IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 10.14

Why Voice Quality Stinks… And Why It Doesn't Have To P.30

IEEE

5.0 NOW AVAILABLE
comsol.com/release/5.0

FROM MODEL

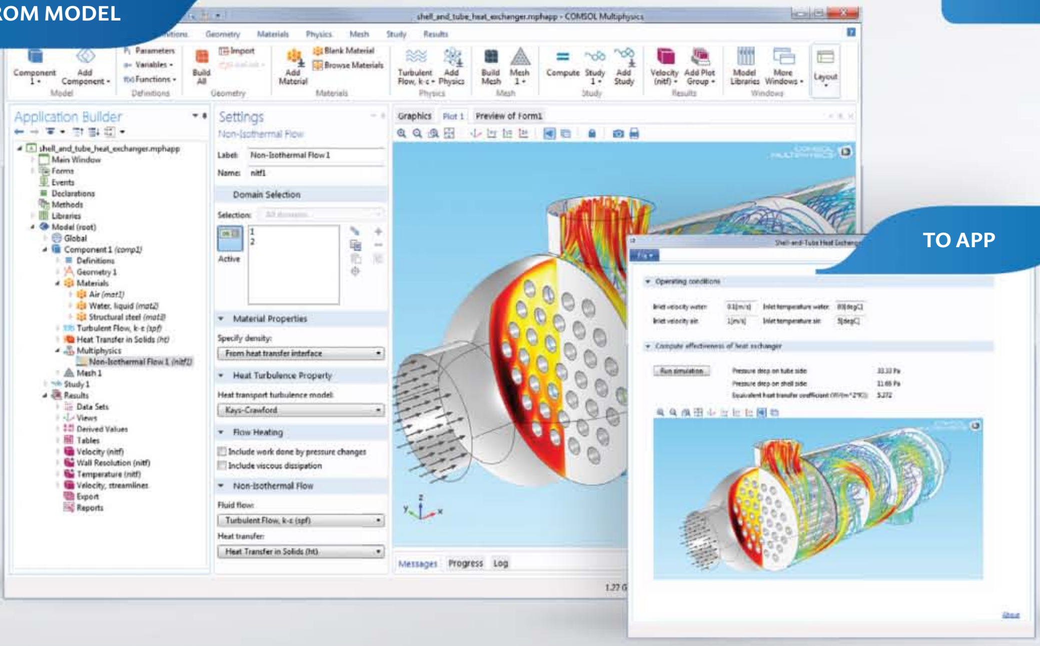

TO APP

Verify and Optimize your Designs with COMSOL Multiphysics®

NOW FEATURING THE APPLICATION BUILDER

The Application Builder provides you with tools to easily design a custom interface for your multiphysics models. Use COMSOL Server to distribute your apps to colleagues and customers worldwide.

Visit comsol.com/release/5.0

PRODUCT SUITE

COMSOL Multiphysics
COMSOL Server

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

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

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

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

MULTIPURPOSE
Optimization Module
Material Library
Particle Tracing Module

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

COMSOL

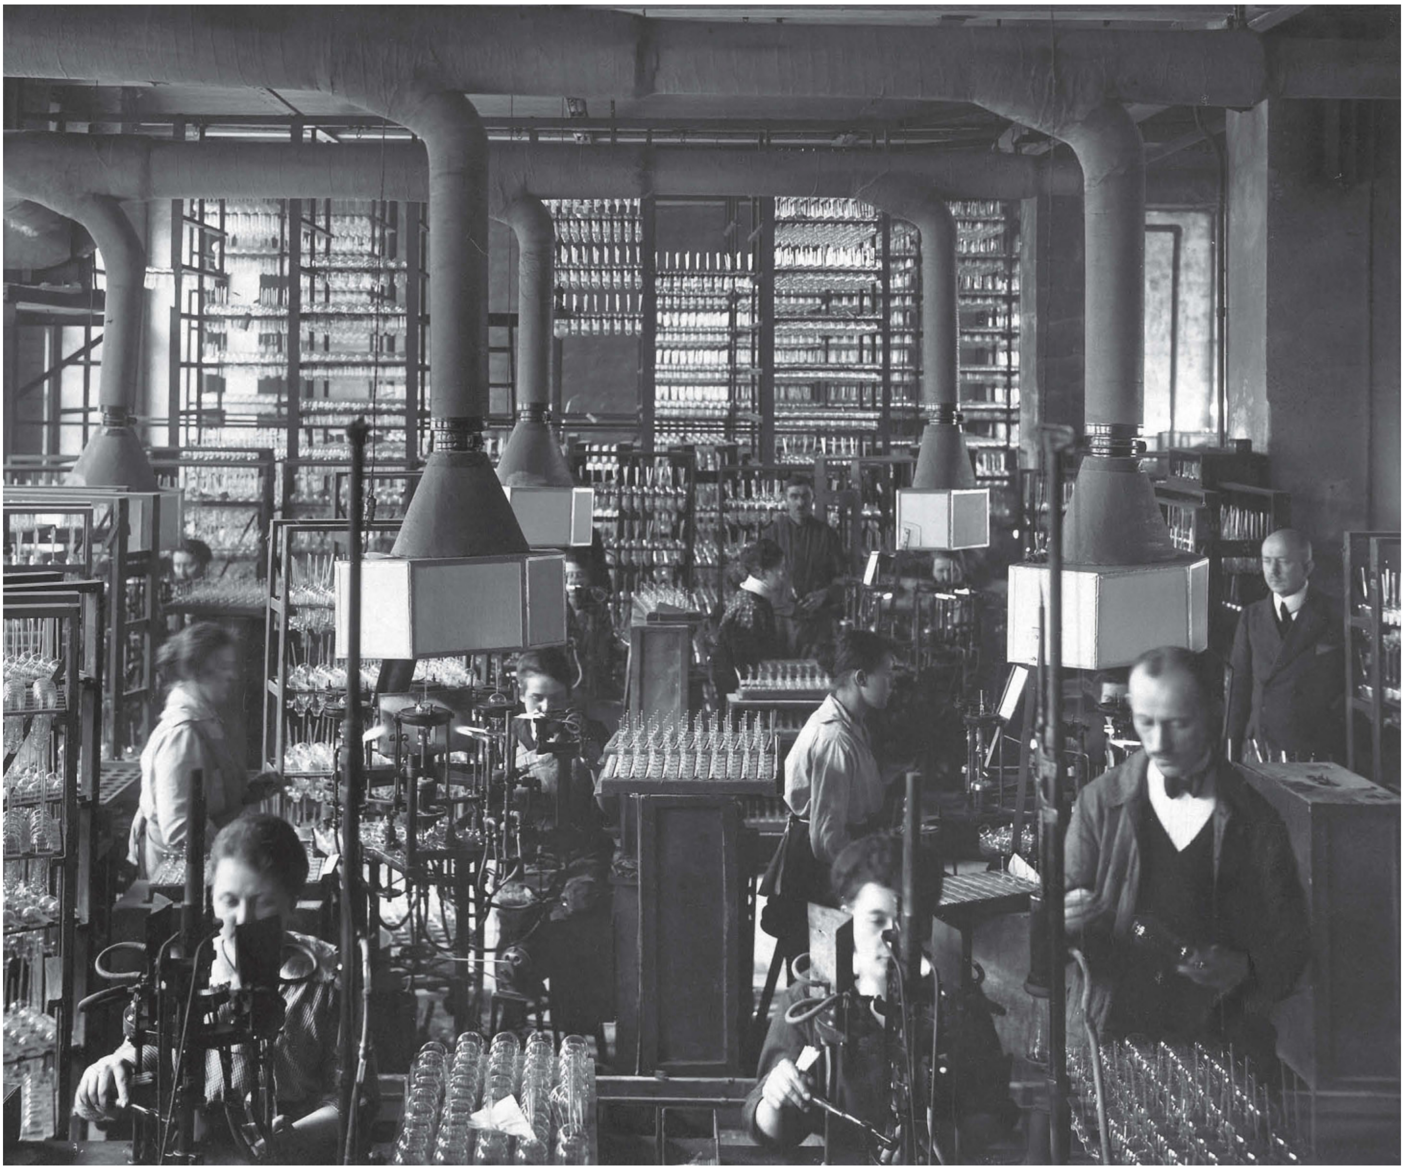

SIEMENS CORPORATE ARCHIVES

On the Cover **Illustration for *IEEE Spectrum* by Serge Bloch**

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

DEPARTMENTS_01.13

IEEE SPECTRUM

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- How Proper Application of Thermally Conductive Materials Will Improve Motor Power Density—2 October
- Designing Building Structures for Protection Against EMP and Lightning—9 October
- The Internet of Things and Open Source—9 October
- Simulation of Accelerator Components—16 October

NEW PRODUCT RELEASE LIBRARY
http://spectrum.ieee.org/static/new-product-release-library

WHITE PAPERS
http://spectrum.ieee.org/whitepapers

The Institute

Available 6 October at theinstitute.ieee.org

MIND-CONTROLLED MACHINES IEEE Senior Member Dean Aslam has come up with a way to control everyday devices just by thinking about them. By inserting a wireless sensor in something that lies on the forehead, say, a baseball cap, the wearer can transmit a signal produced by brainwaves to another sensor in an object, such as a coffee machine, and turn it on or off.

IEEE.TV ADDS NEW CONTENT The IEEE Computer and IEEE Robotics and Automation societies have made hundreds of videos available for free, with thousands more on the way.

HEAVY MACHINERY When she's not at her desk parsing lines of code, Member Amy Jones can be found driving large vehicles like excavators and tree-felling machines to test the software she's developed for them.

FROM LEFT: MOTOROLA SOLUTIONS; RYAN MATTHEW SMITH; GOOGLE

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2014 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 51, 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 the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

BACK STORY_

When Fiction Meets Fact

IN THOMAS PYNCHON'S CULT CLASSIC *Gravity's Rainbow*, an immortal lightbulb named Byron runs afoul of a secretive international industrial alliance known as the Phoebus cartel. When the cartel detects that Byron has exceeded his programmed life span, the Committee on Incandescent Anomalies dispatches a hit man to take Byron out.

Markus Krajewski, like many readers, found the story both "wild and weird." And yet, he says, "I knew that Pynchon's prose style mixes fact and fiction, and so I wondered: Could this be true?"

Turns out, many parts of Pynchon's tale were indeed based on fact: There really was a Phoebus cartel, and it really did target lightbulbs. Krajewski learned that Pynchon had relied on bona fide economic histories in weaving his tale of the cartel, including George W. Stocking and Myron W. Watkins's 1946 text, *Cartels in Action*. Digging deeper, he discovered that the Municipal Archives in Berlin housed corporate records from Osram, a key cartel member. At the time, in the late 1990s, he was studying in Berlin at Humboldt University, so he decided to "order up some files." In this rich trove were letters and reports that documented how the cartel conspired to engineer a shorter-lived incandescent lightbulb.

Now a professor of media studies at the University of Basel, in Switzerland, Krajewski [shown above, at the Berlin archives] says at first he was "quite astonished that the cartel so meticulously tried to control everything connected to such an ordinary object like the lightbulb." But when he considered the huge profits involved, it all made sense.

Still, he says, by reducing the lightbulb's life span, the cartel was essentially working against progress. "William Meinhardt, the head of Osram, always argued that the cartel was for the benefit of the consumer," Krajewski says. And that, he concluded, was the biggest fiction of all. ■

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. 51, no. 10 (INT), October 2014, p. 64, or in *IEEE Spectrum*, Vol. 51, no. 10 (NA), October 2014, p. 76.

10.14

IEEE
SPECTRUM

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

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

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

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

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

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

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

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

EXECUTIVE PRODUCER, SPECTRUM RADIO Sharon Basco
ASSISTANT PRODUCER, SPECTRUM RADIO Francesco Ferorelli, f.ferorelli@ieee.org

ADMINISTRATIVE ASSISTANTS
Ramona Foster, r.foster@ieee.org
Nancy T. Hantman, n.hantman@ieee.org
INTERNS Theresa Chong, Sarah Lewin

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

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

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

EDITORIAL ADVISORY BOARD
Susan Hassler, *Chair*; Gerard A. Alphonse, Francine D. Berman, Jason Cong*, Sudhir Dixit, Limor Fried, Kenneth Y. Goldberg, Bin He, Robert Hebner, Chenming Hu*, Grant Jacoby, Christopher J. James, Norberto Lerendegui, John P. Lewis, Steve Mann, Jacob Østergaard, Umit Ozguner, John Rogers, Jonathan Rothberg, Umar Saif, Gaurav Sharma, Takao Someya, Sergio Verdú, Jeffrey M. Voas, Kazuo Yano, Larry Zhang*, Yu Zheng, Kun Zhou
* Chinese-language edition

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

PHOTOGRAPH BY **Christian Werner**

CONTRIBUTORS_

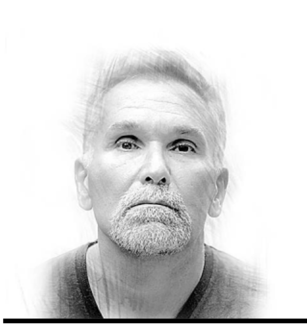

Ozalp Babaoglu

Babaoglu and coauthor Moreno Marzolla are computer scientists at the University of Bologna, in Italy. Babaoglu became intrigued by the kind of peer-to-peer technology they describe in "The People's Cloud" [p. 44] after applying biologically inspired techniques to computing. Marzolla says his interest in cooperative computing was sparked by a sci-fi novel in which insectlike micromachines "managed to wipe out all other life on their planet. But that was only science fiction, right?"

Theresa Chong

Chong was an editorial intern at *IEEE Spectrum* from April to August 2014. In addition to writing, she has produced videos for *Spectrum* and *Engineering News-Record*. Prior to entering journalism, Chong worked as a field engineer in her native Alberta, Canada, for several years, and then headed off to Northwestern University's journalism school. One of her first projects at school was about wearable technology for firefighters, and she was glad to be able to pick up the story again at *Spectrum* [p. 18].

Daniel Dern

Dern is a Boston-based technology and business writer. In this issue he reports on remote-collaboration tools for engineers [p. 25]. Dern says hardware engineers in smaller and newer groups may find such tools especially useful because they help users work through the complexities of government compliance and international manufacturing. "For hardware, spec errors can mean expensive manufacturing redos, so tools that include some built-in manufacturing know-how make sense," he says.

Jeff Hecht

An IEEE life member, Hecht writes for *Laser Focus World* and *New Scientist* and is the author of 11 books. As a journalist, he's become painfully aware of the deterioration in telephone voice quality, as interviewees increasingly take his calls on cellphones. "If everyone's on a cellphone, how are we going to understand each other?" he wonders. Hecht explores how mobile carriers plan to improve voice service—and why they haven't done so already—in "All Smart, No Phone" [p. 30].

Jérôme Lodewyck

Lodewyck, an associate professor at France's National Center for Scientific Research, enjoys doing battle with atoms, pinning down every factor that might affect their behavior. This compulsion came in handy in 2007, when he took a postdoctoral position at the Paris Observatory and was immediately assigned to build an optical-lattice clock from scratch. In this issue, he describes the promise of these clocks [p. 36], which are now the world's most accurate and could one day redefine the second.

10.14

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 J. Roberto B. de Marca, president@ieee.org
+1 732 562 3928 FAX: +1 732 465 6444

PRESIDENT-ELECT Howard E. Michel
TREASURER John T. Barr **SECRETARY** Marko Delimar
PAST PRESIDENT Peter W. Staecker

VICE PRESIDENTS
Saurabh Sinha, Educational Activities; Gianluca Setti, Publication Services & Products; Ralph M. Ford, Member & Geographic Activities; Karen Bartleson, President, Standards Association; Jacek M. Zurada, Technical Activities; Gary L. Blank, President, IEEE-USA

DIVISION DIRECTORS
Ellen J. Yoffa (I); Jerry L. Hudgins (II); Harvey A. Freeman (III); Jozef W. Modelski (IV); Susan K. Land (V); Bogdan M. Wilamowski (VI); Wanda K. Reder (VII); Roger U. Fujii (VIII); Marina Ruggieri (IX); Stephen Yurkovich (X)

REGION DIRECTORS
Vincent P. Socci (1); Parviz Famouri (2); Mary Ellen Randall (3); Karen S. Pedersen (4); J. Derald Morgan (5); Michael R. Andrews (6); Amir G. Aghdam (7); Martin J. Bastiaans (8); Norberto M. Lerendegui (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
PUBLICATIONS Anthony Durniak
+1 732 562 3998, a.durniak@ieee.org
CORPORATE ACTIVITIES Elena Gerstmann
+1 732 981 3452, e.gerstmann@ieee.org
EDUCATIONAL ACTIVITIES Douglas Gorham
+1 732 562 5483, d.g.gorham@ieee.org
ACTING CHIEF INFORMATION OFFICER Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
HUMAN RESOURCES Shannon Johnston, SPHR
+1 732 562 6343, s.johnston@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Eileen M. Lach, +1 212 705 8990, e.m.lach@ieee.org
CHIEF MARKETING OFFICER Patrick D. Mahoney
+1 732 562 5596, p.mahoney@ieee.org
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 Baillieul, Babak Beheshti, Herbert Bennett, Jennifer Bernhard, Stuart Bottom, Jean-Luc Gaudiot, David Allan Grier, Glenn Gulak, Lawrence Hall, Sheila Hemami, David Hodges, Hulya Kirkici, Khaled Ben Letaief, Carmen Menoni, Paolo Montuschi, Lloyd A. "Pete" Morley, William Moses, Michael Pecht, Sorel Reisman, Charles Rubenstein, Tariq Samad, Gaurav Sharma, Curtis Siller, David Tian, Joel Trussell, Leung Tsang, Stephanie White, Timothy 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**

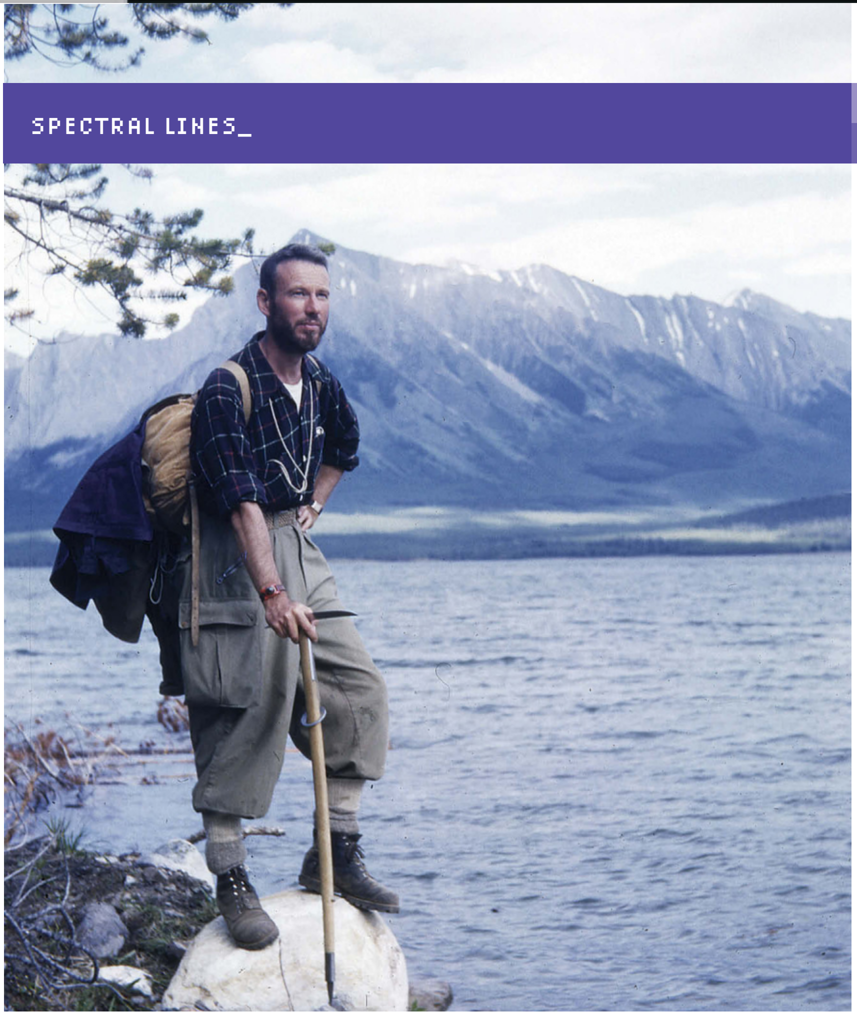

LOFTY AMBITIONS: *IEEE Spectrum*'s third editor, J.J.G. "Jerry" McCue, loved a challenge. An avid hiker and mountain climber, he vacationed at Spray Lake, in Alberta, Canada, in 1952.

When Freedom Rang

The Vietnam war and an unusual engineer prompted *IEEE Spectrum*'s first forays into controversy

AT THE TENDER AGE OF 5, *IEEE Spectrum* had its first identity crisis, and it was a doozy. The year was 1969, and the United States was racked by intense social and political upheaval. The war in Vietnam was prompting huge demonstrations, and as that tumult inevitably seeped into the offices of *Spectrum*, the magazine was soon forced to decide whether it was going to incorporate timely journalism or be another academic journal, disconnected from the important issues of the day.

By shoving the young magazine firmly toward journalism, one extraordinary man precipitated the fledgling magazine's crisis–and bestowed on the publication one of the greatest favors it would ever receive. He was *Spectrum*'s third editor, and his name was J.J.G. "Jerry" McCue.

IEEE SPECTRUM 1964 / 2014

The episode began with a letter published in the April 1969 issue of *Spectrum*. It came from the Union of Concerned Scientists, which had been recently formed at MIT. It started with these words: "Misuse of scientific and technical knowledge presents a major threat to the existence of mankind. Through its actions in Vietnam our government has shaken our confidence in its ability to make wise and humane decisions." The letter was followed two months later in the magazine by an article in which George Wald, a Nobel Prize winner and professor at Harvard University, referred to the Vietnam war as "shameful and terrible."

A mail squall soon pelted the *Spectrum* and IEEE offices. Some of the letter writers applauded the Union of Concerned Scientists letter or Wald's article. Quite a few more criticized them, some thoughtfully, some not. One in the latter category disclosed that he had written to his congressman requesting that McCue be tried for treason.

The 55-year-old electrical engineer at the center of this postal paroxysm was an improbable target. John Joseph Gerald McCue was a respected researcher, university professor, textbook author, and a dutiful IEEE volunteer. He had graduated from high school in New Jersey at age 16 and then worked on a vacuum-tube project at Bell Laboratories in 1930 before going off to college at Harvard and then on to a Ph.D. in physics, under the great Hans Bethe, at Cornell.

As an adult, he loved mountain climbing and what we now call adventure travel. He was also an early and ardent collector of Shaker furniture. Indeed, much of what we know about McCue, who died in 2011, comes from a family in the antiques business. Will and Karel Henry became friendly with McCue and his wife, Miriam, in the 1980s and recently sold much of McCue's huge Shaker collection. The Henrys' son, Tyler Henry, interviewed McCue several years ago and wrote a biographical essay for the catalog of the collection.

That essay, combined with McCue's own writings for *Spectrum*, revealed a complex and nuanced man. A registered Republican, he had worked on a secret radar project at the MIT Radiation Laboratory during World War II. During his editorship of *Spectrum*, he was an engineering researcher at Lincoln Laboratory, a leading defense-

WILLIS HENRY AUCTIONS

Editor's note: In this 50th anniversary year of IEEE Spectrum, *we are using each month's Spectral Lines column to describe some pivotal moments of the magazine's history. Here we describe the surprising and eventful tenure of the magazine's third editor, J.J.G. McCue.*

industry research organization. There he had done pioneering work for the U.S. Navy on the physics of bat echolocation.

McCue, like all editors of *Spectrum* until 1972, was drafted into the role from the ranks of high-level IEEE volunteers. He had not been trained as a writer, and before his service at *Spectrum*, his editing experience had been limited to some associations with academic journals. And yet at *Spectrum* and during those turbulent times, he would show remarkable insight and judgment, and those qualities would prove critical to *Spectrum*'s future.

McCue was deeply interested in not only the larger issues of journalism and publishing and ethics but also in the craft of writing itself. His Spectral Lines columns are invariably thoughtful, well structured, and written with imagination and discipline. Here's a passage from a 1969 column: "In the past 30 years the groves of Academia have turned into jungles, whose denizens have learned that in the struggle for existence, the writing of review articles has little survival value." He had also published an article and photographs in *National Geographic* magazine and authored a book, *The World of Atoms: An Introduction to Physical Science* (1956). It is unusually well written for a textbook.

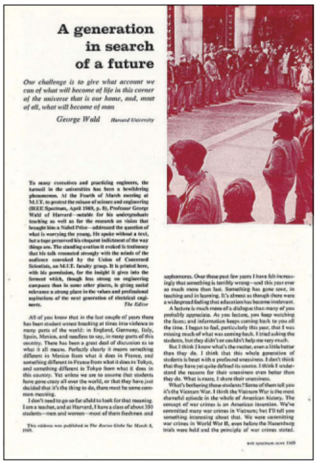
A generation in search of a future

Our challenge is to give what account we can of what will become of life in this corner of the universe that is our home, and, most of all, what will become of man

George Wald Harvard University

KINDLING: An article by George Wald in the June 1969 issue of *Spectrum* fueled a firestorm of protest.

