around the world. Recycled parts, relabeled and sold as new, threaten not only military systems but also commercial transportation systems, medical devices and systems, and the computers and networks that run today's financial markets and communications systems.

USUALLY, IT'S GOOD to recycle electronics products. Ethical recycling companies properly dispose of the toxic ingredients in discarded computers, printers, mobile phones, and other systems and harvest the precious metals they contain. But not everyone in the electronics recycling ecosystem acts ethically.

In theory, with proper screening, electronics components could be safely reused in some low-cost applications, such as calculators and remote controls, where tolerances aren't particularly demanding. But given the low cost of electronics parts for those products, such reuse wouldn't usually be worth the trouble. There's far more economic incentive to use recycled components as replacement parts in more expensive systems with long service lifetimes.

But sometimes old components are misused. Some companies have built a business model based on pulling old parts from cast-off products and reselling them as new. Relabeled or otherwise altered parts masquerading as new can fail prematurely in critical systems, such as those in airplanes and cars, with potentially catastrophic results.

A SAVVY CONSUMER can often tell with a glance whether a designer handbag or a pair of shoes is genuine or counterfeit. But electronics counterfeits, hiding deep within products and systems, are not so easy to detect. Very few of us open up gadgets to inspect the components inside, and even if we did, almost none of us would know how to distinguish good components from bad. Unethical recyclers of electronics parts don't need to fool everyone in a supply chain; often they just need to fool a single company among the many that sit upstream from an end product.

Today's supply chains are built largely on trust. Company A sells a part to Company B and warrants that it will function as specified. Company B may then incorporate that part into a subsystem and sell the subsystem to Company C, in turn warranting that the subsystem will perform as designed. A part can pass through a half dozen or more different intermediaries before it ends up in a finished product. To move undetected through this chain, recycled parts need only function well enough to pass a few tests that are conducted along the way–tests usually designed to weed out accidental defects and design flaws, not to identify parts that counterfeiters have specifically altered to masquerade as something they are not.

Most known cases of electronics parts counterfeiting have been uncovered by military supply chain investigations, sometimes triggered when a key component fails. In part, that's a reflection of the mismatch between broader electronics-industry obsolescence cycles, which may be as short as one or two years, and the life cycles of defense systems, which may sometimes be decades. A market for replacement parts that were last manufactured a decade ago creates ample opportunity for unethical suppliers to relabel and resell scavenged parts.

But that's not the whole story. The P-8A Poseidon aircraft, for example, is a relatively new platform, so it's not just old systems that are targeted. Military-systems manufacturers find counterfeit parts more often than do their commercial brethren, in part because they test more often and more rigorously–they need to do so to make sure the parts will work under extreme conditions. And under a U.S. law in place since the end of 2011, U.S. government contractors are obliged to be more diligent in screening parts destined for military systems, in an effort to identify potential counterfeits.

There's much less screening going on in the commercial community, and the consumers and businesses that purchase electronics devices and systems are often unaware of the potential risks posed by the parts they contain. Recycled parts may get into the consumer-product supply chain as well and just go undetected,

COUNTERFEITS ON BOARD: An investigation by the U.S. Senate Committee on Armed Services identified suspect components in the supply chains for [from left] the CH-46 Sea Knight helicopter, the C-17 military transport aircraft, the P-8A Poseidon sub hunter, and the F-16 fighter.

BELOW, FROM LEFT: ANDREW SCHMIDT; KENN MANN/U.S. AIR FORCE; BOEING; MASTER SGT. ANDY DUNAWAY/U.S. DEFENSE IMAGERY although this is less likely when companies manufacture products in large volumes and purchase parts directly from their manufacturers or their approved distributors.

RECYCLED PARTS ARE just one segment of a vast and fast-growing counterfeit electronics industry. In 2011, the Semiconductor Industry Association pegged the cost of electronics counterfeiting at US $7.5 billion per year in lost revenue and tied it to the loss of 11 000 U.S. jobs. The problem isn't limited to the United States; electronics counterfeiting is a global concern.

Besides improper recycling, counterfeiting also includes parts that are made in an authorized production run but fail testing and are sold anyway instead of being destroyed, excess inventory intended for the scrap heap that isn't disposed of properly, and some parts that are simply phony from the beginning and don't work at all. But the most pervasive problem is recycled parts, accounting for 80 to 90 percent of counterfeit parts in circulation, according to a 2010 estimate by SMT Corp., based in Sandy Hook, Conn.

The U.S. government has long recognized the risks posed by counterfeit parts. A 2010 U.S. Department of Commerce report on counterfeit electronics in the Department of Defense supply chain noted that between 2005 and 2008 the number of companies reporting incidents involving counterfeit chips more than doubled. In February 2012, market-intelligence firm IHS iSuppli stated that "reports of counterfeit parts have soared dramatically in the last two years," with "1363 separate verified counterfeit-part incidents" reported worldwide in 2011.

In some instances, counterfeit parts end up in high-performance military systems like the Boeing aircraft. In another case that's become public, an investigation led by a branch of the U.S. Department of Homeland Security shut down Florida-based VisionTech in 2010 and arrested several of its executives. The prosecution charged certain individuals at VisionTech with importing thousands of shipments of confirmed or suspected counterfeit parts from China that had been improperly marked as military grade, some of which were purchased by Raytheon Missile Systems and installed on circuit boards intended for use in a radar-detection system on F-16 aircraft. Fortunately, Raytheon was able to identify the presence of the counterfeits during testing of the boards and didn't install them in the planes.

In 2011, the U.S. Senate Committee on Armed Services collected data on 1800 cases involving a million suspect parts. The committee investigated about 100 of those cases and found defective parts in the supply chain for the Air Force C-17 military transport, the Marine Corps CH-46 helicopter, an Army missile defense system, and more. The committee asked the Government Accountability Office to set up a fake company to purchase electronic parts for military systems online. The GAO found suppliers willing to sell components having imaginary part numbers, and even when the GAO provided real part numbers, the parts it received were counterfeit.

CLOSE-UP: A specialist in counterfeit electronics at SMT Corp. uses a handheld USB microscope to examine a suspect component.

ABOVE: SMT CORP.

THE COUNTERFEITING SCAM begins at the foot of a mountain of discarded electronics products that grows by some tens of millions of tons annually.

To harvest the components, recyclers heat circuit boards to a temperature high enough—sometimes up to 400 °C—to melt the solder that attaches them to the boards. Unethical recyclers with little concern for how components will later be used may then bang the boards repeatedly against a hard object to dislodge the parts, which they clean and then sort by size, package style, number of pins, part number, and manufacturer name.

Parts that emerge from this process with their labels and pins undamaged go promptly back into the supply chain, and a dishonest components distributor can quickly sell them as new. If however, extracting the components damages the labels or pins, there's more work to do. Pins are simple to straighten, and labels can easily be changed. Unethical distributors often alter even undamaged labels to falsely indicate that a part has a higher performance or operating range than it actually does.

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

TO CATCH A COUNTERFEIT

Investigators at SMT Corp. use a variety of signs to spot a counterfeit component. Some counterfeits can be detected by just a visual inspection, but spotting some sophisticated alterations requires other tests.

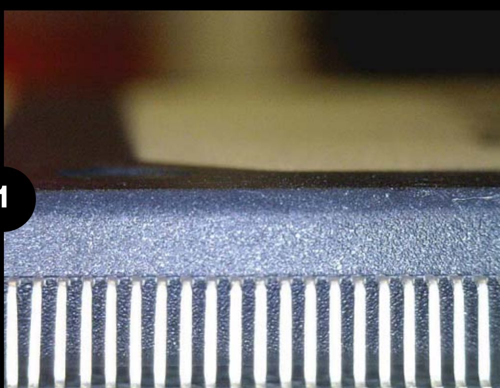

1

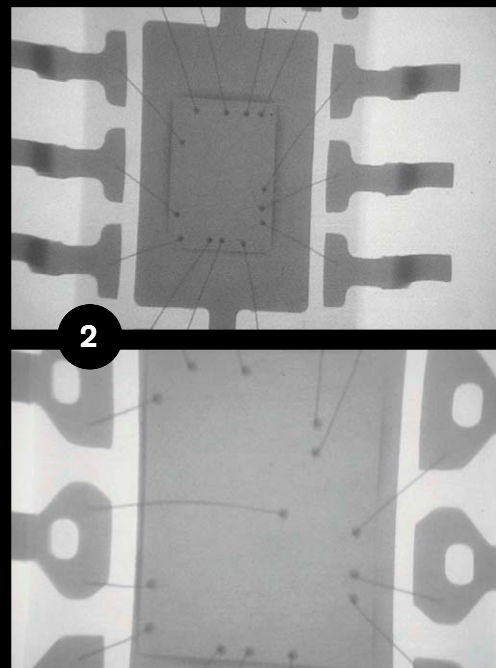

2

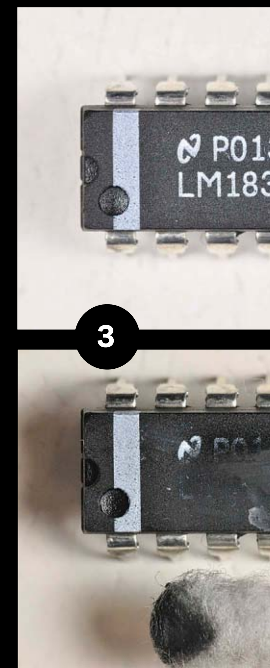

3

To relabel a part, the distributor sands off the original markings and applies a black coating—a "blacktopping"—that's almost indistinguishable from the original topcoat. Traditionally, recyclers used blacktopping materials that resembled paint; these could be detected using a simple test such as swabbing the part with acetone, which would dissolve the coating. In 2009, however, SMT discovered that counterfeiters had developed sophisticated blacktop materials that were chemically more similar to the original coatings and therefore much harder to detect. Once a new blacktop coating is in place, the counterfeiter relabels the parts by printing new text onto the blacktop coating, and the recycled parts are ready to sell.

Parts that emerge from this process pose many risks. First, to state the obvious, they are no longer new. This means that expectations regarding performance and lifetime are off. Second, the removal process itself, which exposes parts to high heat, water, chemicals, physical impacts, and other stresses, can damage them. Third, fraudulently labeled parts may end up in the hands of unsuspecting buyers, who use them in unsuitable operating environments where excessive temperature, humidity, or vibrations could lead to premature failure.

TECHNOLOGY OFFERS SOME solutions to the counterfeiting problem. Companies can scrutinize packages for signs that pins have been straightened or indications that labels have been sanded or repainted. They can also perform more detailed analyses, using X-ray, scanning electron, or acoustic imaging to look inside a package for things that might be amiss, like the improper placement of a chip within its package.

Another telltale feature is a chip's electrical behavior. Age makes transistors sluggish, so signals can take longer to pass from one portion of a chip to another. Statistical analysis of these path delays and other electrical attributes can play an important role in counterfeit detection.

Applied DNA Sciences, based in Stony Brook, N.Y., is taking another step to make it easier to detect counterfeits: It's mixing DNA synthesized from gene sequences found in plants into the ink used to stamp chip packages. The company keeps the specific gene sequences used to mark chips a secret so that few people can make or verify the marks.

EVEN THE MOST SOPHISTICATED technological countermeasures face challenges. As Tom Sharpe, vice president of SMT, observes, "Our industry has gotten much better at detecting counterfeits—but the counterfeiters are getting better as well."

In this cat-and-mouse game, the economic equation gives an advantage to the unethical distributors and the recyclers who supply them with parts. They buy discarded components for a song and then sell them at full price. They can then use some of the profits to finance the R&D that makes them ever better at hiding evidence of recycling. Legitimate component buyers and sellers, by contrast, often have razor-thin profit margins. Testing adds to their costs, which may put buyers who perform careful counterfeit-detection testing at a competitive disadvantage compared with buyers who do little or no testing. It would be even more costly, and in most cases impractical, for companies that buy subsystems to test every component within them.

It's bad enough that these recycled components are increasing the dangers for those who operate military airplanes, helicopters,

SMT CORP. (7)

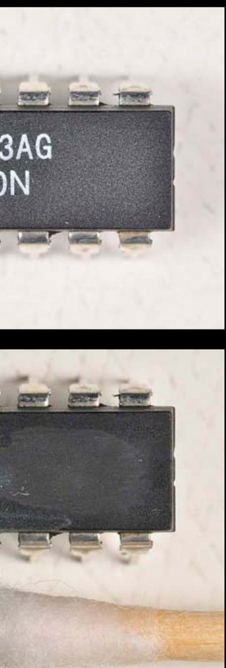

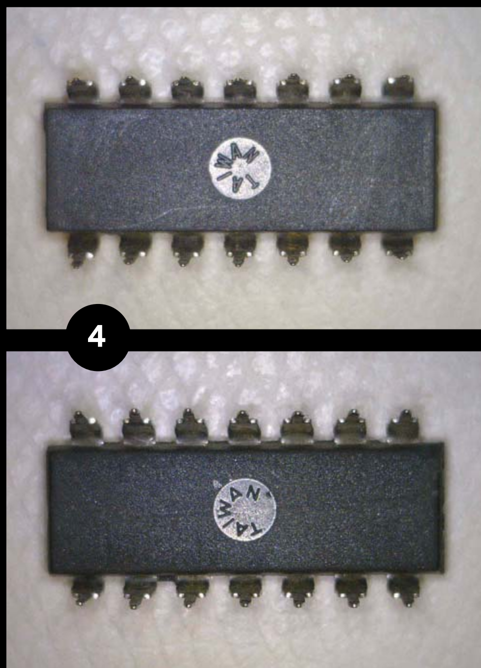

4

1] AT THE MARGINS: An uneven band of color along the top and side surfaces of a component package indicates the presence of resurfacing material.

2] X-RAY VISION: An X-ray of a component can reveal irregularities in the size and shape of the die, the routing of the wires, and other internal construction details. The device in the top image is authentic; the bottom image shows a suspected counterfeit with the same part number.

