IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 08.14

COMPUTING'S BIG BREAKTHROUGH? Silicon will unleash quantum's power **P. 32**

MERCEDES'S CAR WILL BEAT GOOGLE'S Robocar advantage is cheap, clever radar **P. 40**

THE FORGOTTEN ROCKET SCIENTIST Frank Malina hot-wired the space age **P. 46**

SUMMER SCIENCE FICTION SPECIAL A preview of our upcoming e-book at **SPECTRUM.IEEE.ORG**

THE COLLAPSE OF Internet Privacy

Browser fingerprinting lets advertisers track your every online move **P. 26**

IEEE

POWER SUPPLY COOLING: Simulation of the thermal behavior of an enclosed computer power supply unit.

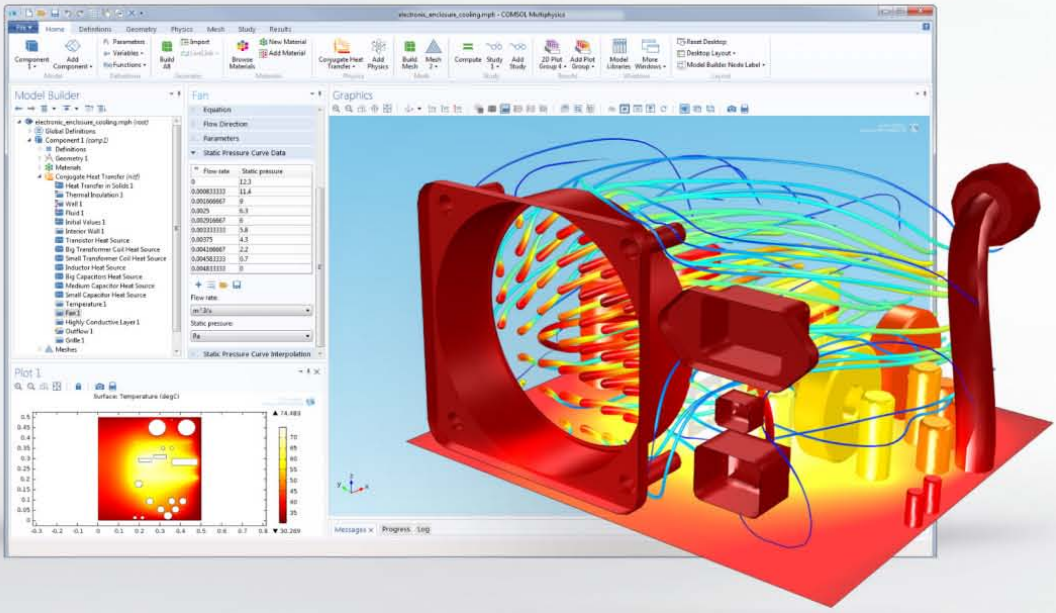

VERIFY AND OPTIMIZE YOUR DESIGNS WITH

COMSOL MULTIPHYSICS®

Multiphysics tools let you build simulations that accurately replicate the important characteristics of your designs. The key is the ability to include all physical effects that exist in the real world. To learn more about COMSOL Multiphysics, visit www.comsol.com/introvideo

Product Suite

COMSOL Multiphysics

ELECTRICAL

AC/DC Module
RF Module
Wave 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
ECAD Import Module
LiveLink™ *for* SolidWorks®
LiveLink™ *for* Inventor®
LiveLink™ *for* AutoCAD®
LiveLink™ *for* Creo™ Parametric
LiveLink™ *for* Pro/ENGINEER®
LiveLink™ *for* Solid Edge®
File Import *for* CATIA® V5

FEATURES_08.14
IEEE SPECTRUM

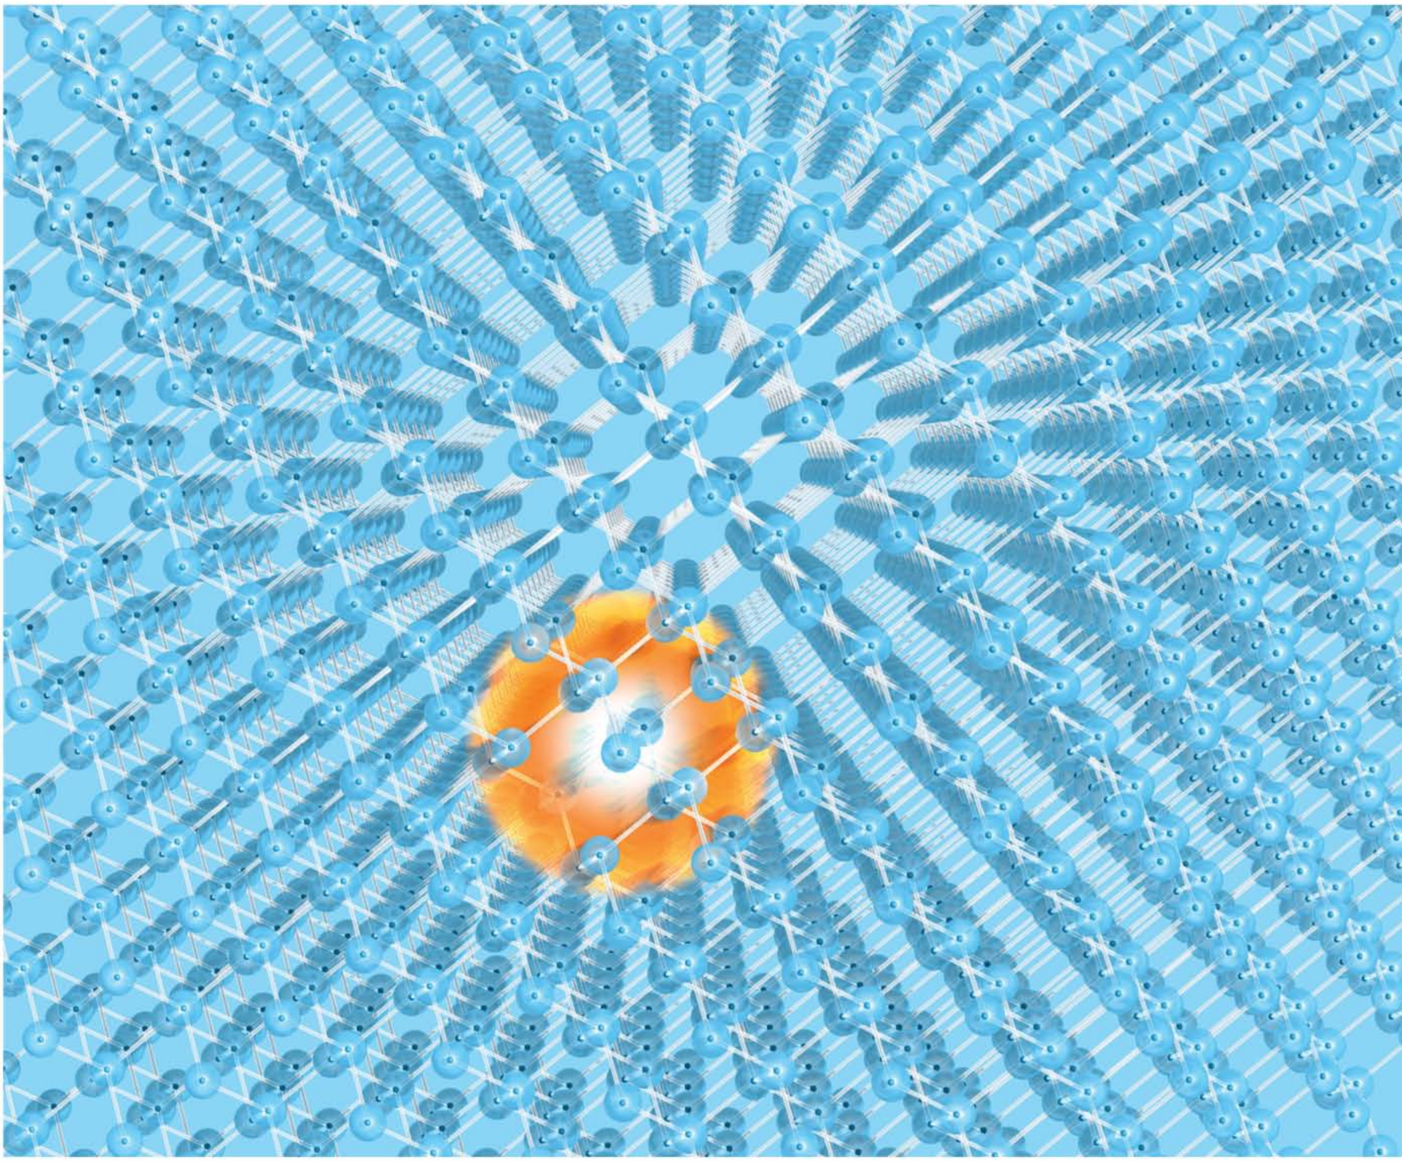

32

SILICON'S SECOND ACT

Silicon took the electronics industry by storm. Could it do the same for quantum computing?

BY CHEUK CHI LO & JOHN J.L. MORTON

26

Browse at Your Own Risk

Browser fingerprinting is tracking you right now. And there's nothing you can do about it.

By Nick Nikiforakis & Günes Acar

40

Making Bertha See

Mercedes-Benz's self-driving car has many sensors, but radar is the critical one.

By Jürgen Dickmann, Nils Appenrodt & Carsten Brenk

46

America's Forgotten Rocketeer

Frank Malina's motley crew of engineers and enthusiasts legitimized rocket science.

By James L. Johnson

COVER, PROP STYLIST: DOMINIQUE BAYNES

On the Cover **Photograph for *IEEE Spectrum* by Dan Saelinger**

ILLUSTRATION BY **Bryan Christie Design**

Infinite Designs, One Platform
with the only complete system design environment

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

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

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

800 453 6202

NATIONAL INSTRUMENTS™

DEPARTMENTS_08.14

IEEE SPECTRUM

09

News

Holographic 3-D Displays for Smartphones
At least two companies are close to commercializing the ultimate in 3-D.
By Sarah Lewin

11 Mightier Magnetic Memory
12 Storing Wind Energy Underwater
14 The Practical Polariton Laser
16 The Big Picture

19

Resources

Antigravity Glitter
Build a levitating disco-ball night-light; bell-bottom pants optional.
By W. Wayt Gibbs

21 Patents: Choose the Filing Strategy That's Right for You
22 Careers: Where the Jobs Are
23 Start-ups: Brewbot, Bonaverde
56 Dataflow: Phone-Snooping Laws

06

Opinion

Spectral Lines
What do Edward Teller and William Shockley have in common? They both wrote articles for *IEEE Spectrum*.
By Glenn Zorpette

04 Back Story
05 Contributors
25 Technically Speaking

Online

Spectrum.ieee.org

Coming Soon Enough
As part of its 50th anniversary celebrations, *IEEE Spectrum* is releasing an e-book anthology of six science fiction stories. Nancy Kress, Greg Egan, Brenda Cooper, and others present their vision of the future of technology. The e-book is sold through Amazon.com and iTunes; a preview is available online.

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Modeling Batteries and Fuel Cells Using COMSOL Multiphysics—14 August

- **IBM RATIONAL SMARTER PRODUCTS**
  http://spectrum.ieee.org/static/ibm-rational-smarter-products
- **IBM RATIONAL INDUSTRY SOLUTIONS LIBRARY**
  http://spectrum.ieee.org/static/ibm-rational-industry-solutions-library
- **NEW PRODUCT RELEASE LIBRARY**
  http://spectrum.ieee.org/static/new-product-release-library
- **WHITE PAPERS**
  http://spectrum.ieee.org/whitepapers

The Institute

Available 8 August at theinstitute.ieee.org

- **SAVING MOTHERS** Read about the latest life-saving medical devices aimed at improving maternal health in developing regions.

- **EFFECTIVE CAREER NETWORKING** Not all contacts are equally valuable, and not every interaction is immediately rewarding. Learn how to increase the value of each networking interaction so you're more likely to land a job.

- **PREDICTING THE FUTURE** The future of humans, health care, networks, and energy are among the topics that will be covered at this year's IEEE Technology Time Machine event, to be held 21 to 22 October in San Jose, Calif.

FROM LEFT: OSTENDO TECHNOLOGIES; W. WAYT GIBBS; MARTIN ANSIN

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. 8, 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_

A Rocketeer to Remember

THE FIRST TIME HE HEARD the name Frank Malina, James L. Johnson was taking an aerospace history class at the University of North Dakota. The professor mentioned how Malina and his colleagues had pioneered U.S. rocketry. "I knew about Wernher von Braun and the German engineers," Johnson recalls. "But this guy [Malina] sounded really interesting, and I wondered why I hadn't heard about him before."

Years later, as a graduate student in history at Case Western Reserve University, in Cleveland, Johnson decided to write his thesis about Malina. Johnson's research included going through all of Malina's correspondence—he was a prolific letter writer—at the U.S. Library of Congress. Upon learning that Malina had been investigated as a suspected communist, Johnson also filed a Freedom of Information Act request, through which he gained access to Malina's voluminous FBI file. Given Malina's present-day obscurity, Johnson was shocked to learn that during and shortly after World War II, Malina was considered the leading U.S. rocket scientist.

Johnson is now vice president of operations for a company that makes hydraulic seals, and in his spare time he continues to research this forgotten hero. One place where Malina hasn't been forgotten, Johnson notes, is in Brenham, Texas, Malina's hometown. Two years ago, the local museum held a special exhibition about the rocketeer. Johnson was in the area on business and made a point of stopping by [see photo].