The upshot is that at that pivotal time in *Spectrum*'s history, the magazine had as its leader precisely the person it needed: a man of impeccable engineering credentials who could think like a magazine editor and who also had courage and vision. There weren't many people like him, then or now.

McCue's widow, Miriam, says that her husband tried to keep his work at *Spectrum* secret from his employers at Lincoln Laboratory, because any such moonlighting was officially forbidden there. But McCue knew that what he was doing at *Spectrum* was important. "He'd come home from work at 5:30 or 6:00, then after dinner he'd go up to his office and start working again," she recalls. Sometimes "he'd be up at 1:00 or 2:00 or 3:00 a.m., working for *Spectrum*," she adds.

To the onslaught of hysterical letters following the Concerned Scientists letter and the Wald article, he responded with equanimity and by publishing many points of view, even ones he surely found abhorrent. In fact, he printed many more angry letters than supportive ones. When several readers complained that the letter from the Union of Concerned Scientists was unsigned, he responded by printing the names of the people involved. The list included such distinguished EEs as Mildred Dresselhaus, Murray Eden (who would later serve on *Spectrum*'s advisory board), Samuel J. Mason, Marvin Minsky, Arthur R. von Hippel, and Joseph Weizenbaum.

To those who insisted that *Spectrum* stop covering controversial topics, he had a succinct response: "In the plan for *Spectrum* originally laid down by the Board of Directors in 1963, one of the functions of *Spectrum* was to be the publication of 'News of political and social interest to the profession.' Until the past few months, this policy has been pretty much of a dead letter. If it is taken seriously in these times of questioned values, it must result in the treatment of controversial themes."

Some suggested that *Spectrum* should be permitted to cover a controversial topic—but only if it gave space simultaneously to "the other side." It's an idea that seems logical at first glance. But it is actually impossible to implement without incurring costs that no good magazine would ever accept, namely, being forced to throw out many excellent articles or publish them months after they had ceased being timely or relevant. McCue had encountered this problem firsthand, and he sought to explain it to readers by citing one of his own recent experiences.

It involved a draft article that sought to explain why the Tennessee Valley Authority, a U.S. federal power generator, was buying equipment from non-U.S. suppliers. One reviewer demanded that the story be accompanied by another one arguing the other side—that "the U.S. government should buy domestic products even when foreign ones cost less," McCue wrote. "I invited him to write such an explanation, but he said no. Nor would he name anybody who might write it."

Faced with a choice between publishing a timely, relevant article with a point of view or shelving that article indefinitely—and probably forever—McCue chose to publish. As any good magazine editor would do.

Regarding his own role—and this very column—he wrote: "In this column, over his name, the Editor may properly make comment that is colored by his political or ideological bias, but he is expressing his own opinion, not that of the Institute, and if he permits his bias to affect his selection of material to be published, he is not doing a good job."

Meanwhile, the IEEE's president, F. Karl Willenbrock, was surfing another tsunami of letters, some of them angrily demanding McCue's immediate resignation. Willenbrock, a Brown University graduate, found that when it came to controversial subject matter in *Spectrum*, the IEEE "did not have a defined policy in this area." The IEEE Publications Board soon drafted one. Published in the September 1969 issue of *Spectrum*, it was accompanied by a letter from Willenbrock asking for members to read it and send him their opinions.

The proposed policy had four provisions, and they were explicitly declared to apply not just to *Spectrum* but to all material published by the IEEE. The most notable was an exhortation that "every effort should be made to provide for adequate and

timely presentation of differing viewpoints." A key word here is "timely." It suggests that the board understood the ludicrousness of attempting to force a magazine to always provide simultaneous opposing views on any matter that might at any time (and by anyone) be deemed controversial.

> One letter writer disclosed that he had written to his congressman requesting that McCue be tried for treason

Another provision called for subject matter to "have some relevance" to electrical engineering "or to the relationship of that work to the needs of society." There were also stipulations that authors be "responsible or representative spokesmen for the viewpoint they espouse" and that it be clear that they were giving their own opinion, and not the IEEE's.

Those 1969 provisions are basically still in effect. Today, they are incorporated into the IEEE Publication Services and Products Board Operations Manual as section 8.1.3, "Presentation of Nontechnical Material." And after 45 years, you might think the fact that *Spectrum* covers controversial subjects would no longer itself be a controversial topic, but sometimes it is. However, every time IEEE members have ever been asked whether they want *Spectrum* to be free to wade into controversy, or to always shy away from it, an overwhelming majority have said they want us to keep wading.

Spectrum's freedom to publish controversial articles is what sets us apart from most other association publications. It's a major reason why we have won about 150 significant journalism honors over the past decade: This freedom enables us to attract top talent to our staff and to unleash them to do first-rate work. Among the best of this work are articles and multimedia segments whose investigative revelations served IEEE members and also society at large–thereby fulfilling the IEEE's core mission of supporting and explaining technology for the benefit of humanity.

In recent years, these included hard-hitting articles on such topics as fiascoes in big-government IT projects; efforts to maintain a power plant in the Gaza Strip; the Fukushima Daiichi reactor meltdowns and their aftermath; the massive failures in electrical reconstruction in Iraq and Afghanistan; the supposed crisis in STEM education; and the often-overlooked environmental costs of electric vehicles.

Though J.J.G. McCue was an engineer, he instinctively grasped something that all journalists know and that the writer Albert Camus expressed most concisely. "A free press can of course be good or bad," he said in 1955. "But, most certainly without freedom it will never be anything but bad." –GLENN ZORPETTE

NEWS

US $1,000: VALUE OF UNCLAIMED CHAMPAGNE OFFERED TO ANY IBM ENGINEER WHO COULD FIND A BUG IN THE TRUENORTH CHIP

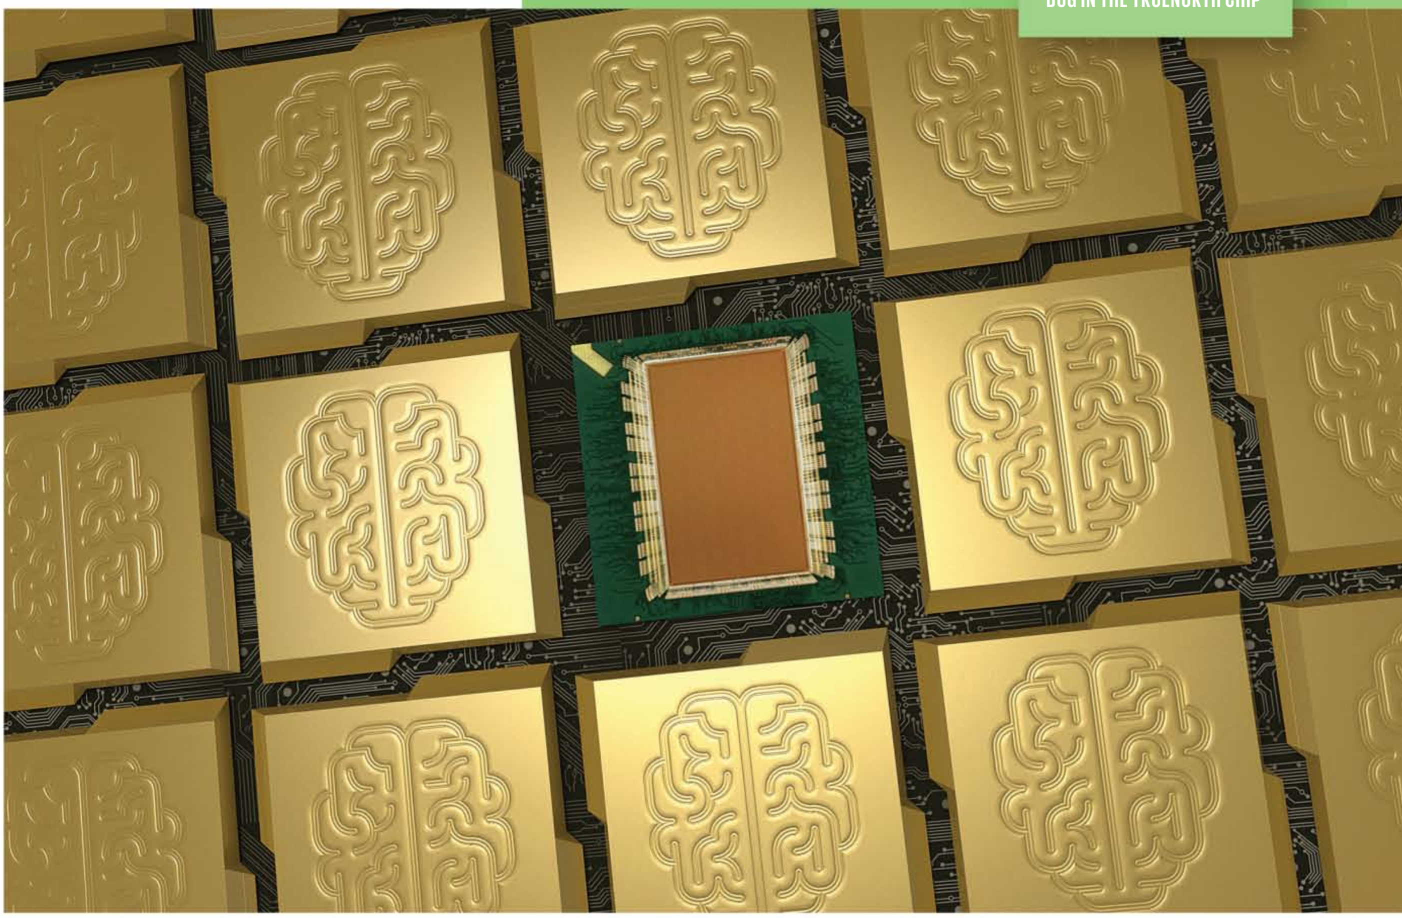

IBM'S NEW BRAIN

The TrueNorth neuromorphic chip takes a big step toward using the human brain's architecture to reduce computing's power consumption

IBM

Neuromorphic computer chips meant to mimic the neural network architecture of biological brains have generally fallen short of their wetware counterparts in efficiency—a crucial factor that has limited practical applications for such chips. That could be changing. At a power density of just 20 milliwatts per square centimeter, IBM's new brain-inspired chip [above] comes tantalizingly close to such wetware efficiency. The hope is that it could bring brainlike intelligence to the sensors of smartphones, smart cars, and—if IBM has its way—everything else.

The latest IBM neurosynaptic computer chip, called TrueNorth, consists of 1 million programmable neurons and 256 million programmable synapses conveying signals between the digital neurons. Each of the chip's 4,096 neurosynaptic cores includes the entire computing package: memory, computation, and communication. Such architecture helps to bypass the bottleneck in traditional von Neumann computing, where program instructions and operation data cannot pass through the same route simultaneously.

"This is literally a supercomputer the size of a postage stamp, light like a feather, and low power like a hearing aid," says Dharmendra Modha, IBM » fellow and chief scientist for brain-inspired computing at IBM Research-Almaden, in San Jose, Calif.

Such chips can emulate the human brain's ability to recognize different objects in real time; TrueNorth showed it could distinguish among pedestrians, bicyclists, cars, and trucks. IBM envisions its new chips working together with traditional computing devices as hybrid machines, providing a dose of brainlike intelligence. The chip's architecture, developed together by IBM and Cornell University, was first detailed in August in the journal *Science*.

"The impressive aspects of TrueNorth are the integration density—a million neurons on a single, admittedly very big, chip—and the very low power consumption for this many neurons," says Steve Furber, a professor of computer engineering at the University of Manchester, in England, who is behind a competing effort [see "To Build a Brain," *IEEE Spectrum*, August 2012].

With a total of 5.4 billion transistors, the computer chip is one of the largest CMOS chips ever built. Yet it uses just 70 mW in operation and has a power density about 1/10,000 that of most modern microprocessors. That brings neuromorphic engineering closer to the human brain's marvelous efficiency as a grapefruit-size organ that consumes just 20 W.

IBM minimized power usage in several ways. For one, it traded the traditional processor's clock—used to trigger and coordinate computational processes—for a more biological concept called event-driven computing. TrueNorth's digital neurons can work together asynchronously without a clock by reacting to signal spikes, which are the output of both real neurons and silicon ones.

IBM also saved on power through the design of an on-chip network that interconnects all the chip's neurosynaptic cores instead of using extra power to communicate with off-chip memory. And finally, it made the chip using a process technology meant for producing low-power mobile processors.

One brainlike feature that IBM did not mimic to reduce power consumption was to make TrueNorth's neurons analog instead of digital. The choice to go all-digital led to a number of advantages: First, IBM dodged the problem of slight differences in the manufacturing process or temperature fluctuations that have an outsize effect on analog circuits.

BIG-BRAIN CHIPS

TrueNorth IBM, DARPA SyNAPSE

Project features Low-power neuromorphic chip designed for applications in mobile sensors, cloud computing, and so on.

Chip specs 1 million neurons; 256 million synapses

Largest current configuration 16 chips: 16 million neurons; 4 billion synapses

Power density 20 mW/cm²

Final configuration 10 billion neurons; 100 trillion synapses

Neurogrid Stanford University

Project features Simulates biologically realistic brain models; team has moved on to autonomous robot applications.

Chip specs 65,536 neurons; 375,000 synapses

Largest current configuration 16 chips: 1 million neurons; 8 billion synapses

Power density 50 mW/cm²

Final configuration Info not available

HRL neuromorphic chip HRL Laboratories, DARPA SyNAPSE

Project features Emphasizes biological realism with learning capabilities on chip; has flexible, programmable routing approach.

Chip specs 576 neurons; 73,000 synapses

Largest current configuration 4 chips: 2,304 neurons; 292,000 synapses

Power density 120 mW/cm²

Final configuration Info not available

SpiNNaker University of Manchester, Human Brain Project

Project features Enables low-power, large-scale digital model of brain; helps improve models of brain diseases.

Chip specs Up to 16,000 neurons; 16 million synapses

Largest current configuration 1,152 chips: Up to 20 million neurons; 20 billion synapses

Power density 1,000 mW/cm²

Final configuration Up to 1 billion neurons; 1 trillion synapses

HICANN Heidelberg University, Human Brain Project

Project features Simulates brain activity at 10,000 times normal speed to compress a day into 10 seconds; focuses on studying brain learning and plasticity.

Chip specs 512 neurons; 144,688 synapses

Largest current configuration 6 wafers: 1.2 million neurons; 300 million synapses

Power density 3,000 mW/cm²

Final configuration 5,000 wafers; 5 billion neurons; 1.3 trillion synapses

SOURCES: HEIDELBERG UNIVERSITY, HRL LABORATORIES, IBM RESEARCH-ALMADEN, STANFORD UNIVERSITY, UNIVERSITY OF MANCHESTER

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

Second, the lack of analog circuitry allowed the IBM team to dramatically shrink its hardware. Many experimental neuromorphic chips still use analog circuits that must be built using a process that on the Moore's Law curve is more than a decade behind the process used today, Furber explains. By comparison, IBM fabricated its chip using Samsung's 28-nanometer process technology–typical for manufacturing chips for today's mobile devices.

And finally, the digital design enabled TrueNorth's hardware to become functionally equivalent to its software–a factor that allowed the IBM software team to build TrueNorth applications on a simulator before the chip itself had been built.

The chip represents the culmination of a decade of Modha's personal research and almost six years of funding from the U.S. Defense Advanced Research Projects Agency (DARPA). Modha continues to lead DARPA's SyNAPSE project, a global effort that has committed more than US $100 million since 2008 to making computers that can learn. "Our long-term end goal is to build a 'brain in a box' with 100 billion synapses consuming 1 kilowatt of power," Modha says.

But the goal of that brain is different from the goals of other, similar projects. TrueNorth's digital circuits are designed with commercial applications in mind. Other projects have the goal of better understanding how the brain works.

SpiNNaker, a neural network based on digital circuits and led by Furber, uses a general-purpose parallel computing system tuned to run neurons based in software rather than hardware. The SpiNNaker team, which soon plans to simulate 100 million neurons using 100,000 chips, sacrificed the efficiency of dedicated hardware to gain the flexibility of software: "Our primary goal is to understand biology, so the flexibility is important in understanding the biological brain," Furber says. For applications-focused IBM, "the efficiency and the density of their chip is perhaps more important than the flexibility they've retained." –JEREMY HSU

BRETT HOKR

NEWS

IDENTIFYING EXPLOSIVES AT A DISTANCE

The random Raman laser is the latest technology to detect explosives and other nasty stuff from a safe vantage

Being standoffish is usually frowned upon—that is, unless what you're standing off from might be an explosive or a cloud of anthrax spores. That's why efforts have accelerated to develop standoff detection techniques that use lasers to identify chemicals and biological substances from a safe distance.

The newest entry in the field is called random Raman spectroscopy. Shine a laser beam into a loose material–say, a powder–and if the density is right, the photons will bounce around among the powder's particles until they stimulate a new laser emission. Such a random laser, as it is known, works much the same way as a more traditional laser cavity, only without mirrors.

Normally, about 1 in 10 million photons undergoes a process called spontaneous Raman scattering, in which it drops to a lower frequency determined by the particular molecule it's bouncing off. The random laser enhances this Raman scattering, producing a signal strong enough for a detector to pick up at a distance. By measuring the shift in frequency, scientists can tell the chemical makeup of the powder.

Marlan Scully and Vladislav Yakovlev of Texas A&M University, in College Station, demonstrated such a setup. Scully says they can perform spectroscopic analysis of a material at a distance of a kilometer, and that 10 kilometers should be possible. That would be useful for, say, a drone flying over an

RANDOM RAMAN LASER LIGHT: A laser beam fired at a powder causes the powder itself to become a laser, beaming out information about the material's molecular structure.

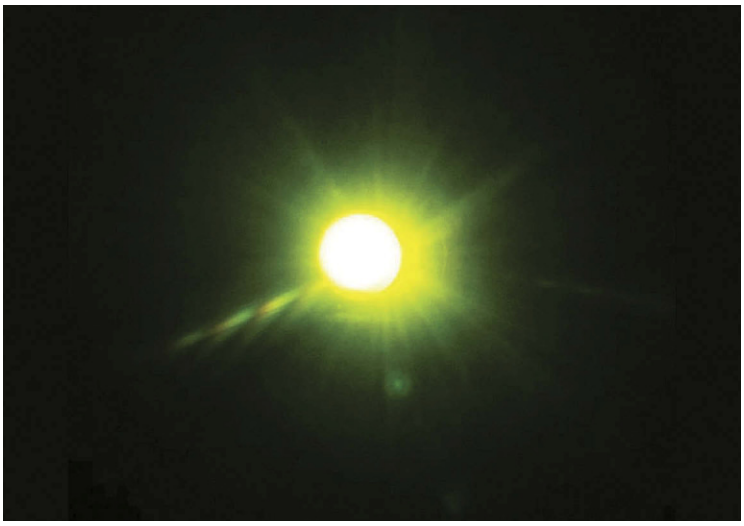

area where explosives might be hidden, or an airplane measuring the quality of soil on a farm.

Another approach to spectroscopy uses terahertz waves, or T-rays. T-rays have the advantage of being able to penetrate many substances, without the ionizing radiation of X-rays. But they have a downside. "The terahertz wave does not travel easily through the atmosphere because of water absorption," says Xi-Cheng Zhang, head of the University of Rochester's Institute of Optics, in New York.

One way around the problem is what he calls terahertz-radiation-enhanced emission of fluorescence, which is designed to detect trace gases emitted by an explosive. He focuses laser beams at two wavelengths on a point in the air, where they interact to create a plasma filament that fluoresces in the ultraviolet. He also emits a terahertz pulse. The T-rays interact with the material being studied to provide the spectroscopic information and then interact with the plasma field to alter its fluorescence. That encodes the spectroscopic information onto the ultraviolet radiation, which is easily picked up by a photodetector. Zhang says the challenge of doing this increases with distance, but he's already demonstrated detection at 10 meters.

Fow-Sen Choa, a professor of computer science and electrical engineering at the University of Maryland, Baltimore County, uses a quantum cascade laser to do photoacoustic spectroscopy. Heating a material with a modulated laser beam causes it to expand and create a pressure wave, as if it were a tiny audio speaker. Microphones pick up the sound wave and identify the material based on its frequency. "Whether it's TNT or fertilizer, you can tell pretty easily," Choa says.

Most of the development of this technique is focused on the accurate detection of the sound and elimination of noise, Choa says. "Distance is not yet the focus," he says. "The issue is how accurate you will be."

There's no ideal distance for how far the detector should stand off, says James Kelly, senior scientist at Pacific Northwest National Laboratory, in Richland, Wash., other than far enough to be safe. The distance, and the most effective technology, really depends on the particular requirements of an application.

Kelly's team is working on being able to measure a substance at a dosage of 1 milligram per square centimeter on a surface. That's about what you'd get if someone handling explosives had then touched something and left a fingerprint. The team would like to be able to use an eye-safe system such as hyperspectral imaging to scan vehicles coming to a checkpoint or parking at a stadium, for instance, to see if there are any traces of explosives on them. Because it can be challenging to tease out such a signal from those given off by the paint and other coatings on the surface of a car, researchers at PNNL and other teams are using an eye-safe tunable laser to do reflectance-based hyperspectral imaging, in which multiple images of the surface are taken at different wavelengths under the laser's illumination. Two substances that might be indistinguishable at one wavelength can look very different at another.

For that application, which could be used by the United States' Transportation Security Administration or border patrol, a distance of 50 to 100 meters might be desirable, Kelly says. A drone surveying a war area would probably require detection distances in the kilometer range.

For a lot of the techniques being developed, it's not so much the detection technology itself that's the bottleneck but the analysis of the signal, Kelly says. Finding trace amounts of explosives does little good at a checkpoint if it takes several minutes of computer processing to identify them.

In the end, no one technology is likely to win out, researchers say. More probably, the one that is used will be the one best suited to a particular need. "There's a whole gamut of techniques people are looking at," Kelly says. —NEIL SAVAGE

1 milligram per square centimeter is what you get if a person handling an explosive touched something and left a fingerprint

A BETTER TEST THAN TURING

The Turing test is a flawed metric for human-level AI. Can Winograd schemas do better?

In June 2014, a computer program named Eugene was able to convince 33 percent of the humans it chatted with that it was a real 13-year-old Ukrainian boy named Eugene Goostman. The experience was part of an artificial intelligence contest held by the University of Reading, in England. Eugene was declared to be the first AI to have passed the Turing test–Alan Turing's 1950 attempt to provide a framework for determining whether machines can think.

The Turing test is simply a short conversation between a human judge and either another human or a computer program. If the judge can't tell within 5 minutes whether he or she is talking to a human or a computer more than 70 percent of the time, Turing suggested, we should be comfortable saying that the computer program is capable of thinking (or has achieved a specific kind of human-level intelligence) because it's consistently able to make us believe that it's thinking.

But there has been plenty of criticism about the value of the Turing test, as well as over the implementation of the University of Reading contest. "Chatbots like Eugene Goostman get away with changing the subject, not giving direct answers, and joking around," says Leora Morgenstern, chair of the executive committee at Commonsense Reasoning, an organization focused on enabling artificial intelligence systems to solve problems the way humans do. "But that doesn't make them intelligent in any real sense. What does fooling a human judge really have to do with intelligence?"

Morgenstern believes there are more effective ways of determining whether a computer program has the ability to demonstrate intelligence: by seeing whether it is able to use a large body of knowledge to correctly reason about an intentionally ambiguous statement. Such statements are called Winograd schemas, named for Stanford computer science professor Terry Winograd, who

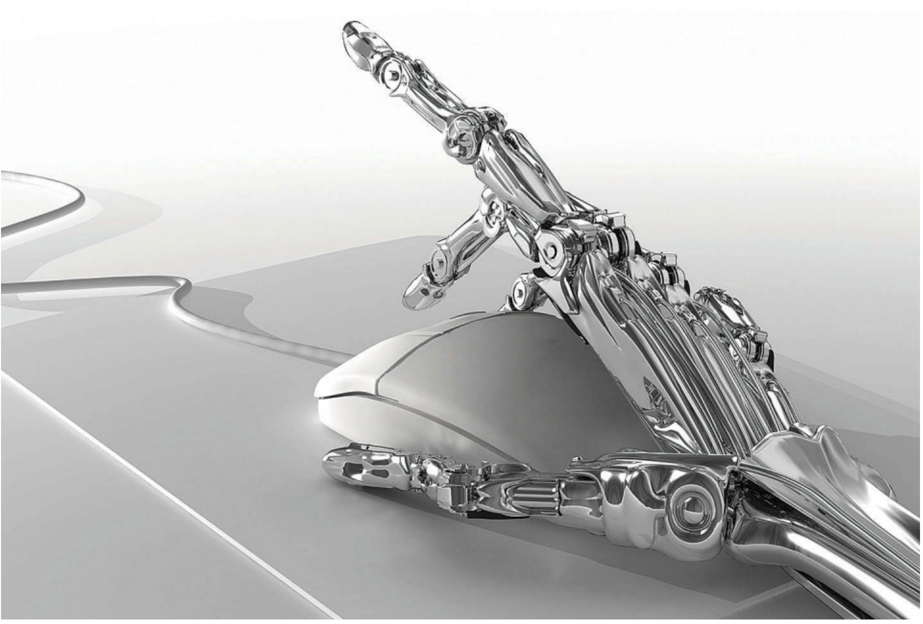

KONSTANTIN INOZEMTSEV/GETTY IMAGES

first proposed their use. This is what they look like:

Pete envies Martin because he is very successful. Who is very successful?