3] BAD BLACKTOP: Counterfeiters sometimes use resurfacing material to hide the original identification numbers on a component, but this coating, unlike the original black surface, dissolves in acetone. Here is a suspect component before and after an acetone swipe test.

4] MISMARKED: The variation in the appearance of the country of origin marks indicates that these samples came from different lots, but counterfeiters have marked both with the same lot number.

and ships. But from the standpoint of a counterfeiter, the commercial supply chain is in some ways a much more attractive target, because the market is much larger and more diversified, the level of testing is lower, and product life cycles are often shorter. This gives counterfeit parts less time to fail and provides counterfeiters with more opportunities to sell their wares.

The prospect of a commercial electronics supply chain laden with counterfeit parts is sobering: Commercial parts are used in routers, servers, storage hardware, and many other electronics systems that enable the communications, financial, transportation, power, and other critical infrastructure to run smoothly. While it is clearly not practical to subject every commercial chip to the level of testing used for flight-critical military systems, a well-designed, widely used system of randomized testing explicitly designed to detect counterfeit parts in the commercial supply chain could help dramatically. In addition, with such testing we might be able to picture the true scope of the counterfeit electronics problem. That picture, though, is likely to be a grim one.

Technological solutions are indeed vital for detecting counterfeit electronics. But detection technologies can identify problematic components only after they have entered the supply chain. Truly solving the problem will require making sure that far fewer counterfeit parts enter the supply chain in the first place. That means making it harder for counterfeiters to operate.

To put the squeeze on counterfeiters, we need a better way to track parts as they move through the supply chain and then use that information to call out unethical suppliers. Counterfeiters stay in business because they suffer few or no consequences for their actions. More-effective systems to log, report, and share information about unethical suppliers could put a big dent in their business. Some valuable industry and government counterfeit-reporting programs are already in place; some government agencies collect information about counterfeit incidents and report it to their suppliers. In addition, the Independent Distributors of Electronics Association, for example, keeps a list of suspected counterfeit incidents. But as important as these efforts are, they likely capture information on just a small fraction of the global trade in counterfeit electronics. We need more companies–including those not directly involved in manufacturing systems for government customers–to be aware of the problem and to improve their efforts at detecting, tracking, and reporting counterfeit parts.

We must also address this problem on an international level, and in doing so we must acknowledge the geopolitical realities of the counterfeit-parts industry. Counterfeiters stay in business because they are at the upstream end of an opaque, loosely regulated global market for parts. Trade negotiators need to talk more openly about the risks posed by counterfeit electronic components. Governments can help by closing regulatory loopholes that allow counterfeiters to operate and by increasing whistle-blower protections to encourage reporting. Together, these steps could make it much harder for recyclers and other counterfeiters to operate.

Now that electronic products are everywhere, the threat of recycled electronics parts is everywhere as well. Although we will never be able to eliminate the threat of recycled components completely, we can and should reduce the risks they pose. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/chopshop1013

Electronic Treasures of The Sarnoff Collection

A NEW EXHIBITION OFFERS A FASCINATING WINDOW INTO A VANISHED CORPORATE PAST

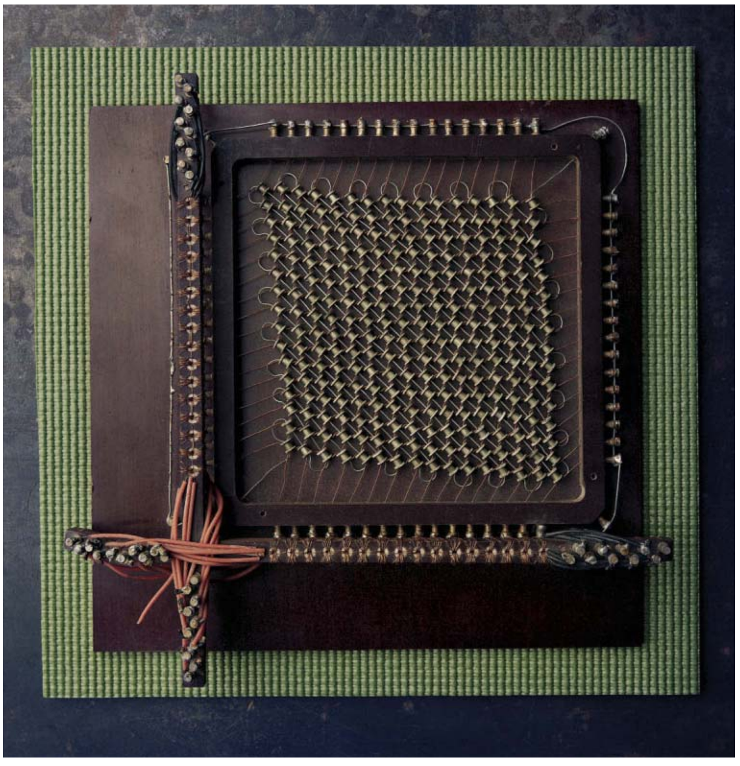

256-bit magnetic core memory (c. 1952) Slow data retrieval and storage speeds limited the utility of early computers. RCA researcher Jan Rajchman's solution was a memory array consisting of a wire matrix with doughnut-shaped magnetic cores at each intersection. By applying a current to a given set of horizontal and vertical wires, you could select a specific core and quickly change the direction of its magnetic field.

The history of the Radio Corporation of America is in many ways the history of 20th-century American innovation. From the company's founding in 1919 to its sale in 1986, the RCA name was synonymous with products that shaped how Americans lived and worked. Long before the rise of Silicon Valley, RCA Laboratories, in Princeton, N.J., was at the center of the nation's consumer electronics industry, harnessing the creative impulses of thousands of scientists, engineers, and technicians to systematize the invention of new technologies.

This month, a new exhibition highlighting RCA's rich history opens at the College of New Jersey, in Ewing. It draws from the more than 6000 artifacts that the college inherited after the David Sarnoff Library–RCA's main technical archive and museum–closed in 2009. (The IEEE Foundation funded a new study center connected to the exhibition.) The installation covers the development of radio, television, and broadcasting, as well as RCA's work in liquid-crystal displays, electron microscopy, solid-state physics, and computers. The stories it reveals highlight the challenges of managing complex technical projects and the effects of social, economic, and political trends on industrial research and development.

TEXT BY **BENJAMIN GROSS**
PHOTOGRAPHY BY **SUZANNE KANTAK**

Model 100 Radiola speaker (1925) Edwin Armstrong's ingenious superheterodyne circuit made it easier for radios to detect and amplify high-frequency signals. RCA licensed the technology from Armstrong and used the superheterodyne in its popular Radiola series of receivers, sold in the 1920s and early '30s.

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

David Sarnoff's garrison cap, RCA employee ID badge, military decorations, appointment orders »
Neither a scientist nor an inventor, RCA chairman David Sarnoff nevertheless oversaw the development of the technologies that came to define the information age. Under his leadership, RCA organized the first radio broadcasting network, perfected black-and-white and color television, and established a research center in Princeton, N.J., that made crucial contributions to digital computing, integrated circuitry, and flat-panel displays.

Born in what is now Belarus in 1891, Sarnoff was a fervent patriot of his adopted country. Upon hearing of the Japanese attack on Pearl Harbor, he sent a telegram to the White House: "All our facilities are ready and at your instant service. We await your commands."

RCA scientists and engineers went on to make major advances in military radar and sonar, as well as mobile broadcasting equipment. Dwight D. Eisenhower recruited Sarnoff to coordinate all radio traffic for the D-Day invasion of Normandy in June 1944. In December of that year, Sarnoff was made a brigadier general, as indicated by the sterling silver star on his garrison cap. From then on, he would be known as "General Sarnoff" or "The General."

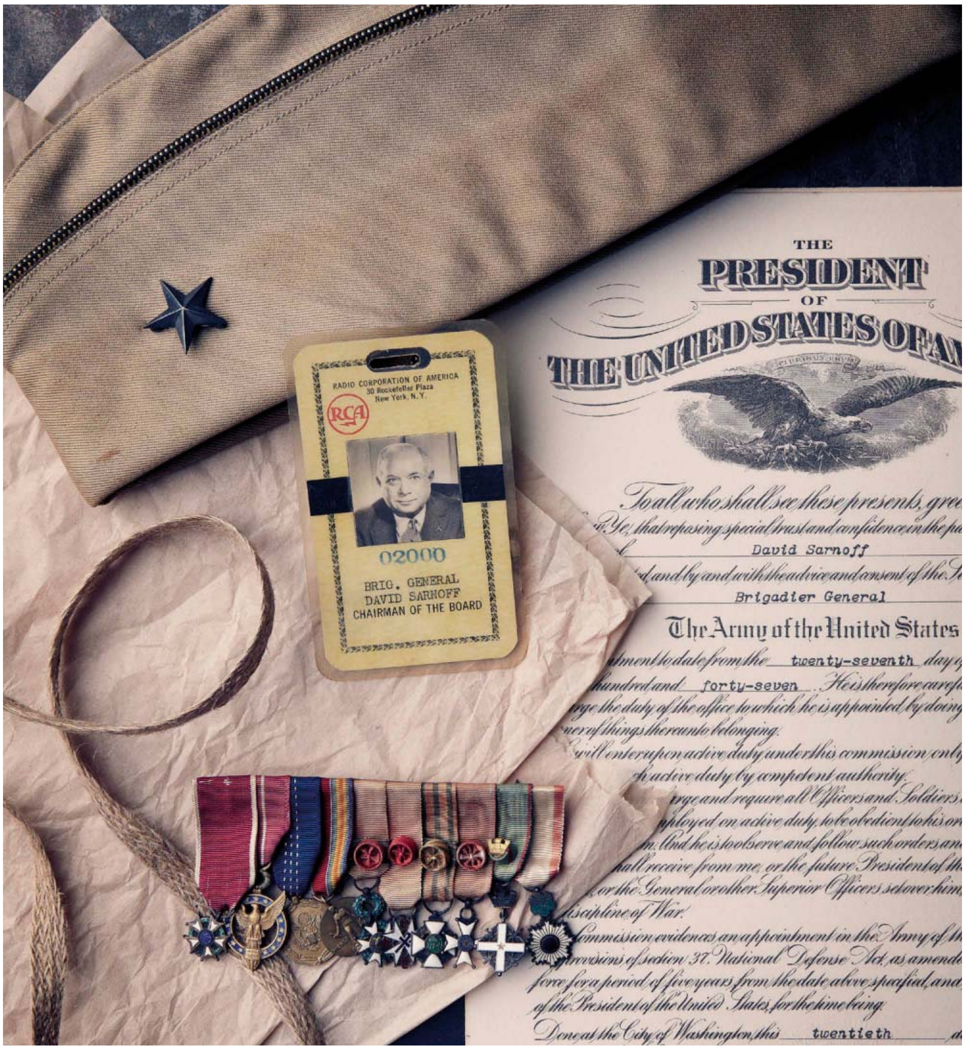

Telegraph key (c. 1912) »
Even as a young man, Sarnoff capitalized on every opportunity. A skilled telegraph operator, he was assigned in 1912 to manage Marconi's Wireless Telegraph Co.'s station at the Wanamaker department store in New York City. On the evening of 14 April 1912, the grand ocean liner RMS *Titanic* struck an iceberg and began transmitting distress signals. Sarnoff came on duty soon after and stayed at his post for three straight days, using this telegraph key to contact rescue ships and compile the names of survivors. He later referred to the incident as a turning point in his career, noting that "the *Titanic* disaster brought radio to the front, and incidentally me."