"It's ironic that we celebrate von Braun, the former Nazi, while we completely overlook Malina, who was just as talented, just as passionate, and whose contributions were just as important to the field of rocket science," says Johnson. With his article "America's Forgotten Rocketeer," in this issue, he hopes to help set the record straight. ■

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. 8 (INT), August 2014, p. 56, or in *IEEE Spectrum*, Vol. 51, no. 8 (NA), August 2014, p. 64.

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
Joshua J. Romero (Interactive), j.j.romero@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
Ariel Bleicher, a.bleicher@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
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

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

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

ANN NEFF

CONTRIBUTORS_

Jürgen Dickmann

Dickmann started his career working on high-frequency radar systems that help airplanes sense trouble in the heavens; he then moved to Daimler, in Germany, to teach cars to do the same on the ground. He says it's a harder job: "Signals from a car or pedestrian must be distinguished from stray reflections" off lampposts, manhole covers, the grass. In this issue, he and Daimler colleagues Nils Appenrodt and Carsten Brenk write about how it's done in an experimental self-driving Mercedes [p. 40].

W. Wayt Gibbs

Gibbs, who brewed his own conductive ink in his last *IEEE Spectrum* how-to, returns to create a magnetic floating disco ball [p. 19]—which he uses every evening when he dims the lights to watch TV. Gibbs says DIY projects present a unique challenge to science writers, who normally deal in concepts at a higher level: "There's no way to BS yourself through a project like this. You have to understand it all the way down to the bottom."

John J.L. Morton

Nanoelectronics professor Morton and research fellow Cheuk Chi Lo, both at University College London, describe how silicon could be used to build powerful quantum computers [p. 32]. Recent worldwide advances in this area have been staggering, Morton says. "Quantum bits now live longer than the time it takes to make a cup of tea, and you can watch the telltale signal on an oscilloscope of a single quantum bit flipping back and forth," he says. "It's an exciting time to be working in silicon."

Nick Nikiforakis

Nikiforakis, who joins the faculty of New York's Stony Brook University in September, works at the University of Leuven, in Belgium, where coauthor Günes Acar is a Ph.D. student. In "Browse at Your Own Risk" [p. 26], they explain how advertisers can track your online activities without using cookies. Nikiforakis was raised in Greece, where the Muppets are somewhat obscure, but he sometimes refers to the creepy new technique as a "cookieless monster"—an apt label if you value your privacy.

Andrew Zbihlyj

Zbihlyj, who created intricate maps for *Spectrum*'s 2010 water and energy report, returns to illustrate "Making Bertha See," about a self-driving Mercedes-Benz [p. 40]. Imagining the scene from the sensors' perspective, he drew a photorealistic car surrounded by people made of sketched radar blips. Zbihlyj loves illustrating unfamiliar concepts: "To be honest, I'm not that big a car person," he admits. "I don't own a car, but I can appreciate that this is exciting technology."

PHOTO-ILLUSTRATIONS BY **Gluekit**

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
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
CORPORATE STRATEGY Matthew Loeb, CAE
+1 732 562 5320, m.loeb@ieee.org
CHIEF MARKETING OFFICER Patrick D. Mahoney
+1 732 562 5596, p.mahoney@ieee.org
CHIEF 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

FELLOW TRAVELERS: John R. Pierce [left] and Rudolf Kompfner perfected the traveling wave tube at Bell Laboratories in the 1950s. Both also wrote memorable feature articles for *IEEE Spectrum*.

YALE JOEL/TIME & LIFE PICTURES/GETTY IMAGES

Famous Bylines in *IEEE Spectrum*

How I lost an article by Bill Gates

A MAGAZINE, LIKE A POLITICIAN, needs friends. And the more powerful, the better. Throughout its first 50 years, *IEEE Spectrum* has benefited in no small measure from the support of the great, the famous, and the well-connected.

Sometimes they agreed to be interviewed about a sizzling topic of the day; sometimes they shared a personal recollection of a landmark breakthrough; sometimes they pulled strings behind the scenes for us. Every now and then they tipped us off to an upcoming event of seismic consequence. From time to time, they even wrote articles for us. Indeed, the roster of dazzling bylines in *Spectrum* is long.

IEEE SPECTRUM 1964 / 2014

Members of the trio associated with the establishment of Intel–Robert Noyce, Gordon Moore, and Andrew Grove–are listed as authors or coauthors on four different *Spectrum* articles. (Moore, however, contributed his most famous magazine article–the 1965 piece that introduced his eponymous law–to *Electronics* magazine, our archrival in those days.) William Shockley, one of the three inventors of the transistor and a former boss of Noyce and Moore, wrote a turgid feature we published in 1966 called "Articulated Science Teaching and Balanced Emphasis." The article is interesting today mainly for a passage in which Shockley recalls a protracted argument, at age 15, with a science teacher concerning the best way to calculate the forces on a rowboat.

The other two inventors of the transistor, John Bardeen and Walter H. Brattain, never wrote for us. But Bardeen shared a touching anecdote with my colleague Karen Fitzgerald in 1988. He told her about the time he disclosed to his wife, in four words, the achievement of which he was part, which happened to be the greatest invention of the 20th century. As his wife prepared dinner one evening in 1947, their three children scampering around them, Bardeen said, "We discovered something today."

"That's nice, dear," his wife replied.

In 1973 we published a feature article by Dr. Strangelove himself. Edward Teller, the man behind the hydrogen bomb, contributed an article with the clunky title "A Future ICE (Thermonuclear, That Is!)." Based on a conference talk, the story sketches out the basic physics of controlled thermonuclear fusion, for both power generation and space propulsion. In an odd but revealing sidebar, Teller renders this opinion: "I reject any expressions like 'human engineering.' These two words in this context don't go together. Problems in engineering are solved by really different processes–often by logical processes. Whereas human problems are very often solved, and rightly solved, by illogical processes." You read it here first, folks.

John R. Pierce, an early pioneer of pulse-code modulation and colleague of the trio at Bell Labs that invented the transistor, contributed no fewer than four feature articles to *Spectrum*. Pierce had one of the most brilliant and productive engineering

careers of the 20th century. He improved on the traveling-wave tube, which had been invented in Britain during World War II by Rudolf Kompfner, who himself published a feature article in *Spectrum* in 1967 titled "Electron Devices in Science and Technology." With writer Arthur C. Clarke, Pierce is credited with conceiving the communications satellite, and he later led the project that built the first commercial one, Telstar 1. What's more, Pierce published science fiction novels and, at the end of his career, was a music professor at Stanford, where he discovered a new kind of musical scale.

Pierce was also known for his wit: "Funding artificial intelligence is real stupidity," was one of many memorable quotes. An article he wrote for *Spectrum* elaborated on that very notion. "Many early computer enthusiasts thought that computers should resemble human beings and be good at exactly the tasks that human beings are good at," he wrote. He likened that approach to "designing an airplane that will light on a tree. It is facing the future with one's back squarely towards it." Nowadays, with enthusiasm high for the technological singularity and its promise of machine consciousness, I long for Pierce's authority, clear-eyed realism, and pungent humor.

Luminaries continue to appear in the pages of *Spectrum*. Best-selling author and Harvard Business School professor Clayton M. Christensen contributed to two memorable articles, one on the microprocessor business and the other on how the Toyota production method could be applied to semiconductor manufacturing. In 2008, not long after Hewlett-Packard researcher R. Stanley Williams and his team succeeded in building the long-hypothesized memristor, Williams described that work in a feature article for us. Nathan Myhrvold, a former chief technical officer of Microsoft and a cofounder of the huge patent holding company Intellectual Ventures, has contributed to two *Spectrum* features. The first, published in our June 2013 issue, described efforts to build a food compositor, which would assemble edible "elements" into delectable custom-made meals. His second article, published just last month, described why even costly modern ovens perform poorly and how they could be improved.

And then there was the article that got away. In 2006, a publicist for Microsoft Corp. sent us a manuscript and informed us that the byline could list as authors Bill Gates and one other Microsoft executive. However, when we showed the executive an edited version of the article, which was about a robotics initiative, the company balked and politely withdrew the draft. We next saw the piece in the pages of *Scientific American*, some months later. Predictably, it did not seem to have been substantially edited.

During my own career at *Spectrum*, I've been grateful many times for the help I received from the powerful or famous. I have interviewed U.S. senators, congressmen, and military generals; British diplomats; and government ministers in India and Iraq. I spent a blustery day sailing around San Francisco Bay and drinking ancient port with Silicon Graphics founder and Netscape cofounder Jim Clark, in his prebillionaire days. I had dinner with Jon Rubinstein around the time he was leading the design of the iPod, and I've spoken with Gordon Bell, best known for his pioneering work on minicomputers at Digital Equipment Corp., more times than I can remember. In Afghanistan, I interviewed the governor of Kandahar province in his palace, days after riots a few miles away killed 16 people.

At the fabled Washington, D.C., eatery Duke Ziebert's, I ate alligator with an astronaut: Owen Garriott, who was trained as an electrical engineer and once held the record for the longest stay in space. Ben Rich, the legendary engineer who oversaw the design of the F-117A stealth fighter-bomber and helped design the U-2 spy plane, regaled me with dirty jokes during a freewheeling interview a year or so before he died.

The list goes on. I don't have the space (and you don't have the time) for an exhaustive acknowledgment of all the stellar and powerful friends who've supported *Spectrum* over the years, but I will pull a few more names out of the ether. Alfred N. Goldsmith, the radio pioneer and educator, and Bernard Oliver, who became R&D chief at Hewlett-Packard, helped us get off the ground in 1964. Jack Kilby and J. Fred Bucy, of Texas Instruments, were also reliable supporters in the 1960s and 1970s, as was Clarence Lester "Les" Hogan, another electronics-industry pioneer who combined engineering and executive brilliance. Michiyuki Uenohara, senior executive vice president of NEC, and Tsuneo Nakahara, vice chairman of Sumitomo Electric Industries, were our steadfast champions in East Asia; in a culture where personal connections are crucial, they introduced us to sources we otherwise had no way of reaching. Eberhardt Rechtin, CEO of the Aerospace Corp., Norman R. Augustine, CEO of Lockheed Martin, and Malcolm Currie, CEO of Hughes Aircraft, enabled us to bypass the bureaucracy and get aerospace and military stories we wouldn't have gotten otherwise. Lewis Terman, who was a leader in IBM's research establishment and the son of Frederick Terman, continues to advise us to this day.

Trying to make computers resemble human beings was "facing the future with one's back squarely towards it," Pierce wrote

All of these people went out of their way to help us when they surely had plenty of other claims on their attention. The only possible explanation for their generosity is that they believed in us and in our mission. They expected nothing in return. But, hopefully, this column is better than nothing. —GLENN ZORPETTE

NEWS

2,500: NUMBER OF PERSPECTIVES FROM WHICH YOU CAN VIEW A 3-D IMAGE USING OSTENDO'S QUANTUM PHOTONIC IMAGER CHIP

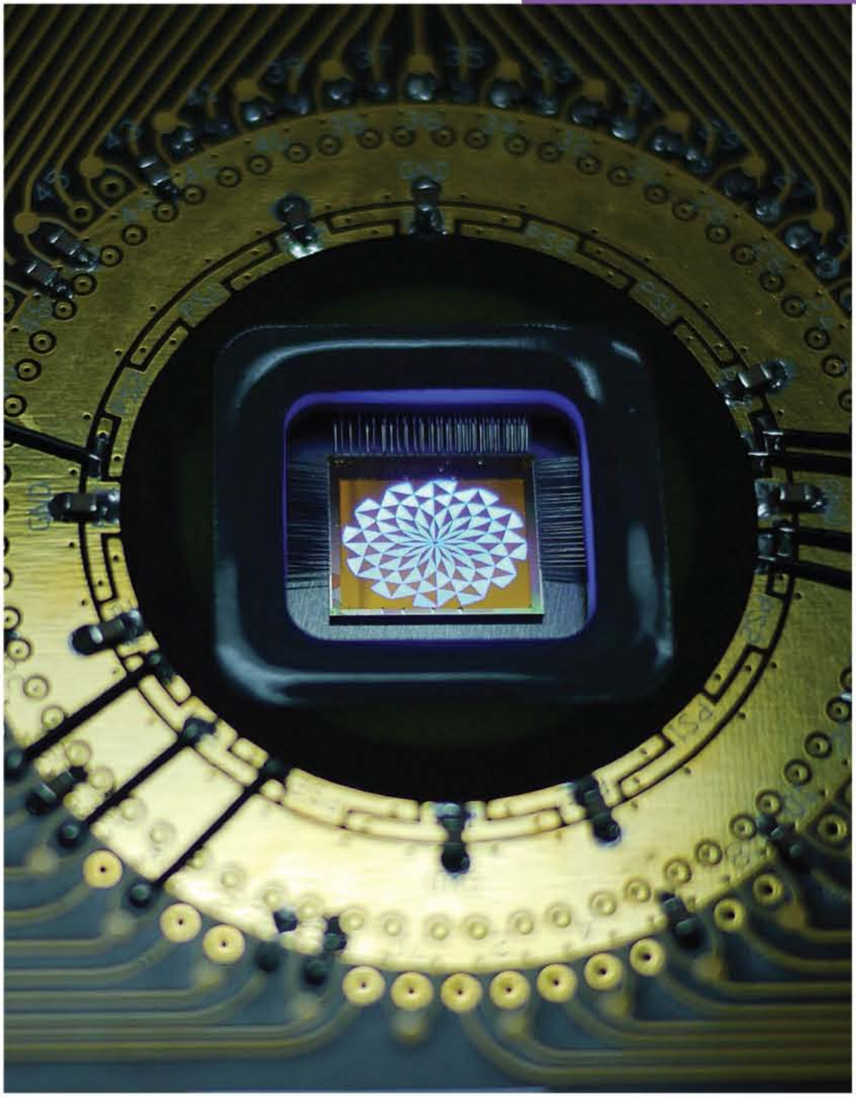

HOLOGRAPHIC DISPLAYS COMING TO SMARTPHONES

Light-field displays for mobile devices might be only a year away

Looking at a stereoscopic 3-D display takes some mental gymnastics. When you look at objects in real life, your brain expects the area of focus to be the same as where your eyes need to converge. But in order to see in stereoscopic 3-D–in which a different image is presented to each eye–you focus on the screen but your eyes converge where the image appears to be. For some, this is a headache-inducing dilemma.