Answering this question is as simple as determining which person the pronoun *he* refers to. It's straightforward for humans because we have a large body of knowledge to reason from that includes what *envy* and *success* mean and how they apply to interactions between people. But as Morgenstern explains, "Envy and success are sophisticated concepts that most computers don't even represent in any meaningful way, let alone reason with."

Morgenstern and two colleagues first proposed using Winograd schemas as an alternative to Turing tests in 2012. Critically, a computer program's explanation of how it solves a Winograd schema can be used to assess how "deep" its thinking process really is. For example, consider:

The man couldn't lift his son because he was so heavy. Who was heavy?

One way an AI might solve this problem is by searching the Web to find similar phrases that might include clues, or even answers. Many programs that do well at chat-based Turing tests use this method. But a program that solves the problem through search would not be thinking nearly as "deeply" as a program that instead reasons using the knowledge that heavy objects are hard to lift. And even if searching solved the first case, Winograd schemas could still sniff out the dull bots. In each schema there is a "special word" that if replaced with another word maintains the sense of the sentence but changes the answer. Substituting *weak* for *heavy* above would befuddle a search-based bot, even if it managed to get the original version right.

Nuance Communications, in Burlington, Mass., and Commonsense Reasoning are cosponsoring the Winograd Schema Challenge, scheduled to take place in 2015. The challenge will consist of at least 40 Winograd schemas that each computer program can answer any way it sees fit. The organizers, including Morgenstern, are hoping that programs that use commonsense knowledge and inference will come out on top. "It will have to know about physical properties of materials, goals and plans, emotions, [and] personal interactions," she says. It will have to "reason from all that knowledge to get to the correct answer. If it can do that, it's going to be able to do a whole lot of other commonsense reasoning as well."

Even if a computer program wins the Winograd Schema Challenge next year, saying that artificial intelligence has reached a human level would still be nearly as hasty as making the same claim for a chatbot that passes a Turing test. But the ability to reason and infer using a large database of commonsense knowledge has the potential to make software much better at both understanding what humans want and communicating with us as we do with one another.

—EVAN ACKERMAN

A version of this article appeared online in July.

ANATOMY OF AN AI TEST

Winograd schemas might be a better test of human-level artificial intelligence than the Turing test because they require reasoning about a broad body of knowledge. Each schema has four requirements.

1 Two parties (males, females, groups, objects) are mentioned in a sentence.

THE TROPHY DOESN'T FIT IN THE BROWN SUITCASE BECAUSE IT IS TOO BIG/SMALL.

WHAT IS TOO BIG/SMALL?

2 A pronoun or possessive adjective is used in the sentence to refer to one of the parties, but that word could also refer to the second party.

3 The question involves determining the referent of the pronoun or possessive adjective.

4 The sentence, and possibly the question, includes a "special word." Replacing the special word with a designated alternate word results in a sentence that still makes sense but also produces a different answer to the question. (Here, replace "big" with "small.") An AI should be able to answer either version correctly.

NEWS

NEWS

FIREFIGHTERS GET LIFE-SAVING WEARABLES

Futuristic firefighter gear integrates vital-sign sensors and indoor tracking

In 2012, the U.S. Fire Administration reported that nearly 50 percent of U.S. firefighter deaths were caused by heart attacks. Wearable health-monitoring devices and other personal tech could help prevent some of these attacks, say experts, but for a variety of reasons–technological, economic, and social–they haven't been adopted. Now might be the time, say several technology suppliers. Like all consumers, firefighters are starting to adopt wearable tech in their personal lives, and that's paving the way for wearables in their work.

"We're basically seeing an explosion of wearables, not just in consumer markets but also in our core markets, like public safety," says Bert Van Der Zaag, senior manager of interaction design at Motorola Solutions, in Ft. Lauderdale, Fla.

Possibly the first to make inroads will be Globe Manufacturing Co.'s Wearable Advanced Sensor Platform. WASP is a flame-resistant T-shirt with an embedded adjustable strap that has a removable physiological sensor. Secured to a waist belt is a device (made by TRX Systems) that uses accelerometers, a compass, and other sensors to locate firefighters relative to a fixed point, even inside buildings. Mark Mordecai, director of business development at Pittsfield, N.H.-based Globe, says that this fall three firefighter-training academies in the United States will start instructing students in the use of the WASP system. He estimates that within the next year, thousands of firefighters will be training with WASP.

"If we really want to make a difference, we have to go beyond just creating the protective envelope around the firefighters and really take it to the next level by monitoring how they're doing and where they are," says Mordecai.

INFERNO INTELLIGENCE: Globe Manufacturing Co. is training firefighters to use its Wearable Advanced Sensor Platform, which tracks the firefighter's vital signs and location.

GLOBE MANUFACTURING CO.

Motorola had that in mind when its team of engineers envisioned a high-tech firefighter suit, which is part of a prototype concept called Next Generation Fireground Communications. It incorporates a host of wearable technologies, including a helmet-mounted camera, a head-up display on the breathing mask, an environmental sensor, a strap that monitors vital signs, indoor location tracking, and a rugged radio.

For Motorola, keeping the system hands free was critical. "If you look specifically at firefighters, when they run into a building, they never run into a building empty handed," says Van Der Zaag, whose team runs into blazing buildings alongside firefighters during training drills. So if they want to access critical stats, such as a heart rate, "they can't really pull out a smartphone in the middle of a building to get that information." The head-up display shows firefighters when battery levels, oxygen supplies, and even their own heart rates are reaching critical levels.

Although the visual display helps firefighters keep their eyes up and hands free, there's a fine line between providing vital information and mucking up the ability to navigate through dangerous obstacles. "You don't want to put anything in there that you don't need at a moment's notice," Van Der Zaag says.

It will take some time before it's clear whether firefighters really warm to wearables. Even a few years ago, such systems never made it past the prototype phase. ProeTEX, which started as a multi-university research project, also integrated heart rate, breathing rate, and internal temperature sensors into a firefighter's inner garment. But according to Nicholas Walker, the former project manager of ProeTEX, not much has happened with the technology since the last prototype was completed in 2010.

What has happened since then is that the market and culture for wearable technology has shifted. "Sensors are getting smaller. Things like head-up displays are becoming accessible," says Van Der Zaag. "So you see firefighters experimenting with it. You see them creating their own apps to solve some of the problems that they're having. So it has really accelerated the adoption of and acceptance of these types of technologies." –THERESA CHONG


WE DON'T JUST FILL YOUR POSITIONS, WE FULFILL YOUR VISIONS

PHOTOGRAPH BY Kenneth Armstrong/Reuters/Corbis

ROADSIDE ROBOT

AT THE END OF JULY, this conglomeration of parts—which at first glance has the look of a discarded sixth-grade science project—was placed on the curb in front of Canada's Nova Scotia College of Art and Design, in Halifax. But that represented the beginning, not the end, of its useful life. The hitchBOT, as it's called, took a 6,000-kilometer cross-country trip this summer, thumbing its way to Victoria, B.C., by the end of August. The crude-looking robot, designed and built in Canada by an interdisciplinary team of researchers from McMaster University and Ryerson University, is actually quite sophisticated. Advanced conversation software let it ask drivers how far they could take it and whether they'd mind letting it recharge via their cigarette lighters. Once inside the vehicles, hitchBOT took pictures of the countryside, played music, and live tweeted its experience the whole way.

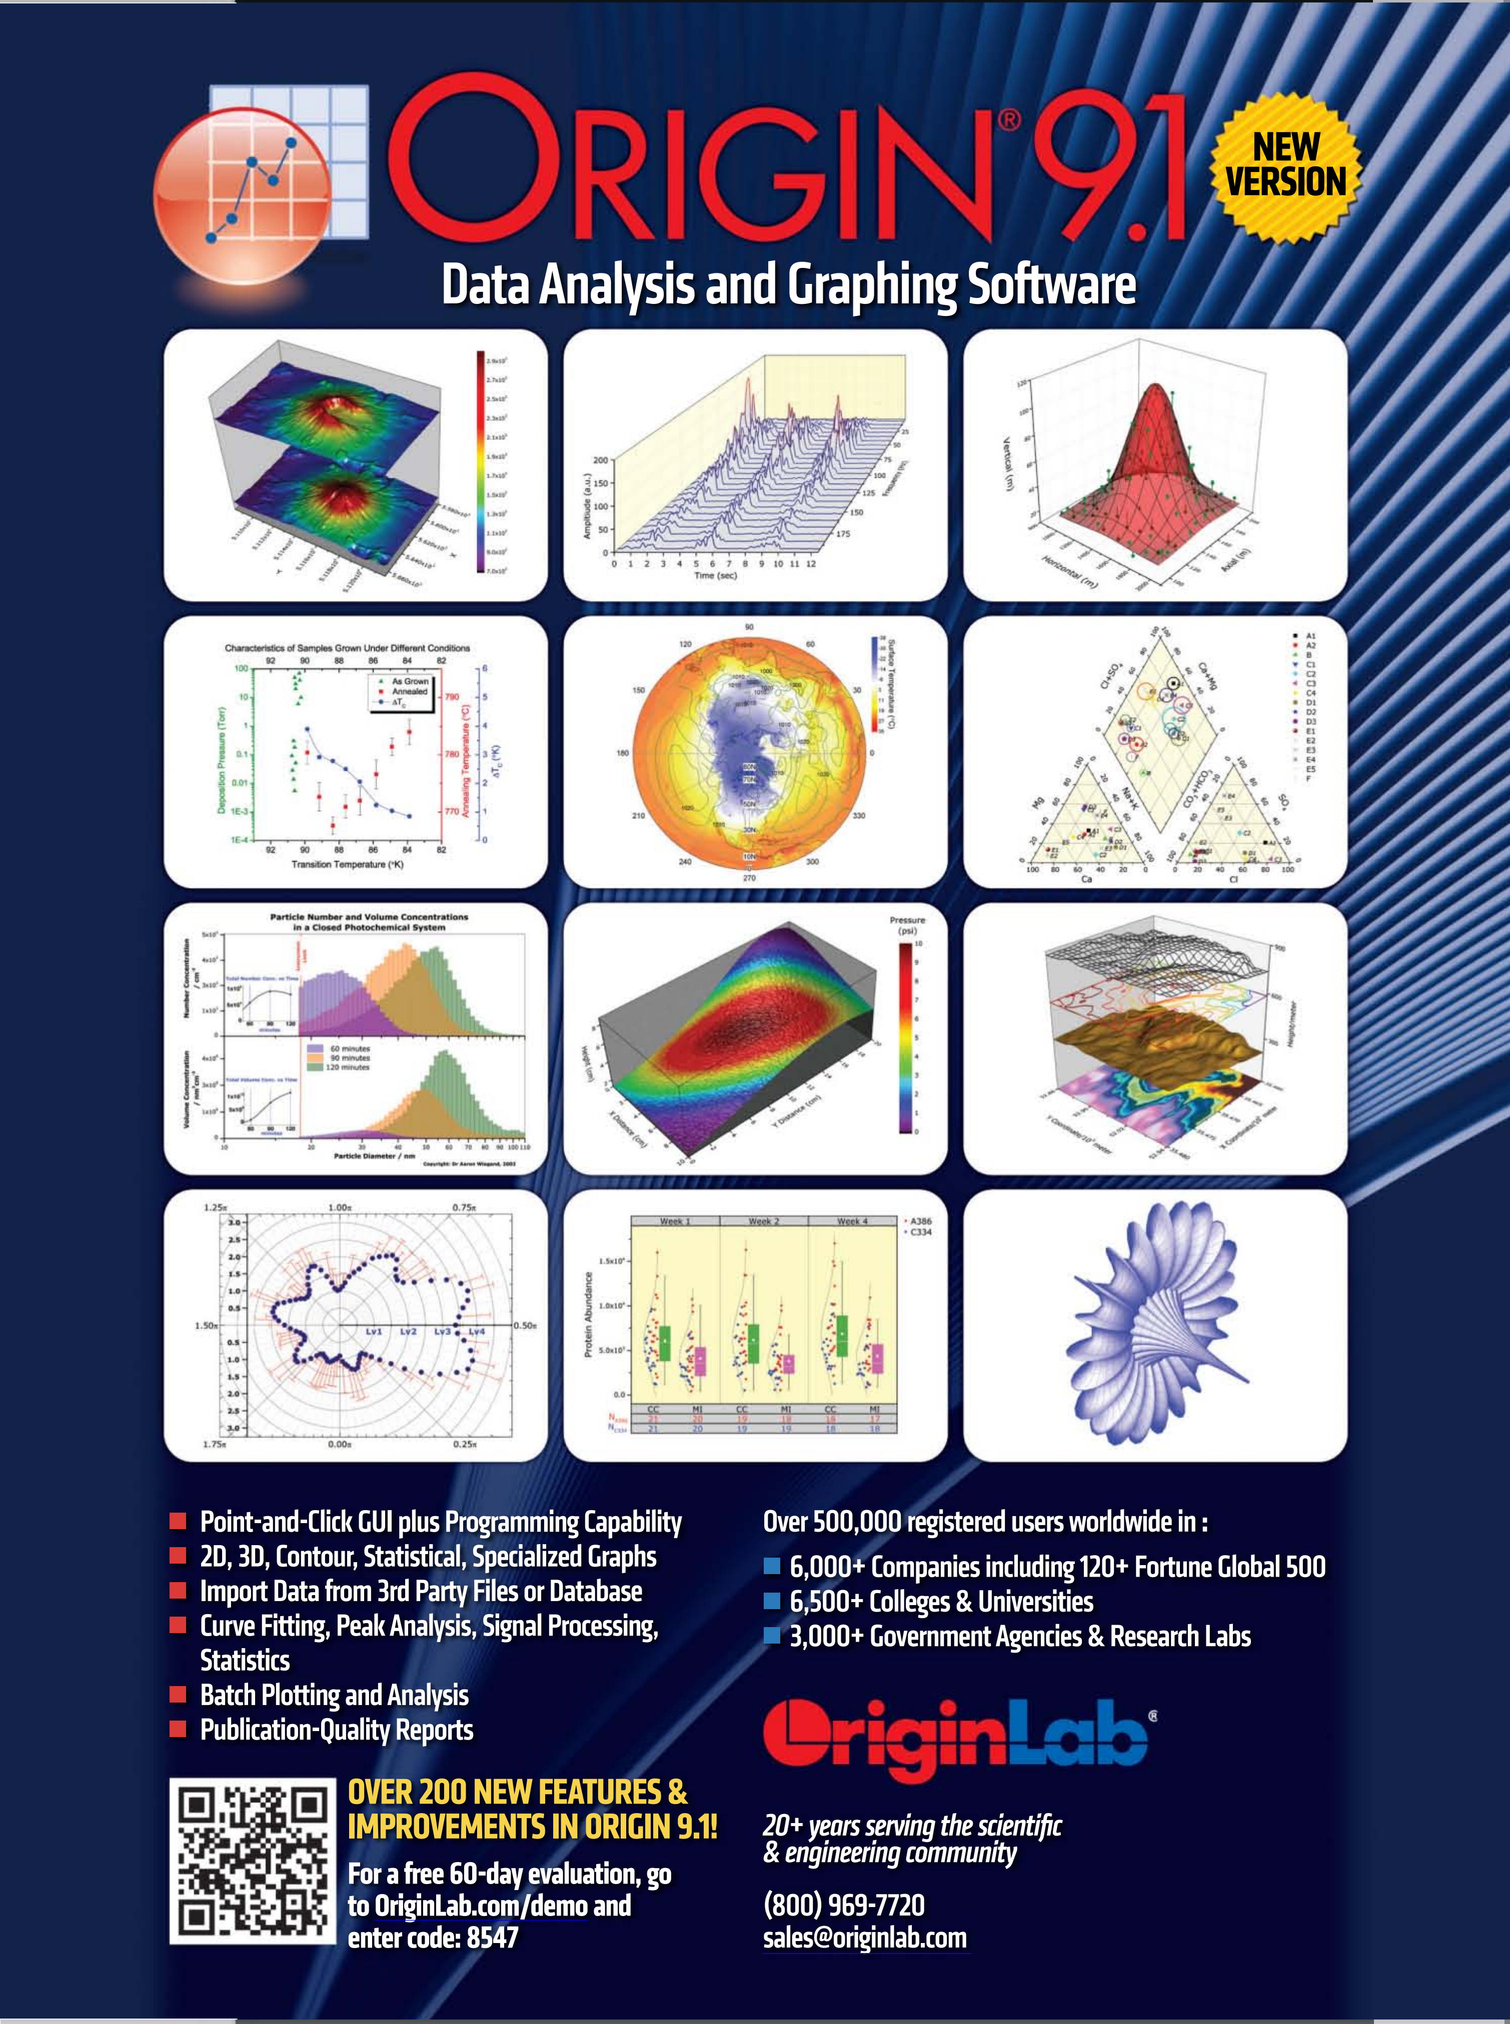
ORIGIN® 9.1
NEW VERSION
Data Analysis and Graphing Software
Point-and-Click GUI plus Programming Capability
2D, 3D, Contour, Statistical, Specialized Graphs
Import Data from 3rd Party Files or Database
Curve Fitting, Peak Analysis, Signal Processing, Statistics
Batch Plotting and Analysis
Publication-Quality Reports
Over 500,000 registered users worldwide in :
6,000+ Companies including 120+ Fortune Global 500
6,500+ Colleges & Universities
3,000+ Government Agencies & Research Labs
OVER 200 NEW FEATURES & IMPROVEMENTS IN ORIGIN 9.1!
For a free 60-day evaluation, go to OriginLab.com/demo and enter code: 8547
OriginLab®
20+ years serving the scientific & engineering community
(800) 969-7720
sales@originlab.com

RESOURCES

1,000 YEARS: THE LENGTH OF *LONG-PLAYER*, A MUSICAL COMPOSITION THAT COMPUTERS AND HUMANS STARTED PERFORMING NEARLY 15 YEARS AGO

HANDS-FREE SHEET MUSIC

NEVER LOSE YOUR PLACE, WITH A FLAT-PANEL MONITOR AND FOOT PEDALS

As an old-school writer, I prefer my reading matter printed and bound. But as an avid piano player, I am sick to death of music books. The music stand on our Baldwin grand piano holds a few pages laid side by side just fine. But most piano music is sold in songbooks: double-sided pages bound to inflexible spines. These books close up on you midsong when they are new; they fall to pieces when they get old. And playing complex music without interruptions is virtually impossible unless you have an assistant to turn pages for you. • Tinkerers since the early 1900s have patented myriad mechanical page turners to address this problem, but to this day they are generally considered a pain to load, noisy, and unreliable. • After years of growling at my misbehaving music, I finally resolved to create my own solution. I set the bar high. The system should work with all of my music books, allow annotations, and turn pages while my hands remain securely on the keys. It must be big enough to show two pages of notation sharply at full size, yet compact enough to fit on my piano's music stand. I wanted something self-contained that would work with a digital keyboard when I play with a band. And I didn't want to spend more than a few hundred bucks. • Clearly, the key was ditching print in favor of pixels. I lugged my boxes of books down to a copy center, and for US $30 they ran them through a guillotine to slice off the bindings. ▸

PHOTOGRAPH BY Ryan Matthew Smith

I then used a Fujitsu ScanSnap document scanner (modified to accept large pages) to convert the books into PDFs. Calibre, an open-source e-book library organizer, let me tag and sort scores by genre, difficulty, composer, and so on.

It was less obvious how best to display the scanned scores. Many people use tablets such as the iPad for this purpose, but none are big enough to display two typical 23-by-30-centimeter (9-by-12-inch) pages at full size. All-in-one computers are either too bulky to fit on the piano or exceeded my budget. What I need, I thought, is a thin, no-frills 24-inch monitor—and a cheap computer.

It occurred to me that a $35 Raspberry Pi might have just enough computing power and that I could tuck both the Pi and its cables out of sight behind the monitor, which I acquired for $170. Assembling the pieces was straightforward: The Pi, its case, a wireless mouse receiver, and a USB receiver for a pair of $80 wireless foot pedals all fit neatly onto the back of the display, held in place with Velcro strips. I ran a small bundle of power and Ethernet cables to a power-line network adapter, along with an audio cable to a separate amplifier for playing MP3 backing tracks.

The Pi runs Linux, so I had reams of open-source software to draw on. Only an obsolete version of Calibre had been released for the Pi in a precompiled package, so I had to download the latest source code and compile it, a tricky process that took several hours. I discovered that although Calibre is very good at organizing PDFs, it renders them too slowly when running on the Pi. So I turned to another open-source program called qpdfview, a streamlined app that prerenders upcoming pages.

I loaded the Duke Ellington song "It Don't Mean a Thing (If It Ain't Got That Swing)" and began playing. A bar before the first page turn, I tapped on the foot control. The next two pages appeared nearly instantaneously. I paged back with another tap—it was working!

But two problems emerged. Each time I finished a song, I had to pick up a wireless keyboard to close the PDF and return to the Calibre score library. And sometimes when I tapped the foot pedal, it sent two forward or back commands rather than just one, a classic case of contact bounce. I needed to filter out page-turn signals that arrived in quick succession.

MUSIC MACHINE

After cutting musical scores from their bindings, I fed them into a document scanner whose page guides I had removed [top]. I connected a pair of foot pedals [second from top] to a Raspberry Pi [second from bottom] via a wireless dongle. I then attached the Pi to the back of the monitor used to display scores.

To tackle the first issue, I dug through Linux software libraries and discovered voicecommand, a utility that hooks into Google's speech recognition server to allow the Pi to interpret spoken commands and utter responses. I hooked up a USB webcam for an audio input, and before long I was having conversations with the piano.

"Baldwin?" I would say.

"Hello," it would acknowledge in a pleasant voice.

"Change songs."

After pondering this for about 10 seconds, voicecommand complied by executing a command I had preprogrammed—sometimes. Just as often, it replied, "I do not understand."

This was no good, and the futility of this approach became clear as I launched into Gershwin's jazzy "I Got Rhythm" and voicecommand, mistaking the music for speech, chirped along, "I do not understand…I do not understand…"

More searching led me to the Brightside and xdotool utilities, which offer a better way to eliminate the keyboard. Combining them, I set up the system to load a new song by executing a sequence of commands whenever I moused the pointer into the upper right corner of the screen.

Debouncing the pedal signals was more straightforward: I got the source code for qpdfview, modified it to accept only one page-change command every second, and recompiled.

Now, a run-through of "I Got Rhythm" zips along interruption-free. And as an added bonus, the Pi can accompany me by playing prerecorded bass and drum parts on certain songs—and act as a Web server through which I can access any piece of sheet music from my library when I am away. It's music the way I like it: unbound.

—W. WAYT GIBBS

PHOTOGRAPHS BY Ryan Matthew Smith

REMOTE-COLLABORATION SOFTWARE FOR ENGINEERS

HARDWARE MANUFACTURERS NEED MORE THAN GENERAL-PURPOSE TOOLS

Hardware development, like many engineering projects, is typically a group effort, with contributors increasingly working from home or across multiple buildings, organizations, states, countries, and continents rather than face-to-face. And as a new wave of hardware start-ups are finding out, coordinating those contributors often requires more sophisticated tools than just e-mail and phone calls.

Many general-purpose remote-collaboration platforms are currently available and in use by engineers, such as Dropbox for file sharing, Basecamp for project management, Google Docs for collaborative document sharing and editing, and WebEx for screen sharing. But the nature of engineering often requires more specialized systems, such as the Git platform, which is designed for managing changes to source code and documentation.

Originally developed for managing the source code of Linux, Git has found broader use, often in conjunction with other tools. Matt Wood, principal engineer in charge of hardware and firmware development at FanWise, which makes mobile-commerce hardware modules for vending machines, arcade games, and other devices, says, "We use [commercial] tools like Atlassian Stash to do private Git repositories…for hardware files like schematics, board layouts, [bills of materials], and firmware."

But Git isn't always well suited to hardware. "Our private Git repository lacks a sign-off procedure—we create a tag for released files, but the system doesn't require any formal review by the team to do that," says Wood. "So we have to use custom tools to make sure we check off against the specs. You have to know how you're handling formal review and sign-off so a mistake isn't caught after

a design has been released to manufacturing and 10,000 bad boards have been made."

One company that's trying to provide better support for distributed hardware tools for EEs is Altium. "The old methods of copy and paste tend to be messy and error prone and require lots of cleanup time, which may prove to lose what should otherwise have been gained. Having a formal collaboration mechanism and work-in-progress version-control system breaks those problems down," says Max Clemons, an application engineer at Altium.

"Hardware development follows specific rules and sequences to produce high-quality, stable products in the quantity you need," agrees Lucas Wang, CEO of HWTrek, another company founded to develop software for hardware collaborators. "There are several steps you can't skip or detour before shipping a product." International collaboration using mainstream tools can cause particular challenges, says Wang, giving the example that "Google Docs doesn't function well in China, and there's a version-control issue."