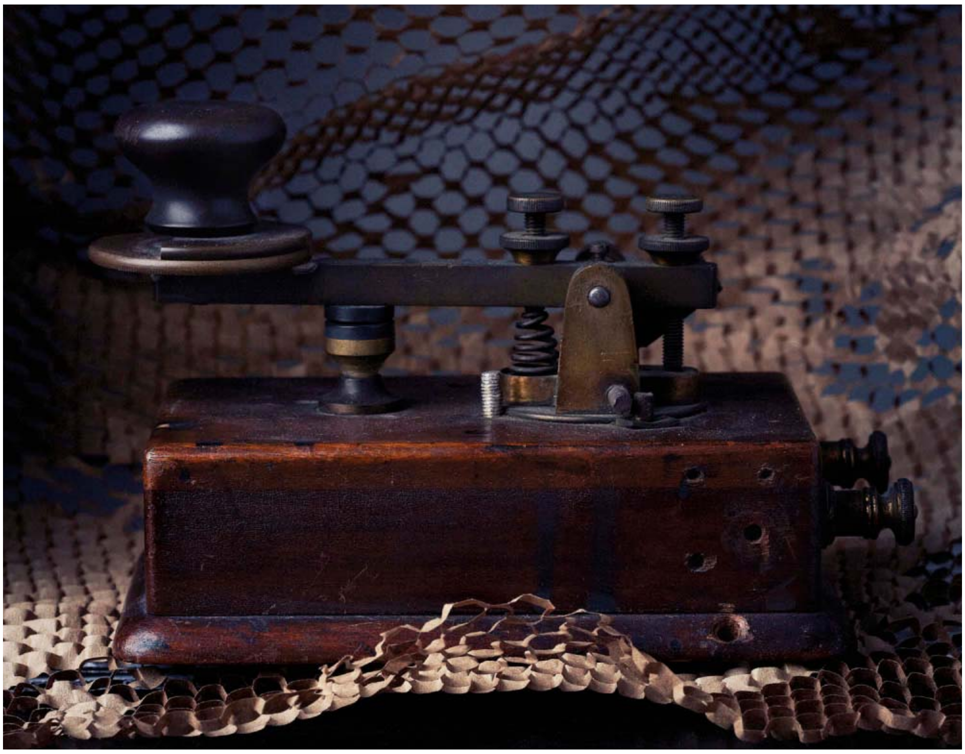

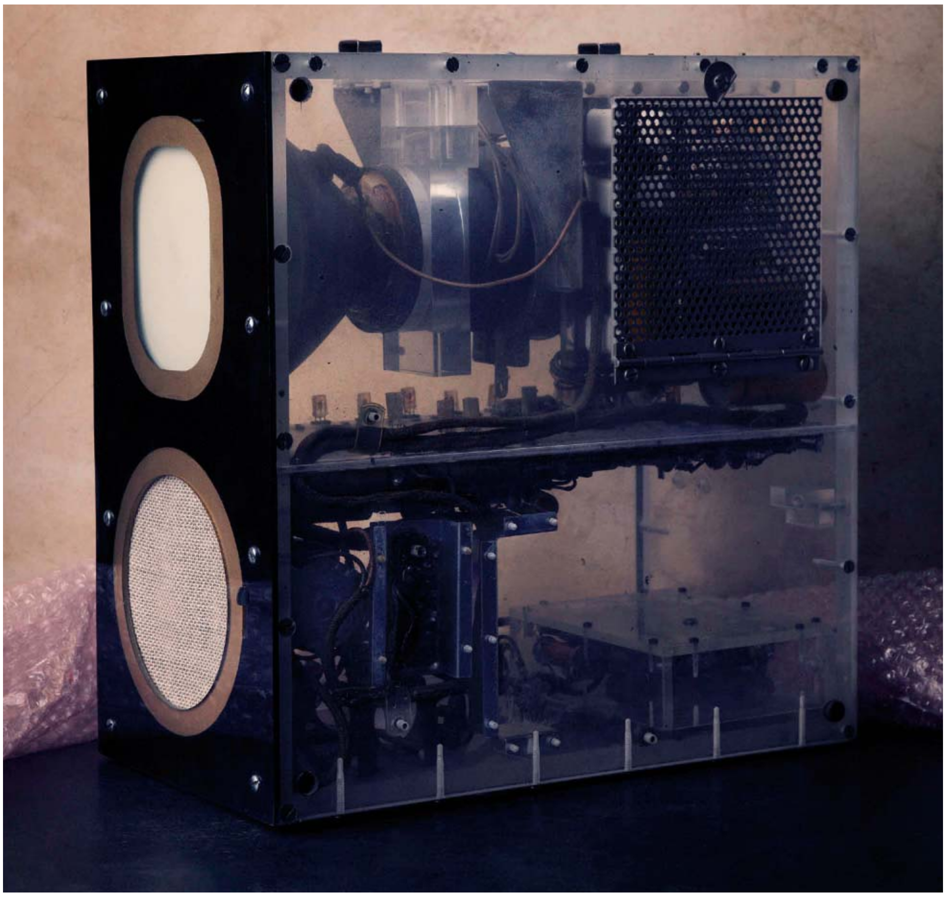

« Transistor TV prototype (1952)
In November 1952, RCA demonstrated the world's first transistorized television. Developed by George Sziklai, Robert Lohman, and Gerald Herzog, this 5-inch set contained 37 transistors, weighed 12 kilograms, and had only one channel.

Miniature image orthicon tube (late 1940s; bottom left)
During World War II, RCA's work on a television-guided bomb culminated in the image orthicon, a camera hundreds of times as sensitive as existing products. A miniaturized version was mounted in an experimental missile but never used in combat. The "immy" later became the standard camera in TV broadcasting and the namesake of the Emmy Award.

First color TV set, CT-100 (1954)
In the 1950s, RCA vied with the Columbia Broadcasting System to introduce a color TV system. CBS's offering had a superior picture but wasn't compatible with monochrome sets. So RCA refined its system, and in December 1953 the Federal Communications Commission endorsed its compatible color standard. RCA's first color TV, the CT-100 sold for US $995—five times the price of a black-and-white set.

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

Prototype propaganda phonograph and Cold War memorandum (1955) »
During the Cold War, David Sarnoff devised a new weapon: a hand-powered plastic phonograph intended to be air-dropped behind the Iron Curtain, along with records made of cardboard. Not only would this device cost less to manufacture than a radio, but its prodemocracy messages would be impossible for enemies to jam. Sarnoff floated the idea to President Eisenhower, and the United States Information Agency expressed an interest, but it was never produced in significant numbers.

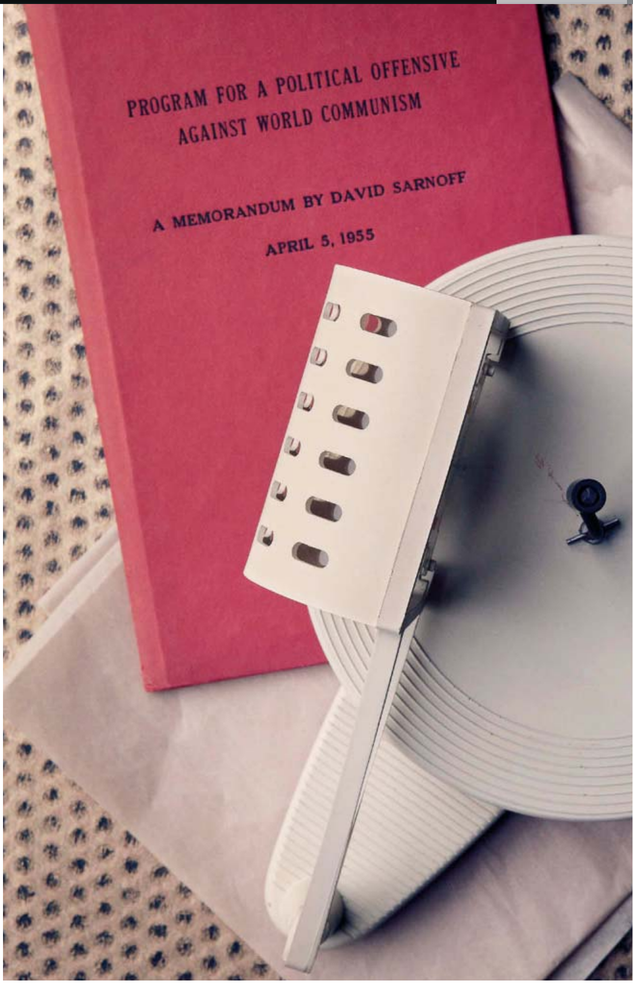

Gallium arsenide ingot (1950s) and gallium nitride blue light-emitting diode (1972)
Beginning in the 1950s, RCA researchers experimented with silicon alternatives, including gallium arsenide, an ingot of which is shown above and which became the basis of some of the earliest light-emitting diodes. Later, in 1972, RCA scientists Herbert Maruska and Jacques Pankove used gallium nitride to create a blue-violet LED, which their boss, James Tietjen, hoped could be used in a flat-panel display. Although the blue LED project was canceled in 1974, LEDs prepared in a similar fashion have since found a place in lightbulbs, televisions, and high-definition DVD systems.

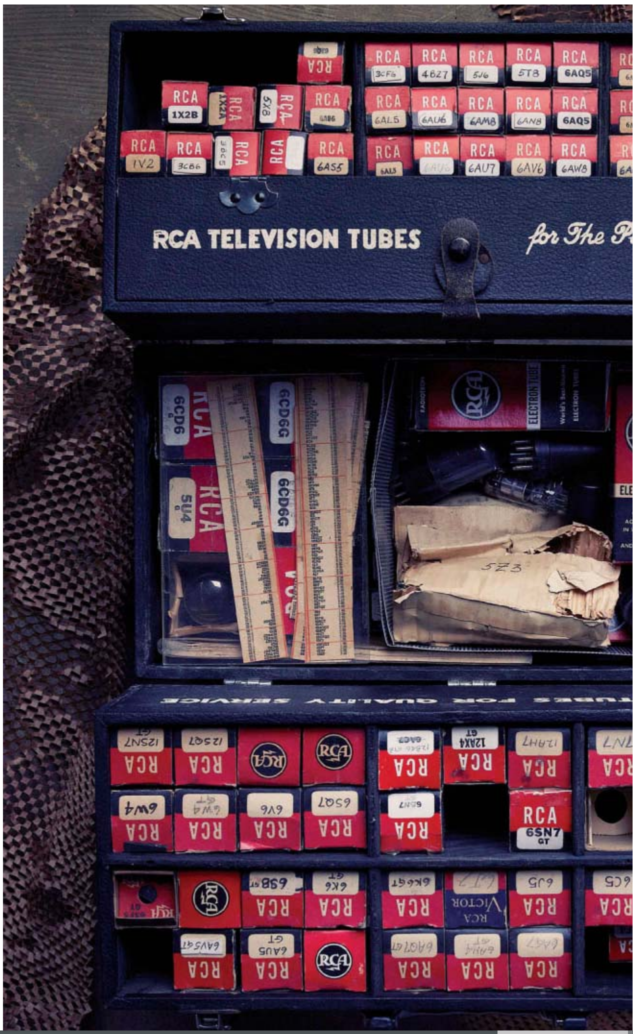

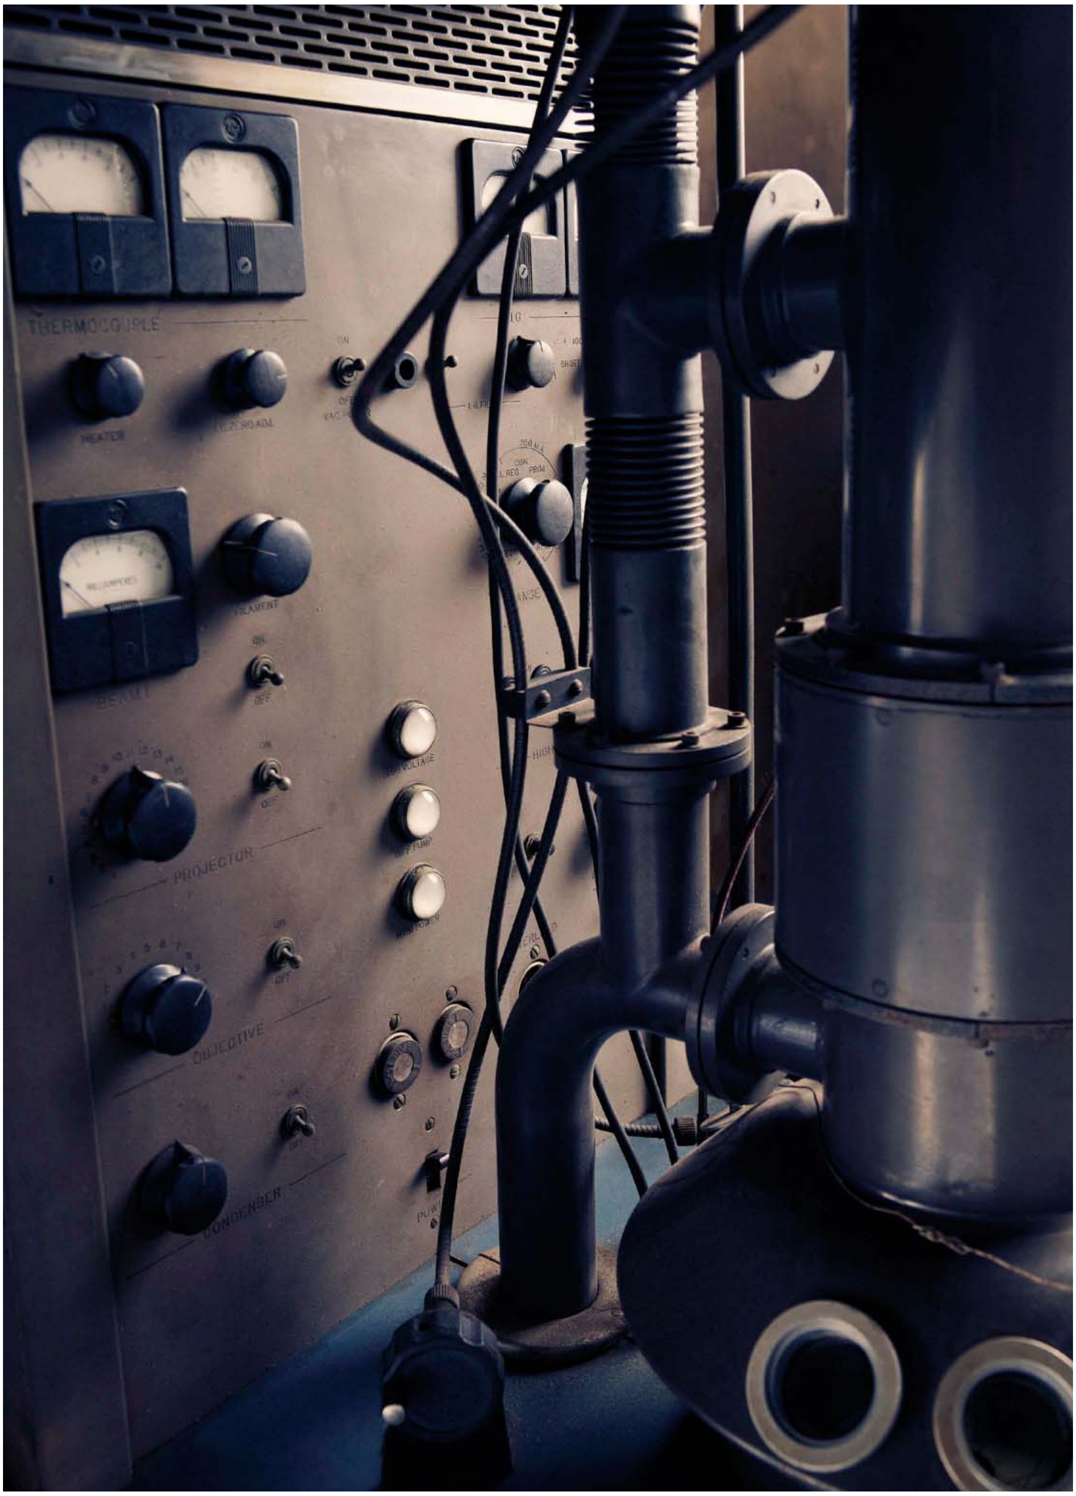

Model EMB-4 electron microscope (c. 1942)
RCA television expert Vladimir Zworykin's interest in the ability of electromagnetic fields to alter the trajectory of charged particles led him to work on a new type of microscope, a machine that used electron beams to scan samples at a much higher resolution than traditional optical instruments could. In 1940, James Hillier, a physicist in Zworykin's lab, had demonstrated the Model B, one of the world's first commercial electron microscopes. Electron microscopy would transform how scientists examine otherwise invisible specimens and bolstered the company's reputation for technical achievement.