Holograms get around that by projecting light right to the spot where your eyes would focus: The light beams travel through that point and hit your eyes just as if they'd come from an object that was actually there. Even better, holograms work from any angle and don't require glasses. Up until now this type of display has been a weighty affair, requiring large projectors and screens or a very restricted viewing angle. But two companies, Ostendo Technologies and Hewlett-Packard spin-off Leia, promise to put such holographic displays–more properly called light-field displays–in your pocket within a year or two. It might not be Princess Leia projected from an astromech droid, but it's close. »

QUANTUM PHOTONIC IMAGER: Ostendo's light-field display produces 3-D images using light-emitting pixels and piles of pixel-level processing.

OSTENDO TECHNOLOGIES

NOT-SO-FUZZY DICE: An array of eight of Ostendo's Quantum Photonic Imager chips made a 3-D image of a die seem to fly. Such light-field displays for smartphones could arrive in a year or two.

OSTENDO TECHNOLOGIES

At Display Week in June, Ostendo demonstrated the culmination of nine years of work, an array of eight Quantum Photonic Imager (QPI) chips in a grid projecting three spinning green dice—one seemingly floating behind the display, one at chip level, and the third in front of the chips.

"Almost every display you see emits light that goes everywhere," says Hussein El-Ghoroury, Ostendo's CEO. "In contrast, the QPI collimates the light to a very narrow angle before emitting it, so you can emit different images in different directions." Ostendo's 3-D images are viewable from 2,500 perspectives.

Each of the 1 million pixels on Ostendo's little chip consists of a layer each of red, green, and blue micro-LEDs (or lasers, in some iterations) sitting on top of its own small silicon image processor. The pixels are between 5 and 10 micrometers on a side. By modulating the power to the individual layers, each pixel can send out any color of light in a thin, focused beam. Multiple vertical waveguides carry the light out from the layers and modulate its direction—although company representatives won't specify exactly how—and an array of microlenses focus and direct the beam further. Having an image processor under each pixel saves power and lightens the overall computational load, which is considerable for complex images because they must be simultaneously rendered for viewing from thousands of different perspectives.

"You need to pack a lot of pixels into a small area to create this light-field effect," says Martin Banks, a professor of optometry and vision science at the University of California, Berkeley, who uses light-field displays in his research on vision. (Banks helped with Ostendo's application for a U.S. government grant by evaluating the QPI chip's capabilities.) "That's what's promising about their technology—small, low-power-consumption devices that can generate a lot of light with very tightly packed pixels."

Banks says that along with the high computational load for such a display, another challenge is the geometry of the display itself. Manufacturing microlenses for placement in front of light-field display pixels is difficult because the shapes and positions must be just right to steer the beams at the correct angles. And with so many viewing perspectives to produce, resolution can be a real problem, says Gordon Wetzstein, a researcher at the MIT Media Lab, who is working to address the issues of data volume and resolution that affect all small-scale, glasses-free 3-D technology. "It's really hard to give multiview or light-field images at high resolution," he says. "If you want to have 10 views, each of the views is 10 times lower resolution than the original display." Wetzstein is developing software to weave different views together so that 3-D displays will be less computationally intense and won't have to sacrifice as much resolution.

"Everybody wants to put 3-D on a smartphone," says David Fattal, founder of Leia, a light-field display start-up. "And customers aren't going to want to compromise between a holographic 3-D phone that has mediocre 2-D performance and a normal phone. They're going to want the best of the best."

Leia's 3-D display works by putting a grid of gratings behind an ordinary LCD. The gratings point the light beneath them in different directions, creating up to 64 different viewing angles for a 3-D image or video. Fattal's goal is a system that would be easy to scale up and integrate with existing screens or even with transparent displays. The company is planning its first commercial product for 2015.

The light-field display "is going to be the next big revolution in displays," says UC Berkeley's Banks. "But the public won't accept it if it doesn't have good color, doesn't look high resolution. Right now, it's not feasible with cost-effective equipment. But computers are getting faster, and people are getting smarter and learning shortcuts, so someday that will be achievable."

"That day will come," says El-Ghoroury. "Maybe two to three years, not much longer."

—SARAH LEWIN

It's really hard to achieve multiview or light-field images at high resolution. If you want to have 10 views, each view has one-tenth the resolution of the original display

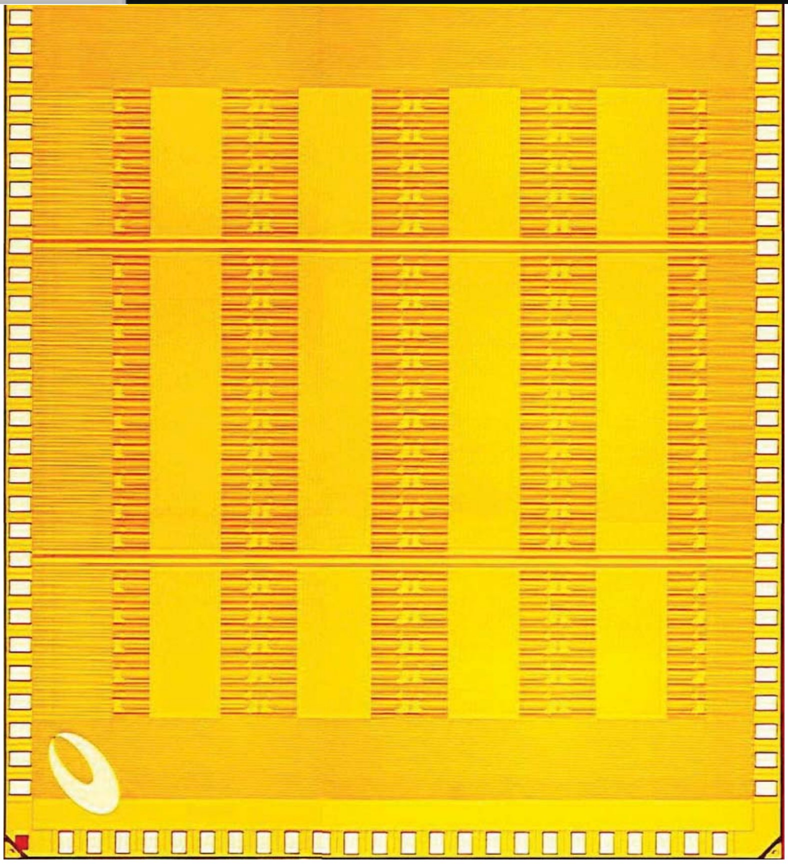

SPIN'S GOLDEN MOMENT: TDK-Headway's 8-megabit test chip stores bits in microscopic magnetic pillars and can be written to at record speeds.

TDK-HEADWAY TECHNOLOGIES

SPIN MEMORY SHOWS ITS MIGHT

Spin-transfer-torque MRAM could edge out some mainstream memories

The read head of a hard-disk drive might seem an unlikely place to hunt for the future of memory technology. But TDK-Headway Technologies, in Milpitas, Calif., is betting that the lowly magnetic tunnel junction—the device it makes to read data off hard-disk platters—could be redesigned and repackaged to create a new way of storing information.

Magnetoresistive random-access memory, or MRAM, has undergone a few incarnations already. But TDK-Headway and a number of other companies are now converging on a scheme they say could upend the memory business. Dubbed spin-transfer torque (STT) MRAM, it promises speed and reliability comparable to that of static random-access memory, or SRAM—the quick-access memory embedded inside microprocessors—along with the "nonvolatility" of flash, the storage of smartphones and other portables.

In June, at the 2014 Symposia on VLSI Technology and Circuits in Honolulu, Guenole Jan, who heads up characterization and magnetic design for TDK-Headway's MRAM team, presented results on the company's new 8-megabit STT-MRAM test chip. The capacity of the chip itself—roughly what would be expected from a small patch of embedded flash—was not intended to turn heads.

Instead, the team had a few other things to show off. One was the write speed, which stands at 1.5 nanoseconds. That's fast enough, Jan says, to compete with the SRAM that takes up most of the memory space on a modern microprocessor: the level-three cache. A second item was the team's manufacturing process, which is used to make an array of memory cells that can withstand some of the critical last steps of the chipmaking process, which demand temperatures as high as 400 °C.

This is quite a feat for STT-MRAM, in which each cell is made from a magnetic tunnel junction, a delicate pillar of layered materials. The junction usually consists of more than 10 different layers, says Gill Lee, a senior director in the silicon systems group at Applied Materials, in Santa Clara, Calif.; some of them can be less than a nanometer thick. When these layers are etched into pillars, Lee notes, some material can be ejected and then redeposited along the sides of the pillar, shorting out or damaging the device.

STT-MRAM is a lower-energy alternative to an earlier version of MRAM, which used current flowing through wires to create a magnetic field that could set the magnetization of a bit in either one direction or another. In STT-MRAM, the bit flipping is done by passing a current through the magnetic tunnel junction. Applying voltage draws current through a "reference layer," which is often made up of two oppositely magnetized layers. The electrons flowing through the stack align their spins to match the direction of magnetization of each layer. The spin-aligned electrons can then tunnel through an insulating barrier into a "free layer." There, they "torque" the spins of the free layer's electrons, imparting their spin to them and magnetizing that layer. A similar process of spin selection occurs when the voltage is reversed; electrons enter from the free-layer side of the stack, bounce off the insulating layer, and become polarized, causing spins in the free layer to point in the other direction. The two different states can be detected by running current through the device.

The stack will have low resistance when the magnetization of the free layer and topmost part of the reference layer point in the same direction and high resistance when they point in opposite directions.

Since SRAM takes six transistors to make a cell and STT-MRAM needs just one, it could potentially be used to make a more compact working memory. But the technology is likely to be used earlier as an energy-saving alternative to small-capacity embedded flash, which is used to store vital information such as network keys on mobile chips and chip-based sensors.

The mobile chip giant Qualcomm has also been working on STT-MRAM, and in a recent test of TDK-Headway's chips, it found no errors in data retention after 528 hours at 150 °C. Seung Kang, who leads R&D on advanced memory technologies at Qualcomm's headquarters in San Diego and reported the result at the symposia, says his company is less interested in replacing existing memories than it is in revamping the basic memory structure of the system-on-a-chip. "We call it nonvolatile working memory," Kang says. "That's kind of a new class." One potential application of the technology, he says, is as a form of both storage and working memory in future wireless sensors, for which energy consumption and longevity are priorities.

Other companies are targeting stand-alone memory chips. At the same symposia in Honolulu, Samsung reported success in making cells that could potentially be fabricated using a 15-nanometer manufacturing process, the state of the art today. That could put STT-MRAM in line as a replacement for dynamic RAM computer memory.

In fact, STT-MRAM is already making small inroads on that front. Everspin Technologies, a spin-off of Freescale Semiconductor based in Chandler, Ariz., has shipped samples of a 64-Mb chip that uses a common DRAM interface.

Everspin's first STT-MRAM chip is built with 90-nm-node technology, which puts it generations behind cutting-edge memory chips in terms of density. But Joe O'Hare, the company's director of marketing, says that Everspin thinks its chip could find a home as a form of "persistent DRAM" that could afford greater protection against power outages to data centers and storage devices. The company expects to ramp up sales early next year.

Everspin hopes to boost the capacity of its next chips, and it is looking into switching to magnetic tunnel junctions in which the magnetization runs perpendicular to the layers instead of parallel to them. This architecture, which is now being pursued by Samsung, TDK-Headway, and others, has been shown to be more energy efficient and easier to miniaturize. "There's a lot on our to-do list," says O'Hare.

STT-MRAM may be claiming some of the enthusiasm once reserved for other alternative memories, such as ferroelectric RAM, phase-change memory, and resistive RAM. But its success will come down to manufacturing technology and how well it can compete on cost.

Hard-disk read-head tools make just one magnetic tunnel junction at a time and so must be adapted for high-throughput operation. "The issue is making a million or a billion of these devices on a chip and having good yields and identical properties on all devices," says TDK-Headway's Jan.

"I think [STT-]MRAM has all the desired features—speed, nonvolatility, low power, low energy, almost unlimited endurance," says Lee of Applied Materials. But, he adds, "there's still work to do."

—RACHEL COURTLAND

STT-MRAM offers the speed and reliability of SRAM and the nonvolatility of flash. Some engineers call it nonvolatile working memory

STASHING ENERGY IN UNDERWATER BAGS

Submerged bags of air could turn wind and solar power into round-the-clock resources

With the worldwide proliferation of wind- and solar-generated power, the fickleness of these renewable sources is a problem crying out for a good solution. A Canadian start-up called Hydrostor thinks it has an answer: air-filled bags.

This month, the Toronto company plans to sink several large balloonlike bags into Lake Ontario, and then, using electricity from Toronto Hydro's grid to run a compressor, it will fill the bags with air. Later, when the utility needs electricity, the air will be emptied from the bags and run through a turboexpander, which uses the expanding air to drive a turbine. The result will be the world's first commercial facility for underwater compressed-air energy storage.