Part of the challenge is a lack of domain knowledge, adds Wang: "Large buyers like Dell and HP know how to work with larger manufacturers like Foxconn and have project managers to handle communication, scheduling, and following details. But many hardware developers today are small start-ups, working on smaller projects—and don't have the in-depth knowledge of manufacturing, logistics, or other details involved in producing hardware."

To address this, HWTrek has been developing HWTrek Project Development Hub (PDH), a software-as-a-service tool to guide hardware product development from build to delivery. "PDH incorporates a lot of data, like hardware certification information and [step-by-step] schedules" that help get projects ready for manufacturers, says Wang.

Adam Kell, founder of FlameStower, a company that makes a gadget that can recharge USB devices by tapping heat from campfires, began using PDH after having difficulties getting the company's recharger built. He says, "Our main office is in San Francisco, our sales office is in Boulder, Colo., and our contract manufacturers are in Shenzhen, China.… One of the main pain points we had when working with our initial manufacturers was communicating when one of the team members wasn't on the ground in China." Kell adds that PDH is improving communication and scheduling among the involved parties. "Skype is helpful, but having a higher-level schedule to guide the progress of deliverables makes a much bigger difference."

—DANIEL DERN

ILLUSTRATION BY Mark Allen Miller

REVIEW: WIKIMUSICAL
THIS SHOW PUTS THE COLLABORATIVE NATURE OF INTERNET MEDIA ON STAGE

TREADING THE BOARDS: Trey Harrington [left] and Perry Sherman performing the number "Search Engine Crash" from *WikiMusical.*

BILLY BUSTAMANTE

Twenty years ago only a quarter of U.S. homes had PCs, and the biggest digital media decision facing a child was choosing between a Sega Genesis or a Super Nintendo games console. Now, digital media and the Internet are a constant presence in the lives of millions—but can you sum up their essence? Blake J. Harris thinks he can—with musical theater.

Harris and Frank Ceruzzi coauthored *WikiMusical,* which debuted at the Pearl Theatre Company during the 2014 New York Musical Theatre Festival, in July.

Harris is also author of the recently published book *Console Wars*, which looks at the 1990s battle between Sega and Nintendo. He cowrote the lyrics and script, which is a meandering journey through the sites, memes, and trends that make up the Web. (The music was composed by Trent Jeffords.)

It's silly and surreal, and hidden inside is an exploration of the ways we're making the Internet our home.

"The story of the tech age is such an interconnected story," says Harris. "It touches upon a lot of different threads. One of the great things about the Internet is that there's a seemingly infinite number of rabbit holes you can go down. The story is an attempt to make a narrative out of all of it."

The show starts in a simpler time, with siblings based on Harris and his brother getting a Gateway computer from an eccentric Santa Claus. It quickly jumps to the present: The grown brothers are pulled bodily into the modern Internet, where they must defeat the sinister and seductive Spam King to save the Web and return home, mending their relationship along the way. On their journey, they meet a blogger on a quest and a cast of unlikely online characters—including the cats that (evidently) invented the Internet, the video-game characters Mario and Luigi, and Morgan Freeman.

The festival production of *WikiMusical* had strong acting from the cast, and the songs were often catchy, although the narrative thread was sometimes hard to follow among all the different encounters and gags. In fact, narrative is one of the themes of the show: how to pull the universal tropes of a hero's journey from the nebulous, multifaceted, constantly evolving Internet.

According to Harris, "We live in such a niche age, whether it is websites or television channels, that having flagship websites or central hubs is more important than ever to [creating] shared communal experiences." Using well-known sites along with familiar memes and Web personalities as landmarks, Harris takes the audience on a State of the Internet tour, in an attempt to capture the zeitgeist of what we all experience when we experience the Internet and the communities we've created.

WikiMusical is at its strongest when it explores the Internet's collaborative nature and makes us stop and think about what we're all building together. Harris sees Wikipedia as a prime example of that world under construction: *WikiMusical*'s plot originally focused just on the editable encyclopedia—hence the show's title—before ballooning to encompass the larger online world.

"Wikipedia is basically this giant Ouija board that we all put our hands on and try to create an information network together," he says. "It represents the best and worst of the technological age." —**SARAH LEWIN**

A version of this article appeared online in July.

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

RESOURCES_TOOLS

THE BEAN
BOOSTING HEARING WITHOUT A PRESCRIPTION

For those over 40, it's not un- usual to begin to find it difficult to see nearby objects. But many can avoid—at least for a while—the hassle and expense of prescription lenses by buying a pair of cheap off-the-rack reading glasses from a drugstore. Hearing declines with age as well, but unlike with vision there hasn't traditionally been good intermediate technological assistance between simply turning up the TV volume and getting a hearing aid, which can cost US $4,000 and up.

In recent years, a number of companies have tried to step in to fill this niche with personal in-ear digital sound amplifiers that selectively boost the volume for the audio frequencies used in speech and do some ambient noise cancellation. There are also smartphone apps that use the phone's microphone and send enhanced sound to your earbuds. Now Chicago-area Etymotic Research has entered the market with a different approach.

Etymotic's Bean Quiet Sound Amplifier is the size and shape of a wireless earbud and costs $375 for one or $700 for a pair. It pays attention to volume as well as frequency, boosting both soft sounds and high-pitched sounds (human hearing inevitably degrades over time, with less and less sensitivity to high frequencies). Other sounds are passed through without amplification; very loud sounds (above 120 decibels) are not passed through at all. It can handle sounds from 50 hertz to 16 kilohertz, a wider range than most hearing aids.

CECI N'EST PAS UNE HEARING AID: Designed for occasional use rather than the treatment of hearing impairments, the Bean selectively amplifies quiet sounds and voices.

The Bean is classified as a personal sound amplifier and not a hearing aid by the U.S. Food and Drug Administration, allowing it to be sold directly to consumers—the difference between the two being that personal sound amplifiers are intended for occasional use and not to treat a hearing impairment.

I'm not in the target market, at least not yet. But I was able to test a pair—and notice a difference—in several situations. The Bean's plastic casing makes it light but also makes it feel fragile, although I didn't have any durability problems. A switch that changes the maximum level of amplification from 15 dB to 23 dB was inaccessible when the gadget was in my ear—that is a bit frustrating. And the on-off toggle is a little too basic: To turn it off, you open the battery door, which can accidentally snap shut if you're carrying it in a pocket.

Each Bean comes with a selection of tips for seating it properly in the auditory canal. The standard ones fit me fairly comfortably, and I'm someone who is rarely comfortable with earbuds. It was a little tricky to get them inserted at just the right depth: A tad too shallow and the devices whistle; a little too deep and the sound seems muffled.

The Beans' effect was startling sometimes. I first put the pair in while walking through town on a busy street. Traffic noise was unchanged, but the sounds of unseen birds jumped dramatically: It felt like I had entered an aviary. In a movie theater, they were less useful; I had to terminate my test when the sounds of people four or five seats away munching popcorn turned out to be extremely annoying. In a noisy bar though, they lowered the overall din and helped me hear my friends' conversations a little better when their volume dropped off and everybody else at the table started asking "What?"

I found the gadgets most useful when watching TV late at night. I try to avoid disturbing my noise-sensitive kids down the hall, and I was able to turn the TV down to a murmur and still easily follow the dialogue. There's no standard for TV volume, but, for comparison, during the day I might have the volume control on my bedroom TV, about 5 meters from my head, turned to level 15; normally at night I might get it down to 11 or 12. With the pair of Beans it was fine at 6 or 7. Maybe I'm not so far off the target market as I thought. —TEKLA S. PERRY

ETYMOTIC RESEARCH (2)

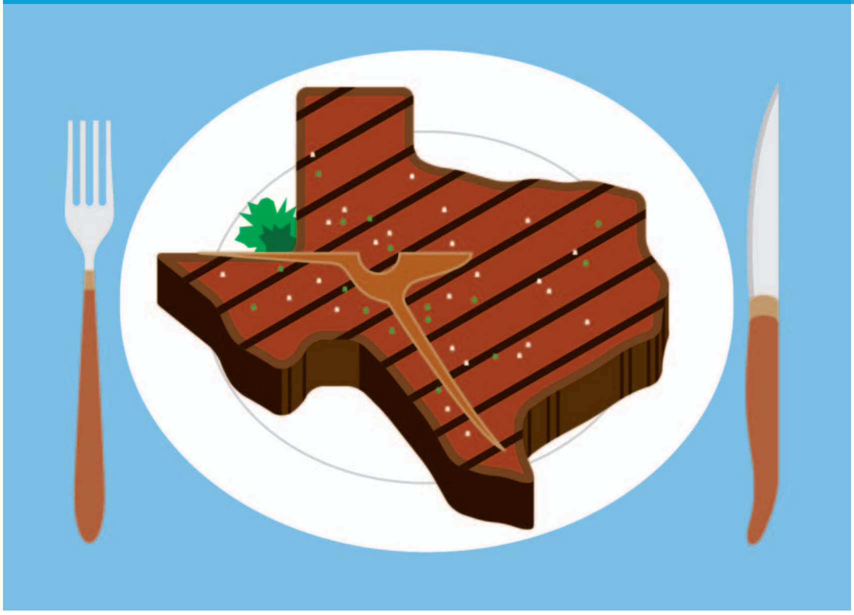

MAPHEADS AND ROADGEEKS: THE NEW CARTOGRAPHY

It's never enough just to *be* at a place—anyone can do that. The trick is to know where you are.
—Ken Jennings, maphead

WHEN GOOGLE LAUNCHED its Maps service in early 2005, it didn't include an application programming interface (API), but that didn't stop Paul Rademacher from figuring out how to use Maps to display markers indicating available apartments in the San Francisco Bay Area. This was not only the first **mashup** (information created by combining data from multiple sources) but also the unofficial beginning of **neogeography** and **neocartography**. • Neogeography is the practice of combining online maps with data–such as blog posts, websites, and annotations–related to locations on those maps. It's a subset of neocartography, also called **citizen cartography**, which is mapmaking as practiced by nonprofessional cartographers like Paul Rademacher and, nowadays, just about everyone else. Services such as Google Maps and OpenStreetMap, as well as the availability of massive location-based data sets, have made **neogeographers** of many of us. Great chunks of the population have been revealed as **mapheads**, people who are passionate about maps and cartography. • This **cartophilia** takes many forms, but one of the strangest (and hardest to pronounce) is **cartocacoethes** (kart-oh-kak-oh-EE-theez), the tendency to see random patterns as maps. This mouthful of a word combines the prefix *carto-*, "maps," and the word *cacoethes*, "an itch or compulsion." It has led to the fun disciplines of **accidental cartography** and **found cartography**, where everyday objects bear an uncanny resemblance to maps. The opposite, in a sense, is **counter cartocacoethes**, where maps are concealed from prying eyes by making them look like something noncartographic. • **Geonerds** combine their passion for maps and their **topophilia** (the love of landscape) into new hobbies. One of the most popular is **geocaching**, a type of scavenger hunt in which participants are given the geographical coordinates of a cache of items and then use their GPS devices (smartphones, nowadays) to locate the cache. **Geocachers** fall into a variety of categories. The **megacachers** are the alphas of the geocaching world because they've found the most caches (numbering in the thousands), while **power cachers** enjoy the challenge of finding as many caches as possible in a set period of time (often using **power trails**, paths that yield lots of caches). **Extreme cachers** take only the most dangerous **geotrails**, while **puzzle cachers** have to work out a riddle or similar puzzle before they can locate a cache. At the bottom of the hierarchy are the **nongeocachers**, who may accidentally happen upon a cache and are sometimes called **muggles** (because they're the geocaching equivalent of nonwizards in the Harry Potter books). Particularly prized are **virgin caches**: Getting to one of these before anyone else earns the geocacher a coveted **FTF** ("first to find").

Another favorite pastime of the **geoworld** is the **confluence hunt**, in which participants seek **degree confluence points**: locations where the latitude and longitude coordinates are perfect integers. **Confluence hunters** find the exact spots, snap pictures, and then upload them to the Degree Confluence Project.

Why this newfound **geojoy**? Perhaps it's because maps are a built-in feature of our brains. As we find our way in the world, we construct **cognitive maps**, mental representations of the real world, so is it any wonder that physical and digital maps resonate with us in powerful ways? Some believe that using GPS devices prevents us from creating these cognitive maps because we let the device do all the work. We no longer ask for directions because we're never lost, and we no longer truly experience a place because we are focused on following a line on a screen. These are serious concerns, but my own belief is that GPS-enabled devices can, if used in the right spirit, make us *more* likely to get lost and to experience where we are. We just have to remember to put our phones away and wander. The maps will always be there when we need them. ■

ILLUSTRATION BY L-Dopa

AFTER SEVERAL RINGS, John Beerends picks up my call on his cellphone. Beerends, a senior researcher at the Netherlands Organization for Applied Scientific Research, in Delft, is one of the world's top experts on sound perception, and I've called from Boston to ask his opinion on the quality of audio on mobile phones. But the connection keeps cutting out, and what I can hear is almost unintelligible. I must sound just as bad, because he asks me to dial him back on his landline. This time, his voice is much clearer. And he immediately confirms what now seems glaringly obvious: Despite their ubiquity and decades-long existence, cellphones still make for pretty poor phones.

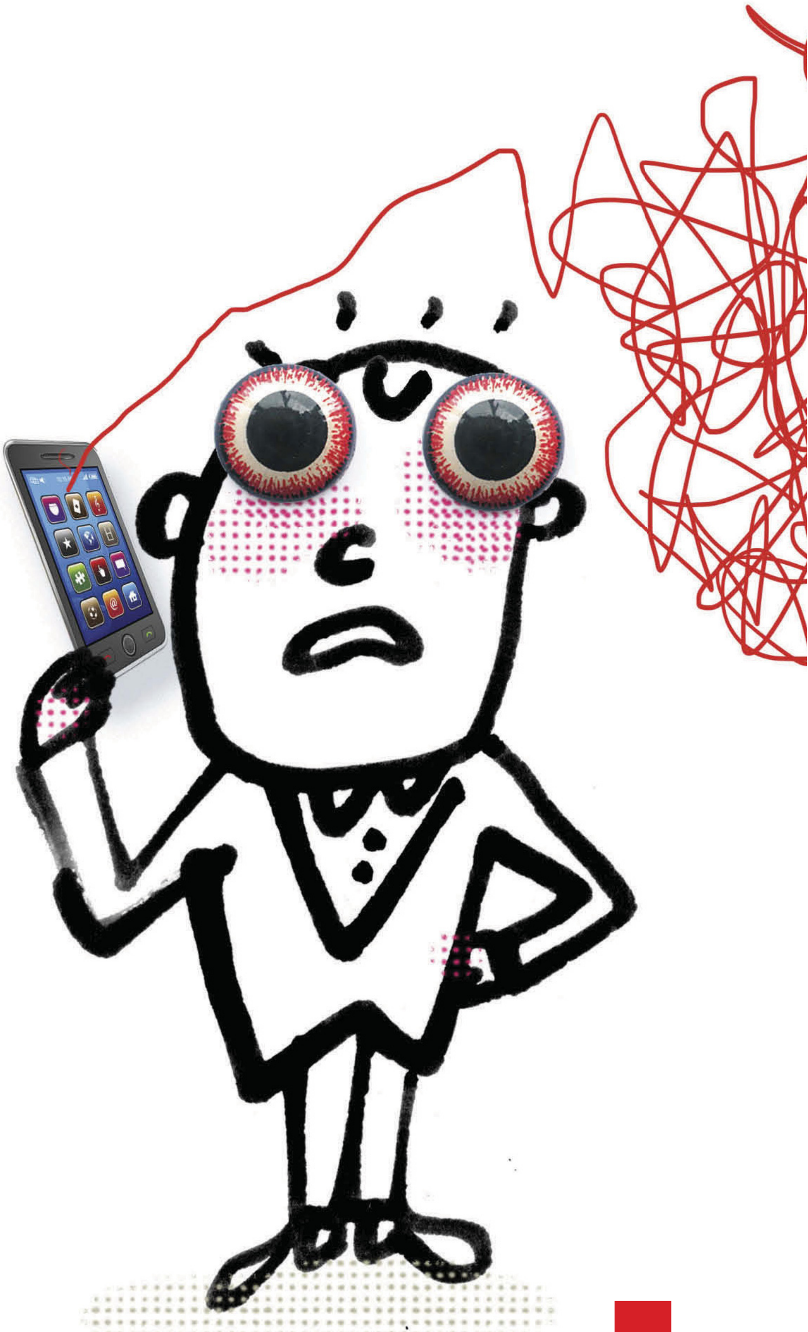

All Smart,

CELLULAR CARRIERS ARE DRAGGING THEIR HEELS

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

How can that be? After all, today's smartphones are incredible feats of engineering. Packing the processing power of a mid-1980s supercomputer into a sleek, pocket-size slab, they can take photographs, play music and videos, and stream tens of megabits of data to the palm of your hand every second. But try calling your boss in rush-hour traffic to say you're running late, and there's a good chance your message won't get through. "Mobile companies have rather lost the focus on a smartphone also being a telephone," says Jeremy Green, now a tech-industry analyst at Machina Research, in Reading, England–on a cell connection that keeps dropping words. ››

By **Jeff Hecht**
Illustrations by Serge Bloch

No Phone

OVER TECHNOLOGY TO IMPROVE VOICE QUALITY

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

BOTTLENECK #1

Handset design

Once upon a time, cellphones were just phones. The rise of smartphones has seen microphones take a back seat to other design considerations.

Solution: Using multiple microphones can clean up a distorted voice.

Laboratory tests by Broadcom confirm that it's not just my aging ears: Even in the best conditions, including a quiet environment and a strong wireless signal, users consistently rate voice quality lower on a cellphone than on a landline. Weaken the cellular link or add background noise, such as from wind or street traffic, and callers' opinions of the experience drop dramatically.

For example, engineers at Nokia found that when they compressed voice data to 5.15 kilobits per second, which cellphones do automatically when a tower connection is weak, user ratings fell from "good" to "fair." When the engineers decoded and then recompressed the data, which happens when a call travels through the backbone network to another cellphone, the ratings dipped lower still.

So why do we mobile subscribers—all 4.5 billion of us—put up with such crummy voice service? In the early days of cellphones, their fickleness wasn't such a big deal. Back then, mobility was a luxury, a handy supplement to a dependable wired line. But now, more and more people are cutting the cord—or never installing one. Today in the United States, for instance, 40 percent of homes rely exclusively on mobile phones for making and receiving calls. In Africa, cellular subscriptions outnumber landlines 52 to 1, according to the International Telecommunication Union.

At the same time, U.S. wire-line carriers, which by law must provide service to everyone at reasonable prices, are pushing the government to let them phase out the aging public switched telephone network. Rather than replace decaying or damaged copper circuits, they plan to deliver Internet-based service, known as Voice over IP (VoIP), using cable or fiber lines. In remote areas where wired broadband is too costly, however, customers could be forced to rely on wireless links. On Fire Island, in New York, for example, Verizon refused to rebuild the copper infrastructure destroyed by Hurricane Sandy in 2012, offering a fixed wireless service called Voice Link in its stead. Although Verizon eventually promised to install a fiber network on the island after customers protested, other communities may not be so lucky in the future.

Upgrading cellular networks to provide high-quality voice service is now more important than ever. So what's the holdup?

One explanation is that there's no silver-bullet fix. Most cellular voice traffic today passes through a patchwork of diverse systems, each exchange point an opportunity for degradation and delays. But lest you despair, here's some good news: Solutions to many of these impediments are in the works or already available, including a standard known as HD voice or wideband audio, and an add-on to today's 4G LTE systems called Voice over LTE. Although many operators, particularly in the United States, have been slow to roll them out, these technologies could very well prove to be game changers.

THE FIRST OBSTACLE to a good-quality voice connection on today's mobile phones is their design. Handsets have evolved considerably since Motorola debuted the original "brick" phones, made famous by Michael Douglas's suave character Gordon Gekko in the 1987 movie *Wall Street*. With its ear-size speaker and microphone pointed directly at his mouth, the monstrosity Gekko used was clearly constructed with voice calls in mind. Modern phone makers have taken a new tack. The smartphone's form "is driven by industrial design and not voice quality," says Chris Kyriakakis, founder and chief technology officer of Audyssey, a Los Angeles-based acoustic design company.

For example, to create an elegant, palmable chassis for watching videos and thumbing through music playlists, smartphone designers shrink and flatten speakers and sometimes even cover them in plastic, Kyriakakis explains. By themselves, small, compressed speakers damp down low frequencies, causing Darth Vader to sound like Tiny Tim. So

BOTTLENECK #2

Sharp background noise

Your phone can filter out background noise, such as traffic. But when someone honks, noise-canceling software may not be nimble enough to silence the interruption.

Solution: Specialized headsets with noise-canceling technology can silence everything but what you want to hear.

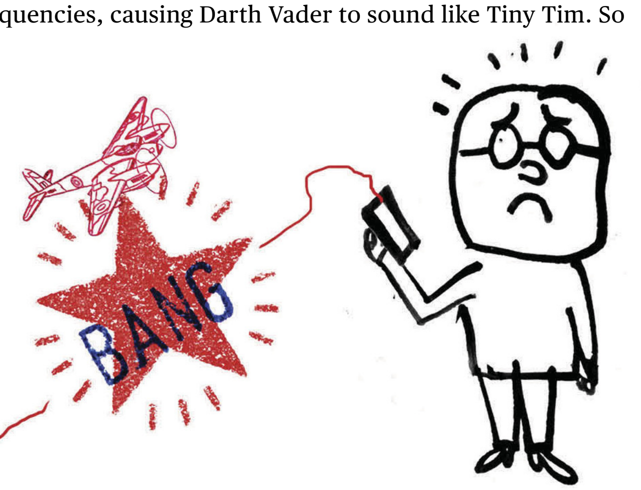

smartphones use software to lessen such distortions, making voices sound more realistic.

Your smartphone's puny microphone is similarly problematic. And the farther it is from your mouth, the more unwanted noise it picks up. Many high-end smartphones address this problem by using multiple microphones, typically three. With one microphone situated as close as possible to the user's lips and the additional ones set farther away—at the opposite end, for example—a smartphone can compare the different incoming signals to better filter out background sounds.

But noise-cancellation algorithms aren't a sure-fire fix, because they can take a few seconds to recognize a noise. So while they're quite good at removing consistent sounds such as the thrum of a leaf blower or the hiss inside a passenger jet, they do a poor job of eliminating sudden or irregular disruptions, like a baby crying. Voice echoes are especially difficult to weed out because the algorithms must also preserve speech, says Jari Sjoberg, an audio expert at Microsoft and Nokia. Too much noise suppression removes much of the natural acoustic variation in human speech, making it sound robotic.

BOTTLENECK #3

Phone-to-tower connection

Distance from a cellular tower can degrade the quality of a voice call or result in the conversation being dropped altogether.

Solution: Cellular providers offer signal boosters that can improve call quality if you're not near a tower.

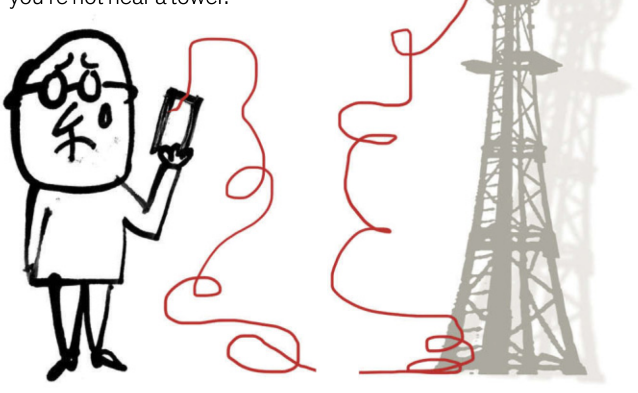

And you can pretty much forget about getting a clear voice signal through a Bluetooth ear clip or the speakerphone in your car. These setups put the microphone closer to the speaker than to your mouth, causing callers' voices or reflections within the car to echo back at them.

Yet even if your cellphone distills crisp, noise-free speech, there's no guarantee it will arrive at the listener intact. The next threat comes when the phone transmits the call to a base station. Modeled after standard wire-line phones, most mobile phones today digitize audio frequencies from 300 to 3,400 hertz. But unlike landlines, which provide each caller with a dedicated, full-capacity channel, cellphones must share a limited amount of wireless spectrum. So they compress the voice data to let more users connect.

Standard compression rates vary from 12.2 kb/s to 4.75 kb/s, depending on the volume of voice traffic and the strength of the wireless signal. Calls compressed to speeds as low as 7.95 kb/s can still sound almost as good as a landline connection. But beyond that, "you start to hear compression artifacts," including missing syllables and distortions such as ringing or warbling, says Jerry Gibson, a wireless-engineering expert at the University of California, Santa Barbara.

If you're making a local call to a mobile user on your own carrier network, count yourself lucky. The compressed data will likely travel to the receiving cellphone without further manipulation, and so voice quality may not be half bad. But say you're talking to someone across the country or on a different carrier. In those cases, your local network will typically direct the call into the backbone telephone network, which was designed to carry landline traffic at 64 kb/s. So transcoding equipment at the exchange point must convert the mobile voice data to the higher wire-line rate.