RCA repairman's case (c. 1960)
Although the transistor was invented in 1947, vacuum tubes remained at the heart of most consumer electronics until well into the 1960s. RCA repairmen often carried kits filled with replacement vacuum tubes to fix broken radio and television sets.

Autographed model of Apollo 15 camera (1971) »
During the 1960s, RCA continued to develop smaller, lighter-weight color TV equipment. Among the beneficiaries of this work was NASA, which mounted an RCA camera on the lunar rover used during the Apollo 15 mission. The camera could be operated remotely by mission control personnel in Houston. NASA was so pleased with the system's performance that it asked RCA to supply cameras for the Apollo 16 and 17 missions. The model shown here was signed by the Apollo 15 crew and presented to RCA chairman and CEO Robert Sarnoff; the original camera remains on the moon.

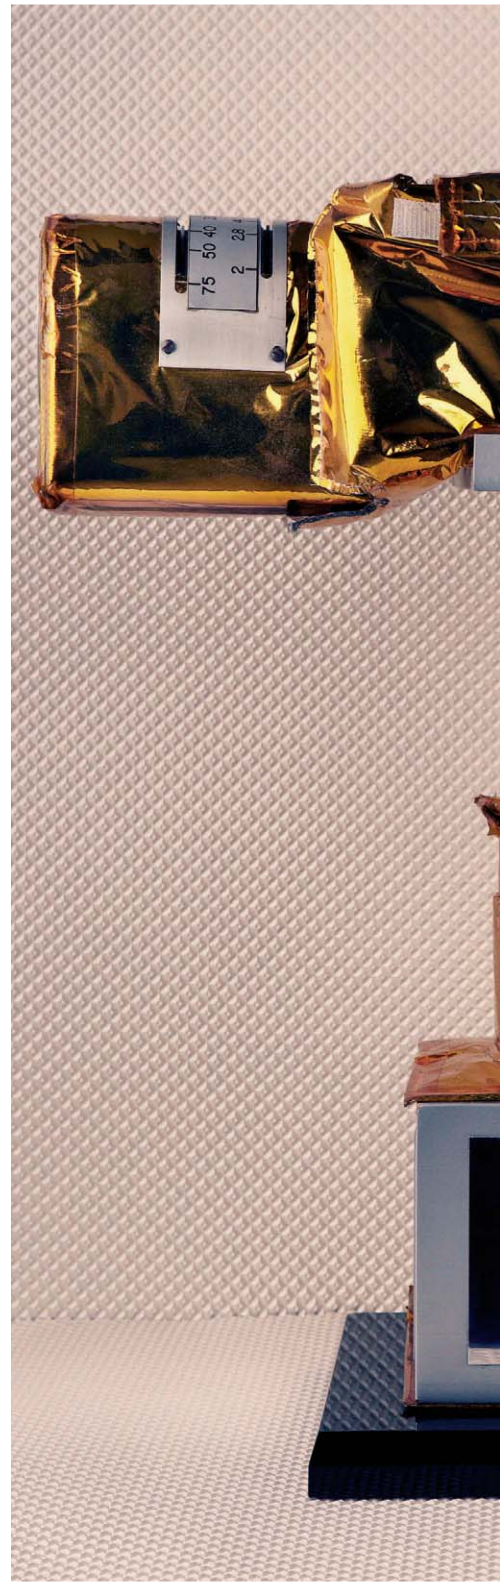

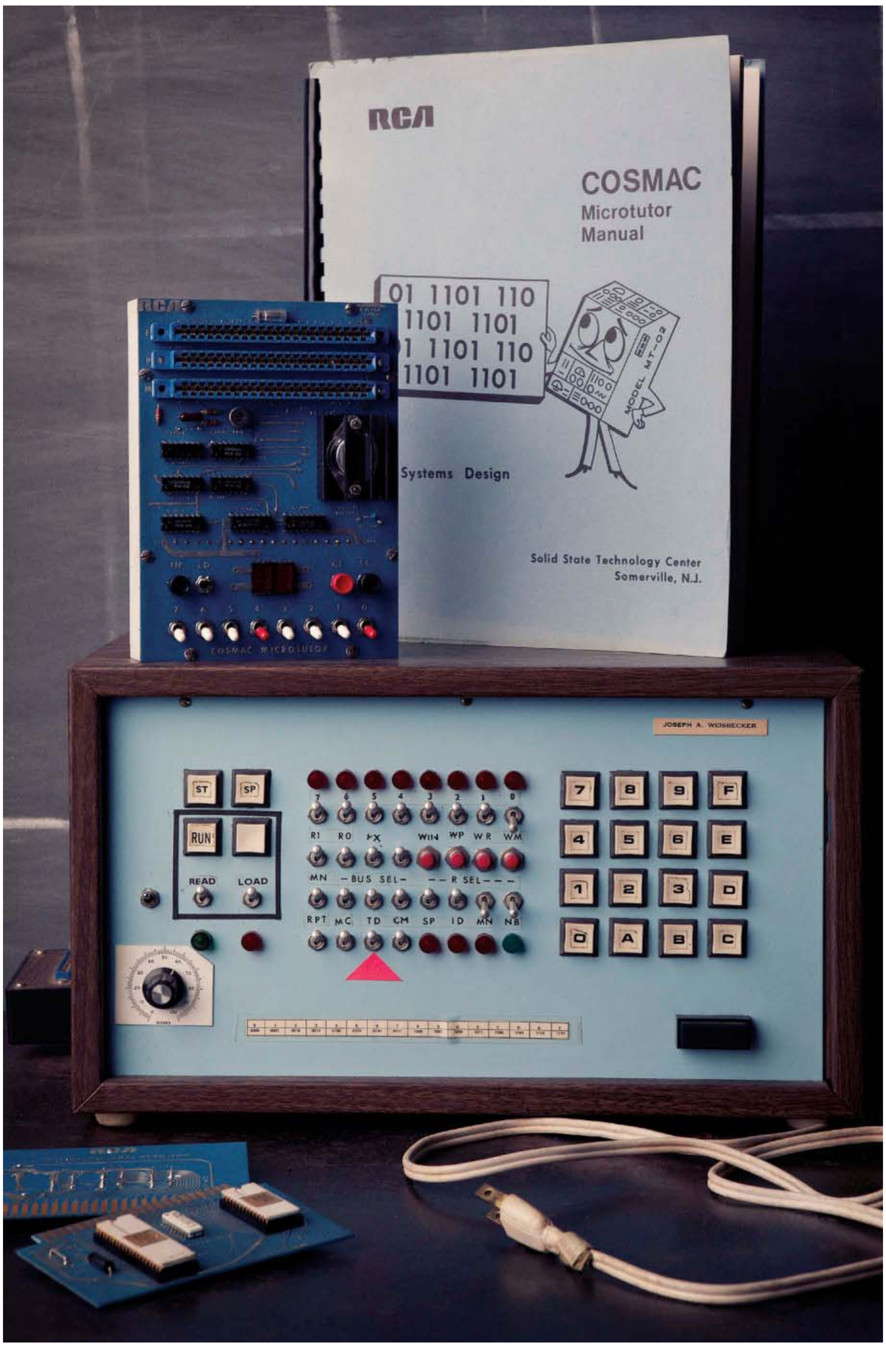

Model 00 personal computer (1972) and COSMAC Microtutor (1976)
After David Sarnoff stepped down as RCA chairman in 1970, his son Robert launched an ambitious campaign to challenge IBM in the realm of computing. The move backfired, and the following year RCA sold off its computer division, taking a US $490 million write-off, the largest in U.S. history up to that point. Among the few insiders who still saw a future for RCA in computing was Joseph Weisbecker, who thought the company should target home users. To demonstrate the feasibility of this idea, he developed the Model 00 [bottom], a $975 computer that could run simple games when connected to a TV set. For those on a tighter budget, there was the $350 Microtutor, which used a new CMOS (complementary metal-oxide semiconductor) microprocessor and could be programmed using a row of eight toggle switches.

Optel LCD watches (c. 1971)
Although RCA scientists and engineers made the first key breakthroughs in liquid-crystal displays, they failed to commercialize the technology [see "How RCA Lost the LCD," *IEEE Spectrum*, November 2012]. But several members of the company's liquid crystal research group eventually left to set up their own LCD companies. One such firm, Optel Corp., would create the first wristwatches with dynamic scattering displays.

SFT-100 VideoDisc player (1981) and VideoDiscs »
RCA spent more than a decade and $200 million developing a home video player, which, much like a phonograph, used a stylus to play prerecorded vinyl discs. Released in 1981, the product bombed with U.S. consumers, who preferred videocassettes that let them both play and record programs. The VideoDisc was RCA's last major commercial venture. In 1986, General Electric bought RCA for $6.3 billion.

(Good-Enough Computing)= <We could **save energy** in everything from smartphones to super-computers by letting them make **m1stake5**;

by ADRIAN SAMPSON, LUIS CEZE & DAN GROSSMAN

Illustrations by JUDE BUFFUM

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

Voltaire once wrote, "The best is the enemy of the gOod."

That's just one of the many ways of saying that perfectionism is not always a virtue. Yet computers today are relentless perfectionists, because even rare flaws in their output can be expensive, if not disastrous. A bug in Intel's P5 Pentium chip, for example, cost the company almost half a billion U.S. dollars, even though it caused only slightly inaccurate answers in about one in 9 billion computations.

The extraordinary reliability of computers makes them useful for many tasks, especially ones that we error-prone humans would never be able to do well. But this wonderful quality comes at a price–energy. You see, to ensure that errors are so rare that you can safely assume they never happen, computers consume gobs of energy.

That near-perfect correctness was a foregone conclusion by 1971, when the first commercial microprocessor, the Intel 4004, hit the market. For the next three decades, the main goal of microprocessor designers was to maintain that attribute while stuffing ever more transistors onto each chip; energy efficiency was, if anything, an afterthought. Today, however, energy use is front and center in chip designers' minds, for several reasons.

First off, the processors that go into the mobile computers we've all grown so fond of, like smartphones and tablets, must be energy efficient to conserve battery life. Those found in supercomputers and data centers need to be similarly thrifty with power, because the electricity to run these facilities often costs their owners millions of dollars a year. And in contrast to the situation in 1971, designers can no longer count on advances in semiconductor technology to bring drastically better per-transistor energy efficiency year after year.

Indeed, they now face the "dark silicon" problem: While Moore's Law remains intact, consistently delivering more transistors per chip with each new generation of manufacturing, microprocessors can't use all these additional transistors at the same time. Only a portion of them can be powered up before the chip becomes impossible to cool. So chip designers must arrange things to leave much of the microprocessor unpowered, or "dark," at any given moment–the rolling blackouts of the silicon world.

In this energy-constrained era, computer scientists of all stripes need to reexamine each of the things computers devote power to–including error-free operation. Absolute correctness was a great attribute when energy didn't matter. But now it's time to embrace the occasional slipup.

Compromising on correctness may seem to be a dangerous strategy. After all, when computers make mistakes, the results can be catastrophic. So if you let them make more mistakes, wouldn't they run the risk of becoming useless? Not really. The trick is controlling when the goofs can happen. Sometimes they won't pose a problem, because many of the things we use computers for today don't demand strict correctness.

Let's say you're on a long flight, watching a movie on your tablet computer. When the original movie file was compressed, the software that encoded it threw away unimportant details from each frame to produce a file that's a lot smaller. If the software decoder running on your tablet messes up a few pixels while playing the movie back, you probably won't object–especially if it means more battery life remains by the time the credits roll.

This situation is not unique to media players. For many kinds of software–speech recognition, augmented reality, machine learning, big-data analytics, and game graphics, to give a few examples–perfection in the output isn't the goal. Indeed, completely correct answers to the problem at hand may be impossible or infeasible to compute. All that's desired is an approximate answer. And when today's perfectionist computers execute these fundamentally approximate programs, they squander energy.

To seize the opportunity this extravagance presents, computer designers could build machines that can switch into an energy-saving–albeit somewhat error-prone–mode on demand. The machine might turn down the CPU voltage, for example, at the expense of causing arithmetic errors every once in a while. Reducing the refresh rate on dynamic random-access memory (DRAM) chips could also save energy, with the trade-off being a few unwanted bit flips. And wireless devices could cut back on the power they draw if some communication errors were allowed.

Many researchers are working on such energy-saving hardware modifications. We have been exploring how programmers could make use of them. To that end, our research group at the University of Washington, in Seattle, has developed a computer language that lets benign errors occur every now and then while preventing the catastrophic ones. We call it EnerJ. It's our contribution to a new approach for boosting energy efficiency called approximate computing.

The key difficulty with approximate computing is that even an approximate program sometimes needs to produce absolutely correct results. Consider a photo-manipulation program. While a handful of incorrect pixels won't mar a large image, a single wrong bit in the file's JPEG header could render the output useless.

This dichotomy is common to many kinds of software: Some parts of a program can tolerate occasional mistakes or imprecision, while other parts must always be executed precisely and without error. Approximate computers need to support both modes, and approximate programs must be written to use the energy-efficient mode wherever possible while avoiding errors that would lead to catastrophic failures.

But how does a computer distinguish between the parts of a program that can tolerate approximation and those that can't? At this stage at least, the programmer needs to instruct the computer to do that, using a language that offers some mechanism for making the distinction.