Using compressed air to store energy is not a new idea. The first such systems emerged back in the 1870s, and these days compressed air is stored in underground caverns, in pipes, and even in small tanks for powering cars and locomotives. Variants of the underwater storage idea have also been floated, so to speak, since at least the 1980s, says Seamus Garvey, a professor of dynamics at the University of Nottingham, in England. Garvey, who's not affiliated with Hydrostor, designed an underwater storage system using Thin Red Line Aerospace's bags and deployed a prototype off Scotland's Orkney Islands in 2012. "The idea is to put the storage where it matters most, which is where the intermittent energy is being generated from offshore wind," Garvey says.

Hydrostor CEO Curtis VanWalleghem says his company began looking at the technology four years ago as a side project to a wind farm it wanted to develop. At first, the company planned to use pumped hydro storage, in which water is pumped uphill and then released later to reclaim the electricity. Pumped hydro can have efficiencies of 80 percent or higher, but it works only in certain geographies and isn't economical on a small scale. "So we thought, if lifting a cubic meter of water into the air is the best way to store energy, maybe the

TOP: KEITH THOMPSON/THIN RED LINE AEROSPACE; BOTTOM: BRIAN CHEUNG

IN THE BAG: Energy bags like this 5-meter-diameter one, from Thin Red Line Aerospace, of Canada, could be used to store electricity underwater as compressed air. Engineers hope the technology could one day smooth out the intermittency of electricity produced by offshore wind farms and other renewable energy sources.

reverse would also work—submerging air below water," says VanWalleghem.

The concept is simple enough: When the energy bag is anchored underwater—at least 25 meters deep and ideally 100 meters or more—the weight of the water naturally pressurizes the air, allowing more air, and thus energy, to be stored in a given volume. (The pressure increases roughly 1 atmosphere, or about 100,000 pascals, every 10 meters.) At depths greater than 500 meters, says Garvey, "the cost of the containment becomes negligible compared with the costs of the power-conversion machinery."

In the Toronto system, the bags (or "flexible accumulators," as Hydrostor calls them) will be deployed at a depth of 80 meters, and they should be able to supply about a megawatt of electricity for 3 hours or so. The company will also be testing fixed-wall accumulators, in which the compressed air will displace water inside the vessel. "This is the smallest size we would contemplate," says VanWalleghem. A more typical capacity, he says, would be 20 to 30 megawatts that can be discharged over 10 to 20 hours. Eventually, the company will aim for an efficiency of about 60 to 70 percent. The technology easily scales up, he adds. "We just make the air cavity bigger, so there really is no upper limit." By year's end, the company plans to build a bigger and deeper underwater energy storage facility in Aruba.

One key challenge Hydrostor faced was how to capture the heat given off when air is compressed and then use it later to warm the air as it cools during expansion. "The air can reach temperatures of 650 °C during compression, so if you're not judicious in capturing that waste heat, you lose efficiency," explains Rupp Carriveau, an associate professor at the University of Windsor, in Ontario, Canada, who advised Hydrostor early on. The solution they settled on was an off-the-shelf heat exchanger coupled with an insulated water bath.

DRY RUN: In 2011, Toronto start-up Hydrostor tested its underwater compressed-air energy-storage system in Lake Ontario. This month, it plans to deploy a commercial version, the world's first.

In fact, Hydrostor has tried to use existing components wherever possible. VanWalleghem explains why: "Reliability is very important for utilities. You need them to be comfortable with the technology." Off-the-shelf equipment already has rated life spans in the field, which Hydrostor's partners and investors found reassuring. "The downside is that you have to live with what's available," VanWalleghem says. "But it's worth it in terms of speed to market and not having to design and build everything from scratch." Hydrostor wouldn't disclose the exact cost of the Toronto system, but VanWalleghem says it's "in the numerous millions of dollars."

Although the technology is still new, the need for this kind of energy storage is obvious, says Carriveau. Much of the world's population lives near a coast, he notes. "So that's your load. And because of the losses you get during transmission, it follows that you want to keep your energy generation and your storage as close as possible to your load."

Garvey sees the underwater storage as part of a holistic system. "An offshore wind farm should not simply be a subsystem that produces electricity when the wind blows. It should be a system which takes energy from the wind and does whatever is needed to deliver energy to shore as that [energy] is needed."

The energy bags, he says, "are one very possible step toward that utopian view."

—JEAN KUMAGAI

THE PRACTICAL POLARITON LASER

New design could be key to on-chip optical interconnects

Lasers have largely replaced the copper wiring of old to speedily stream bits among servers in large data centers. Engineers have long dreamed that they could use lasers to do the same within the servers' data-congested processors, too. One problem standing in their way is that even the most efficient lasers need an awful lot of current before they light up. But a new kind of device, the polariton laser, could light the way using merely a trickle of electrons. Until recently, however, the device has worked only in an impractical environment.

Now, researchers at the University of Michigan have built the first polariton laser that both runs on electricity and operates at room temperature. It emits coherent light when provided with a current of only 169 amperes per square centimeter. A similar structure operating as an ordinary laser would take more than 250 times as much current, and even the best gallium nitride laser, enhanced with quantum dots, requires at least 1000 A/cm^2 before it starts lasing.

The difference comes from the way the device works. In a conventional laser, incoming current raises electrons to a higher energy state; once a majority are in that state, they can suddenly drop to a lower state and emit photons, an action known as stimulated emission. In contrast, a polariton laser works via the stimulated scattering of polaritons—quasiparticles that combine an electron, a hole, and a photon and that can exist only within a crystal. When energy is pumped into the system, the polaritons absorb it, then quickly release it as photons. Unlike with electrons in conventional lasers, the majority of polaritons don't have to be excited before lasing can begin, so the so-called lasing threshold—the energy required to start the process—is much lower.

"For 50 years, we've had one kind of coherent light source—the laser," says Pallab Bhattacharya, professor of electrical engineering at Michigan, who reported the work in June in *Physical Review Letters*. "This coherent light source is based on a completely different physical principle."

Polariton lasers were first described in 1996. In the years that followed, researchers developed several versions that used the light from other lasers as an energy source. Last year, Bhattacharya's group and a separate team of researchers—from Stanford; the University of Würzburg, in Germany; and other schools—both showed they could make electrically pumped versions, an important step in turning any laser from a laboratory curiosity into a practical device. Both those devices, however, operated only at extremely low temperatures, about 30 kelvins. This new laser is the first to run on electricity and at room temperature.

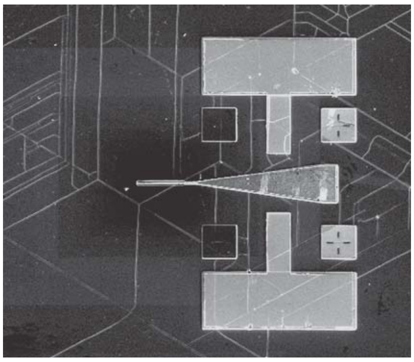

REFLECTED GLORY: University of Michigan engineers built a polariton laser [triangle] that runs on electricity and operates at room temperature. The key was to lower the resistance of the device's reflectors.

UNIVERSITY OF MICHIGAN

"Nothing prevents commercialization of polariton lasers now," says physicist Alexey Kavokin, head of the nanophysics and photonics department at the University of Southampton, in England, who was not involved in the work.

The difference in the physics also means the device can be switched on and off much faster than conventional lasers. That capability for fast modulation, plus the low power requirements, points to a use in optical communications, such as across a computer chip, or as part of an optical memory system.

Bhattacharya says that success with this laser came from a fundamental change to the design. Previous attempts used a vertical structure, in which the lasing material was sandwiched between two sections of distributed Bragg reflectors, with the light eventually emerging from the top of the laser. The reflectors were alternating layers of material with different indices of refraction that reflected photons back into the lasing cavity. The structure required 20 to 25 pairs of alternating layers, however, which increased electrical resistance. Pumping in more energy to overcome the resistance disrupted the polaritons and destroyed the lasing effect.

This time, Bhattacharya used a horizontal lasing cavity 690 nanometers long, not much bigger than the wavelength of the light it would emit. The cavity is capped at either end with distributed Bragg reflectors built of alternating layers of silicon dioxide and titanium dioxide, five on one end and six on the other. The gallium nitride lasing material sits on top of a layer of indium aluminum nitride, which acts as a barrier, preventing photons from leaking out of the bottom; air acts as a barrier on the top. Instead of the kilo-ohms of resistance in the previous design, this device has only 10 ohms. "That is really the turning point," Bhattacharya says.

The new laser is very simple to fabricate, according to Bhattacharya, which he says could speed commercial versions. He's experimenting with other materials to get emissions at longer wavelengths; this laser emits in the ultraviolet, whereas most optical communications rely on the near infrared. But, he says, for chipmakers faced with a data bottleneck in processors and other complex chips, any working device would prove valuable, no matter what the wavelength. "If you put a laser on silicon, anybody will take any color they can get." —NEIL SAVAGE