A standard landline phone can decode that signal without losing more information. But if your call is sent to another cellphone, voice quality will take another nosedive when the base station serving the phone recompresses the data to fit into a cramped wireless channel.

Other parts of the telephone network may require additional conversions, which can further degrade quality. For instance, international carriers sometimes compress voice data to stuff more calls through subsea cables rather than pay for additional capacity. The extra compression cycles "can explain the very poor international voice call quality that we sometimes experience," says Jan Derksen, head of technical marketing at Ericsson, in Stockholm.

A couple of technologies already exist that can circumvent these choke points—or at least lessen the damage. Many new smartphones have one or both built in. But for you to use these enhancements to their full potential, carriers will have to make major network upgrades, which will take time and money.

One solution is HD voice. This transmission standard more than doubles the range of audio frequencies that represent speech, letting phone systems collect and relay signals from 50 to 7,000 Hz. At their healthiest, normal human ears can perceive frequencies as low as 20 Hz and as high as 20,000 Hz. But early telephone networks had limited bandwidth, and engineers decided that frequencies between 300 and 3,400 Hz would be adequate for conveying intelligible speech.

By the 1980s, however, acoustic researchers had demonstrated that people need to hear a wider range of wavelengths to fully understand speech. Frequencies above 3,400 Hz, for example, help listeners distinguish between some consonants. "If I said 'fox' and 'socks' in isolation, you couldn't tell them apart over standard telephone bandwidth," says Mark Clements, a signal-processing expert at Georgia Tech. Like-wise, names like Jeff and Jess sound the same on the phone. Believe me, I speak from experience.

After the International Telecommunications Union standardized HD voice more than a quarter century ago, radio broadcasters were among the first adopters. Today, they still use the technology to transmit remote interviews from a sports stadium or another studio over high-speed digital telephone lines. In my experience, the improvement in sound quality is striking. And others seem to agree. In laboratory tests at Nokia, for instance, users rated HD voice calls nearly a full point higher than standard voice calls on a five-point scale. "It really is better," says Machina Research's Green.

In September 2009, in an unlikely test ground sandwiched between Romania and Ukraine—the country of Moldova—the wireless carrier Orange became the first company to launch HD voice on a cellular network. The technology has since spread around the world. According to the latest count, 329 smartphone models support the standard, and 109 mobile operators offer service in 73 countries. "The widest deployments are in Asia and Europe, where it's really taken off," says Alan Hadden, president of the Global mobile Suppliers Association (GSA), which tracks these numbers.

If you're wondering why you haven't noticed these changes, that's because HD voice equipment still typically defaults to standard "narrowband" service. Even if your phone is HD compatible, for instance, you won't hear an improvement unless the person you're talking with is also on an HD phone and all of the networking equipment in between supports the technology. But that situation almost never happens because the circuit-switched backbone still uses standard voice technology. Until it's upgraded, you won't be able to roam among different HD networks or place an HD call to someone on a different carrier.

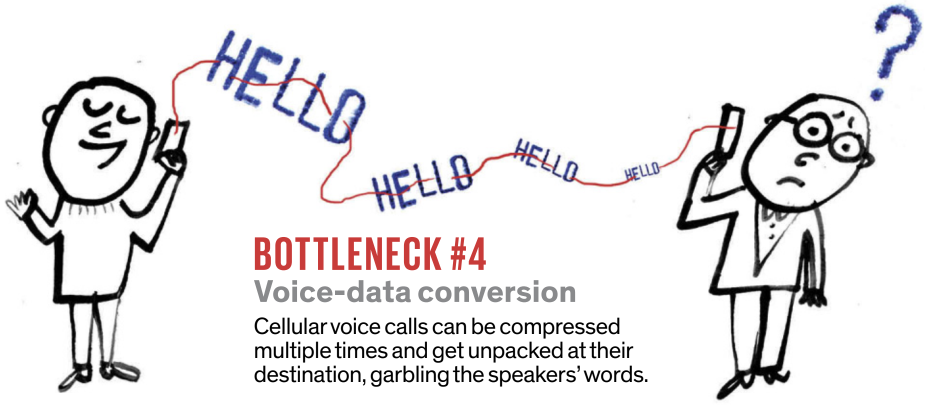

BOTTLENECK #4
Voice-data conversion
Cellular voice calls can be compressed multiple times and get unpacked at their destination, garbling the speakers' words.

Solution: VoLTE technology can ensure your call has to be compressed only once, minimizing voice degradation.

A promising fix for this problem is the second technology capable of boosting cellular voice quality: Voice over LTE (VoLTE). Today, the majority of mobile calls, including HD traffic, are carried on a 2G or 3G network despite the widespread deployment of LTE technology. This newest cellular standard is the first generation to ferry data via Internet-style packet switching.

VoLTE lets mobile carriers deliver voice traffic just like regular data—a characteristic it shares with other VoIP services, such as Microsoft's Skype and Verizon's FiOS. By compressing a voice call into a series of standardized packets that can travel between carriers and across national borders over an IP backbone, VoLTE eliminates the need to convert the data into different formats for different parts of the system. So no information is lost. "The same bits that leave one phone will enter the other phone without any changes," Ericsson's Derksen explains. And because LTE is designed to deliver any data packet regardless of its content, VoLTE networks can support HD voice right out of the box.

But most LTE carriers don't yet offer VoLTE. "LTE was originally designed without a native voice service," says Peter Carson, senior director of technical marketing at Qualcomm. That's because traditional packet switching doesn't ensure good voice quality. By treating all packets on a first-come, first-served basis, LTE carriers can't guarantee that voice packets will arrive at their destination in a timely manner. Packets can be lost or delayed, for example, when a network is busy, creating unintentional silences that can garble speech or cause callers to talk over one another. This unreliability helps explain why VoIP calls can sometimes sound great one minute and poor the next.

In general, the quality of VoIP services has gotten better in recent years as broadband speeds

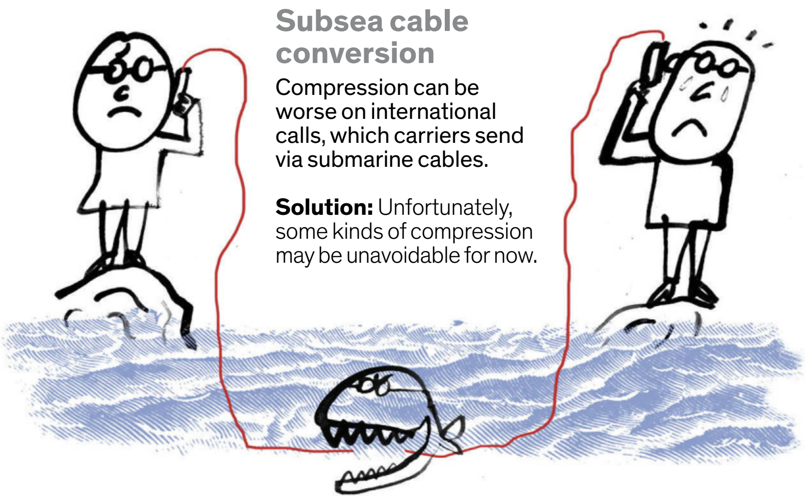

BOTTLENECK #5
Subsea cable conversion
Compression can be worse on international calls, which carriers send via submarine cables.

Solution: Unfortunately, some kinds of compression may be unavoidable for now.

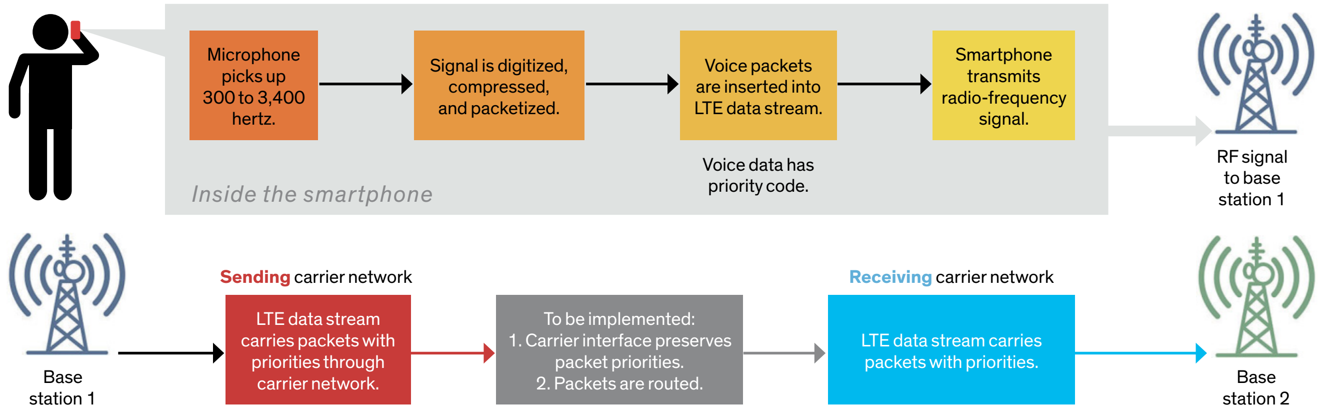

VOICE OVER LTE (VoLTE) lets cellular carriers send voice calls the same way they send other data, like bits of a streaming video. As VoIP services do, VoLTE compresses and digitizes your voice, sending it as packets of data over the LTE network. The packets are sent with priority codes to ensure they arrive in order and don't garble your message. LTE is now widely available, but carriers still have work to do to ensure that all networks respect one another's priority codes.

have increased. And some users have configured their private home or business networks to prioritize voice traffic over other data. But when VoIP packets enter the Internet or a cellular network, they're handled as "best-effort traffic" along with other data, Carson says. So VoIP providers can't promise that sound quality will always be adequate.

Mobile operators want to control quality so that their customers will keep paying premium prices for cellular voice service. VoLTE lets LTE carriers manage voice traffic using a software platform called the IP Multimedia Subsystem, or IMS. This control layer essentially acts as a traffic cop, opening fast lanes for voice data and other time-sensitive streams, such as video calls and online gaming.

The IMS prioritizes some types of traffic over others by assigning each data connection a single-digit code, called the QoS (quality-of-service) class identifier, or QCI. This number, which is stored in a routing table, describes the transmission requirements for the link, including the maximum packet latency, acceptable number of losses, and whether the network will guarantee a given bit rate. Voice calls, for example, get a QCI of 1, which ensures that 99.99 percent of packets will arrive at their destination within 100 milliseconds, even during peak use times. By comparison, typical Internet traffic, such as e-mail and browsing data, receives the lowest-priority QCIs: 8 and 9. Each router along the way can now usher packets into different transmission queues depending on their QCIs, preventing VoLTE packets, for example, from getting stuck in a Netflix traffic jam.

For years, VoLTE deployments lagged far behind HD voice over 2G and 3G networks. But now, carriers are finally committing to the upgrade, says GSA's Hadden. More than 100 smartphone models include the technology, and although the service is currently available from only 10 carriers, in Hong Kong, Japan, Singapore, South Korea, and the United States, others are quickly getting on board. As of July, 56 more operators had announced plans to test or commercialize VoLTE in 35 countries around the globe, including Algeria, Germany, Kazakhstan, and Russia.

Meanwhile, carriers are expanding their IP infrastructures, including backbone networks and local broadband links, which will let VoLTE packets flow seamlessly between mobile handsets and other IP phones, including computers and landlines. Voice quality should continue to improve as more networks support priority protocols and callers move onto the same packet-based system.

Eventually, if the carriers get their way, the old circuit-switched networks will go dark. But that transition will take time as companies invest in new equipment and regulators work to ensure that some customers aren't left with shoddy voice service—or none at all.

There was a time when telephone operators took pride in their voice networks. When Sprint built the first nationwide fiber footprint in the late 1980s, for example, it ran commercials boasting that the system was "so incredibly quiet that you could actually hear a pin drop."

But since the arrival of the smartphone era, carriers have been strangely—and frustratingly—mum about sound quality. Could the tide be turning at last? Even the second-, third-, and fourth-largest U.S. carriers—AT&T, Sprint, and T-Mobile—which long shied away from discussing voice quality with the public, announced plans for VoLTE rollouts this year. Verizon, the top U.S. carrier, which plans to start deploying the technology before year's end, calls it "the next evolution in wireless calling."

If that's really true, it's reason for me and other voice customers to be optimistic. But I've experienced too many lousy connections to take these promises at face value: I'll believe it when I hear it. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/voicequality1014

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

An Even Better Atomic Clock

By Jérôme Lodewyck

Optical-lattice clocks could redefine the way we mark time

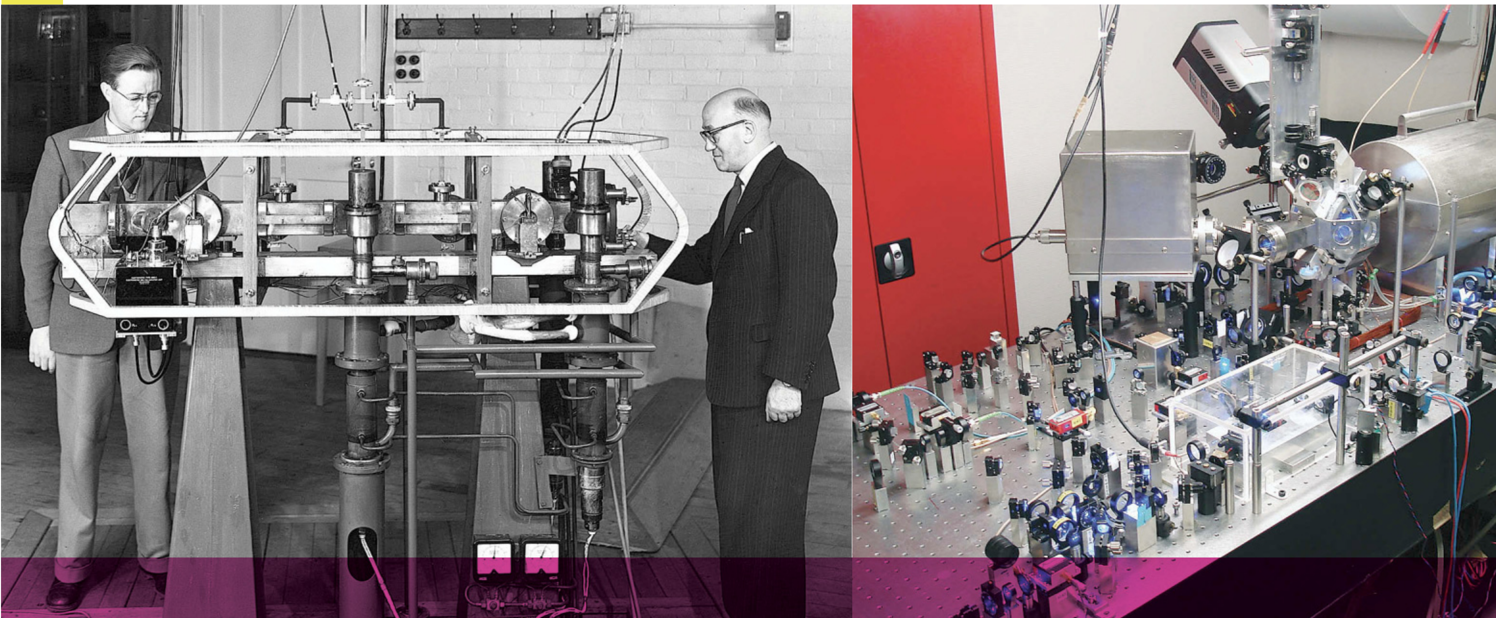

TIME TRANSFORMED: In the photo at left, John V.L. Parry [left] and Louis Essen stand with the first cesium atomic clock, in 1956, at the United Kingdom's National Physical Laboratory. The instrument paved the way for a redefinition of the second in 1967. At right is one of two modern optical-lattice clocks that have been built at the Paris Observatory.

In 1967, time underwent a dramatic shift. That was the year the key increment of time—the second—went from being defined as a tiny fraction of a year to something much more stable and fundamental: the time it takes for radiation absorbed and emitted by a cesium atom to undergo a certain number of cycles.

This change, which was officially adopted in the International System of Units, was driven by a technological leap. From the 1910s until the mid-1950s, the most precise way of keeping time was to synchronize the best quartz clocks to Earth's motion around the sun. This was done by using telescopes and other instruments to periodically measure the movement of stars across the sky. But in 1955, the accuracy of this method was easily bested by the first cesium atomic clock, which made its debut at the United Kingdom's National Physical Laboratory, on the outskirts of London.

Cesium clocks, which are essentially very precise oscillators, use microwave radiation to excite electrons and get a fix on a frequency that's intrinsic to the cesium atom. When the technology first emerged, researchers could finally resolve a known imperfection in their previous time standard: the slight, irregular speedups and slowdowns in Earth's rotation. Now, cesium clocks are so ubiquitous that we tend to forget how integral they are to modern life: We wouldn't have the Global Positioning System without them. They also help synchronize Internet and cellphone communications, tie together telescope arrays, and test fundamental physics. Through our cellphones, or via low-frequency radio synchronization, cesium time standards trickle down to many of the clocks we use daily.

The accuracy of the cesium clock has improved greatly since 1955, increasing by a factor of 10 or so every decade. Nowadays, timekeeping based on cesium clocks accrues errors at a rate of just 0.02 nanosecond per day. If we had started such a clock when Earth began, about 4.5 billion years ago, it would be off by only about 30 seconds today.

But we can do better. A new generation of atomic clocks that use laser light instead of microwave radiation can divide time more finely. About six years ago, researchers completed single-ion versions of these optical clocks, made with an ion of either aluminum or mercury. These surpassed the accuracy of cesium clocks by a full order of magnitude.

Now, a new offshoot of this technology, the optical-lattice clock (OLC), has taken the lead. Unlike single-ion clocks, which yield one measurement of frequency at a time, OLCs can simultaneously measure thousands of atoms held in place by a powerful standing laser beam, driving down statistical uncertainty. In the past year, these clocks have managed to surpass the best single-ion optical clocks in both accuracy and stability. With further development, they will lose no more than a second over 13.8 billion years—the present-day age of the universe.

So why should you care about clocks of such mind-boggling accuracy? They are already making an impact. Some scientists are using optical-lattice clocks as tools to test fundamental physics. And others are looking at the possibility of using them to better measure differences in how fast time elapses at various points on Earth—a result of gravity's distortion

LEFT: NATIONAL PHYSICAL LABORATORY; RIGHT: JÉRÔME LODEWYCK

of the passage of time as described by Einstein's theory of general relativity. The power to measure such minuscule perturbations may seem hopelessly esoteric. But it could have important real-world applications. We could, for example, improve our ability to forecast volcanic eruptions and earthquakes and more reliably detect oil and gas underground. And one day, in the not-too-distant future, OLCs could enable yet another shift in the way we define time.

According to the rules of quantum mechanics, the energy of an electron bound to an atom is quantized. This means that an electron can occupy only one of a discrete number of orbiting zones, or orbitals, around an atom's nucleus, although it can jump from one orbital to another by absorbing or emitting energy in the form of electromagnetic radiation. Because energy is conserved, this absorption or emission will happen only if the energy corresponding to the frequency of this radiation matches the energy difference between the two orbitals involved in the transition.

Atomic clocks work by exploiting this behavior. Atoms–of cesium, for example–are manipulated so that their electrons all occupy the lowest-energy orbital. The atoms are then hit with a specific frequency of electromagnetic radiation, which can cause an electron to jump up to a higher-energy orbital–the excited "clock state." The likelihood of this transition depends on the frequency of the radiation that's directed at the atom: The closer it is to the actual frequency of the clock transition, the higher the probability that the transition will occur.

To probe how often it happens, scientists use a second source of radiation to excite electrons that remain in the lowest-energy state into a short-lived, higher-energy state. These electrons release photons each time they relax back down from this transient state, and the resulting radiation can be picked up with a photosensor, such as a camera or a photomultiplier tube.

If few photons are detected, it means that electrons are largely making the clock transition, and the incoming frequency is a good match. If many photons are being released, it means that most electrons were not excited by the clock signal. A servo-driven feedback loop is used to tune the radiation source so its frequency is always close to the atomic transition.

Converting this frequency reference into a clock that ticks off the seconds requires additional steps. Generally, the frequency measured in an atomic clock is used to calibrate other frequency sources, such as hydrogen masers and quartz clocks. A "counter," made using basic analog circuitry, can be connected to a hydrogen maser to convert its electromagnetic signal into a clock that can count off ticks to mark the time.

The most common atomic clocks today use atoms of cesium-133, which has an electron transition that lies in the microwave range of the electromagnetic spectrum. If the atom is held at absolute zero and is unperturbed (more on that in a moment), this transition will occur at a frequency of exactly 9,192,631,770 hertz. And indeed, this is how we define the second in the International System of Units–it is the time it takes for 9,192,631,770 cycles of 9,192,631,770-Hz radiation to occur.

In actuality, cesium-133 isn't so perfect a pendulum. Atoms experience various forms of perturbation because of their imperfect environment. For example, an atom's motion through space, which in the laboratory can easily be as fast as 100 meters per second, can shift the frequency of an electron transition by means of the Doppler effect. This is the same phenomenon that affects the pitch of ambulance sirens and other sounds as the source of the sound moves relative to the listener. Interactions with the electron clouds of other atoms can also alter the energies of electron states, as can stray external electromagnetic fields.

Perturbations decrease a clock's accuracy: how much the atom's average frequency is shifted from its natural unperturbed value. A number of these offsets can be accounted for, and changes in clock design have helped minimize these shifts. Indeed, one of the most dramatic such improvements occurred in the early 1990s, when physicists developed the fountain clock. This clock uses a laser to launch cooled cesium atoms upward, as if they were water droplets from a fountain, so that the Doppler shift caused by the upward motion cancels out nearly all of the shift that occurs as they fall.

But nowadays cesium clocks can't be improved much more. Tiny gains are increasingly difficult to achieve, and any gains we try to make now will take a long time. That's because cesium clocks are pushing the limit of the other key metric we use to evaluate clocks: the stability of their frequency.

Frequency stability characterizes how clock frequency fluctuates over time. The bigger the frequency instability, the greater the frequency noise, so the clock frequency will sometimes be a bit higher and sometimes a bit lower than its average value.

Careful engineering can minimize most sources of frequency noise. But there's a fundamental source of instability that is very difficult to overcome, because it comes from the probabilistic nature of

MAGIC TRANSITION

In an optical-lattice clock, an electron [yellow dot] can absorb electromagnetic radiation to jump from a lower- to a higher-energy orbital around a clock atom's nucleus [center]. Light used to trap the atom can shift the natural energy of each orbital [dotted lines] down in energy [solid lines]. This would ordinarily change the energy associated with the jump. But for a "magic wavelength" of trapping light, the energy shift of each orbital will be identical, and the frequency of the transition will remain the same.

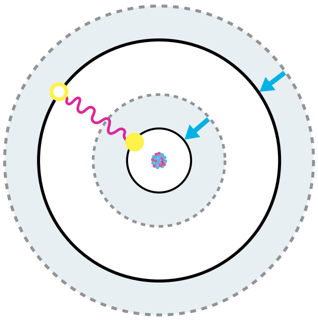

FINDING THE FREQUENCY

The fundamental frequency of a clock-atom transition (f_0) is often offset by some amount (Δf). To better detect fluctuations around the resulting transition frequency ($f_0+\Delta f$), a clock laser's frequency is set slightly off (f_{lock}). This lowers the probability that an electron will make a transition and results in unneeded corrections to the clock laser frequency that yield noise called quantum projection noise. QPN decreases with number of atoms (*N*).

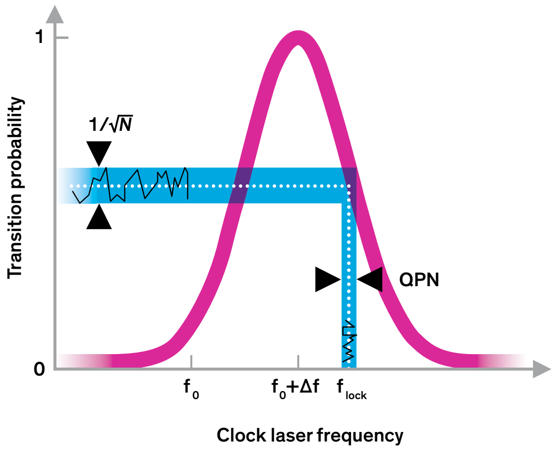

quantum mechanics. To understand it, let's go back to the basic operating principle of an atomic clock.

We typically excite the electrons in an atomic clock with radiation whose frequency doesn't quite match the transition frequency. That's because the probability that an electron will be excited follows a bell-curve-like distribution. On the sides of the bell curve, it's easier to see whether a small change in frequency has occurred because it produces a more detectable effect, more dramatically increasing or decreasing the likelihood that an electron is excited [see illustration, "Finding the Frequency"]. Because of this, during the ordinary operation of an atomic clock, the clock radiation is set so that it has only a 50 percent probability of getting any given atom to make the clock transition. But even if the clock radiation frequency is set precisely at that point, an electron will be in either an excited or an unexcited state after it's measured. The servo loop will then wrongly assume that the clock radiation frequency is either too high or too low and will introduce an undue frequency correction.