Our prototype language, EnerJ, works by letting the programmer mark data as either approximate or precise. Note that we're using those terms somewhat loosely here. The two kinds of data could differ in just numerical precision—that is, in the number of bits devoted to holding a value. Or they could differ in how prone they are to errors: An "approximate" datum would have a small but non-negligible chance of being garbage, whereas a "precise" datum can for all practical purposes be considered error free. Some of the hardware we've been investigating mix these two energy-saving approaches, allowing occasional errors in the least significant bits while making sure that the most significant ones are always correct.

EnerJ is an extension of the Java programming language, but its overall design can be applied to most languages with data types that the programmer declares explicitly. Such declarations are used to indicate whether a data element is meant to hold a Boolean (true/false) value, a byte, a 32-bit integer, a 64-bit integer, a 32-bit floating-point number, a 64-bit floating-point number, or various other possibilities. Everywhere the programmer declares a Java data element, he or she can mark it as approximate by writing "@Approx" before the declaration. Data types declared without such an annotation are implicitly made precise. In other words, the system uses approximation only where the programmer has specifically allowed it.

As a concrete example, imagine that you want to write some code that calculates the average shade of a black-and-white image that measures 1000 by 1000 pixels. Your code might start with the first pixel of the first row and note its value. It would then go on to the second pixel in the first row and add its value to the first. Then it would do the same with the third pixel, and so on, each time adding the pixel's value to a running sum it's maintaining in some variable—let's call it *TOTAL*. When it gets to the 1000th pixel in row one, the program starts all over on the second row of pixels and continues in that manner until it finishes up with the 1000th pixel of the 1000th row. At the end, it divides *TOTAL* by 1 million to calculate the average pixel value. Easy enough.

In addition to the variable *TOTAL*, the code requires two counter variables: one to keep track of the row number and one for the column number of the pixel it's adding to the sum. After the column counter gets to 999, it resets to 0 and progresses to the next row. When both row and column counters reach 999, the program can divide by a million to get the average.

If along the way one of the million pixel values is a little off or doesn't get added correctly to the sum, well, it's no big deal: The answer will be affected, but not by much. If, however, one of the row or column counters doesn't increase correctly, the program might throw an error, terminate prematurely, or even go into an infinite loop.

With EnerJ, the programmer simply marks *TOTAL* and the array holding the pixel values with "@Approx" when he or she declares them. The two counter variables—let's call them *I* and *J*—remain precise.

The yet-to-be-built energy-saving computer this EnerJ program runs on would then be allowed to use approximate computation indiscriminately–as long as it didn't affect the program's precise data. For example, the computer might store the individual pixel values in unreliable, low-refresh-rate DRAM, but it would have to use the normal, reliable part of its memory to store *I* and *J*. Or the addition operations that add the pixel values to *TOTAL* could be run at a lower voltage level because both of the operands are approximate, but the operations that increase the two counter variables would have to be run at the normal voltage so that they were always computed exactly. That is, just by annotating the type declarations in the code, the programmer specifies where approximate storage and operations can be used and where such approximation is forbidden.

EnerJ also helps programmers avoid bugs that would let approximate computations contaminate data that need to stay precise. Specifically, it forbids the program from putting approximate data into precise variables. Continuing with our example, this assignment would be illegal:

```
I = TOTAL;
```

Because *I* is precise, EnerJ guarantees that no approximation can be used to compute it. *TOTAL* is approximate, so copying its value into *I* would introduce approximation into a part of the program that the programmer wants to keep precise and free of errors. On the other hand, the opposite assignment has no such issues:

```
TOTAL = I;
```

Precise-to-approximate assignments like this one do not violate any guarantees of precision because they affect only approximate data.

By allowing one kind of assignment while prohibiting the other, EnerJ ensures that data move in just one direction. This constraint acts like a one-way valve: Precise data can flow freely into the approximate part of the program, but not vice versa.

The researchers who pioneered systems for enforcing unidirectional data movement, called information-flow tracking, were not interested in saving energy. They were concerned about software security. They wanted their programs to distinguish between low-integrity (possibly compromised) data and data deemed to be of high integrity.

Information-flow languages can, for example, prevent a kind of computer attack called SQL injection. (Structured Query Language, or SQL, is used for database management.) Here, the attacker provides the system with a database command when it's expecting just some ordinary input data. The system then carries out that command–perhaps doing great mischief in the process. Information-flow tracking can prevent such attacks by assuring that no inputs from the user flow to SQL statements. EnerJ borrows this idea to keep imprecision or rare errors from contaminating good data.

There's another way, though, that approximate data can affect precise data: through control-flow statements. Consider this code:

```
if (TOTAL > 0.0)
    I = 1;
else
    I = 2;
```

There are no illegal assignments, but the approximate value *TOTAL* still influences the precise variable *I*. This situation is called implicit information flow and must also be avoided.

EnerJ prevents implicit information flow simply enough: It forbids approximate conditions in any control-flow statement. The expression "*TOTAL* > 0.0" is approximate because one of the operands is approximate. So the compiler–the software that translates EnerJ statements into the low-level code that runs on whatever machine you're using–issues an error when it sees that expression as the condition in an "if" statement.

The strict isolation that EnerJ enforces between the approximate and the precise can in some cases be overly limiting, though. Sometimes programs produce approximate data that then need to move into a precise part of the program–say, for checking or output. For example, we experimented with a smartphone app for reading bar codes that used several approximation-tolerant image-analysis algorithms to produce a string of bits from a photo of a QR code. It then calculated a "checksum"–a number used to ensure the integrity of some other block of data.

Checksums are very handy for detecting errors. Here's a simple example: Say you've got a million bytes (a byte is a group of 8 bits) that you want to transmit in a way that is less than perfectly reliable. You might start by adding those million bytes together and then throwing away all but the least significant 8 bits of the answer. The result is the checksum, which you send after you send the data. This single-byte checksum can now be used to judge data integrity after transmission. Just compute the checksum again at the far end and compare it with the value that was sent. If even a single bit of the data got corrupted during transmission, the original and recalculated checksums will not match.

For our bar-code-reading app, we wanted to keep the checksum computation precise to catch any errors that might crop up because of the approximations used elsewhere. But that required the output of an approximate computation (the bar-code-reading algorithm) to flow into a precise computation–the checksum. Normally, that wouldn't be possible. But EnerJ provides an escape hatch in the form of something called endorsements–markers that indicate that the programmer gives permission to break the strict isolation EnerJ otherwise enforces.

EnerJ's data-type annotations, endorsements, and the isolation guarantees they provide make it possible to write programs that take advantage of approximate computation while retaining enough precise computation to stay reliable. With them, programmers can safely trade precision and strict correctness for increased energy efficiency. Just as programmers are accustomed to optimizing software for performance, EnerJ lets them optimize for energy savings.

How much savings? It's a little hard to say, because nobody has yet built the kind of two-mode, energy-saving computers EnerJ needs to run on. But our group has written many EnerJ programs and run them in an environment that simulates hardware with various energy-saving features in its mem-

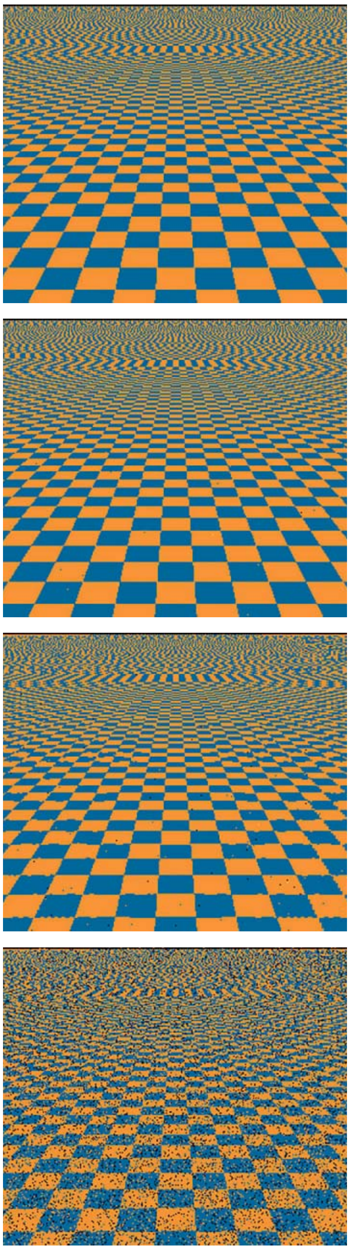

ADRIAN SAMPSON

A BIT OFF: This illustration shows the output of a 3-D-rendering program run on simulated hardware. The top panel is the result of exact calculations, with increasing levels of energy-saving approximation shown in the panels below.

ory or central processing. Depending on what kind of program it is and how aggressively the simulated hardware seeks to reduce energy consumption, EnerJ saves anywhere from 10 to 40 percent.

We were happy that our newly minted computer language did so well in these trials. We were also pleasantly surprised to discover the great variety of programs that have parts where approximation makes sense. Image rendering and image recognition are both natural fits for EnerJ. Signal-processing algorithms and geometry problems in 3-D gaming were less obvious candidates, but they too could be readily split into approximate and precise components. Indeed, when you have EnerJ or another approximate-computing hammer at your disposal, every program starts to look like a nail.

We now need realistic CPU designs that can save energy when running EnerJ programs. With our former University of Washington colleague Hadi Esmaeilzadeh, who is now at Georgia Tech, and Doug Burger at Microsoft Research, we have begun designing processors that can run at two different voltage levels. When a program needs error-free behavior, it uses the higher voltage. For calculations that can tolerate the occasional glitch, the processor switches to the lower voltage and saves energy. We've also been exploring some more radical designs for hardware that can act as dedicated approximate coprocessors.

To make approximate computing blossom, programmers will also need tools to help them understand the trade-offs involved in sacrificing precision for energy savings. Programming-language researchers, including those in our group, are beginning to design tools that are analogous to today's debuggers and performance profilers. With them, a programmer can find the answer to a question like "What will happen to the playback quality of a video file if I mark this part of my media player as approximate?" The right tools should let the programmer save quite a bit of energy without the user even noticing a difference.

Indeed, the energy that computers expend avoiding mistakes–even the rarest of small errors–looks increasingly wasteful when so much of today's software already incorporates imperfections in the form of error-prone data or approximate output. To rectify things, the computing industry needs some new approximate programming languages, some new approximate hardware designs, and tools to help programmers understand how best to use them. Together, they should help computers become a little more like human brains, which are less than perfect but amazingly energy efficient. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/enerj1013

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

The Tunneling Transistor

CONTINUED FROM PAGE 34 | germanium have indirect bandgaps, which means that in order to transition from one band to another, electrons must also absorb some extra energy from vibrations in the crystal lattice that makes up the material. This extra hurdle significantly lowers the probability that charge carriers will make the leap. As a result, the current-carrying capacity of silicon and germanium TFETs is only a trickle compared with that of today's transistors.

That might be a stumbling block for adoption by the industry. However, there are a range of direct-bandgap materials, based on a mix of elements picked from columns III and V of the periodic table, that have considerably higher tunneling probabilities. These materials have yet to make it into mass production in logic chips, but work on incorporating them into traditional MOSFETs is already gearing up [see "Changing the Channel," *IEEE Spectrum*, July 2013]. The notion that they might emerge in logic chips in the foreseeable future is not nearly as far-fetched as it would have seemed just a few years ago.

Research into TFETs made from III-V materials has also been advancing rapidly in recent years. Suman Datta and his colleagues at Pennsylvania State University were the first to demonstrate III-V TFETs, in 2009. They used channels made of a mix of indium, gallium, and arsenic and immediately set a record, with an "on" current that was 50 times as high as with the best germanium TFET.

Since then, the Penn State team and my group at the University of Notre Dame, in South Bend, Ind., have both shown even higher currents in TFETs made from a mix of two compounds: aluminum gallium antimonide and indium arsenide. The former material has bands that can be shifted up or down by tuning the ratio of aluminum to gallium. This lets us create tunnel junctions that have a natural overlap between bands, which means less voltage is needed to turn them on. And because the barrier can be quite thin–just a single atom or so wide–they permit more current. The devices we have built perform well at just 0.5 V and can carry nearly 200 microamperes across a 1-micrometer-wide channel, comparable to what can be accomplished with a state-of-the-art MOSFET.

The one caveat is the subthreshold swing of these "heterojunction" TFETs, which so far hasn't been able to beat the 60-mV-per-decade limit for the MOSFET. Many research groups are now struggling with this challenge. The main culprit is defects–many of which arise from dangling chemical bonds–at the interface between the semiconductor and gate oxide. These defects trap and immobilize charges, leaving fewer charges available for conduction. This means we have to apply a greater voltage to the gate to induce charge carriers in the channel.

That said, there is reason for optimism. Groups based at Intel, in Hillsboro, Ore., and Hokkaido University, in Sapporo, Japan, have demonstrated III-V TFETs with subthreshold swings of less than 60 mV per decade. And simulations from Intel suggest that it's possible to drive down subthresh-


We're making waves
With the world's most advanced EM and RF simulation software and measurement systems, it's no wonder we're getting noticed.
From concept to design to fielded solutions, Delcross has the tools to help you ride the wave of success. Our Savant software rapidly solves installed antenna problems that other electromagnetic solvers cannot address. With GPU acceleration in Savant, the power of a cluster can be realized for the price of a video card. Using our EMIT software and AMS measurement system, you can extend your analysis beyond the antenna ports to tackle your toughest cosite interference and link margin analysis problems. We provide unique solutions that get you ahead.
www.delcross.com
Savant
EMIT
AMS
DELCROSS TECHNOLOGIES
For more information, please contact your local CST AG representative (www.cst.com).

old swing even further without monumental changes in materials, simply by scaling down the transistors they have already built. In principle, devices with subthreshold swings of around 20 mV per decade appear possible; the ultimate limit will be set by thermal vibrations in the crystal, which make the edges of the conduction and valence bands less sharp.

MUCH AS IT WOULD HAVE BEEN HARD to predict the MOSFET's ultimate capabilities 50 years ago, it's difficult to say exactly what may ultimately be achieved with the TFET.