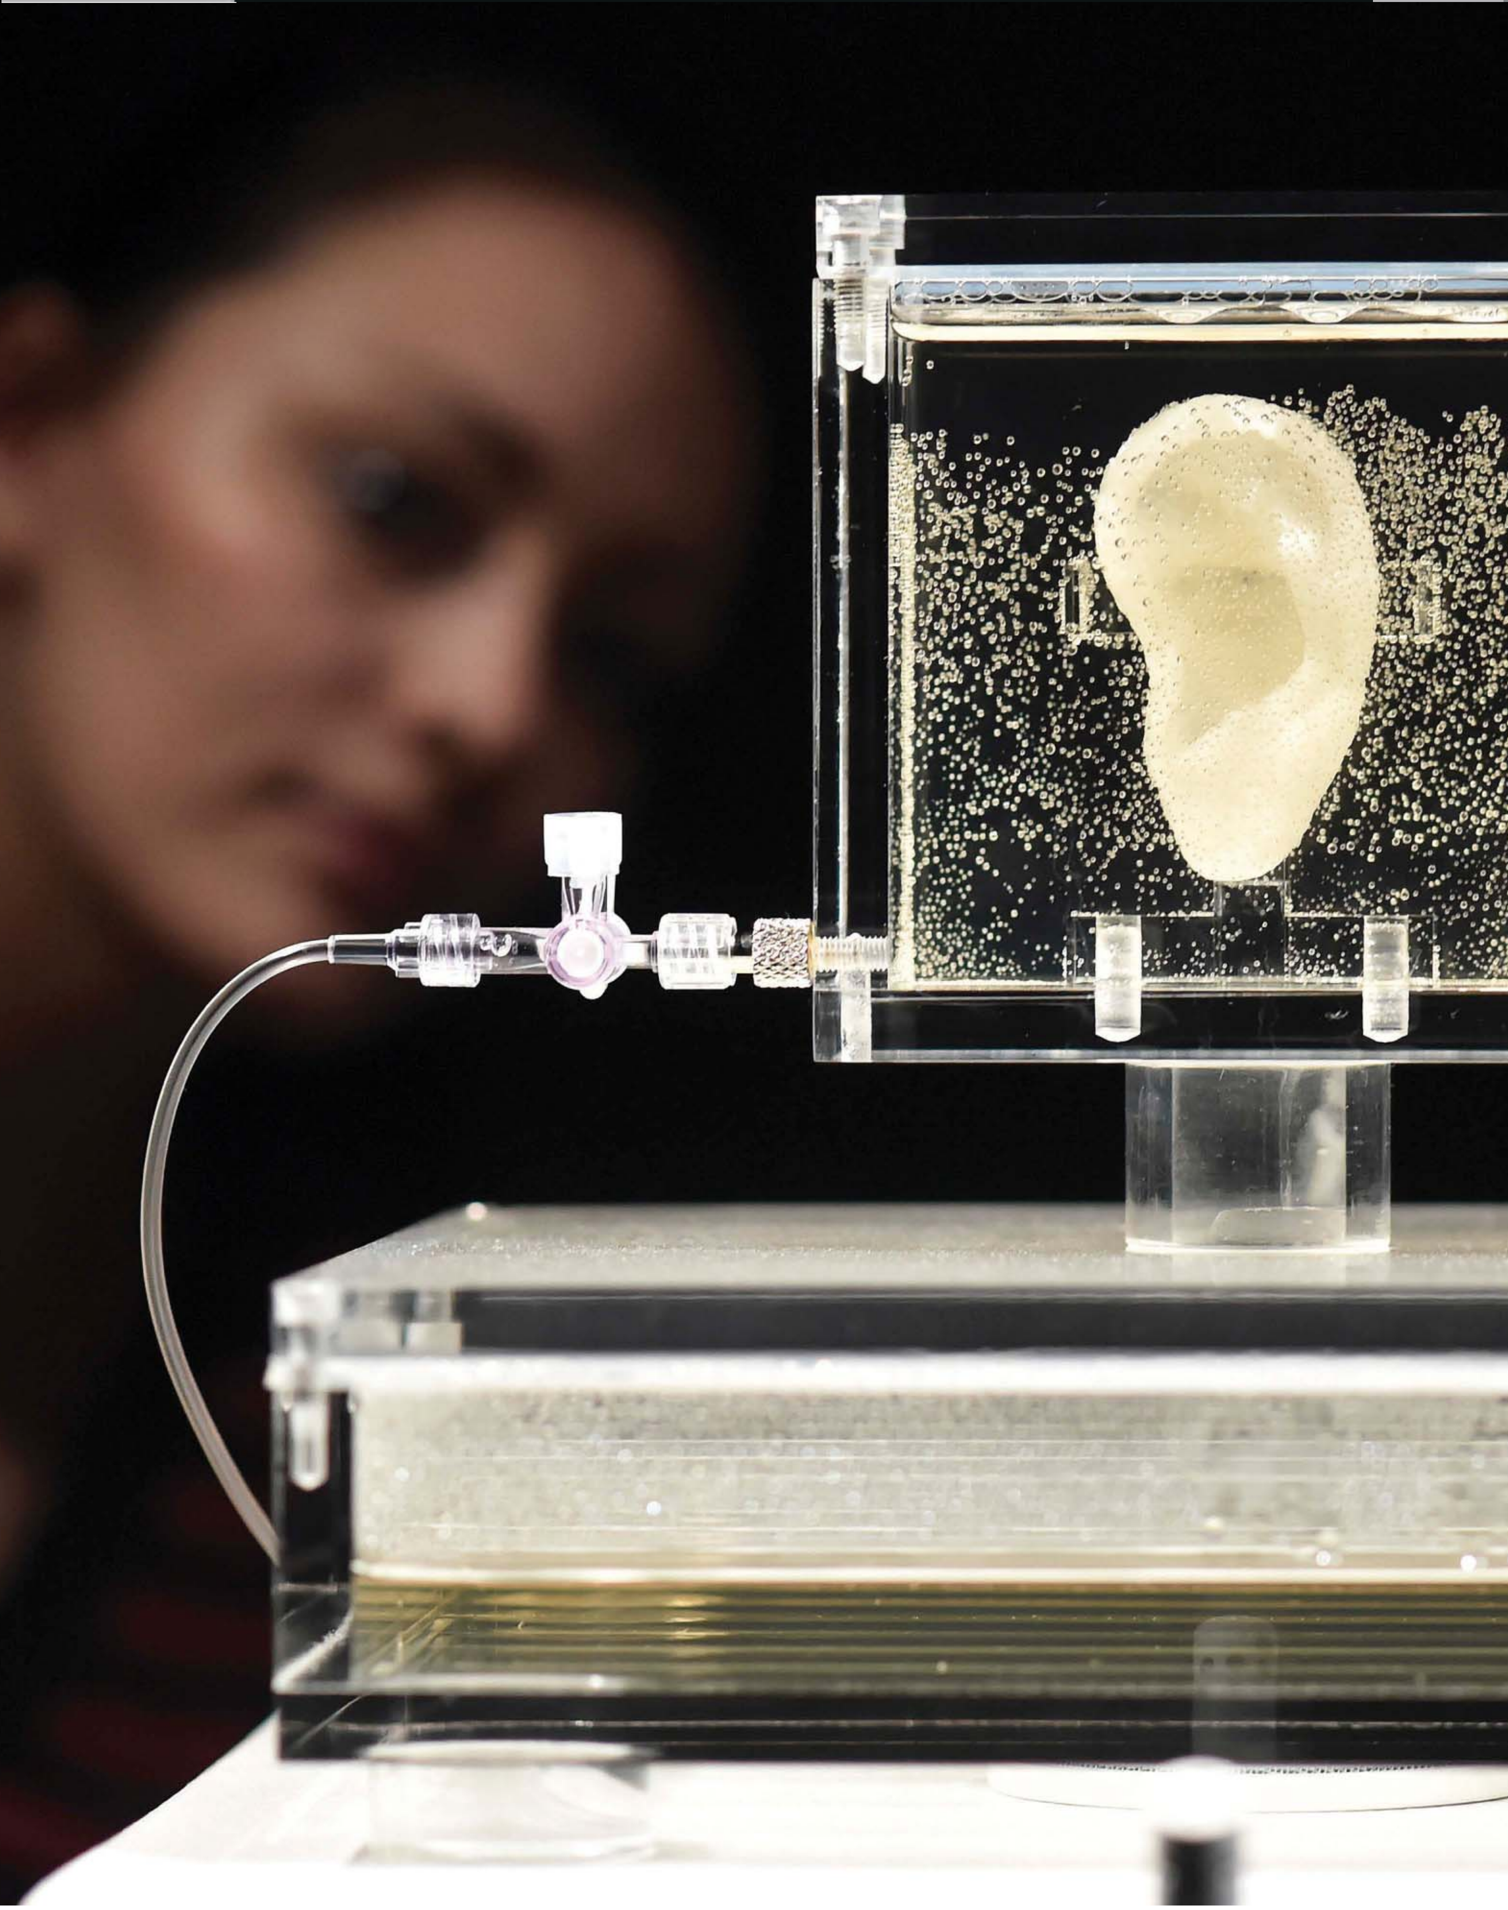

PHOTOGRAPH BY Uli Deck/AP Photo

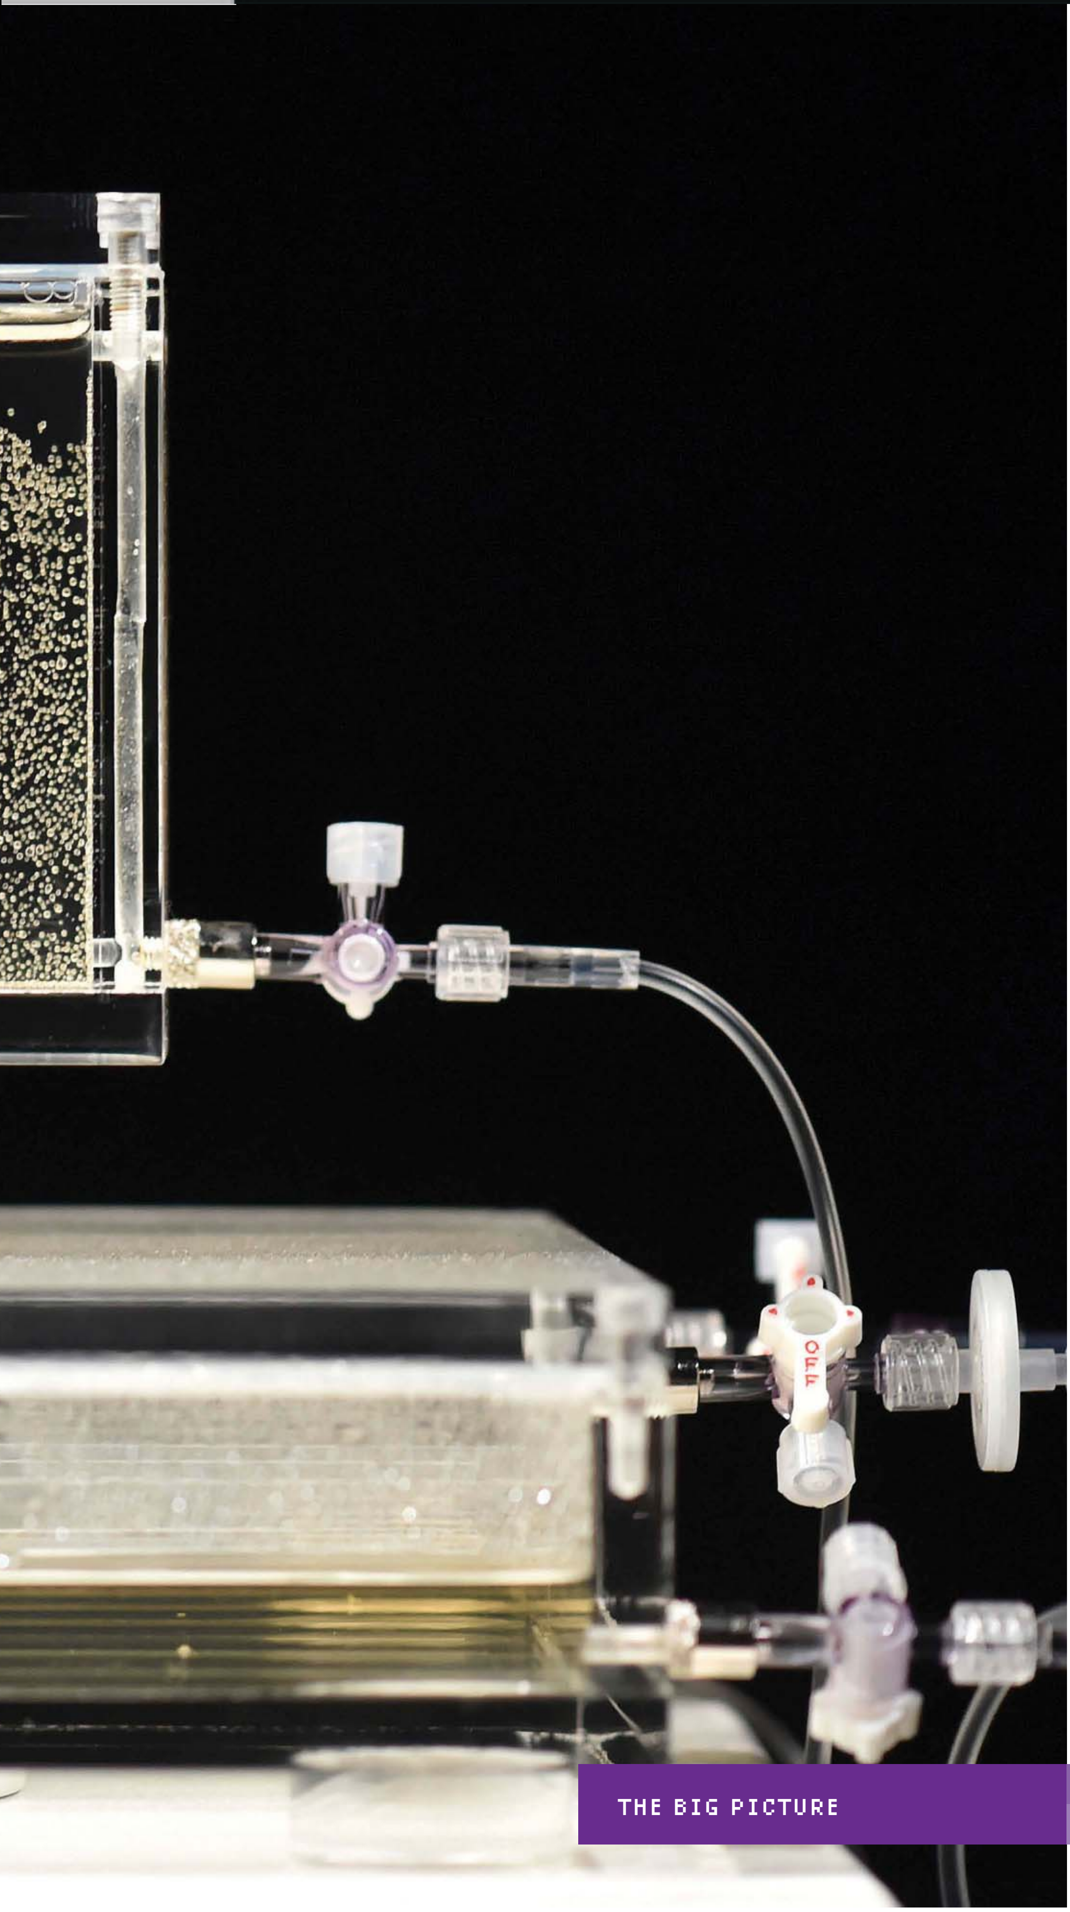

VAN GOGH'S NEW EAR

DUTCH ARTIST Vincent van Gogh is remembered nearly as much for (reportedly) severing his own left ear as for his dazzling paintings such as *The Starry Night* and *Still Life: Vase With Twelve Sunflowers*. Continued fascination with the missing appendage, combined with the latest in bioengineering, has led to another work of art: a 3-D printed replica of van Gogh's ear generated from tissue taken from a descendant of van Gogh's brother Theo. Artist Diemut Strebe featured the ear at an exhibition that ran through 6 July at the Center for Art and Media in Karlsruhe, Germany. The replica can even hear; it picks up sound with a built-in microphone and software that simulates auditory nerves.