These miscorrections yield additional noise in the clock that we call quantum projection noise (QPN), and they are the main source of frequency instability in the best cesium clocks. Like many random sources of noise, the average level of QPN decreases with time. The longer you observe the clock, the more often the random upward shifts in frequency cancel out the downward shifts, and the noise eventually becomes negligible.

The catch is that this takes a long time in cesium: It takes about a day for the stability of the best cesium clocks to reach 2 parts in 10^{16} –their steady-state accuracy level. (Metrologists commonly measure quantities such as accuracy and stability in fractional units. For a cesium clock with a frequency of 9.2 gigahertz, an accuracy of 2×10^{-16} translates to an uncertainty of 1.8 microhertz in the frequency.)

You could run a series of experiments to make cesium clocks more accurate. But each measurement would have to consist of a lot of data taken over a very long time in order to minimize random fluctuations from measurement to measurement. In a series of experiments designed to push the clock accuracy down to 1 part in 10^{17}, a 20-fold improvement, it could take an entire year just to make a single measurement.

Fortunately, there are other ways to minimize QPN. The noise is the same regardless of frequency, but its relative impact decreases the higher in frequency you go. And just as the average QPN decreases the longer you observe a clock, increasing the number of atoms you interrogate at the same time will boost the signal-to-noise ratio. The more you can sample in one go, the less uncertainty you'll have in the number of atoms that made the clock transition.

Moving to higher frequencies is what motivated work on the optical atomic clock. The first of these clocks was developed in the early 1980s, and nowadays they can be built from any of a number of neutral or ionized versions of elements, including mercury, strontium, calcium, ytterbium, and aluminum. What they all have in common are relatively high resonance frequencies, which lie in the optical spectrum around several hundred thousand gigahertz–10,000 times cesium's frequency. Using a higher frequency lowers the QPN, and it also lowers the relative impact of several factors that can shift the clock frequency. These include interactions with external magnetic fields coming from Earth or nearby metal (or, in Paris, the Métro lines). As an added bonus, if an optical clock is built with ions, those charged atoms can easily be trapped in an oscillating electric field that will cancel out most of their motion, effectively eliminating the Doppler effect.

But optical clocks have limitations of their own. If all other aspects of a clock are the same, the move to optical frequencies should lower the QPN to 0.01 percent of what it is in cesium. But many optical clocks are made with ions instead of neutral atoms, such as those used in cesium clocks. Because they're charged, ions are fairly easy to trap, but they also easily push on one another when placed close together, creating motion that's hard to control and causing a Doppler frequency shift. As a result, such clocks tend to use just one ion at a time and so are only about 20 times as stable and 25 times as accurate as the best cesium clocks, which can easily contain a million atoms. To get closer to the factor-of-10,000 boost in stability promised by optical clocks, we must find a way to boost the number of atoms in the optical clock, simultaneously interrogating many atoms so that the QPN averages out. And with the optical-lattice clock, researchers realized they could go quite big, measuring not just a handful of atoms but 10,000 or more at the same time.

It certainly isn't easy. To build a clock out of 10,000 atoms, you must find a way to make an atomic ensemble that is both tightly confined (to minimize the Doppler effect) and very low in density (to minimize electromagnetic interactions among the atoms). The atoms in a typical crystal move too fast and interact too strongly to work, so the best way to proceed is to produce an artificial material with a lattice of your own creation.

To build an optical-lattice clock, we start much the same way we do in many cold-atom experiments, with an ensemble of slow-moving, laser-cooled neutral atoms. We send these into a vacuum vessel containing a single laser beam that has been reflected back on itself. An interference pattern arises in the areas

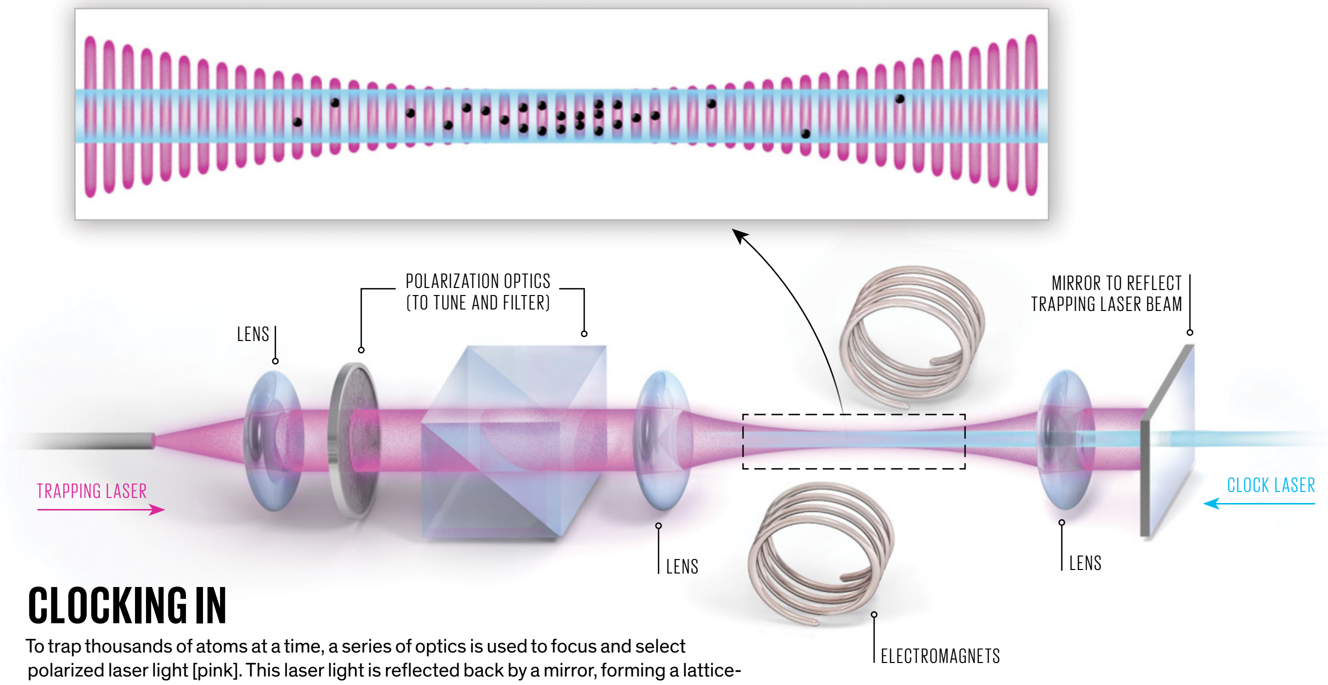

CLOCKING IN

To trap thousands of atoms at a time, a series of optics is used to focus and select polarized laser light [pink]. This laser light is reflected back by a mirror, forming a lattice-like pattern of regions of high and low intensity as the laser beam interferes with itself [top illustration]. Atoms are attracted to thousands of bright areas in this optical lattice. Three pairs of electromagnets [one pair shown] are tuned to cancel external magnetic fields. A clock laser beam [blue] is used to excite electrons in the trapped atoms.

where the beam overlaps with itself, creating an optical lattice made of thousands of small "pancakes" of light. The atoms fall into the lattice like eggs into an egg carton because of a force that draws each of them toward a spot where the light intensity is at a maximum. Once the atoms are in place, we use a separate "clock laser" to excite the atoms so that we can measure the frequency of the clock transition.

The difficulty is that the clock atoms aren't so easy to coerce into this lattice. Inexpensive lasers have outputs in the milliwatts. To create a lattice strong enough to trap and hold a neutral atom, you need several watts of light. Such a powerful laser beam, however, can shift energy levels in clock atoms, pushing their transition frequency far from their natural state. The amount of this shift will vary with the intensity of the trapping light, and that intensity is hard to control. Even with very careful calibration, this large frequency shift would render the clock much more inaccurate than even the very first cesium clocks.

Fortunately, physicist Hidetoshi Katori conceived a workaround in the early 2000s. When atoms are hit with the trapping light, the energy associated with each electron orbital decreases. Katori, then at the University of Tokyo, noted that each orbital will respond differently, with an energy shift that will depend on the wavelength of the trapping light. For a specific, "magic" wavelength, the shift of both orbitals will be identical, and so the energy difference between the two orbitals will be unchanged. This magic wavelength, where the clock frequency stays the same whether the atoms are trapped or not, is different for each element. For strontium, it's 813 nanometers, in the infrared part of the spectrum. Ytterbium's magic wavelength is 759 nm; mercury's is in the ultraviolet part of the spectrum, at 362 nm.

When Katori made his proposal, my group at the Paris Observatory's Systèmes de Référence Temps-Espace (LNE-SYRTE) department, which is responsible for maintaining France's reference time and frequency signals, had already been investigating the use of strontium for optical clocks. We set to work almost immediately to see if we could make an optical-lattice clock using strontium, competing at first with just two other groups that had long-standing experience working with cooled strontium: Katori's team in Tokyo and Jun Ye's group at JILA, in Boulder, Colo. A decade and many projects later, other groups have built lattice clocks using strontium and ytterbium. More experimental projects using mercury or magnesium, which require still higher-frequency and less-well-developed lasers, are also in the works.

ILLUSTRATION BY Emily Cooper

TIME MARCHES ON

Atomic clocks have made great strides since their start in 1955, improving in accuracy by a factor of 10 or so each decade. Cesium clocks [green], which employ microwave radiation to interrogate ensembles of cesium atoms, were the first. These were surpassed in accuracy in the 2000s by optical clocks [pink], which use laser light and often just a single ion. This year, optical-lattice clocks, which incorporate thousands of atoms [blue], became the most accurate atomic clocks. The symbol by each optical-lattice point denotes the atomic species used in the clock: strontium (Sr), mercury (Hg), and ytterbium (Yb).

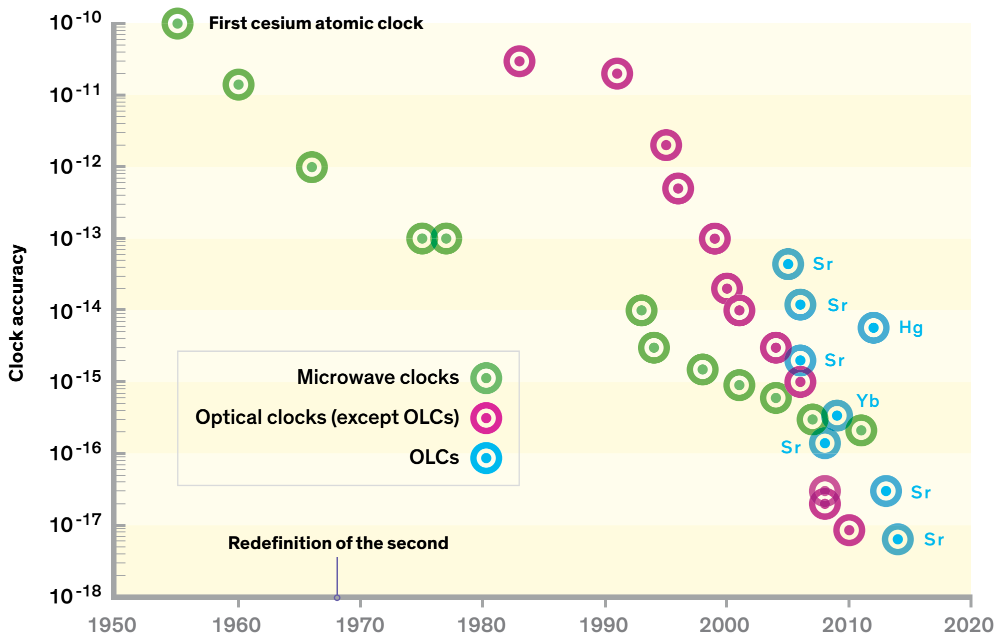

One of the key factors in making optical-lattice clocks more accurate over the past few years has been the development of clock lasers with very narrow spectra—essentially just a small spike at one particular frequency. We need these to effectively explore the region around the transition frequency of the clock, to see in fine detail how a slight shift in the clock frequency affects the transition probability.

The best way to make narrow-lined laser light is to feed it into a mirrored chamber called a Fabry-Pérot cavity. After bouncing back and forth up to a million times inside this cavity, light of any arbitrary wavelength will have interfered with itself and canceled itself out. Only laser light with a wavelength that is a unit fraction of the length of the cavity emerges.

While the cavity helps to filter out natural fluctuations in the frequency of a laser source, the technique isn't perfect. The frequency of the clock laser that emerges from the cavity can wobble around because of thermal fluctuations that cause the cavity to slightly expand or contract.

But over the past few years, researchers have found ways to help mitigate this effect. Cavities were made longer, so the relative impact of a small change in length is smaller. Vibrations were damped. The cavities were also cooled to cryogenic temperatures, to limit tiny expansions and contractions due to thermal energy.

The net result was much more stable clock lasers. Nowadays, over the few seconds it takes to prepare and probe clock atoms, a 429-terahertz clock laser might drift in frequency by just 40 millihertz or so. For a typical cavity, with a length of a few dozen centimeters, that amounts to changes in its length of no more than a few percent of the size of a proton for the several seconds it takes to prepare and probe the atoms in the optical clock.

Largely due to this effort, the stability reached within one day with cesium clocks, or within a few minutes with optical-ion clocks, can now be reached in 1 second with an optical-lattice clock, close to the QPN limit. This improved stability makes the clock itself a tool. The less time you need to gather data to measure an atomic clock's frequency with precision, the faster you can use the clock to run experiments to explore ways to make it better. Indeed, just three years after the first frequency stability improvements were demonstrated in optical-lattice clocks at the U.S. National Institute of Standards and Technology, these clocks took the lead for accuracy. The published record is now held by one of the strontium OLCs at JILA, which boasts an estimated accuracy of 6.4 parts in 10^{18}.

A clock is only so good on its own. Evaluating one clock requires another, comparable clock to serve as a reference. When OLCs were first developed a decade ago, the initial comparisons were done between strontium OLCs and cesium clocks. These measurements were enough to establish the early promise of OLCs. But to truly ascertain the accuracy of an atomic clock, it's crucial to directly compare two clocks of the same type. If they are as accurate as advertised, their frequencies should be identical.

So as soon as we had finished building one strontium optical-lattice clock in 2007, we began work on a second. We finished the second clock in 2011, and set to work making the first comparison between two optical-lattice clocks in order to directly establish their accuracy, without relying on cesium clocks.

Once a second clock is built, previously undetectable problems soon become apparent. And indeed, we soon uncovered flaws that had been overlooked. One was the influence of static electric charges that had become trapped on the windows of the vacuum chamber. We had to shine ultraviolet light on the windows to efficiently dislodge the charges.

In a paper that appeared last year in *Nature Communications*, we showed that our two strontium OLCs agree down to the 1 part in 10^{16} level, a solid confirmation that these clocks are more accurate than the best cesium clocks. Earlier this year, Katori's team at the research institution Riken, in Wako, Japan, reported an agreement of a few parts in 10^{18} in similar clocks, this time enclosed in a cryogenic environment.

Incidentally, the frequency of an optical clock is so fast that no electronic device could possibly count its ticks. These sorts of clock comparisons rely on a new technology that's still very much in development: the frequency comb. This instrument uses femtoseconds-long laser pulses to create a spectrum that consists of coherent, equally spaced teeth that span the visible and infrared spectrum. In effect, it acts like a ruler for optical frequencies.

The ability to perform comparisons between OLCs pushes us further along the road to redefining the second. Before a redefinition in the International System of Units can take place, a large number of laboratories must demonstrate not only that they can implement the new standard but also compare their measurements. Consensus is needed to establish that all the laboratories are on the same page. It is also necessary to ensure that the world can literally keep time: Coordinated Universal Time, the time by which the world's clocks are set, and the International Atomic Time it's derived from, are created by making a weighted average of a large number of microwave clocks around the world.

Cesium clocks are "networked" using the signals emitted by satellites and are compared by microwave transmission. This is good enough for microwave clocks but too unstable for distributing more-accurate optical-lattice clock signals. But soon, international comparisons of optical clocks will reach a new milestone. New fiber connections, built with dedicated phase-compensation systems that can cancel small timing shifts introduced by the lines, are now being constructed.

By the end of this year, thanks to a number of national and international projects, we expect to be able to start using such connections to make the first comparisons between optical-lattice clocks based at LNE-SYRTE in Paris and the Physikalisch-Technische Bundesanstalt, Germany's national metrology center, in Braunschweig. A link to the National Physical Laboratory, in London, which has strontium- and ytterbium-ion clocks, is also set to be completed early next year. These efforts will pave the way for an international metrology network that could enable a new standard for the second.

In the meantime, scientists have already begun using optical-lattice clocks as a tool to explore nature. One focus has been on measuring the frequency ratio between two clocks that use different types of atoms. This ratio depends on fundamental physical constants, such as the fine-structure constant, which could reveal new physics if it turns out to vary in time or from place to place.

Astronomers may also benefit from optical clocks. Atomic clocks are used as a time reference in radio astronomy, allowing astronomers to combine the light collected by telescopes separated by hundreds or thousands of kilometers to produce a virtual telescope, with an angular resolution equivalent to that of a single telescope spanning that entire distance. As optical atomic clocks mature, they could enable a similar feat for optical telescopes.

And it's not hard to imagine that optical-lattice clocks could offer new insight into the world beneath our feet. According to Einstein's theory of general relativity, a clock sitting on a denser part of Earth will tick slower relative to one situated on a part that's less dense. Although gravimeters can be used to measure gravitational force at any one point, measuring gravitational potential—which could shed light on different, deeper structures inside Earth—must be done by integrating the measurements of gravimeters at different points around Earth's surface or by measuring the orbits of satellites. Metrologists and geodesists are now teaming up to understand what optical-lattice clocks will be able to offer. It's possible that they could be used at different points around Earth to assist with oil detection, earthquake monitoring, and volcano prediction.

In the meantime, there is still work to be done to keep improving the stability and accuracy of OLCs. Recently, a large effort has been made to fight the effect of black-body radiation. This radiation is unavoidably emitted by any physical body with nonzero temperature, including the vacuum chamber that surrounds the clock atoms. When it interacts with the atoms it shifts the energy levels of the clock transition. This shift can be corrected after the fact, but a precise knowledge of the temperature and emissivity of the vacuum chamber must be acquired. It is also possible to enclose the atoms in a cryogenic environment or use an atomic species that is inherently less sensitive to black-body radiation, such as mercury, a route that our group is exploring.

Before the end of the decade, new generations of ultranarrow lasers are also likely to help push stabilities below 1 part in 10^{17} after a single second of data gathering. That will make it practical for us to achieve an accuracy below 10^{-18}—more than 100 times the precision of cesium clocks. As OLCs become more accurate, the scope of applications will continue to expand.

Even if OLCs are wildly successful, we won't abandon the cesium clock, which will remain more compact and less expensive to build. And in the future, OLCs may be supplanted by clocks of even higher frequencies that rely on energy transitions inside the atom's nucleus instead of among the electrons in orbit around it. These nuclear transitions are mostly out of reach of current laser technology, although researchers are starting to explore them.

But before long we will see yet another time standard that could significantly influence the way we relate to our universe. Just as surely as time keeps on ticking, improvements in our ability to measure it will go on. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/clock1014

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

The People's Cloud

Peer-to-peer cloud computing could free us from the tyranny of data centers

NOT LONG AGO, any start-up hoping to create the next big thing on the Internet had to invest sizable amounts of money in computing hardware, network connectivity, real estate to house the equipment, and technical personnel to keep everything working 24/7. The inevitable delays in getting all this funded, designed, and set up could easily erase any competitive edge the company might have had at the outset. • Today, the same start-up could have its product up and running in the cloud in a matter of days, if not hours, with zero up-front investment in servers and similar gear. And the company wouldn't have to pay for any more computing oomph than it needs at any given time, because most cloud-service providers allot computing resources dynamically according to actual demand.

By **OZALP BABAOGLU** & **MORENO MARZOLLA**

Illustrations by **ROB WILSON**

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

Scattering the Clouds

With the right software, geographically distributed hardware can provide a unified cloud-computing resource

With the computing infrastructure out of sight and out of mind, a start-up can concentrate its attention on launching and improving its product. This lowers the barriers to entry, letting anyone with an Internet connection and a credit card tap the same world-class computing resources as a major corporation. Many of the most popular and successful Internet services today, including Netflix, Instagram, Vine, Foursquare, and Dropbox, make use of commercial clouds.

These clouds might seem a bit ethereal to their end users, but they in fact require some very down-to-earth facilities. Their stadium-size data centers are immensely costly to construct, and, not surprisingly, most are run by giant corporations like Amazon, Google, and Microsoft. Each offers a variety of service models, depending on exactly how the customer interacts with its cloud-computing environment.

The lowest-level model is known as infrastructure as a service, or IaaS, which outfits each customer with one or more virtual machines running on the cloud provider's physical equipment. One actual computer might, for example, simulate five different virtual computers, each leased to a different customer. In addition to leasing such virtual machines, an IaaS provider may include a choice of operating systems to run on them. Notable examples of such IaaS clouds include Google's Compute Engine and Amazon's Elastic Compute Cloud.

At a higher level of abstraction are platform-as-a-service, or PaaS, clouds. These include an environment for developing the online applications that are to run on the provider's equipment. Customers don't have to manage virtual machines. They just create their applications using various software libraries, application-programming interfaces, and other software tools such as databases and middleware, and then one or more virtual computers are spun up automatically as needed to run all of this. Examples of PaaS clouds are Amazon's Elastic Beanstalk, Google's AppEngine, Microsoft's Azure, and SalesForce's Force.com.

At a still-higher level are software-as-a-service, or SaaS, clouds. Their customers know nothing of the underlying infrastructure or computing platform: They simply use some Web-based application or suite of applications to handle the task at hand. This is probably the model of cloud computing that most people are familiar with. It includes services like Apple iWork, Gmail, Google Docs, and Microsoft Office 365.

But is this the only way cloud computing can work? At the University of Bologna, in Italy, we've been investigating a very different strategy to do cloud computing without those giant centralized facilities at all–using peer-to-peer technologies of the kind sometimes associated with shady file-sharing operations. Their use here, though, could help democratize cloud computing. Our prototype software is still at a very early stage, but its development and similar successes by researchers elsewhere show real promise.

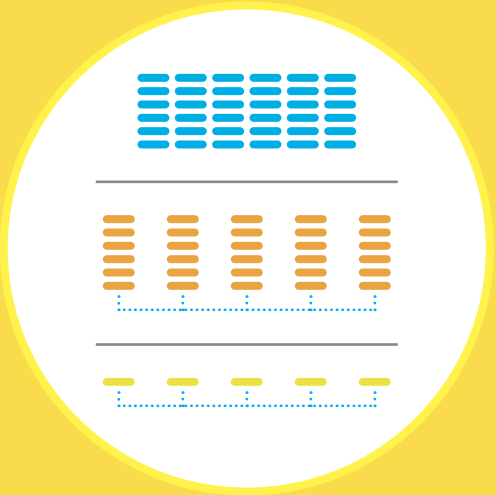

SOME CLOUD-COMPUTING PROVIDERS put all their hardware eggs in one data-center basket [blue]. Others employ multiple data centers networked together [orange]. The logical extension is a peer-to-peer cloud made of individual computers [yellow].

PLACING THE PHYSICAL INFRASTRUCTURE for a cloud-computing operation where it's usually found, in a single massive data center, has definite advantages. Construction, equipment procurement, installation, and maintenance are all simplified, and economies of scale reduce costs. On the other hand, a single large data center consumes an enormous amount of electrical power, often comparable to what you'd need to

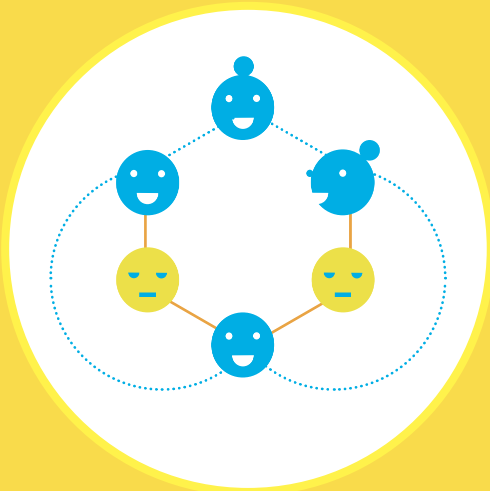

THE COMPUTERS in some P2P networks resemble gossiping office workers. Gossip-based protocols allow information to flow reliably, even if some computers leave the system and break previously established links [orange lines].

GOSSIP-BASED PROTOCOLS are used to maintain an unstructured peer-to-peer network of individual computers, some of which do the work of one customer [orange] while other combinations serve different customers [blue and yellow].

run a small town, and dissipating the waste heat it generates is usually a big headache.

Perhaps the most serious shortcoming, though, is that a centralized cloud-computing facility can end up being a single point of failure, no matter how cleverly it is designed. Redundant power supplies, backup power generators, and replicated network connections help, but they can't protect absolutely against catastrophic events such as fires, hurricanes, earthquakes, and floods.

Another drawback of centralized clouds is that their geographic location, which may be best for the owners, may not be best for the customers. This is the case, for example, when governments place restrictions on sensitive data crossing national borders. A data center located in one country may then be off-limits to customers in some other countries.