One uncertainty is the maximum current a TFET can carry when it's on. On-current is what ultimately determines the maximum speed of circuits, and for a long time, researchers thought it might be fairly low. But in 2010, Siyu Koswatta at IBM showed in simulation that gallium antimonide and indium arsenide could potentially carry 1.9 μA per micrometer of channel width when supplied with just 0.4 V. If such a device could be built, it would compete directly with the MOSFET in high-performance applications. The International Technology Roadmap for Semiconductors targets a current of 1.685 μA per micrometer of channel width, at a voltage of 0.73 V.

We will also have to tackle the issue of current leakage when the TFET is in its off state. As the channel gets shorter and shorter, it will be easier for electrons to tunnel directly from the source into the drain.

Figuring out the ultimate limits of the device will depend on such factors as electronic structure, defects, and performance requirements. Fortunately, computational tools developed over the past five years at Purdue University and at ETH Zurich now allow researchers to simulate entire devices, atom by atom and bond by bond, to predict device behavior. This activity is helping to guide experiments.

While the TFET's electrical characteristics look promising, there are also quite a few practical things we must tackle before we can start building chips with these transistors. Researchers have been focusing most of their energy on developing *n*-channel TFETs. *P*-channel TFETs—and a complementary process technology that could pair the two transistor types to make circuits—are still on the drawing board.

And chipmakers still have to find ways to address the problem of variability. As MOSFETs shrink, the placement and concentration of dopants and the roughness of interfaces can lead to significant variability in electronic properties. TFETs—which will likely be even smaller than MOSFETs when they're introduced—won't be immune to this problem. As with the MOSFET, we will have to develop other approaches in parallel, such as redundancy and error correction, to address this issue.

Still, I am optimistic that there are more promising results to come. It took just 10 years to get from the first silicon MOSFET to the first CMOS microprocessor. The jump to the TFET is arguably a much bigger challenge. But with more than half a century of experience with semiconductors under our belts, it might come about quicker than we think. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/tunnelfets1013

a place of mind
THE UNIVERSITY OF BRITISH COLUMBIA

UNIVERSITY OF BRITISH COLUMBIA
DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
TENURE TRACK FACULTY POSITIONS

ece Electrical and Computer Engineering

The Department of Electrical and Computer Engineering (ECE) at The University of British Columbia (UBC) invites applications for tenure-track positions at the rank of Assistant or Associate Professor in the area of Multicore Computing. We are particularly interested in expertise at the intersection between software systems (programming languages, compilers, runtimes, operating systems) and computer architecture.

Applicants must have either demonstrated or possess a clear potential and interest in achieving excellence in research and teaching. Successful applicants will preferably have relevant industrial experience and be active in enhancing educational and research links within the community. All faculty members are expected to teach at both undergraduate and graduate levels, supervise graduate students, develop a sponsored research program, collaborate with other faculty and be involved in service to the community and profession. Applicants should have received a Ph.D., or equivalent, in electrical engineering, computer engineering or an appropriate area. Registration as a Professional Engineer in British Columbia is required within five years of the appointment. These appointments are expected to commence as early as 1 July 2014. In addition to start-up funds from the University, significant start-up funding to new faculty may be available through the Canada Foundation for Innovation (CFI) and other sources.

UBC is rated among the top 40 research-intensive universities worldwide. The campus is surrounded by parks and water, and is located on an attractive peninsula in what the Economist recently rated one of the most liveable cities in the world – Vancouver. The Department currently comprises just over 50 faculty members, nearly 1000 undergraduates and more than 400 graduate students, and has the largest graduate program on campus with a strong interdisciplinary research culture. The Department is an active participant in the UBC Strategic Initiative of the "Campus as a Living Lab". This initiative promotes interdisciplinary research, the demonstration of new technologies, and industrial partnerships towards innovation and commercialization. Clean energy, smart integrated systems, health and sustainability are UBC research thrusts or priorities. In addition, UBC is committed to a novel research partnership paradigm known as Campus as a Living Lab, offering a number of opportunities for ECE faculty members to participate and lead. Additional information is available at http://www.ece.ubc.ca/.

The nature of an appointment as an Assistant or Associate Professor and the criteria for achieving tenure are described at the University's faculty relations site: http://www.hr.ubc.ca/faculty-relations/.

Review of applications will begin 1 November 2013 and continue until the positions are filled. UBC hires on the basis of merit and is strongly committed to equity and diversity within its community. We especially welcome applications from visible minority group members, women, Aboriginal persons, persons with disabilities, persons of minority sexual orientations and gender identities, and others with the skills and knowledge to productively engage with diverse communities. All qualified candidates are encouraged to apply; however Canadians and permanent residents will be given priority.

To apply, please submit your cover letter and CV online at **http://hr.ubc.ca/careers/faculty**.

Applicants are also required to provide three reference letters. Please request that referees send electronic copies of reference letters to references@ece.ubc.ca. Please note, that we are unable to accept e-mails from Hotmail, Yahoo, Gmail, MSN or other free e-mail accounts for referees.

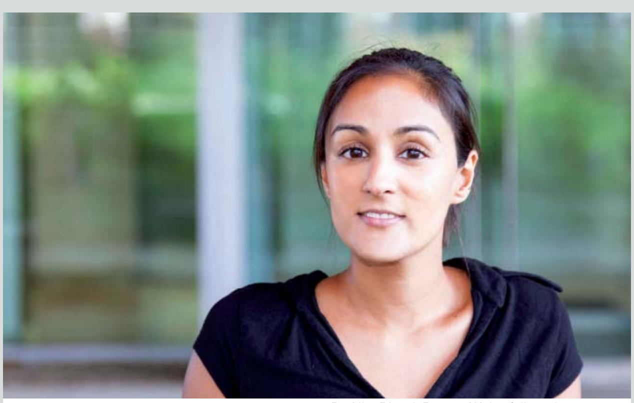

Dr. Nita Bharti, Branco Weiss fellow since 2012

Know what science will look like tomorrow? Apply today.

The magnitude of challenges we face today requires people with fresh thinking and novel approaches. To help find new ways forward, Society in Science - The Branco Weiss Fellowship gives extraordinary postdocs and engineers a generous grant to pursue an unconventional project for up to five years anywhere in the world. Have an idea that could change tomorrow? Get in touch with us today!

www.society-in-science.org

society in science
The Branco Weiss Fellowship

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Preston M. Green Department of Electrical & Systems Engineering
Tenured/Tenure-Track Faculty

The Preston M. Green Department of Electrical & Systems Engineering at Washington University in St. Louis invites applications for a faculty position at the assistant or associate professor level for fall 2014. Well qualified candidates at the full professor level will also be considered. Candidates should be exceptionally strong, possess novel and creative visions of research, and commit gladly to teaching at both the undergraduate and graduate levels. Candidates should have an earned doctorate in Electrical Engineering, Applied Physics, Systems Engineering, Mathematics, Operations Research or related fields.

Technical areas of interest include, but are not limited to, signal processing and imaging, information theory, applied physics, control systems, operations research, optimization, applied mathematics, and applied statistics. Applications include biomedicine, energy, the environment, robotics, financial engineering, network sciences and modeling of physical and complex systems. Successful candidates are expected to teach, conduct research, publish in peer-reviewed journals, and participate in department and university service.

Applications will be accepted immediately, and interviews will begin after January 1, 2014. The details of the application process and necessary documents are found at **http://ese.wustl.edu/aboutthedepartment/Pages/faculty-openings.aspx.**

Washington University in St. Louis is a medium-size private university, which is 14th in U.S. News & World Report's national university ranking.

Washington University in St. Louis is an Equal Opportunity and Affirmative Action employer, and invites applications from all qualified candidates. Employment eligibility verification required upon employment.

THE UNIVERSITY of TENNESSEE
KNOXVILLE

FIVE FACULTY POSITIONS

IN THE DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

The Department of Electrical Engineering and Computer Science (EECS) at The University of Tennessee (UT) is seeking candidates for five faculty positions which are further specified as follows. Successful candidates will be expected to teach at both undergraduate and graduate levels, to establish a vigorous funded research program, and to have a willingness to collaborate with other faculty in research.

EECS is housed in a new $37.5 million teaching and research facility completed in 2012. The department currently has an enrollment of more than 550 undergraduate and 220 graduate students, with a faculty of 40, and research expenditures that exceed $13 million per year. The department was awarded an NSF Engineering Research Center in 2011 focusing on Wide Area Control of Power Transmission and also houses the Center for Information Technology Research, which is one of nine UTK Centers of Excellence. The College of Engineering is a leading research institution with strong research partnerships with organizations such as the nearby Oak Ridge National Laboratory.

ERICSSON-HARLAN D. MILLS CHAIR OF SOFTWARE ENGINEERING

Applications and nominations are invited for the Ericsson-Harlan D. Mills Chair of Software Engineering. Associate Professors and recently promoted Full Professors are strongly encouraged to apply, although consideration will be given to all outstanding candidates. Applicants should have demonstrated an outstanding independent research and teaching program in software engineering or in a related software-encompassing field, and should be capable of working with other faculty to further strengthen our existing research. Applicants should have an earned Ph.D. in Computer Science or equivalent and be eligible for appointment at the rank of full Professor.

COMMUNICATIONS AND CONTROLS

Applications are sought for a tenure-track junior-level faculty position in Electrical Engineering with a concentration in controls or communication systems. The research areas include (but are not limited to) wireless communications, optical communication systems, information theory, communications signal processing, coding theory, communication hardware implementations, control over communication networks, adaptive control and decision making for coordinated autonomous systems, decentralized/distributed control of spatially distributed systems, robust and adaptive control of nonlinear processes, and control of complex, multi-scale, highly uncertain systems. Applicants should have an earned Ph.D. in Electrical Engineering or related field.

COMPUTER ENGINEERING AND COMPUTER SCIENCE

Applications are sought for two tenure-track junior-level faculty positions in Computer Science and in Computer Engineering. Applicants should have an earned Ph.D. in Computer Engineering, Computer Science, or a related field. Individuals with expertise in computer architecture, high performance computer systems, parallel and distributed computing, performance modeling and simulation, computing theory, cloud computing, data analytics, digital systems, VLSI, computational accelerators, dependable and secure systems, software engineering, intelligent systems and robotics are encouraged to apply, although qualified applicants from all areas of computer science and computer engineering will be considered.

POWER ENGINEERING

Applications are sought for a tenure-track junior-level faculty position in Power Systems. Applicants should have an earned Ph.D. in Electrical Engineering or equivalent. UTK has one of the largest power programs in the U.S. with state-of-the-art laboratories for power systems and power electronics and is the lead institution of an NSF/DOE Engineering Research Center, the Center for Ultra-wide-area Resilient Electric Energy Transmission Networks (CURENT). Information about CURENT can be found at **http://curent.utk.edu**.

The University of Tennessee welcomes and honors people of all races, genders, creeds, cultures, and sexual orientations, and values intellectual curiosity, pursuit of knowledge, and academic freedom and integrity. Interested candidates should apply through the departmental web site at **http://www.eecs.utk.edu** and submit a letter of application, a curriculum vitae, a statement of research and teaching interests, and contact information for three references. Review of applications will begin on December 15, 2013, and continue until the positions are filled.

The University of Tennessee is an EEO/AA/Title VI/Title IX/Section 504/ADA/ADEA institution in the provision of its education and employment programs and services.

FACULTY SEARCH

ShanghaiTech University

School of Information Science and Technology

The School of Information Science and Technology (SIST) in the new ShanghaiTech University invites applications to fill multiple tenure-track and tenured positions. Candidates should have an exceptional academic record or strong potential in frontier areas of information sciences.

ShanghaiTech is founded as a world-class research university for training future scientists, entrepreneurs, and technology leaders. Besides keeping a strong research profile, successful candidates should also contribute to undergraduate and graduate education within SIST.

Compensation and Benefits:
Salary and startup fund are highly competitive, commensurate with academic experience and accomplishment. ShanghaiTech also offers a comprehensive benefit package which includes housing. All regular faculty members are hired within ShanghaiTech's new tenure-track system commensurate with international practice and standards.

Academic Disciplines:
We seek candidates in all cutting edge areas of information science and technology. Our recruitment focus includes, but is not limited to: computer architecture and technologies, nano-scale electronics, high speed and RF circuits, intelligent and integrated signal processing systems, computational foundations, big data, data mining, visualization, computer vision, bio-computing, smart energy/power devices and systems, next-generation networking, as well as inter-disciplinary areas involving information science and technology.

Qualifications:
- Well developed research plans and demonstrated record/potentials;
- Ph.D. (Electrical Engineering, Computer Engineering, Computer Science, or related field);
- A minimum relevant research experience of 4 years.

Applications:
Submit (all in English) *a cover letter, a 2-page research plan, a CV including copies of 3 most significant publications, and names of three referees* to: sist@shanghaitech.edu.cn.

Deadline: December 31th, 2013 (or until positions are filled). We have 10 positions for this round of faculty recruitment.

For more information, visit http://www.shanghaitech.edu.cn.

NUS National University of Singapore

School of Computing

HEAD, DEPARTMENT OF COMPUTER SCIENCE

The National University of Singapore (NUS) (www.nus.edu.sg) invites applications and nominations for the position of Head, Department of Computer Science.

NUS is a highly ranked research university with intensive international collaborations, broad research funding avenues, and excellent facilities. The Departments of Computer Science and Information Systems constitute the School of Computing (SOC) which is one of the 16 faculties in NUS.

The CS Department (www.comp.nus.edu.sg/cs/) has around 70 faculty members representing all the major areas of computer science. The faculty's research is well represented at prestigious international conferences and journals, and they serve on many program committees and editorial boards. The department also has a thriving graduate school and it attracts the best students (undergraduate and graduate) in the region.

The University seeks to appoint a distinguished scholar with significant academic leadership record as the Head of Department. The accomplishments and stature of the candidate should entail her/him to a tenured professorship at NUS. The Head will lead the Department in all administrative and academic matters including strategic planning and fostering inter-faculty academic programs and research.