RESOURCES

9.98 METERS: DIAMETER OF THE WORLD'S LARGEST DISCO BALL—UNVEILED IN MOSCOW IN 2012—ACCORDING TO GUINNESS WORLD RECORDS

RESOURCES_HANDS ON

DISCO FEVER

FEEDBACK AND ELECTROMAGNETISM KEEP THIS MIRROR BALL ALOFT

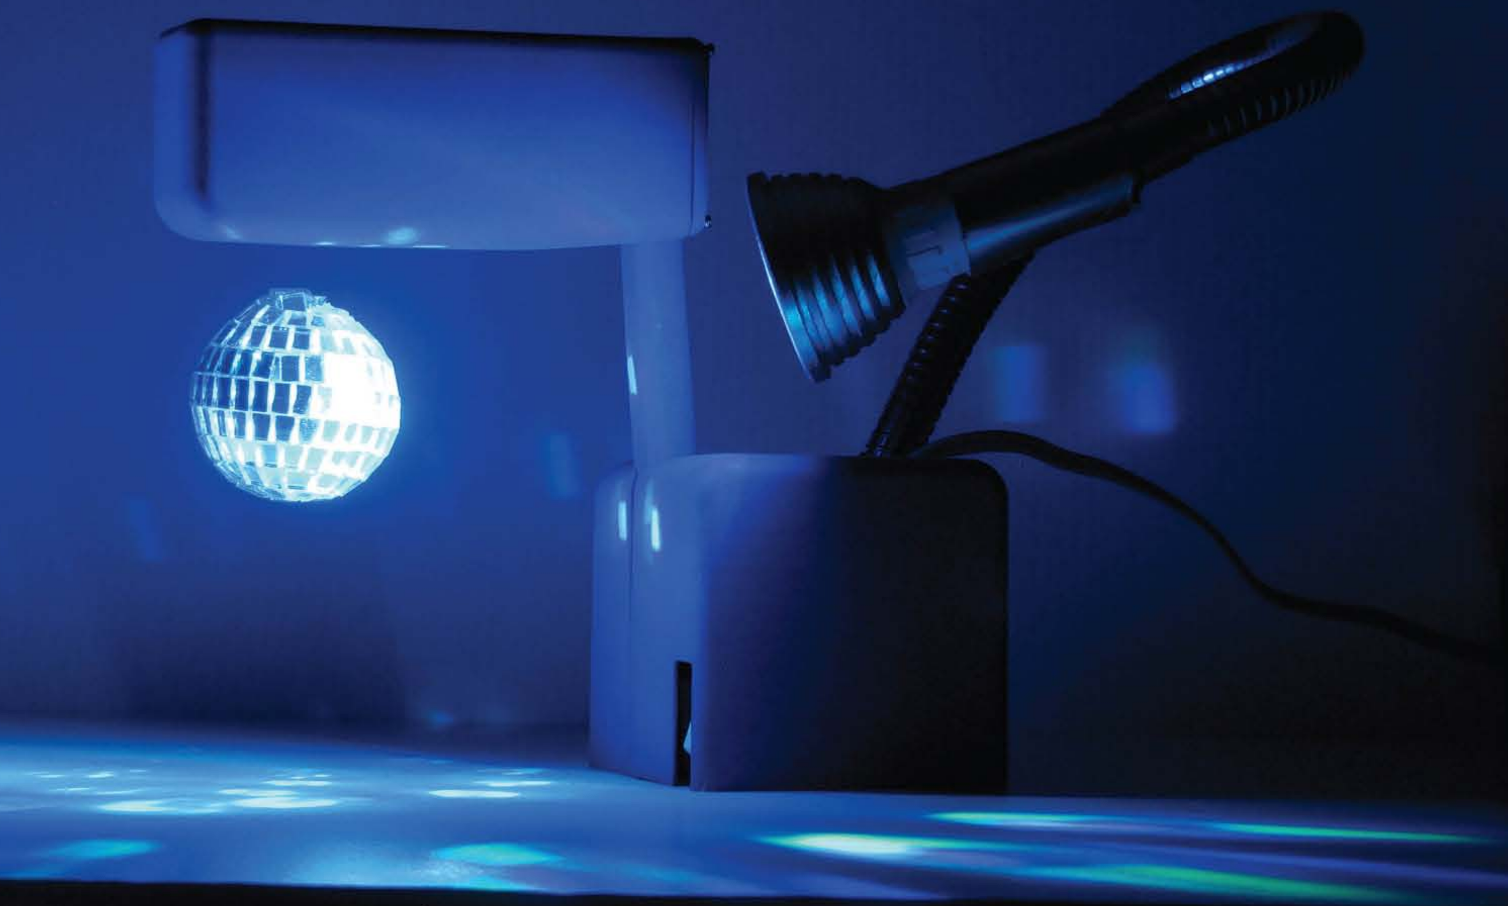

STAYIN' AFLOAT: Illuminated by a programmable multicolor LED, a miniature disco ball is magnetically levitated to create a kaleidoscopic night-light.

IN 1979, WHEN I WAS 11 AND ENTHRALLED BY ***Star Wars*****, magnetism held a special appeal. It was the closest thing in the real world to "the Force." So for our fifth-grade science project, a friend and I wrapped my dad's iron chisel with wire to make an electromagnet. When connected to a lantern battery, it grabbed, repulsed, and spun bar magnets like magic. But when we got curious and plugged it into a wall socket, the resulting pop and puff of smoke produced an unforgettable jolt. ● The experience sparked a lasting fascination with magnetism. So when my son Liam, contemplating what to make for his own fifth-grade science fair, suggested making a magnetic levitation gizmo, I was irresistibly drawn to the idea. ● The project turned out to be surprisingly challenging. Initially, we envisioned some simple circuitry that would make a miniature space ship from** ***Star Trek*** **(actually a Christmas-tree ornament) hover with no strings attached. How hard could it be?**

I found an electromagnet rated for 6 watts at an online store for US $15; a set of 10 neodymium button magnets was $10. I figured we'd just hang the electromagnet from a beam, glue one of the rare earth magnets to the center of gravity of the toy ship, and then manually adjust the current flowing to the electromagnet's coil until the levitating attraction between the magnets exactly balanced the pull of gravity.

While Liam and a friend built a gallows-like support structure out of Lego bricks, I began breadboarding the circuitry. It dawned on me that I ought to do a quick Web search for levitation circuits. And that's when I bumped my head against Earnshaw's theorem.

About 175 years ago, Samuel Earnshaw proved that wherever objects are producing static inverse-square forces—as with magnetism—there is no point in space where the forces balance out in a stable way. The tune-it-and-forget-it system I had envisioned would be impossible.

Fortunately, clever electrical engineering has triumphed where brute physics fails. More searching turned up designs for levitation circuits, all of them variations on a theme: Rather than fixing the power to the electromagnet, you use a position sensor to rapidly adjust it. As the hovering magnet nears the coil, the power sags; as it drops away, the electromagnet's strength is increased.

For our first attempt, we settled on a design that measures the proximity of the permanent magnet to the electromagnet by using Hall effect sensors mounted on the top and bottom of the electromagnet's coil. The electromagnet is powered by a pulse-width-modulated current. An op-amp chip subtracts the output of the two sensors to obtain a voltage that measures the strength of the field originating from the permanent magnet floating below the coil. A second op-amp compares that signal to a reference voltage that sets the distance at which the magnet should hover. Any difference between the signal and reference voltage is sent as an error signal to a KA7500C pulse-width modulator, which adjusts the power output to the coil at about 10 kilohertz to counter the error.

With the system built, we set up our table at the science fair. The coil was too weak to suspend anything other than the magnet itself (so no *Star Trek* spaceship), but it worked—almost. If I kept one finger gently touching the side of the permanent magnet, it hovered for a minute at a time. But without that stabilizing touch, tiny wobbles quickly became big wobbles and shook the magnet loose.

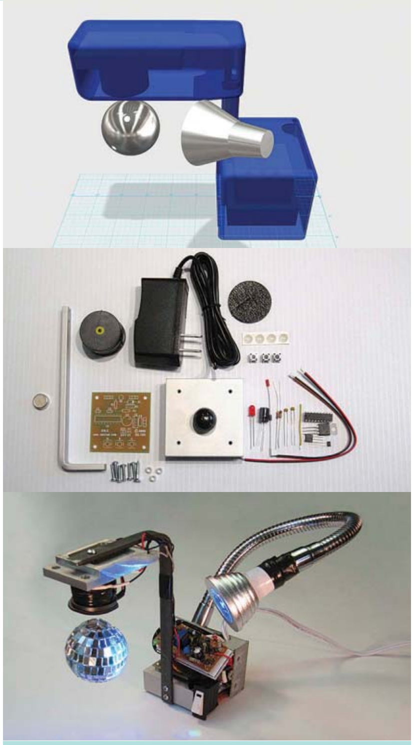

PUTTING IT TOGETHER: The night-light case was designed using computer-aided design software and 3-D printed [top]. The electromagnet that levitates the miniature disco ball is controlled by a kit [middle]. A microphone gooseneck was repurposed to support a multicolor LED that provides illumination [bottom].

Back home, I resolved to make it work sans stabilizing finger—and moreover to make something useful out of it. Other flashes of memory from the 1970s provided inspiration: I would make a far-out, levitating, color-changing, disco night-light. Yeah, baby!

A mini disco ball was easy to find (there seems to be no object that people won't shrink and hang on a Christmas tree). The mirror tiles are glued onto a Styrofoam sphere; I simply pried off some tiles, drilled out the core, inserted three neodymium magnets, and then replaced the tiles. A multicolored LED, complete with a wireless remote for selecting any of four groovy color-transition programs, was $4 (including shipping from Hong Kong). I mounted it to a microphone gooseneck and aimed it at the spot where the ball would hover.

None of the home-brew circuits I found online was strong, stable, and efficient enough to keep a 40-gram disco ball suspended all night

PREVIOUS PAGE AND THIS PAGE: W. WAYT GIBBS (4)

long without overheating or daily adjustment. But I found success with the Maglev Plus System, a kit sold for $73 by Zeltom.

Unlike the designs I had built from scratch, Zeltom's circuit uses only one Hall sensor, and it relies on a preprogrammed PIC microprocessor to dynamically adjust the coil power, rather than tying together a couple of op-amps to a pulse modulator. The PIC has built-in flash memory; a trio of push buttons lets you tune the circuit to the particular mass and distance you want to levitate and then store that setting. I worked through the day and into the night soldering everything onto the circuit board, hooking up the lamp, power supply, and a power switch, and using tin snips to fashion a rough case for the gadget out of thin sheet metal.

Around 2 a.m., I plugged it in and flipped the switch. Holding the mirror ball about 2 centimeters below the coil, I held my breath. I could hear Obi-Wan over my shoulder: "Use the Force, Luke. Let go." So I did.

It was like flipping on a funk generator. The room filled with swirling splashes of magenta shifting to tangerine and then cerulean. The effect was hypnotic.

When Liam padded into the kitchen, bleary-eyed, the next morning, the ball was still hanging in free space, gently twisting to and fro. Nothing was overheating, but the ugliness of the case bothered me. So I downloaded AutoCAD's free 123D Design program and sketched out a suitably high-tech-looking case. I uploaded the file to 3-D printing firm Shapeways, and two weeks later a sleek, blue, $140 plastic enclosure arrived on my doorstep. Like magic.

—W. WAYT GIBBS

PATENT TACTICS

PICK THE FILING STRATEGY THAT'S RIGHT FOR YOU

ISTOCKPHOTO

Typically, the first step in turning an idea into money is to patent it. Starting this process raises the questions "How much does a patent cost?" and "How long will it take?" Unfortunately, there are no fixed answers to these questions. Some factors beyond your control, such as bureaucratic backlog and the complexity of a patent, will affect cost and review time. But your choice of a filing strategy is one thing that will have a predictable and immediate impact.

As a business evolves, so does the weight given to competing concerns regarding cost, timing, and international patent protection. In the United States, this weighting guides the strategic choice between a provisional or nonprovisional patent application.

For some organizations—such as universities that produce many inventions or start-ups—delaying costs while exploring the market is paramount. A provisional patent application establishes the priority of an inventor's claim without requiring complete disclosure about its details, and it costs considerably less than a nonprovisional application. To maintain patent protection, the inventor must apply for a nonprovisional patent within one year. An unprofitable idea can be abandoned.

Going straight to a nonprovisional patent application is a more attractive option for companies with capital on hand, such as a start-up that has secured funding, as it reduces overall costs and the time to receive a patent.

To see how this works, imagine a small company that files a provisional application, and let's use some reasonable estimates for the attorney and drafting fees involved. The up-front cost is US $3,630 ($3,500 in attorney fees and $130 in patent office fees). One year later, the company feels the idea is viable, so it files a nonprovisional application at a cost of $9,730 ($8,500 in attorney fees, $500 for formal patent drawings, and $730 in patent office fees). After 2.5 years from the provisional application filing date, the patent office typically issues an "office action" outlining its opinion about which claims in the patent are allowable and which are not. The applicant pays $2,500 for an attorney to prepare a response. If the response satisfies the patent office, the application is approved at about the 3-year mark, requiring a final $350 in attorney fees and $480 in patent office fees, for a grand total of $16,690.

If the same company had filed a nonprovisional application (also with an approval after one response to an office action), the initial $3,630 for the provisional application wouldn't apply, so the total cost is reduced to $13,060, although most of this would have to be paid up front. In addition, the patent is issued at the 2-year mark following the initial application submission.

Provisional versus nonprovisional considerations also factor into international patent protection. For those seeking protection in just two countries, filing directly with national patent offices will result in faster grants and lower lifetime cost. This approach is typically taken by multinational corporations that know which territories they are interested in before filing any application.

Otherwise, the Patent Cooperation Treaty (PCT) application process is preferable. A qualifying application, such as a U.S. provisional patent application, filed in one country is treated as an application in all PCT member countries, which include most developed nations. It can be a relatively inexpensive way to obtain international protection while gauging which national markets are worth pursuing. However, use of the PCT process extends the time it takes for the final patent to be issued in all jurisdictions.

In one scenario, the company pays an initial $3,630 for the provisional application as described above. One year later, the company deems the idea viable, so it files a PCT application at an additional cost of $11,503 ($8,500 in attorney fees, $2,503 in filing fees—which can vary depending on the precise details of the application—and $500 for drawings). At the 2.5-year mark, the applicant must select the territories it wishes to pursue protection in. The company selects the United States, Canada, and Europe (the European Patent Office can issue a regional patent covering about 40 European countries). The fees break down to $1,200 for the attorneys, $740 in U.S. filing fees, and approximately $2,000 per additional jurisdiction. At the 4-year mark, there are the inevitable office actions, as each patent office typically issues at least one. The cost to respond to the U.S. patent office is $2,500, and the cost to respond to the others is $4,000 each (including attorney fees and government fees). At the 4.5-year mark, the patent is issued, having cost about $830 in attorney and filing fees for the U.S. patent office and $2,000 for each of the other offices, for a total of $36,403.

—NATE BAILEY & ANDY PITCHFORD

The authors are patent attorneys at Waddey Patterson, in Nashville. Opinions are their own.