Cloud-service providers have increasingly addressed these concerns by using not just one but several far-flung data centers connected through fast private networks. Doing so not only protects against local catastrophes, it also provides customers with more options for locating their data.

What would happen if you took this trend in geographically distributing cloud infrastructure to its logical conclusion? You'd end up with clouds made up of millions of individual computers distributed across the globe and connected through the Internet. We would call this a peer-to-peer (P2P) cloud because it shares many of the characteristics of various P2P systems developed for file sharing, content distribution, and the payment networks of virtual cryptocurrency schemes such as Bitcoin.

In principle, a P2P cloud could be built using the ordinary computing, storage, and communication equipment found now in people's homes, with essentially zero initial investment. Broadband modems, routers, set-top boxes, game consoles, and laptop and desktop PCs could all contribute. The challenge is to turn this motley collection into a coherent and usable cloud infrastructure and offer its services to customers. You also have to ensure that the salient features of clouds–on-demand resource provisioning and the metering of service–are maintained.

This would surely be tough to do, but think of the advantages. First, there would be no single entity that owns or controls it. As with most other P2P applications, a P2P cloud could be created and operated as a grassroots effort, without requiring the permission or consent of any authority. People would choose to participate in a P2P cloud by installing the appropriate client software on their local machines, and the value of the resulting cloud infrastructure would be commensurate with the number of individuals who are contributing to it.

A second advantage comes from the fact that a P2P cloud's components are small, individually consume little power, and well distributed. This drastically reduces concerns about local catastrophes. It also removes the usual worries about heat dissipation. Although such P2P clouds couldn't provide the quality-of-service guarantees of a Google or an Amazon, for many applications that wouldn't much matter.

THE IDEA OF CREATING a huge computing resource from a large number of loosely coupled machines is not new. This has long been done, for example, with volunteer com-

puting, where people execute applications on their personal computers on behalf of others. Volunteer-computing systems usually require you to install certain software, which then runs when your computer has no higher-priority tasks to do. That application then uses your spare computing cycles to fetch and process data from some central server and upload the results back to the same server when it's done.

This strategy works well for many scientific problems, for which a central controller can farm out pieces of the desired computation to workers that operate independently and in parallel. If one fails to return a result within some reasonable period, no problem: The same task is simply handed out to some other volunteer worker.

The Berkeley Open Infrastructure for Network Computing (BOINC) is a popular volunteer-computing system that can load and run different client programs. Examples of projects found on the BOINC platform include SETI@home (to analyze radio signals from space in the search for extraterrestrial transmissions), Rosetta@home (to calculate how proteins fold), and Einstein@home (to detect gravitational waves).

Another type of volunteer computing is known as a desktop grid. In conventional grid-computing projects, multiple high-performance computers in different locations are harnessed to work together on a single problem. Desktop grids allow people to contribute the processing power of their personal computers to such efforts. BOINC supports desktop grids, as does EDGeS (Enabling Desktop Grids for e-Science), a project of several European institutions that is based on BOINC, and also XtremWeb, a project of the French National Center for Scientific Research, the French Institute for Research in Computer Science and Automation (INRIA), and the University of Paris XI.

THE SUCCESS OF THE MANY volunteer-computing projects demonstrates the extreme scale that a P2P cloud could in principle attain, both in terms of the number of different computers involved and their geographic distribution. Using such a collection would, of course, mean that equipment failures will be common. And besides, the people who contribute their computers to these clouds could turn them on and off at any time, something the people who run P2P networks refer to as "churn."

So the first task for any P2P cloud is to keep track of all functioning and online devices enrolled in the system and to dynamically partition these resources among customers. And you have to do that in a completely decentralized manner and despite churn.

To deal with such challenges, many P2P systems make use of what are called gossip-based protocols. Gossiping, in this context, is when the computers linked together in a large, unstructured network exchange information with only a small number of neighbors. Gossip-based protocols have been extensively studied and have been used to model, for example, the spread of malware in a computer network, the diffusion of information in a social network, even the synchronization of light pulses in a swarm of fireflies. Gossip-based protocols are appealing for P2P clouds because they are simple to implement and yet enable complex global computations to take place efficiently even in the face of churn.

So when we built our prototype system at the University of Bologna, which we call the Peer-to-Peer Cloud System (P2PCS), we included several decentralized gossip-based protocols. They are used for figuring out what equipment is up and connected, monitoring the overall state of the cloud,

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

partitioning the resources available into multiple subclouds, dynamically allocating resources, and supporting complex queries over the set of connected computers (for example, to identify the most reliable ones). Creating those capabilities was an important first step. But there are still many other requirements for a practical system, only some of which we have attempted to tackle.

If all the equipment is owned by a single organization, building a P2P cloud with it should be straightforward, even if the bits and pieces are located in different people's homes, as might be the case with broadband modems or routers operated by an Internet service provider or set-top boxes operated by a cable-television company. The computing devices will be all pretty similar, if not identical, making it easier to configure them into a single computing environment. And because the equipment's one owner presumably installs the P2P-cloud software, you can be reasonably confident that the data and computations will be handled properly and according to the organization's security policies.

This is not true, however, if the P2P cloud is made up of a diverse collection of different people's computers or game consoles or whatever. The people using such a cloud must trust that none of the many strangers operating it will do something malicious. And the providers of equipment must trust that the users won't hog computer time.

These are formidable problems, which so far do not have general solutions. If you just want to store data in a P2P cloud, though, things get easier: The system merely has to break up the data, encrypt it, and store it in many places.

Unfortunately, there is as yet no efficient way to make every computation running on untrusted hardware tamperproof. For some specific problems (such as mining bitcoins), verifying the results is significantly faster than computing them, which allows the client to check and discard faked results. For those problems that do not have an efficient verification procedure, the best way to detect tampering is to compare results for the same calculation coming from independent machines.

Another issue, common to all P2P systems, is that there must be appropriate incentives to get enough people to cooperate and to discourage free riding. Otherwise, the system is bound to degenerate completely. Coming up with incentives would be easy enough for a company that uses its own devices to create a cloud. That company might have a monetary incentive for creating such a cloud, and the people housing the equipment might have an incentive to keep connected to it because they get better service that way.

Volunteer-computing systems don't enjoy the luxury of having such incentives in place. But they typically have such laudable objectives that getting people to contribute their free CPU cycles is not a problem. Who would not want to help make history when SETI@home, which has been around since 1999, detects the first extraterrestrial radio transmission? For volunteer P2P systems of other kinds, though, the incentives have to be carefully worked out.

DEVELOPMENTS ARE ADMITTEDLY at an early stage, but several research projects and a few commercial systems that have hit the market suggest that P2P clouds can indeed be built and used productively, at least for certain purposes.

Our work on the P2PCS, for example, demonstrated that it is possible to use gossip-based protocols to handle the dynamic allocation of resources and the basic monitoring of the system. Other researchers—at the University of Messina, in Italy (Cloud@Home), at INRIA (Clouds@Home), and associated with the European Union's Nanodatacenters project—have been exploring similar concepts.

The Nanodatacenters project is particularly interesting. The researchers involved worked out how to form a managed P2P network from a far-flung constellation of special home gateways controlled by Internet service providers. Because these "nanocenters" are near end users, the network can deliver data much faster than a few large data centers could.

Some commercial distributed-storage solutions are also based on P2P computing principles. An early version of Wuala's cloud backup, for example, allowed users to trade space on their hard disks. Sher.ly offers a similar service but is oriented toward the business sector: It allows companies to use their own machines and infrastructure to create a secure, always-on private cloud to share files internally. There are also a number of open-source P2P systems for distributed file storage (such as OceanStore, developed at the University of California, Berkeley) or computations (such as OurGrid, developed mostly at the Federal University of Campina Grande, in Brazil).

These pioneering experiments are still few and far between compared with traditional cloud environments. But if they succeed, and if researchers can find ways to deal with the hurdles we've described here, you could easily find yourself making use of a P2P cloud in your daily routine. You might not even know you're doing it. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/cloud1014

The PHOEBUS CARTEL engineered a SHORTER-LIVED LIGHTBULB and gave birth to PLANNED OBSOLESCENCE

THE GREAT LIGHTBULB CONSPIRACY

ON 23 DECEMBER 1924, a group of leading international businessmen gathered in Geneva for a meeting that would alter the world for decades to come. Present were top representatives from all the major lightbulb manufacturers, including Germany's Osram, the Netherlands' Philips, France's Compagnie des Lampes, and the United States' General Electric. As revelers hung Christmas lights elsewhere in the city, the group founded the Phoebus cartel, a supervisory body that would carve up the worldwide incandescent lightbulb market, with each national and regional zone assigned its own manufacturers and production quotas. It was the first cartel in history to enjoy a truly global reach.

By MARKUS KRAJEWSKI

Photo-illustration by GLUEKIT

PREVIOUS PAGES: LIGHTBULB: FIN STEWART; THIS PAGE: PHILIPS COMPANY ARCHIVES

The cartel's grip on the lightbulb market lasted only into the 1930s. Its far more enduring legacy was to engineer a shorter life span for the incandescent lightbulb. By early 1925, this became codified at 1,000 hours for a pear-shaped household bulb, a marked reduction from the 1,500 to 2,000 hours that had previously been common. Cartel members rationalized this approach as a trade-off: Their lightbulbs were of a higher quality, more efficient, and brighter burning than other bulbs. They also cost a lot more. Indeed, all evidence points to the cartel's being motivated by profits and increased sales, not by what was best for the consumer. In carefully crafting a lightbulb with a relatively short life span, the cartel thus hatched the industrial strategy now known as planned obsolescence.

Today, with many countries phasing out incandescent lighting in favor of more-efficient and pricier LEDs, it's worth revisiting this history–not simply as a quirky anecdote from the annals of technology but as a cautionary tale about the strange and unexpected pitfalls that can arise when a new technology vanquishes an old one.

IT WASN'T EASY BEING A LIGHTBULB MAKER in the early 20th century. The rapid spread of electrification and the introduction of new forms of lighting like bicycle lamps, car headlights, and streetlights did offer nearly limitless opportunities for inventors and entrepreneurs. But as thousands of manufacturers vied for market share and a technological edge, no single company felt assured of stable sales from one year to the next. That was as true for tiny backroom operations as it was for the giant corporate entities with multinational factories and research laboratories. Immediately preceding the cartel's formation, for instance, Osram experienced a dizzying drop in its German sales, from 63 million lightbulbs in the financial year 1922–23 to 28 million the following year. Not surprisingly, Osram head William Meinhardt was the first to propose the arrangement that eventually became the Phoebus cartel.

Alliances among lightbulb makers were not exactly new. The Verkaufsstelle Vereinigter Glühlampenfabriken, for instance, was a European cartel of carbon-filament lamp manufacturers that formed in 1903 to stabilize industry ties. It was rendered superfluous when in 1906 two European companies introduced a superior lightbulb whose filament was made from tungsten paste. That bulb was itself eclipsed in 1911 by General Electric's metal-filament bulb, which used pure drawn tungsten wire, and in 1913 by GE's gas-filled tungsten bulb. Dubbed the half-watt bulb, the latter was infused with argon or some other noble gas, which preserved the tungsten better than a simple vacuum; it produced five times as much light per watt as its carbon-filament predecessor.

GE's licensing of its basic lightbulb patents gave rise to yet more alliances, most notably the powerful Patentgemeinschaft ("patent pool"), which controlled the GE patent rights in much of Europe up until World War I. Any company seeking to license GE's intellectual property had to abide by a strict production quota. Philips, for example, was given an annual quota of 5.7 million lightbulbs, despite the fact that its Eindhoven facility could easily produce twice that amount. The Berlin-based patent pool fell apart with the geopolitical reshuffling during the war. As soon as hostilities ended and the lightbulb business once again surged, a new cartel, the Internationale Glühlampen Preisvereinigung, sprang up to try to control prices for much of continental Europe.

None of these efforts, though, had quite the reach and ambition of the Phoebus cartel. On paper, it sounded entirely benign. The document that companies signed to join it was called the "Convention for the Development and Progress of the International Incandescent Electric Lamp Industry." According to that document, the organization's chief goals were "securing the cooperation of all parties to the agreement, ensuring the advantageous exploitation of their manufacturing capabilities in the production of lamps, ensuring and maintaining a uniformly high quality, increasing the effectiveness

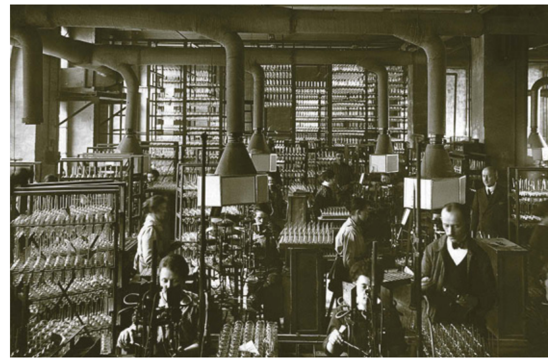

LIGHT BRIGHT: **One of the key achievements of the global lightbulb alliance known as the Phoebus cartel was to engineer shorter-lived lamps. Samples were regularly checked to ensure they conformed to cartel standards, as shown in the photo at left, taken at a facility owned by cartel member Philips, of the Netherlands. Pictured above is a lightbulb factory operated by another key cartel member, Germany's Osram.**

SIEMENS CORPORATE ARCHIVES

of electric lighting and increasing light use to the advantage of the consumer." It covered all electric lightbulbs used for illumination, heating, and medical purposes. In addition to the companies mentioned earlier, its members included Hungary's Tungsram, the United Kingdom's Associated Electrical Industries, and Japan's Tokyo Electric. The U.S. company GE, one of the prime movers behind the group's formation, was itself not a member. Instead it was represented by its British subsidiary, International General Electric, and by the Overseas Group, which consisted of its subsidiaries in Brazil, China, and Mexico. Over the next decade or so, GE would acquire significant stakes in all the member companies that it did not already own.

To oversee national lightbulb markets and their respective development in global trade, Phoebus established a supervisory body, chaired by Meinhardt of Osram. The cartel's other main activities were to facilitate the exchange of patents and technical know-how and to impose far-reaching and long-lived standards. To this day, we still use the screw-type socket—devised by Thomas Edison back in 1880 and designated E26/E27—thanks to the cartel. Most significantly for consumers, Phoebus expended considerable technical effort into engineering a shorter-lived lightbulb.

HOW EXACTLY DID THE CARTEL PULL OFF THIS engineering feat? It wasn't just a matter of making an inferior or sloppy product; anybody could have done that. But to create one that reliably failed after an agreed-upon 1,000 hours took some doing over a number of years. The household lightbulb in 1924 was already technologically sophisticated: The light yield was considerable; the burning time was easily 2,500 hours or more. By striving for something less, the cartel would systematically reverse decades of progress.

The details of this effort have been very slow to emerge. Some facts came to light in the 1940s, when the U.S. government investigated GE and a number of its business partners for anticompetitive practices. Others were uncovered more recently, when I and the German journalist Helmut Höge delved into the corporate archives of Osram in Berlin. Jointly founded in 1920 by three German companies, Osram remains one of the world's leading makers of all kinds of lighting, including state-of-the-art LEDs. In the archives, we found meticulous correspondence between the cartel's factories and laboratories, which were researching how to modify the filament and other measures to shorten the life span of their bulbs.

The cartel took its business of shortening the lifetime of bulbs every bit as seriously as earlier researchers had approached their job of lengthening it. Each factory bound by the cartel agreement—and there were hundreds, including GE's numerous licensees throughout the world—had to regularly send samples of its bulbs to a central testing laboratory in Switzerland. There, the bulbs were thoroughly vetted against cartel standards. If any factory submitted bulbs lasting longer or shorter than the regulated life span for its type, the factory was obliged to pay a fine.

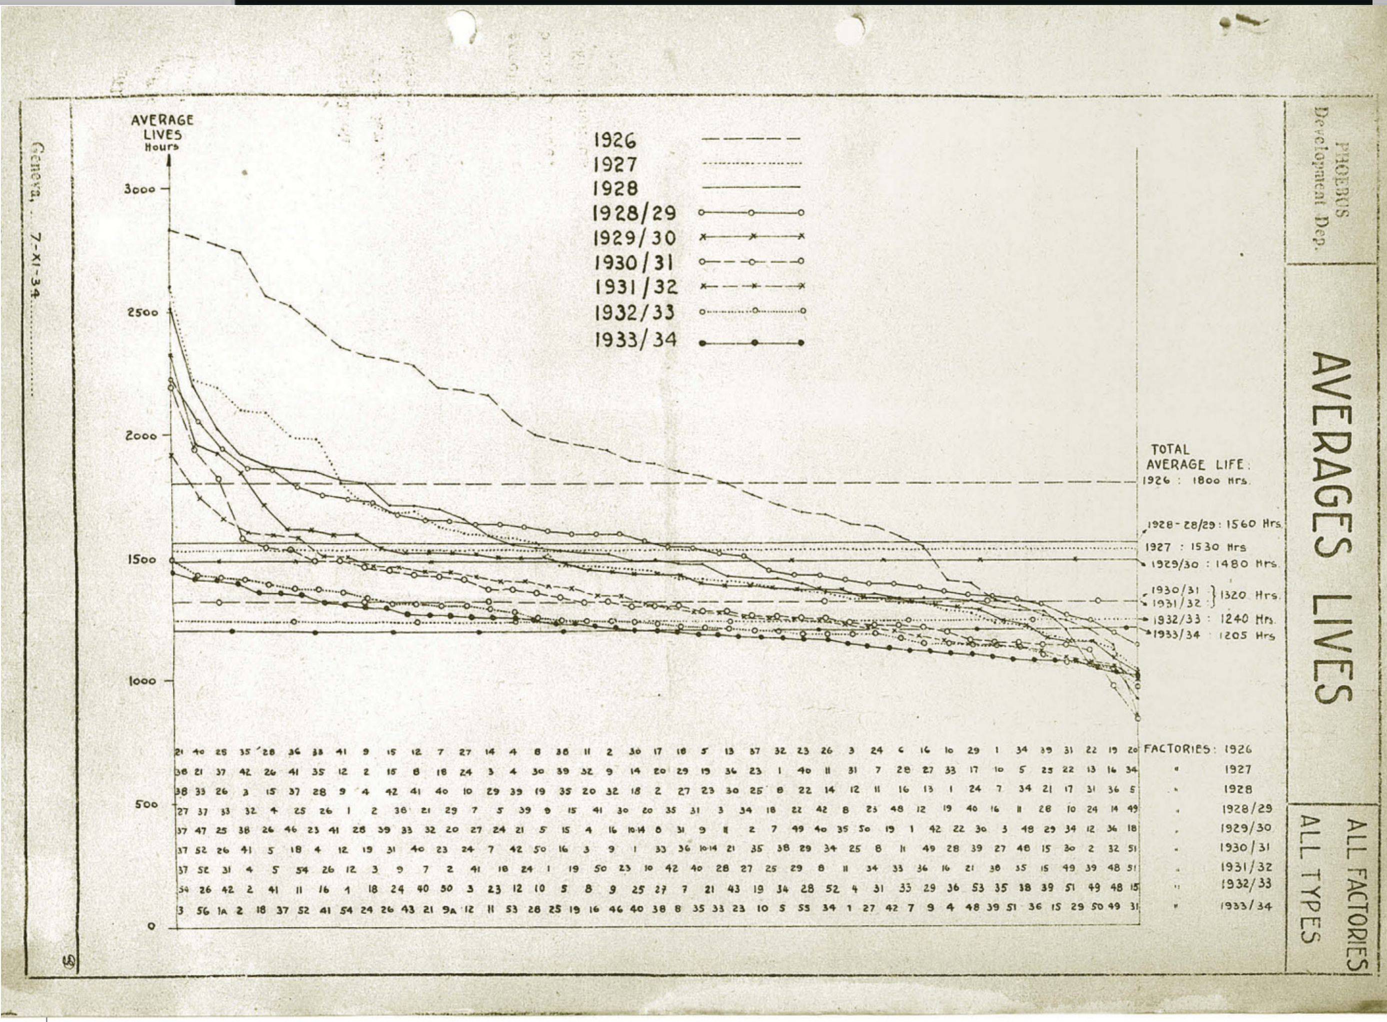

Companies were also fined for exceeding their sales quotas, which were constantly being adjusted. In 1927, for example, Tokyo Electric noted in a memo to the cartel that after shortening the lives of its vacuum and gas-filled lightbulbs, sales had jumped fivefold. "But if the increase in our business resulting from such endeavors directly mean[s] a heavy penalty, it must be a thing out of reason and shall quite discourage us," the memo stated.

There were continual reports of cartel members' attempts to restore the burning time of their bulbs to the old levels in defiance of the watchful eyes of Phoebus. At one point, some members surreptitiously introduced longer-lived bulbs by designing them to run at a voltage higher than the standard line voltage. After the Phoebus development department's customary report of voltage statistics revealed such product "enhancements," Anton Philips, head of Philips, complained to an executive at International General Electric: "This, you will agree with me, is a very dangerous practice and is having a most detrimental influence on the total turnover of the Phoebus Parties.... After the very strenuous efforts we made to emerge from a period of long life lamps, it is of the greatest importance that we do not sink back into the same mire by paying no attention to voltages and supplying lamps that will have a very prolonged life."

As this episode reveals, tweaking a lightbulb's rated voltage was one way to modify the product's life. Another was to adjust the current, as GE engineers did to decrease the life span of its flashlight bulbs. A GE flashlight bulb in the precartel days was designed to last longer than three changes of batteries. This life span was then cut to two battery changes, and in 1932 the GE engineering department proposed that the bulb last no longer than one battery. A GE engineer named Prideaux wrote in a memo, "We would suggest increasing Mazda lamp No. 10 from .27 ampere to .30; and 13.14 and 31 from .30 to .35. This would result in increases of candlepower of 11 and 16 percent respectively." That boost in illumination, he suggested, "would be acceptable to all flashlight users" despite the fact that the higher current would shorten not just the bulb's life but also the battery's.

The cartel's justification for these changes was that at the higher current levels, the bulbs produced more lumens per watt. Alas, more current means not only more brightness but also higher

OUT, OUT, BRIEF BULB: Prior to the Phoebus cartel's formation in 1924, household bulbs typically burned for a total of 1,500 to 2,500 hours; cartel members agreed to shorten that life span to a standard 1,000 hours. Each factory regularly sent lightbulb samples to the cartel's central laboratory in Switzerland for verification. This graph, obtained from the Municipal Archive of Berlin, shows how life spans generally declined over time, from an average of 1,800 hours in 1926 to 1,205 hours in fiscal year 1933–34.

LANDESARCHIV BERLIN

filament temperature and therefore shorter life. Indeed, much of the cartel's life-limiting research focused on the filament, including its material, its shape, and the evenness of its dimensions.

Over the course of nearly a decade, the cartel succeeded in this quest. The average life of a standard reference lightbulb produced in dozens of Phoebus members' factories dropped by a third between 1926 and fiscal year 1933–34, from 1,800 hours to just 1,205 hours. At that point, no factory was producing bulbs lasting more than 1,500 hours.

Of course, given the collective ingenuity of the cartel's engineers and scientists, it should have been possible to design a lightbulb that was both bright and long-lived. But such a product would have interfered with members' desire to sell more bulbs. And sell more bulbs they did, at least initially. In fiscal year 1926–27, for instance, the cartel sold 335.7 million lightbulbs worldwide; four years later, sales had climbed to 420.8 million. What's more, despite the fact that the actual costs of manufacturing were dropping, the cartel maintained more or less stable prices and therefore higher profit margins. From its inception until the end of 1930, the cartel retained its overwhelming share of a growing market. But the good times would not last.

AS THE CARTEL continued its policy of artificially elevated prices, competitors spotted a golden opportunity to sell cheaper, if often inferior-quality, goods. Particularly threatening was the flood of inexpensive bulbs from Japan. Although Tokyo Electric was a cartel member, it had no control over the hundreds of smaller, family-owned workshops that produced bulbs almost entirely by hand. Japanese consumers apparently preferred the higher-quality products sold by the larger manufacturers, and so the majority of these cheap, handmade bulbs were exported to the United States, Europe, and elsewhere, where they sold for a fraction of the price of a Phoebus bulb and well below the average production cost of a cartel bulb, too. From 1922 to 1933, Japan's annual output of incandescent bulbs grew from 45 million to 300 million.

However, as Philips historian I.J. Blanken has noted, these cheap bulbs weren't necessarily a bargain. "Owing to its greater current consumption, the true cost of using one of the poor-quality Japanese lamps, measured over the life of the lamp, was many times greater" than whatever the consumer had saved by buying a cheap lamp rather than a Philips one.

Powerful and influential though it was, the Phoebus cartel was short-lived. Within six years of its formation, the cartel was already starting to struggle. Between 1930 and 1933, its sales volume dropped by more than 20 percent–even as the overall market for lighting was growing. The cartel was also weakened by the expiration of GE's basic lightbulb patents in 1929, 1930, and 1933, by occasional conflicts among its members, and by legal attacks, particularly in the United States. What ultimately killed Phoebus, however, was World War II. As the members' host countries went to war, close coordination became impossible. The cartel's 1924 agreement, which was supposed to last until 1955, was nullified in 1940.