NUS offers highly competitive salaries and generous benefits. Singapore provides a vibrant international environment with world-class health care, excellent infrastructure, and very low taxes.

Please send your applications and nominations by **8 November 2013** to the Chair of the CS Head Search Committee:
Prof. Chua Tat Seng at *cshodrec@comp.nus.edu.sg*

The preferred start date for the appointment is July 1, 2014.

NSIT

Netaji Subhas Institute of Technology (NSIT) invites applications for the positions of Professors, Associate Professors and Assistant Professors in the various specialized areas of Electronics, Communication, Computers, IT, Instrumentation, Manufacturing Processes, Automation and Biotechnology. Candidates having an earned Ph.D. degree with first class Bachelor's and Master's degrees in appropriate branch of Engineering/Technology and possessing relevant Teaching/Industry/Research experience, may send their CV along with the names of three referees, within one month from publication of this advertisement, to: Director, NSIT, Sector 3, Dwarka, New Delhi 110078, India or e-mail to: faculty_appointments@nsit.ac.in

Joint Institute of Engineering

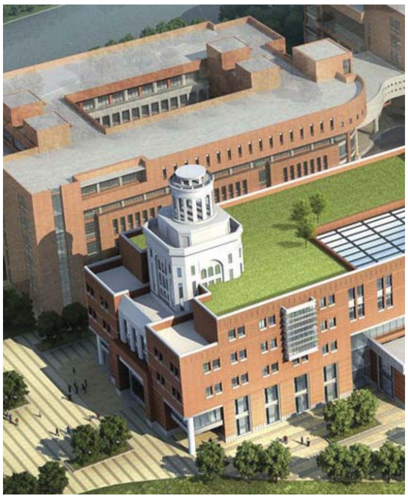

FACULTY POSITIONS AVAILABLE IN ELECTRICAL/COMPUTER ENGINEERING

Sun Yat-sen University & Carnegie Mellon University are partnering to establish the **SYSU-CMU Joint Institute of Engineering (JIE)** to innovate engineering education in China and the world. The mission of the JIE is to nurture a passionate and collaborative global community and network of students, faculty and professionals working toward pushing the field of engineering forward through education and research in China and in the world.

JIE is seeking **full-time faculty** in all areas of electrical and computer engineering (ECE). Candidates should possess a doctoral degree in ECE or related disciplines, with a demonstrated record and potential for research, teaching and leadership. The position includes an initial year on the Pittsburgh campus of Carnegie Mellon University to establish educational and research collaborations before locating to Guangzhou, China.

This is a worldwide search open to qualified candidates of all nationalities, with an internationally competitive compensation package for all qualified candidates.

PLEASE VISIT: **sysucmujie.cmu.edu** for details

SHUNDE INTERNATIONAL

Joint Research Institute

RESEARCH STAFF POSITIONS AVAILABLE IN ELECTRICAL/COMPUTER ENGINEERING

SYSU-CMU Shunde International Joint Research Institute (JRI) is located in Shunde, Guangdong. Supported by the provincial government and industry, the JRI aims to bring in and 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.

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

Candidates with industrial experiences are preferred.

Applications 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.

Please submit the letters of reference and all above materials to the address below.

Application review will continue until the position is filled.

EMAIL APPLICATIONS OR QUESTIONS TO: **sdjri@mail.sysu.edu.cn**

SUN YAT-SEN UNIVERSITY　　Carnegie Mellon University

a place of mind
THE UNIVERSITY OF BRITISH COLUMBIA

UNIVERSITY OF BRITISH COLUMBIA
DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
TENURE TRACK FACULTY POSITIONS

ece Electrical and Computer Engineering

The Department of Electrical and Computer Engineering (ECE) at The University of British Columbia (UBC) invites applications for tenure track positions at the rank of Assistant or Associate Professor in the area of machine learning and control from candidates with a strong fundamental background in machine learning theory for dynamical systems, system identification and control.

Applicants must have either demonstrated or possess a clear potential and interest in achieving excellence in research and teaching. Successful applicants will preferably have relevant industrial experience and be active in enhancing educational and research links within the community. All faculty members are expected to teach at both undergraduate and graduate levels, supervise graduate students, develop a sponsored research program, collaborate with other faculty and be involved in service to the community and profession. Applicants should have received a Ph.D., or equivalent, in electrical engineering, computer engineering or an appropriate area. Registration as a Professional Engineer in British Columbia is required within five years of the appointment. These appointments are expected to commence as early as 1 July 2014. In addition to start-up funds from the University, significant start-up funding to new faculty may be available through the Canada Foundation for Innovation (CFI) and other sources.

UBC is rated among the top 40 research-intensive universities worldwide. The campus is surrounded by parks and water, and is located on an attractive peninsula in what the Economist recently rated one of the most liveable cities in the world – Vancouver. The Department currently comprises just over 50 faculty members, nearly 1000 undergraduates and more than 400 graduate students, and has the largest graduate program on campus with a strong interdisciplinary research culture. The Department is an active participant in the UBC Strategic Initiative of the "Campus as a Living Lab". This initiative promotes interdisciplinary research, the demonstration of new technologies, and industrial partnerships towards innovation and commercialization. Clean energy, smart integrated systems, health and sustainability are UBC research thrusts or priorities. In addition, UBC is committed to a novel research partnership paradigm known as Campus as a Living Lab, offering a number of opportunities for ECE faculty members to participate and lead. Additional information is available at http://www.ece.ubc.ca/.

The nature of an appointment as an Assistant or Associate Professor and the criteria for achieving tenure are described at the University's faculty relations site: http://www.hr.ubc.ca/faculty-relations/.

Review of applications will begin 1 November 2013 and continue until the positions are filled. UBC hires on the basis of merit and is strongly committed to equity and diversity within its community. We especially welcome applications from visible minority group members, women, Aboriginal persons, persons with disabilities, persons of minority sexual orientations and gender identities, and others with the skills and knowledge to productively engage with diverse communities.

All qualified candidates are encouraged to apply; however Canadians and permanent residents will be given priority.

To apply, please submit your cover letter and CV online at **http://hr.ubc.ca/careers/faculty**.

Applicants are also required to provide three reference letters. Please request that referees send electronic copies of reference letters to references@ece.ubc.ca. Please note, that we are unable to accept e-mails from Hotmail, Yahoo, Gmail, MSN or other free e-mail accounts for referees.

M ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
UNIVERSITY of MICHIGAN ■ COLLEGE of ENGINEERING

ELECTRICAL AND COMPUTER ENGINEERING
UNIVERSITY OF MICHIGAN, ANN ARBOR

The Electrical and Computer Engineering (ECE) Division of the Electrical Engineering and Computer Science Department at the University of Michigan, Ann Arbor invites applications for junior or senior faculty positions, especially from women and underrepresented minorities. Successful candidates will have a relevant doctorate or equivalent experience and an outstanding record of achievement and impactful research in academics, industry and/or at national laboratories. They will have a strong record or commitment to teaching at undergraduate and graduate levels, to providing service to the university and profession and to broadening the intellectual diversity of the ECE Division. Although the research areas of particular interest are networks and communications, computer vision, integrated circuits and optics, applications are welcome in all relevant areas of research.

The highly ranked ECE Division (**www.eecs.umich.edu/ece**) prides itself on the mentoring of junior faculty toward successful careers. Ann Arbor is often rated as a family friendly best-place-to-live.

Please see application instructions at **www.eecs.umich.edu/eecs/jobs**

Applications will be considered as they are received. However, for full consideration applications must be received by December 8, 2013.

The University of Michigan is an Affirmative Action, Equal Opportunity Employer with an Active Dual-Career Assistance Program. The College of Engineering is especially interested in candidates who contribute, through their research, teaching, and/or service, to the diversity and excellence of the academic community.

LASSONDE SCHOOL OF ENGINEERING | YORK U UNIVERSITÉ UNIVERSITY

Tenure-Stream Faculty Openings: The Lassonde School of Engineering (LSE) at York University, Toronto, Canada seeks three (3) outstanding candidates in Electrical Engineering and Computer Science at the rank of Assistant or Associate Professor, to commence July 1, 2014, subject to budgetary approval.

One Tier-2 Canada Research Chair with expertise in Big Data Analytics and Scientific Visualization, and emphasis in visual depictions of data from complex systems; scalable analytics; data-enabled methods; or/and applications in data driven design and domain sciences; **2 positions with expertise in the areas of** (i) Power Systems (high power electronics and drives; smart grids; renewable energy resources; storage devices; intelligent control), and; (ii) Medical Devices (Disease diagnostic and treatment technologies; Bioengineering instrumentation for cellular and molecular analysis) or (iii) Electronics (Mechatronics; Nanoelectronics).

The successful candidates must hold PhD degrees in the relevant discipline with outstanding track record of teaching, scholarly research, and professional achievement. These positions will play key roles in the establishment of the Renaissance Engineering initiative at the LSE, an ambitious $250 million dollar hiring 100 new faculty and staff, and expanding the student body by 1500.

Applicants should submit a CV, statements of contributions in research, teaching, and curriculum development, three reference letters, and three sample research publications electronically at *http://lassonde.yorku.ca/new-faculty/* by **October 31, 2013**. For full position details, please visit *http://www.yorku.ca/acadjobs.* York University is an affirmative action employer. The affirmative action program can be found at *http://www.yorku.ca/acadjobs* or a copy can be obtained by calling the affirmative action office at 416-736.5713. All qualified candidates are encouraged to apply; however, Canadian citizens and permanent residents will be given priority.

BAYLOR UNIVERSITY

The Electrical and Computer Engineering Department of Baylor University

seeks faculty applicants for three tenured/tenure-track Faculty Positions at all levels and in all areas of electrical and computer engineering. Desired areas of technical expertise include: embedded systems, cyber-physical systems, computer/network security, software engineering, sensor networks, power, and energy. Applicants seeking a senior position must have an impressive record of scholarship and sustained research funding. All applicants must have an earned doctorate and a record of achievement in research and teaching. The ECE department offers B.S., M.S., M.E. and Ph.D. degrees and is rapidly expanding its faculty size. Facilities include the Baylor Research and Innovation Collaborative (BRIC), a newly-established research park minutes from the main campus.

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

Application reviews are ongoing and will continue until all positions are filled; however, applications received by January 1, 2014 will be assured of full consideration.

Applications must include: 1) a letter of interest that identifies the applicant's anticipated rank, 2) a complete CV, 3) a statement of teaching and research interests, 4) the names and contact information for at least four professional references.

Additional information is available at ***www.ecs.baylor.edu***. Send materials via email to Dr. Robert J. Marks II at Robert_Marks@baylor.edu. Please combine all submitted material into a single pdf file.

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

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Professor/Assistant Professor (Tenure Track) of Power Semiconductors

The Department of Information Technology and Electrical Engineering (www.ee.ethz.ch) at ETH Zurich invites applications for a tenured professorship or tenure-track assistant professorship in power semiconductor devices. The new professor is expected to develop a strong research program in the area of novel power semiconductor devices.

The applicant will be expected to have a PhD degree, as well as an established track record or proven potential in such disciplines as advanced silicon based power semiconductors, wide band gap power semiconductors, power device simulation, packaging, or novel cooling technologies. Furthermore, the applicant should be able and willing to offer courses on power semiconductors in the electrical engineering curriculum (undergraduate level courses in German or English and graduate level courses in English).

Please apply online at www.facultyaffairs.ethz.ch

Applications should include a curriculum vitae, a list of publications, and a statement of your future research and teaching interests. The letter of application should be addressed **to the President of ETH Zurich, Prof. Dr. Ralph Eichler. The closing date for applications is 31 December 2013.** ETH Zurich is an equal opportunity and family friendly employer and is further responsive to the needs of dual career couples. In order to increase the number of women in leading academic positions, we specifically encourage women to apply.

上海交通大學
SHANGHAI JIAO TONG UNIVERSITY

The Department of Computer Science and Engineering (CSE) at Shanghai Jiao Tong University (SJTU) is seeking to fill several tenure-track positions in computer science at the rank of Assistant Professor and above, starting October 2013 and September 2014.

Shanghai Jiao Tong University is one of the oldest and most prestigious universities in China, and CSE is premier in computer science research and education. Candidates for these positions are sought for the well-recognized computer science program (ACM class) at Zhiyuan College of SJTU which provides an outstanding undergraduate education to a select group of 30 research-oriented students. Over the last ten years, students from the ACM class have won five gold medalsin the ACM International Collegiate Programming Contest.

Professor John Hopcroft, 1986 Turing Award recipient, is chairing the committee on curriculum development and faculty recruiting. Since December 2011, he has spent two months a year teaching at Zhiyuan College. In May 2012, he was appointed Special Counselor to President Jie Zhang.

An internationally competitive package for salary and benefits will be offered. Strong candidates in all areas will be considered with special consideration given to systems and networking, architecture, machine learning, theory, and security. In addition to the teaching duties at Zhiyuan College's ACM class, faculty members are required to teach graduate level courses, to supervise Ph.D. students, and to conduct research in the CSE. The overall teaching load is one course per semester.

SJTU makes a great effort to provide opportunities for the development of young faculty,including a startup research grant. There are a number of sources for additional research funding. The positions are provided in strong cooperation with Microsoft Research Asia (MSRA) with opportunities for research collaborations. Candidates are encouraged to apply to the Thousand Talents Program for extra funding and benefit support. Our equal opportunity and affirmative action program seeks minorities, women, and non-Chinese scientists.The criteria for promotion will be professional reputation as judged by international experts in the candidate's field and excellence in teaching.

Applications, including vita and the names of three references, should be sent to Professor John Hopcroft (jeh@cs.cornell.edu) and to Bing Li (binglisjtu@sjtu.edu.cn).

The application deadline is January 31, 2014 for positions starting in September 2014. Applications for starting earlier will be reviewed immediately.