WHERE THE JOBS ARE: 2014

THE MANUFACTURING AND PETROCHEMICAL INDUSTRIES NEED EEs

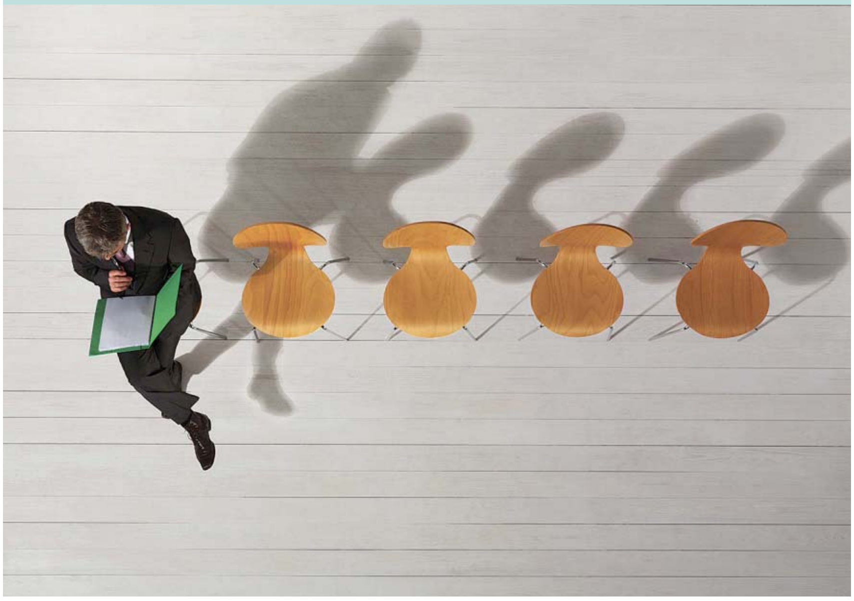

ADIE BUSH/GETTY IMAGES

The slowly recovering economy is yielding gains for electrical engineers: In early 2014, job website CareerBuilder had 3.5 electrical engineering jobs posted for every job-seeking candidate, while U.S. tech recruiting firm Randstad Engineering says demand for electrical engineers is high in semiconductor and telecommunications companies and the energy and transportation sectors. And in manufacturing, "we're seeing an increased demand for EEs, specifically quality and control engineers," says Jay Rogers, vice president of recruiting at Randstad. "There is also a demand for control engineers in the petroleum industry. And salaries in that industry are off the charts." According to the U.S. Bureau of Labor Statistics, the top-paying jobs for electrical engineers are in software development and in oil and gas extraction, with average salaries of US $119,000 and $118,000, respectively.

Many of Randstad's client companies sponsor H-1B work visas for international hires when the supply of U.S. engineers is low, Rogers says. But this year, they are having trouble filling even those work visa positions. The auto industry also has an escalating need for engineers with power electronics as well as computer engineering and software know-how, he adds.

The top 10 jobs in the United States and Canada ranked by highest expected salary gains are in engineering and computer science, according to staffing firm Robert Half. Mobile applications developers, software developers, and software engineers will all see a nearly 8 percent hike in salary this year, more than twice the average salary growth across all fields in 2014. Over the past two years, salaries in the United States for control engineers and power engineers have increased by 10 percent, those for RF engineers by 15 percent, and those for protection and control engineers by 25 percent, according to Randstad.

This high-demand, low-supply scenario applies to Europe as well, Randstad's Rogers adds. The unemployment rate for engineers in the United States is 3 percent, while in Germany it is 2 percent. According to the 2014 European Graduate Career Guide, the U.K. government forecasts engineering job growth in the automotive, aerospace, energy, and medical sectors, and it plans to create 150,000 high-skill jobs in electronics by 2020. The growth of renewables, especially offshore wind in the United Kingdom and Germany, is creating thousands of jobs. But the British manufacturing industry faces a constant dearth of engineers with the right skills, reports the *Financial Times*, with big employers like Dyson and Jaguar Land Rover struggling to fill hundreds of positions.

A recent survey by international employment firm ManpowerGroup found that engineering jobs top the list of hard-to-fill jobs in the Asia-Pacific region. China, the largest oil-product consumer in the Asia-Pacific, plans to reduce dependency on imports by increasing oil refining capacity and establishing shale gas production, so recruitment opportunities are booming, says Joshua Schrijvers of NES Global Talent.

But the petrochemical industry isn't the only hot spot. Thanks to the growth of big data, cybersecurity jobs and income—along with those in mobile and cloud computing—are on the rise. The current average salary for a cybersecurity hire in government is $116,000, while private companies often pay more, says Eugene Spafford, executive director of Purdue University's Center for Education and Research in Information Assurance and Security. Businesses are also creating special positions for computer scientists and math whizzes with degrees in data analysis. Candice Lewis, assistant director of the new M.S. in Business Analytics program at the University of Texas at Austin, says that a job event in the fall brought 53 companies in areas such as retail, consulting, and oil and gas for the school's first crop of 52 graduates.

—PRACHI PATEL

START-UPS: BREWBOT AND BONAVERDE

BETTER BEVERAGES THROUGH NEXT-GEN HOME BREWING

BREW WITH PANACHE: Brewbot's largely automated home beer brewery is monitored and controlled with a smartphone app.

SIMON MILLS

The poster child for the Internet of Things is the smart refrigerator, which keeps track of your groceries, orders new ones, and acts like a digital bulletin board. Alas, grocery management hasn't inflamed much passion among the early adopters crucial to driving new consumer technology. But maybe a new idea for the kitchen of tomorrow can: do-it-yourself beer and coffee.

Two start-ups have recently successfully completed Kickstarter campaigns, one for a smart home brewery and the other for a coffee roaster. The first start-up is Brewbot, based in Belfast, Northern Ireland, whose eponymous product replaces brewing's traditionally hand-tuned processes—adjusting boiling temperatures and tweaking fermentation times—with automated algorithms and smartphone-app control.

"As a group of people, we got together and decided to start brewing," says Chris McClelland, Brewbot's CEO, of his company's origins. "We were worried about plumbing issues and heating elements and. . .[then we thought,] why aren't we looking at our phone and getting indications that say something's happening or not, or when to change something?"

So McClelland and his team of five developers reengineered home brewing for automation. The Brewbot is about the size of a small file cabinet, with tubs and bins inside connected by silicon tubing. An iOS app (an Android version is in the works) keeps tabs on each step, automating some steps and notifying the home brewer when it's necessary to go to the Brewbot to complete other steps. "In the app, you're able to choose the baseline recipe that you want to make, like a pale ale or stout. But you can swap in and out hops and customize the beer to call it your own," says McClelland.

Brewbot is targeting its device at home-brewing hobbyists as well as restaurants and bars that want to have their own line of beers on tap, taking advantage of the current popularity of craft beers. "I think beer is the new wine," McClelland says, "the way people are thinking about it and drinking it and pairing it with food.” Brewbot began shipping units to Kickstarter backers in the second quarter of this year, and the general public can purchase one for about US $2,900.

The second start-up, Berlin-based Bonaverde, is making another eponymous product, but this one is smaller and quicker, and it's aimed toward everyday consumers. Bonaverde communications director Nathalie Sonne says its PC-tower-size appliance began with founder Hans Stier's hatred for the acidic, sour stomach he got from coffee. He came to his eureka moment, like Brewbot's founders, when he realized how much of his beverage's production process could be simplified and automated. And he found he could provide a smoother drink by roasting, grinding, and brewing the beans in the same device.

“The roasting takes 4 to 6 minutes, depending on the roasting degree you've selected,” Sonne says. “The grinding takes around 90 seconds. Then it's the standard brewing process. So it actually takes around 12 to 14 minutes till you have your fresh cup of coffee.”

A BEAN MACHINE: A prototype of Bonaverde's kitchen machine, which does everything required to turn unroasted beans into a hot cup of coffee.

BONAVERDE

Bonaverde has launched a separate Indiegogo campaign for its parallel effort to enable its customers to buy dried (but unroasted) green coffee beans directly from the company's partner farmers in Central America. The company promotes its farmers' Facebook pages and encourages Skype chats and Google Hangouts between farmers and consumers.

“He posts pictures of his farm every day,” Sonne says of one partner in Nicaragua. “We want to make it that transparent, so people can see where their coffee comes from, learn more about their coffee—and make it more like the experience of tasting wine.” The machine will start shipping in December. Though the company's successful Kickstarter campaign priced its machine at the $350 pledge level, Sonne says a retail price for the unit hasn't yet been announced. —**MARK ANDERSON**

Company: Brewbot **Founded:** 2013
Headquarters: Belfast, Northern Ireland
Founders: Jonny Campbell, Kieran Graham Chris McClelland, Alister Sisk **Funding:** US $196,000 **Employees:** 6

Company: Bonaverde **Founded:** 2013
Headquarters: Berlin **Founder:** Hans Stier
Funding: $1 million **Employees:** 6

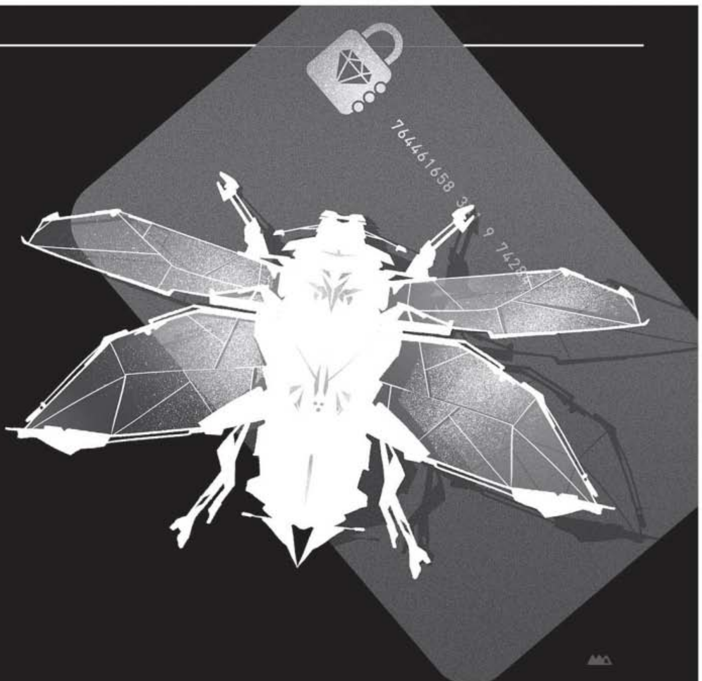

TECHNICALLY SPEAKING_BY PAUL McFEDRIES

OPINION

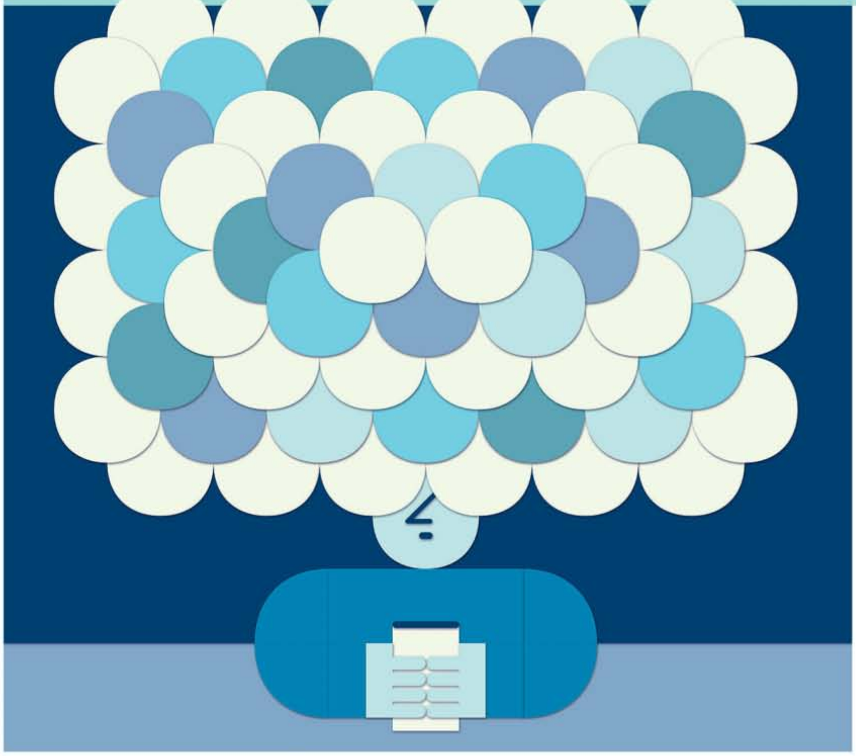

COMPUTING, CONSCIOUSLY

Never before has connectivity offered us so many ways to disconnect. —*Evgeny Morozov*

IN LATE 2010, *New York Times* technology columnist Nick Bilton had an epiphany of sorts: "I've started to notice that while technology allows us to connect to people far away, it can simultaneously disconnect us from people who may be directly in front of us." Okay, perhaps *epiphany* is too strong a word for what might strike many of us as self-evident. So what did he do about it? "My wife and I have started to make a conscious effort to limit the use of our mobile phones during dinner or while spending time with family." Note the use of the word *limit* here in place of, say, *stop* or *ban*. ● I mention this not to rag on Mr. Bilton, who is an excellent columnist and observer, but to illustrate a problem. Our have-device-will-travel lifestyles are combining with sensor-filled cityscapes to create not only **ubiquitous information** but also **ubiquitous distraction**. Our devices are combining with our surroundings to create not only **ambient connectivity** but also *compulsive connectivity*. The result is the scourge of *attention theft* that I described in my June column. If smart people like Nick Bilton *want* to disconnect but *can't*, what is to be done? ● **Disconnectionists** advocate time away from online activities to facilitate mental or spiritual rejuvenation. For example, you could vacation at a **retro retreat** or a **black-hole resort**. I'm sure these are worthwhile diversions, but they're prime examples of **conspicuous austerity** (spending large sums on goods and services that convey an image of simplicity). In any case, once you return from your **digital Sabbath** you're faced with the same distractions. That's okay, reply the **digital detox** disciples, because now your mind is relaxed and fresh, ready for the onslaught! Some **tech-life balance**. ● Other people seek to remove distractions altogether. Some are creating **Walden zones**–rooms without computers or Internet connections. The novelist George R.R. Martin writes using a computer that doesn't have Internet access. The critic Evgeny Morozov, an acknowledged **distraction addict**, purchased a safe with a timer that he uses to lock away his smartphone and Internet cable for days at a time. Too extreme? Perhaps you need some **zenware**, programs that deliberately hide onscreen distractions or block online time sucks. A **distraction blinder** is a utility that removes everything from the screen (including menus and toolbars) except for the program's work area. An **attention shield** is similar, except that it blocks either all incoming online traffic or (if you need the Internet for work) specific sites, such as Facebook and Twitter.