Though long gone, the Phoebus cartel still casts a shadow today. That's true in part because the lighting industry is now going through its most tumultuous period of technological change since the invention of the incandescent bulb. After more than a century of dominance, these bulbs are now being phased out in favor of compact fluorescent and especially LED bulbs.

Consumers are expected to pay more money for bulbs that are up to 10 times as efficient and that are touted to last a fantastically long time–up to 50,000 hours in the case of LED lights. In normal usage, these lamps will last so long that their owners will probably sell the house they're in before having to change the bulbs.

Whether or not these pricier bulbs will actually last that long is still an open question, and not one that the average consumer is likely to investigate. There are already reports of CFLs and LED lamps burning out long before their rated lifetimes were reached. Such incidents may well have resulted from nothing more sinister than careless manufacturing. But there is no denying that these far more technologically sophisticated products offer tempting opportunities for the inclusion of purposefully engineered life-shortening defects. After all, few people will complain, or even notice, if a bulb burns out 9 years after it is installed rather than 14. True, today's lighting industry is much larger and more diverse than it was in the 1920s and '30s, and government monitoring of collusive behavior is more vigilant. Nevertheless, the allure for businesses to cooperate in such a market is strong. And the Phoebus cartel shows how it could succeed. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/phoebuscartel1014

CIRCLING *the* GLOBE

The Phoebus cartel enjoyed a truly global reach. The U.S. company General Electric was itself not a member but was represented through its overseas subsidiaries.

WILLIAM MEINHARDT/ OSRAM
(Germany)

ANTON PHILIPS/ PHILIPS
(Netherlands)

TUNGSRAM
(Hungary)

ASSOCIATED ELECTRICAL INDUSTRIES
(United Kingdom)

INTERNATIONAL GENERAL ELECTRIC

COMPAGNIE DES LAMPES
(France)

OVERSEAS GROUP:

GENERAL ELECTRIC SOCIEDAD ANONYMA
(Brazil)

CHINA GENERAL EDISON

MEXICANA DE LAMPARAS ELECTRICAS

TOKYO ELECTRIC
(Japan)

TOP: SIEMENS CORPORATE ARCHIVES; BOTTOM: PHILIPS COMPANY ARCHIVES

Carnegie Mellon University

College of Engineering

Faculty Positions

Carnegie Mellon University, a world leader in information and communication technology, has extended its global reach into Africa with its Center of Excellence in Information and Communication Technology. Focused on one of the fastest growing economic regions of the world, the Center of Excellence is creating a new generation of graduates able to take advantage of Africa's unique opportunities.

Based in Rwanda, the Center of Excellence benefits from the country's bold ICT strategy and CMU's culture of innovation to provide a platform for students to become technology thought leaders through its Master of Science in Information Technology and Master of Science in Electrical and Computer Engineering.

We are seeking highly qualified candidates who share this vision and wish to join a dynamic world-class faculty. Faculty members are expected to collaborate with regional and multi-national industry in technology innovation and entrepreneurship and deliver innovative, interdisciplinary graduate teaching and research programs for the African context.

Areas of interest include:
•software engineering;
•mobile computing;
•cloud computing;
•big data analytics;
•communications;
•information and cyber security;
•intelligent infrastructures;
•wireless networking;
•mHealth;
•energy systems;
•eLearning; and
•mobile applications.

Candidates should possess a Ph.D. in a related discipline and an outstanding record in research, teaching and leadership. Applications should include a comprehensive resume including a complete list of publications, 3-5 professional references, a statement of research and teaching interests (less than 2 pages each), and copies of 2 research papers (journal or conference papers).

APPLICATIONS SHOULD BE SENT TO:
Director, Carnegie Mellon University in Rwanda
rwanda-coe@cit.cmu.edu

Application review starts Dec. 1

www.cmu.edu/rwanda

Faculty Positions in Robotics and Mechatronics

NAZARBAYEV UNIVERSITY

Nazarbayev University is seeking highly qualified full-time faculty members at Assistant, Associate and Full Professor ranks to join its rapidly growing program in Robotics and Mechatronics in the School of Science and Technology. Successful candidates must have a Ph.D. degree from a reputed research university, a demonstrated ability for research, excellent English-language communication skills 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. The university is committed to be a world-class research university and has made significant investments in research infrastructure, equipment and personnel.

Applications are particularly encouraged from candidates with research interests in the areas of biorobotics, mobile robotics, robot vision, industrial automation, unmanned aerial vehicles, and mechatronic system design. Exceptional candidates with research interests in all topics related to Robotics and Mechatronics are also encouraged to apply.

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

Applicants should send a detailed CV, teaching and research statements, 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 December 15th, 2014. For more information, visit http://sst.nu.edu.kz.

ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

UNIVERSITY of MICHIGAN ■ COLLEGE of ENGINEERING

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. The division invites those with interests in integrated circuits, applied electromagnetics, power electronics, optics, robotics, smart systems, communications, solid-state devices and nanotechnology, biomedical engineering, and all other relevant areas of research to apply.

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, 2014.

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.

a great place to work

KUWAIT COLLEGE OF SCIENCE & TECHNOLOGY
كلية الكويت للعلوم و التكنولوجيا
THE KNOWLEDGE NETWORK

Kuwait College of Science & Technology is seeking highly qualified Vice President – Academic Affairs and Dean of Faculty and full-time faculty members as Assistant, Associate and Full Professors to ranks to join its rapidly growing program in Computer Science and Engineering / Electronics and Communication Engineering. Successful candidates must have a Ph.D. degree from a reputed research university, a demonstrated ability for research, excellent English-language communication skills and a commitment to graduate and undergraduate teaching and program development. Infrastructure of the college is under construction and is expected to be completed in the first quarter of 2015. First academic session of the college is scheduled to start from mid 2015. The college is planned to have a capacity of 1,500 students during the initial period. It will be the first higher education institution in Kuwait to offer the combination of Bachelor of Science in Computer Science - B.Sc.(CS), Bachelor of Technology in Computer Science and Engineering - B.Tech. (CSE) and Bachelor of Technology in Electronics and Communication Engineering - B.Tech (ECE). The College is affiliated with the Indian Institute of Technology, India – IIT. The university is committed to be a world-class research university and will make significant investments in research infrastructure, equipment and personnel. Applications are particularly encouraged from candidates with research interests in the areas of Computer Engineering and Electronics and Communication Engineering. Benefits include an internationally competitive salary, housing allowance (based on rank) and home-leave travel as per policy. ***Applicants should send a detailed CV, teaching and research statements, and list of publications to careers@kcst.net. Review of applications will begin immediately. Applications are to be submitted before October 21, 2014. For more information, visit www.kcst.net***

Vice President for Academic Affairs and Dean of the Faculty

Successful candidate will deputize for the President in all academic and related matters, coordinating all activities at local, national and international levels that contribute to the College's strategic direction. He/she will also oversee the academic activities of the faculties, as well as other independent academic units, including both the appointment of academic staff and the oversight of those aspects of the College's assurance systems designed to underpin and guarantee the high quality of the College's teaching and Research. Candidate provides dynamic and strategic campus leadership for the quality and integrity of all of the College's academic degree programs, assessing student learning outcomes, developing high quality faculty, providing academic support services, and overseeing constituent units that include the library, advising, academic support services, student life, registrar, institutional effectiveness, international studies and career counseling. Substantial involvement in the campus community is essential. The Vice President for Academic Affairs will have an earned doctorate with progressively responsible administrative experience in senior administrative and academic positions; and college-level teaching experience. The candidate is expected to have Ph.D. (CS/IT/ECE); Understanding of best international higher education practice, particularly in the areas of teaching and research with more than 15 years post PhD experience in education, and/or in industry. At least five years as professor. Desirable: Experience as dean or head of department. Experience in Development of new academic and research programs Demonstrable competence in managing sizable groups of academics and professionals. Keen interest in and strong commitment to upholding the academic quality and standard of the College

Faculty - Professors / Associate Professors /Assistant Professors / Lecturers

This position teaches in the Computer Engineering, Computer Science and Electronics and Communication Engineering programs. Duties include teaching computing and engineering courses at undergraduate levels. Duties also include maintaining curriculum, courses of study, and assessment work. Faculty will participate in the process of defining undergraduate computer science and engineering / Electronics and Communication Engineering curriculum. In this capacity, faculties are expected to:

- *Assist in curriculum review and development,*
- *Develop laboratories to support computer science and engineering curriculum,*
- *Create projects designed to enhance lectures,*
- *Read and stay abreast of current topics in computer science and engineering,*
- *Develop syllabus and review textbooks for adoption,*
- *Propose and develop new courses as determined appropriate by the faculty,*
- *Actively participating and support SACS and ABET accreditation preparation.*
- *Providing students with an engaging and challenging environment to learn and develop technology skills*

Kuwait College of Science and Technology

Faculties are required by the College to:

- *Teach contracted load,*
- *Maintain office hours,*
- *Mentor and advise students,*
- *Keep timely records of grades, grade papers, and perform other administrative duties as needed. Faculty will perform other duties as assigned and actively participate in College assemblies and activities.*

Candidates qualifying for Professor position may please indicate whether he/she would accept position of Dean (Faculty) / Dean(Academic) / Dean(Student) /Dean (Research) For the position of Associate Professor, the successful applicant will hold a relevant Ph.D. with 7 years' post PhD undergraduate teaching experience of which at least 2 years as Assistant professor or equivalent. For position of Assistant Professor, the successful applicant will hold a relevant Ph.D. with 2 years' post PhD undergraduate teaching experience. Exceptional candidates, fresh after Ph.D, could be considered for this post. For the lecturer position, fresh PhDs with two years undergraduate teaching experience or M.Sc. with five years undergraduate teaching experience. Industrial experience in research and development is highly desirable. Professional registration is preferred. The successful candidate must be willing and have the ability to pursue a Ph.D. if a terminal degree is not already held. Completion of an acceptable doctoral degree is required for the position to be tenure-track. Rank and salary will be commensurate with experience and credentials. The successful applicant will have strong problem-solving and communication skills and the ability to work in an interdisciplinary environment. A strong commitment to teaching excellence and appropriate publications in leading scholarly journals are expected.

THE UNIVERSITY OF CHICAGO

The Department of Computer Science at the **University of Chicago** invites applications from exceptionally qualified candidates in the areas of (a)systems, and (b) theory of computing for faculty positions at the rank of **Associate Professor**.

Systems is a broad, synergistic collection of research areas spanning systems and networking, programming languages and software engineering, software and hardware architecture, data-intensive computing and databases, graphics and visualization, security, systems biology, and a number of other areas. We encourage applicants working within our strategic focus of data-intensive computing, but also in all areas of systems.

The Theory of Computing ("Theory" for short) strives to understand the fundamental principles underlying computation and explores the power and limitations of efficient computation. While mathematical at its core, it also has strong connections with physics (quantum computing), machine learning, computer vision, natural language processing, network science, cryptography, bioinformatics, and economics, to name just a few areas. We encourage applications from researchers in core areas of Theory such as complexity theory and algorithms as well as in any area with a significant Theory component.

The University of Chicago has the highest standards for scholarship and faculty quality, is dedicated to fundamental research, and encourages collaboration across disciplines. We encourage connections with researchers across campus in such areas as bioinformatics, mathematics, molecular engineering, natural language processing, and statistics to mention just a few.

The Department of Computer Science (cs.uchicago.edu) is the hub of a large, diverse computing community of two hundred researchers focused on advancing foundations of computing and driving its most advanced applications. Long distinguished in theoretical computer science and artificial intelligence, the Department is now building strong systems and machine learning groups. The larger community in these areas at the University of Chicago includes the Department of Statistics, the Computation Institute, the Toyota Technological Institute at Chicago (TTIC), and the Mathematics and Computer Science Division of Argonne National Laboratory.

The Chicago metropolitan area provides a diverse and exciting environment. The local economy is vigorous, with international stature in banking, trade, commerce, manufacturing, and transportation, while the cultural scene includes diverse cultures, vibrant theater, world-renowned symphony, opera, jazz, and blues. The University is located in Hyde Park, a Chicago neighborhood on the Lake Michigan shore just a few minutes from downtown.

Applicants must have a doctoral degree in Computer Science or a related field such as Mathematics, Statistics, etc. Applicants are expected to have established an outstanding research program and will be expected to contribute to the department's undergraduate and graduate teaching programs.

Applications must be submitted through the University's Academic Jobs website.

To apply for the position of ***Associate Professor-Systems***, go to: http://tinyurl.com/pkzpcy5

To apply for the position of ***Associate Professor-Theory***, go to: http://tinyurl.com/kwzb9zu

To be considered as an applicant, the following materials are required:

- cover letter,
- curriculum vitae including a list of publications,
- statement describing past and current research accomplishments and outlining future research plans, and
- description of teaching philosophy, and
- a reference contact list consisting of three people

Review of complete applications will begin January 15, 2015 and will continue until all available positions are filled.

All qualified applicants will receive consideration for employment without regard to race, color, religion, sex, national origin, age, protected veteran status or status as an individual with disability.

The University of Chicago is an Affirmative Action / Equal Opportunity / Disabled / Veterans Employer.

THE UNIVERSITY OF CHICAGO

The Department of Computer Science at the **University of Chicago** invites applications from exceptionally qualified candidates in the areas of (a) systems, and (b) theory of computing for faculty positions at the rank of **Assistant Professor**.

Systems is a broad, synergistic collection of research areas spanning systems and networking, programming languages and software engineering, software and hardware architecture, data-intensive computing and databases, graphics and visualization, security, systems biology, and a number of other areas. We encourage applicants working within our strategic focus of data-intensive computing, but also in all areas of systems.

The Theory of Computing ("Theory" for short) strives to understand the fundamental principles underlying computation and explores the power and limitations of efficient computation. While mathematical at its core, it also has strong connections with physics (quantum computing), machine learning, computer vision, natural language processing, network science, cryptography, bioinformatics, and economics, to name just a few areas. We encourage applications from researchers in core areas of Theory such as complexity theory and algorithms as well as in any area with a significant Theory component.

The University of Chicago has the highest standards for scholarship and faculty quality, is dedicated to fundamental research, and encourages collaboration across disciplines. We encourage connections with researchers across campus in such areas as bioinformatics, mathematics, molecular engineering, natural language processing, and statistics, to mention just a few.

The Department of Computer Science (cs.uchicago.edu) is the hub of a large, diverse computing community of two hundred researchers focused on advancing foundations of computing and driving its most advanced applications. Long distinguished in theoretical computer science and artificial intelligence, the Department is now building strong systems and machine learning groups. The larger community in these areas at the University of Chicago includes the Department of Statistics, the Computation Institute, the Toyota Technological Institute at Chicago (TTIC), and the Mathematics and Computer Science Division of Argonne National Laboratory.

The Chicago metropolitan area provides a diverse and exciting environment. The local economy is vigorous, with international stature in banking, trade, commerce, manufacturing, and transportation, while the cultural scene includes diverse cultures, vibrant theater, world-renowned symphony, opera, jazz, and blues. The University is located in Hyde Park, a Chicago neighborhood on the Lake Michigan shore just a few minutes from downtown.

Applicants must have completed all requirements for the PhD at the time of appointment. The PhD should be in Computer Science or a related field such as Mathematics, or Statistics, etc.

Applications must be submitted through the University's Academic Jobs website.

To apply for the ***Assistant Professor - Systems***, go to: http://tinyurl.com/k3wgaqv

To apply for the ***Assistant Professor - Theory***, go to: http://tinyurl.com/khyc74d

To be considered as an applicant, the following materials are required:
- cover letter,
- curriculum vitae including a list of publications, statement describing past and current research accomplishments and outlining future research plans, and
- description of teaching philosophy, and
- three reference letters, one of which must address the candidate's teaching ability.

Reference letter submission information will be provided during the application process.

Review of application materials will begin on January 15, 2015 and continue until all available positions are filled.

All qualified applicants will receive consideration for employment without regard to race, color, religion, sex, national origin, age, protected veteran status or status as an individual with disability.

The University of Chicago is an Affirmative Action / Equal Opportunity / Disabled / Veterans Employer.

UNIVERSITY OF MINNESOTA
Driven to Discover℠

The **Electrical and Computer Engineering, University of Minnesota – Twin Cities**, invites applications for faculty positions in:

(1) power and energy systems;
(2) biomedical imaging and
(3) control and dynamical systems; robotics and automation; image processing and computer vision; novel sensing and actuation; devices; circuits and systems, to support a University-wide initiative on robotics, sensors, and advanced manufacturing, **http://cse.umn.edu/mndrive**.

Women and other underrepresented groups are especially encouraged to apply. An earned doctorate in an appropriate discipline is required. Rank and salary will be commensurate with qualifications and experience. Positions are open until filled, but for full consideration, apply at:

http://www.ece.umn.edu/

by December 15, 2014. The University of Minnesota is an equal opportunity employer and educator.

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

School of Electrical & Computer Engineering
College of Engineering

The University of Oklahoma School of Electrical and Computer Engineering invites applications for a tenure-track Assistant Professor position. Candidates are sought in the area of electric energy systems and smart grid with a particular focus on the power system. Outstanding candidates in other related areas including control systems, power electronics, renewable energy and distributed generation are also encouraged to apply.

The position requires an earned doctorate in electrical and/or computer engineering, or a closely related discipline. Exceptional candidates with a record of professional achievement sufficient for tenured or tenure-track appointment at the rank of Associate Professor will also be given serious consideration. The successful candidate will be expected to: (a) contribute to research efforts in the development of smart grid technology, control of renewable energy resources and enhancement of grid security; (b) teach at both the undergraduate and graduate levels.

Established in 1890, the University of Oklahoma is a comprehensive public research university offering a wide array of undergraduate, graduate and professional programs and extensive continuing education and public service programs. The newly completed Devon Energy Hall, a 100,000 square foot state-of-the-art building containing research laboratories, offices, and classrooms, is the new home to the School of Electrical and Computer Engineering. Energy has long been a cornerstone of the economy of Oklahoma, and the University has long-term, energy-related research programs ongoing.

Candidates should submit a letter of application, curriculum vitae, teaching and research statements, and the names of three references to ecesearch@ou.edu, or for questions you may contact:

Ms. Cathy Trujillo, Business Manager
OU School of Electrical and Computer Engineering
110 W. Boyd St., Rm. 150
Norman, OK 73019-1102
Voice: 405.325.4723

Electronic submission in PDF format is preferred. Review of applications will begin immediately and continue until the position is filled. Minorities and women are especially encouraged to apply.

The University of Oklahoma is an Affirmative Action/Equal Opportunity Employer and encourages diversity in the workplace.

DATAFLOW_

WHAT ARE VENTURE CAPITALISTS BETTING ON?

OPTIMISM IN INVESTING IS RISING, BUT NOT EVENLY

According to a recent international survey of hundreds of venture capitalists conducted by Deloitte Consulting for the U.S. National Venture Capital Association, overall investor confidence is growing across almost every tech sector. The greatest growth globally is in the clean technology and robotics sectors, but looking at just U.S. venture investors, robotics and consumer businesses led growth. Cloud computing remained the sector inspiring the most confidence, for the second year in a row. —STEPHEN CASS

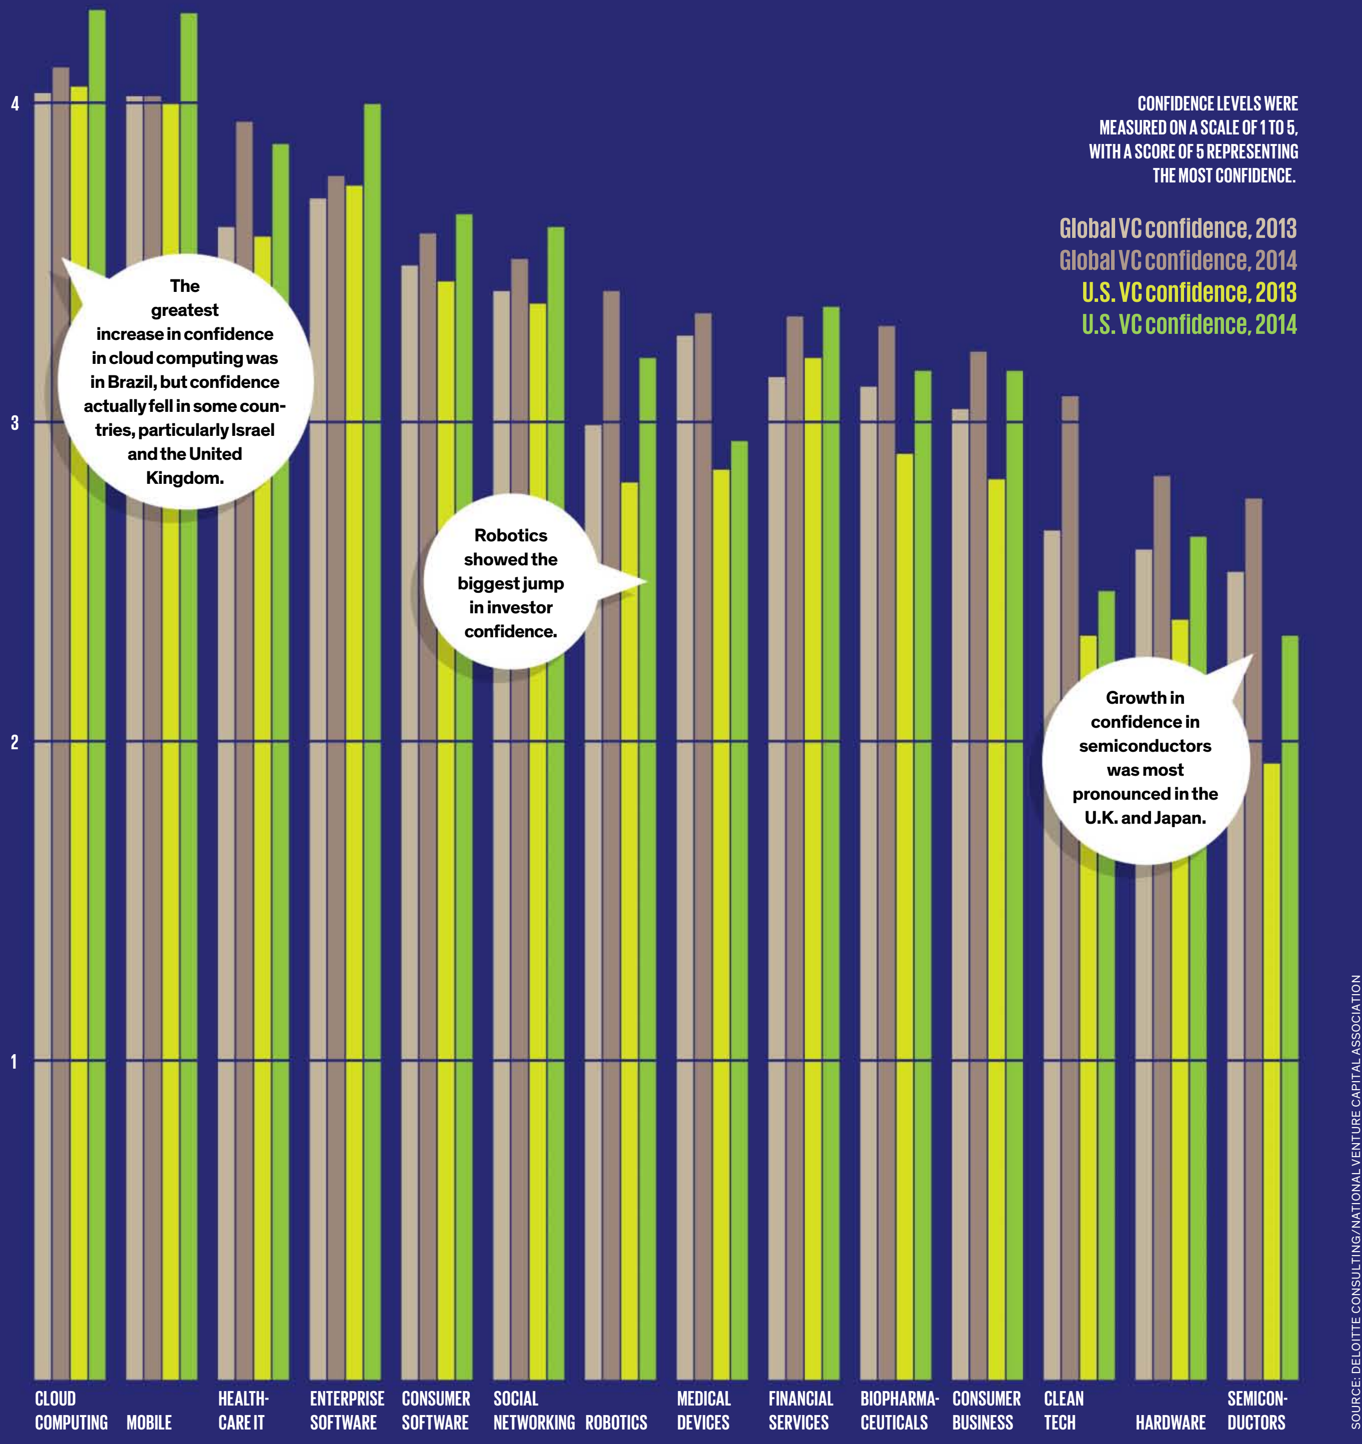

INFOGRAPHIC BY Brandon Palacio

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