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

a place of mind
THE UNIVERSITY OF BRITISH COLUMBIA

UNIVERSITY OF BRITISH COLUMBIA
DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
TENURE TRACK FACULTY POSITIONS

ece Electrical and Computer Engineering

The Department of Electrical and Computer Engineering (ECE) at The University of British Columbia (UBC) invites applications for tenure track positions at the rank of Assistant or Associate Professor in the area of electric power systems engineering. Preferred areas of focus for this position include: Emerging smart adaptive transmission and distribution systems as well as solution methodologies and techniques for analyzing complex power system networks including fast transient solution techniques.

Applicants must have either demonstrated or possess a clear potential and interest in achieving excellence in research and teaching. Successful applicants will preferably have relevant industrial experience and be active in enhancing educational and research links within the community. All faculty members are expected to teach at both undergraduate and graduate levels, and to supervise graduate students, develop a sponsored research program, collaborate with other faculty and be involved in service to the community and profession. Applicants should have received a Ph.D., or equivalent, in electrical engineering, computer engineering or an appropriate area. Registration as a Professional Engineer in British Columbia is required within five years of the appointment. These appointments are expected to commence as early as 1 July 2014. In addition to start-up funds from the University, significant start-up funding to new faculty may be available through the Canada Foundation for Innovation (CFI) and other sources.

The successful candidate will work in close association with renowned UBC power system researchers, as well as other ECE and UBC researchers. The power systems group maintains very strong links with B.C. Hydro and other national and international industry, and participates in highly interdisciplinary research projects with other faculty in ECE, other Departments of Applied Science, and across multiple disciplines within and beyond the University.

UBC is rated among the top 40 research-intensive universities worldwide. The campus is surrounded by parks and water, and is located on an attractive peninsula in what the Economist recently rated one of the most liveable cities in the world - Vancouver. The Department currently comprises just over 50 faculty members, nearly 1000 undergraduates and more than 400 graduate students, and has the largest graduate program on campus with a strong interdisciplinary research culture. The Department is an active participant in the UBC Strategic Initiative of the "Campus as a Living Lab". This initiative promotes interdisciplinary research, the demonstration of new technologies, and industrial partnerships towards innovation and commercialization. Clean energy, smart integrated systems, health and sustainability are UBC research thrusts or priorities. In addition, UBC is committed to a novel research partnership paradigm known as Campus as a Living Lab, offering a number of opportunities for ECE faculty members to participate and lead. Additional information is available at http://www.ece.ubc.ca/.

The nature of an appointment as an Assistant or Associate Professor and the criteria for achieving tenure are described at the University's faculty relations site: http://www.hr.ubc.ca/faculty-relations/.

Review of applications will begin 1 November 2013 and continue until the positions are filled. UBC hires on the basis of merit and is strongly committed to equity and diversity within its community. We especially welcome applications from visible minority group members, women, Aboriginal persons, persons with disabilities, persons of minority sexual orientations and gender identities, and others with the skills and knowledge to productively engage with diverse communities. All qualified candidates are encouraged to apply; however Canadians and permanent residents will be given priority.

To apply, please submit your cover letter and CV online at **http://hr.ubc.ca/careers/faculty**.

Applicants are also required to provide three reference letters. Please request that referees send electronic copies of reference letters to references@ece.ubc.ca. Please note, that we are unable to accept e-mails from Hotmail, Yahoo, Gmail, MSN or other free e-mail accounts for referees.

Faculty positions in Robotics and in Computer Science

NAZARBAYEV UNIVERSITY

School of Science and Technology
Nazarbayev University
Astana, Kazakhstan

Nazarbayev University is seeking highly qualified full-time faculty at all ranks to join its rapidly growing programs in Robotics and Computer Science in the School of Science & Technology. Successful candidates must have an earned Ph.D. degree from an accredited university, excellent English-language communication skills, a demonstrated ability for research, and a commitment to graduate and undergraduate teaching and program development.

Launched in 2010 as the premier national university of Kazakhstan, NU's mandate is to promote the emergence of Astana as the research and educational center of Eurasia. The strategic development of this English-language university is based on the Western model via partnerships with top ranking world universities.

All suitable candidates will be reviewed. Applications are particularly encouraged from candidates with research interests that align with our thrust areas of Embedded Systems, Industrial Automation and Manufacturing **(Department of Robotics)** and Computer Systems, Intelligent Systems, Mobile Computing, Information Security, and Software Engineering **(Department of Computer Science)**.

Benefits include a competitive salary, international health care coverage, housing (based on family size and rank), child educational allowance, and home-leave travel twice per year.

To Apply: applicants should send a detailed CV, including qualifications, experience, and list of publications to **sst@nu.edu.kz**. Review of applications will begin immediately but full consideration will be given to applications submitted no later than **January 15th, 2014**.

FACULTY OF ENGINEERING AND COMPUTER SCIENCE

UNIVERSITÉ Concordia UNIVERSITY

CONCORDIA.CA

The Faculty of Engineering and Computer Science at Concordia University, Montreal, Quebec, Canada, is one of Canada's major engineering schools, offering a wide array of exciting programs, courses and cutting-edge research opportunities in a dynamic setting equipped with state-of-the-art facilities.

For additional information on the Faculty and further details on these postings, visit encs.concordia.ca.

Canada Research Chair — Aerospace Robotics

We invite applications for the position of Tier II Canada Research Chair in Aerospace Robotics. The chairholder will be expected to engage in theoretical and experimental research in such areas as space robotics, on-orbit servicing (OOS), vision-based guidance of robotic systems, unmanned aerial vehicles (UAV's), unmanned space vehicles, autonomous and reconfigurable robotic systems, advanced fault-tolerant sensors/actuators, and related areas. A strong emphasis is placed on fundamental and applied research, interdisciplinary research partnerships, and the ability to establish and develop industrial collaborations as well as international collaborations. Teaching is also an important activity of the chairholder. Shortlisted candidates will be invited to make a presentation on their research. Please apply no later than January 1, 2014, to:

Dr. William E. Lynch, Chair
Department of Electrical and Computer Engineering
blynch@ece.concordia.ca | ece.concordia.ca

Sustainable Energy Engineering

The Faculty of Engineering and Computer Science is seeking excellent tenure-track candidates at the Assistant or Associate Professor ranks in the area of Sustainable Energy Engineering. The successful applicant will be appointed in one of the three following departments, depending on their background: Building, Civil and Environmental Engineering, Electrical and Computer Engineering, or Mechanical and Industrial Engineering. Applicants must possess expertise and research interests in an area of sustainable energy engineering such as, but not limited to 1) application of power electronics, electric machine design, and motor drives to renewable energy systems, electric/plug-in hybrid electric vehicles, and smart/micro grid systems; 2) building-integrated sustainable energy systems, waste-to-energy conversion processes and technologies, innovative energy storage systems, smart building systems; 3) fuel cells and/or hydrogen fuel processes. Please apply no later than January 1, 2014, to:

Dr. Rama Bhat, Associate Dean, Academic Affairs
Faculty of Engineering and Computer Science
rama.bhat@concordia.ca

DATAFLOW_

MORE VIEWING, LESS ATTENTION

TV WATCHERS MULTITASK WITH MOBILE DEVICES

Despite the influx of new media technologies such as smartphones and streaming video services, live television as watched on a big set still dominates the viewing landscape, according to a recent report from the United Kingdom's regulator of the communications industry, Ofcom. But viewers are splitting their attention more often, simultaneously chatting remotely with friends about the show they're watching or doing unrelated activities like checking e-mail. Currently about 16 percent of all U.K. viewers multitask daily while watching TV, and 25 percent do so at least once a week. —STEPHEN CASS

AVERAGE DAILY HOURS SPENT WATCHING TELEVISION PER CAPITA

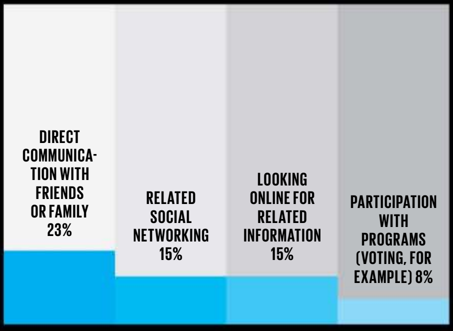

MOST POPULAR MESHING ACTIVITIES

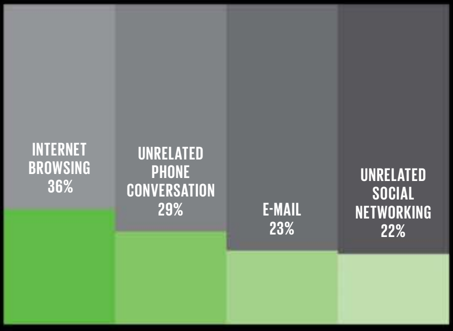

MOST POPULAR STACKING ACTIVITIES

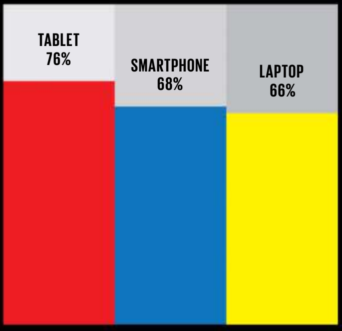

DEVICES USED FOR MULTITASKING
(MULTIPLE DEVICES MAY BE USED)

SOURCE: OFCOM COMMUNICATIONS MARKET REPORT 2013

INFOGRAPHIC BY Brandon Palacio


How can you get your idea
to market first?
On the road to innovation,
speed wins.
Accelerate your R&D and beat the competition
to market. Instant access to over 3 million
top-cited technology research documents can
save hours—and keep your ideas in the fast lane.
IEEE Xplore® Digital Library
Discover a smarter research experience
Request a Free Trial
www.ieee.org/tryieeexplore
Follow IEEE Xplore on
IEEE
Advancing Technology
for Humanity

Find it at
mathworks.com/accelerate
datasheet
video example
trial request

MODEL PHYSICAL SYSTEMS

in

Simulink

with **Simscape™**

- Electrical
- Mechanical
- Hydraulic

and more

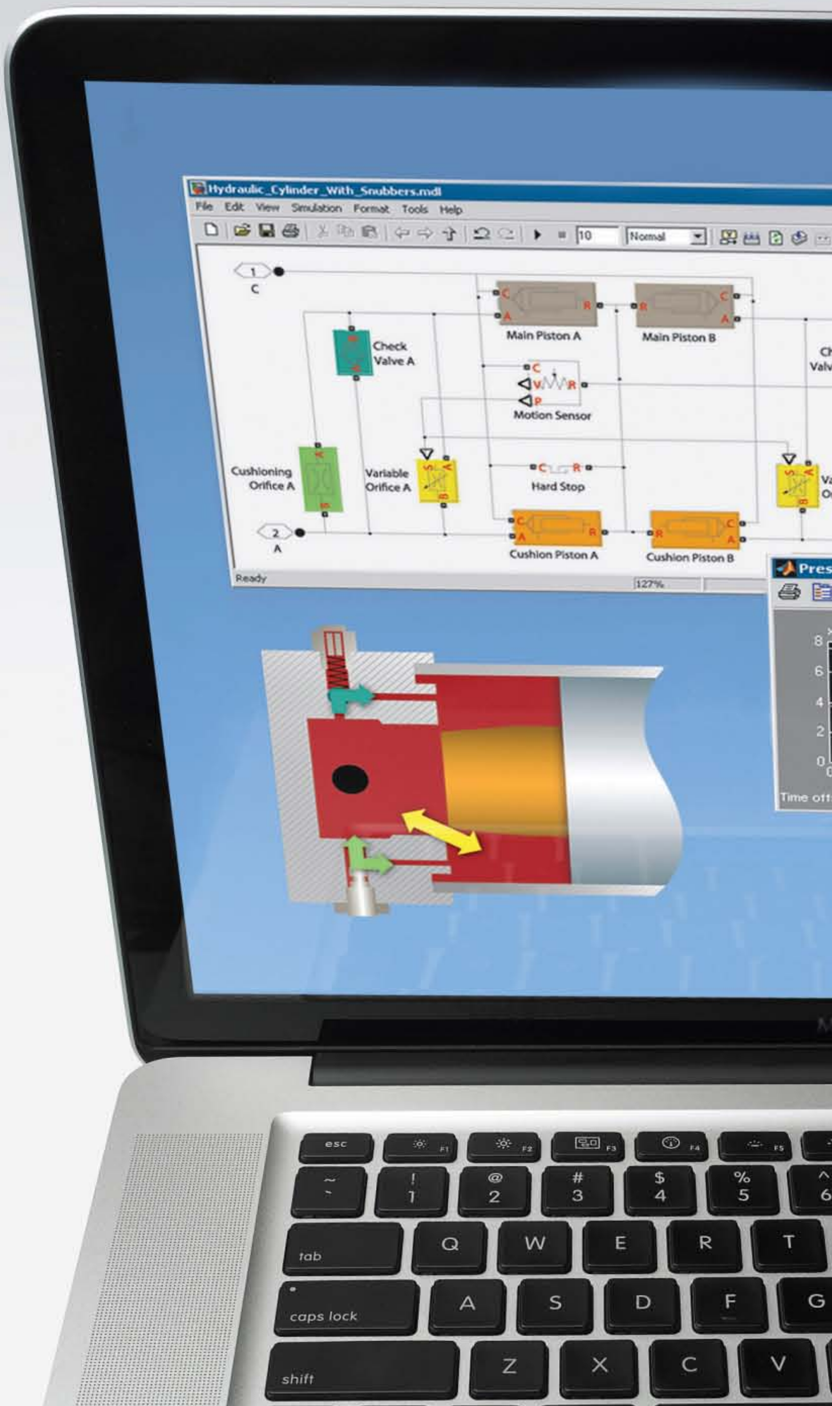

Use SIMSCAPE with SIMULINK to model and simulate the plant and controller of an embedded system. Assemble your model with a graphical interface, or import physical models from CAD systems. Use built-in components or create your own with the Simscape language.

MATLAB® & SIMULINK®

MathWorks®
Accelerating the pace of engineering and science