The problem with solutions such as these is that they put you into a kind of parent-child relationship with yourself, where your conscientious self is the scolding parent and your addicted-to-Twitter self is the naughty child. Telling yourself to "go to your room" is not only weird but also remarkably unhelpful in the long run, because it doesn't treat the underlying problem. So some folks are trying a different approach: **contemplative computing**. Using meditation and mindfulness techniques, they seek not enlightenment but awareness, mental discipline, and increased brain function. As one West Coast meditation teacher put it: "All the woo-woo mystical stuff, that's really retrograde. This is about training the brain and stirring up the chemical soup inside." Former Apple and Microsoft executive Linda Stone would have us go beyond all that and embrace **conscious computing**, where, through the use of tools to monitor telltales such as heart rate or by simply attending carefully to our mental and physical processes, we achieve a state of detachment from our devices and create a healthy distance from our data.

Will any of this help **hyperconnected** people such as Nick Bilton? I can't say, but conscious computing seems like the right approach for the rest of us, at least. After all, our online messages, tweets, and posts are every bit as authentic as our off-line conversations, encounters, and books. The trick lies not in favoring one set of experiences over the other but in finding a personal balance between the two. And, yes, that probably includes not using your mobile device during dinner. ■

ILLUSTRATION BY Greg Mably

Browse at Your Own RISK

EVEN WITHOUT COOKIES, FINGERPRINTING LETS ADVERTISERS TRACK YOUR EVERY ONLINE MOVE

By **Nick Nikiforakis** & **Günes Acar**

Photography by **Dan Saelinger**

IN JULY 1993, THE NEW YORKER PUBLISHED a cartoon by Peter Steiner that depicted a Labrador retriever sitting on a chair in front of a computer, paw on the keyboard, as he turns to his beagle companion and says, “On the Internet, nobody knows you’re a dog.” Two decades later, interested parties not only know you’re a dog, they also have a pretty good idea of the color of your fur, how often you visit the vet, and what your favorite doggy treat is. ● How do they get all that information? In a nutshell: Online advertisers collaborate with websites to gather your browsing data, eventually building up a detailed profile of your interests and activities. These browsing profiles can be so specific that they allow advertisers to target populations as narrow as mothers with teenage children or people who require allergy-relief products. When this tracking of our browsing habits is combined with our self-revelations on social media, merchants’ records of our off-line purchases, and logs of our physical whereabouts derived from our mobile phones, the information that commercial organizations, much less government snoops, can compile about us becomes shockingly revealing. ● Here we examine the history of such tracking on the Web, paying particular attention

to a recent phenomenon called fingerprinting, which enables companies to spy on people even when they configure their browsers to avoid being tracked.

THE EARLIEST APPROACH TO online tracking made use of cookies, a feature added to the pioneering Web browser Netscape Navigator a little over a year after Steiner's cartoon hit newsstands. Other browsers eventually followed suit.

Cookies are small pieces of text that websites cause the user's browser to store. They are then made available to the website during subsequent visits, allowing those sites to recognize returning customers or to keep track of the state of a given session, such as the items placed in an online shopping cart. Cookies also enable sites to remember that users are logged in, freeing them of the need to repeatedly provide their user names and passwords for each protected page they access.

So you see, cookies can be very helpful. Without them, each interaction with a website would take place in a vacuum, with no way to keep tabs on who a particular user is or what information he or she has already provided. The problem came when companies began following a trail of cookie crumbs to track users' visits to websites other than their own.

How they do that is best explained through an example. Suppose a user directs her browser to a travel website–let's call it Travel-Nice-Places.com–that displays an advertising banner at the top of the page. The source of that banner ad is probably not Travel-Nice-Places.com itself. It's more likely located on the Web servers of a different company, which we'll call AdMiddleman.com. As part of the process of rendering the page at Travel-Nice-Places.com, the user's browser will fetch the banner ad from AdMiddleman.com.

Here's where things get sneaky. The Web server of AdMiddleman.com sends the requested banner ad, but it also uses this opportunity to quietly set a third-party cookie on the user's browser. Later, when that same user visits an entirely different website showing another ad from AdMiddleman.com, this ad supplier examines its previously set cookie, recognizes the user, and over time is able to build a profile of that user's browsing habits.

You might ask: If this brings me more relevant online advertisements, what's the harm? True, online tracking could, in principle, help deliver ads you might actually appreciate. But more often than not, the advertisers' algorithms aren't smart enough to do that. Worse, information about your Web browsing habits can be used in troubling ways. A car dealer you approach online and then visit in the flesh, for example, could end up knowing all about your investigations, not only of its inventory but of all the other car-related websites you've been checking out. No wonder such tracking has garnered a reputation for being creepy.

Not long after the use of third-party tracking cookies became common, various media outlets and privacy organizations began questioning the practice. And over the years, people have increasingly come to appreciate that the set of websites they visit reveals an enormous amount about themselves: their gender and age, their political leanings, their medical conditions, and more. The possession of such knowledge by online advertising networks, or indeed by any company or government agency that purchases it from those networks, comes with potentially dire consequences for personal privacy–especially given that users have no control of this very opaque process of data collection.

It should come as no surprise that some of the early news articles about advertisers' use of cookies had headlines announcing "the death of privacy" and made allusions to George Orwell's all-seeing Big Brother. Even the programmers and engineers involved in the development of technical standards got an earful.

In particular, in 1997 a coalition of privacy organizations wrote an open memo to the Internet Engineering Task Force (sending copies to the leading browser

THIS PAGE AND PREVIOUS PAGES, PROP STYLIST: DOMINIQUE BAYNES

developers) that expressed their support for the first cookie standard, RFC 2109, which stated that third-party cookies should be blocked to “prevent possible security or privacy violations.” But advertising companies pushed back harder. And in the end, neither of the two mainstream browsers of that era, Netscape Navigator and Internet Explorer, followed the specification, both allowing third-party cookies.

The winds began to shift in 2005, though, when browser developers started adding a “private browsing” mode to their products. These give users the option of visiting websites without letting those sites leave long-term cookies. Independent developers, too, started producing privacy-preserving extensions that users could add to their browsers.

Today, the most popular extension to Mozilla’s Firefox browser is AdBlock Plus, which rejects both ads and third-party cookies used for tracking. And recently developed tools like Ghostery and Mozilla’s Lightbeam reveal the number of trackers on each website and show how these trackers collaborate between seemingly unrelated sites. Finally, recent studies have shown that a large percentage of people delete their browser cookies on a regular basis, a fact that points to their having at least some understanding of how cookies can compromise privacy online.

But when people started deleting their cookies, the companies involved in tracking didn’t just roll over. They responded by developing new ways of sniffing out users’ identities. Most had one thing in common: They tried to bury the same tracking information found in cookies in some other corner of the user’s browser.

One popular technique was to use Flash cookies. These are conceptually similar to normal cookies, but they are specific to Adobe’s Flash plug-in. In the past, a website could hide information in Flash cookies, which would survive the clearing of normal cookies. The information retained in the Flash cookies would then be used to regenerate the deleted normal cookies. Companies made use of this sneaky tactic for a few years before researchers caught on and started publicizing these shady practices in 2008. Today, most browsers give users the ability to delete all flavors of cookies.

AS YOU MIGHT EXPECT OF THIS long-standing cat-and-mouse game, the advertising networks have not sat idle. In recent years, they have shifted to a form of tracking that doesn’t require Web servers to leave any kind of metaphorical bread crumb on the user’s machine. Instead, these ad networks rely on a process known more generally as device fingerprinting: collecting identifying information about unique characteristics of the individual computers people use. Under the assumption that each user operates his or her own hardware, identifying a device is tantamount to identifying the person behind it.

While this all sounds very sinister, it’s important to realize that such fingerprinting has some very benign, indeed laudable, applications. It can be used, for example, to verify that someone logging into a Web-based service is not an attacker using stolen log-in credentials. Fingerprinting is also helpful for combating click fraud: Someone displays an advertisement on his website in return for payment each time that ad is clicked on—and then tries to run up the bill by having an identity-feigning computer click many times on the ad. The problem is that fingerprinting has become so precise that it makes a sham of browsers’ privacy-protection measures.

In 2010, Peter Eckersley of the Electronic Frontier Foundation showed that tracking various browser attributes provided enough information to identify the vast majority of machines surfing the Web. Of the 470,000-plus users who had participated at that point in his public Panopticlick project, 84 percent of their browsers produced unique fingerprints (94 percent if you count those that supported Flash or Java). The attributes Eckersley logged included the user’s screen size, time zone, browser plug-ins, and set of installed system fonts.

> **TODAY ON THE INTERNET, interested parties not only know you’re a dog, they also have a pretty good idea of the color of your fur**

We have expanded on Eckersley's study by examining not just what kinds of fingerprinting are theoretically possible but, more to the point, what is actually going on in the wilds of the Internet's tracking ecosystem. We started our analysis at the University of Leuven, in Belgium, by first identifying and studying the code of three large fingerprinting providers: BlueCava, Iovation, and ThreatMetrix.

The results were rather chilling. The tactics these companies use go far beyond Eckersley's probings. For instance, we found that one company uses a clever, indirect method of identifying the installed fonts on a user machine, without relying on the machine to volunteer this information, as Eckersley's software did.

We also discovered fingerprinting code that exploits Adobe Flash as a way of telling whether people are trying to conceal their IP addresses by communicating via intermediary computers known as proxies. In addition, we exposed Trojan horse-like fingerprinting plug-ins, which run surreptitiously after a user downloads and installs software unrelated to fingerprinting, such as an online gambling application.

With the information we gathered about these three companies, we created and ran a program that autonomously browses the Web and detects when a website is trying to fingerprint it. The purpose of this experiment was to find more players in the fingerprinting game, ones less well known than the three we studied initially.

We quickly uncovered 16 additional fingerprinters. Some were in-house trackers, used by individual companies to monitor their users without sharing the information more widely. The rest were offered as products by such companies as Coinbase, MaxMind, and Perferencement.

And it seems the companies selling this software are finding buyers. Our results showed that 159 of Alexa's 10,000 most-visited websites track their users with such fingerprinting software. We also found that more than 400 of the million most popular websites on the Internet have been using JavaScript-only fingerprinting, which works on Flash-less devices such as the iPhone or iPad. Worse, our experiment revealed that users continue to be fingerprinted even if they have checked "Do Not Track" in their browser's preferences.

BROWSER FINGERPRINTING IS becoming common, and yet people are mostly in the dark about it. Even when they're made aware that they're being tracked, say, as a fraud-protection measure, they are, in essence, asked to simply trust that the information collected won't be used for other purposes. One of those is targeted advertising, which works even when users switch into their browsers' private mode or delete their cookies. What are those unwilling to go along with this new form of tracking doing about it?

As part of our research on browser fingerprinting, we examined various tools that people are using to combat it. One popular approach is installing browser extensions that let you change the values that identify your browser to the server. Such modifications allow users to occasionally trick servers into dishing out pages customized for different browsers or devices. Using these extensions, Firefox devotees on computers running Linux, for example, can pretend to be Internet Explorer fans running Microsoft Windows. Other extensions go further, reporting false dimensions for the screen size and limiting the probing of fonts.

Our analysis showed that a mildly accomplished fingerprinter could easily overcome any of these supposedly privacy-enhancing browser extensions. That's because modern browsers are huge pieces of software, each with its own quirks. And these idiosyncrasies give away the true nature of the browser, regardless of what it claims to be.

This makes those privacy-protecting extensions useless. In fact, they are worse than useless. Resorting to them is like trying to hide your comings and goings in a small town by disguising your car. If you get a rental, that might work. But if you merely replace the chrome lettering on your Prius with lettering taken from the back of a Passat, not only will your ruse be obvious, you will have now marked your car in a way that makes it easy to distinguish from the many other Priuses on the road. Similarly, installing such a fingerprint-preventing browser extension only makes you stand out more.

GIVEN THAT ADVERTISING IS the Web's No. 1 industry and that tracking is a crucial component of it, we believe that user profiling in general and fingerprinting in particular are here to stay. But more-stringent regulations and more-effective technical countermeasures might one day curb the worst abuses.

Taking Your Print

IN THE PAST, clearing cookies after each session or selecting your browser's "Do Not Track" setting could prevent third-party tracking. But the advent of browser fingerprinting makes it very difficult to prevent others from monitoring your online activities. The diagram at right outlines how an online advertising network can track the sites you visit using fingerprinting.

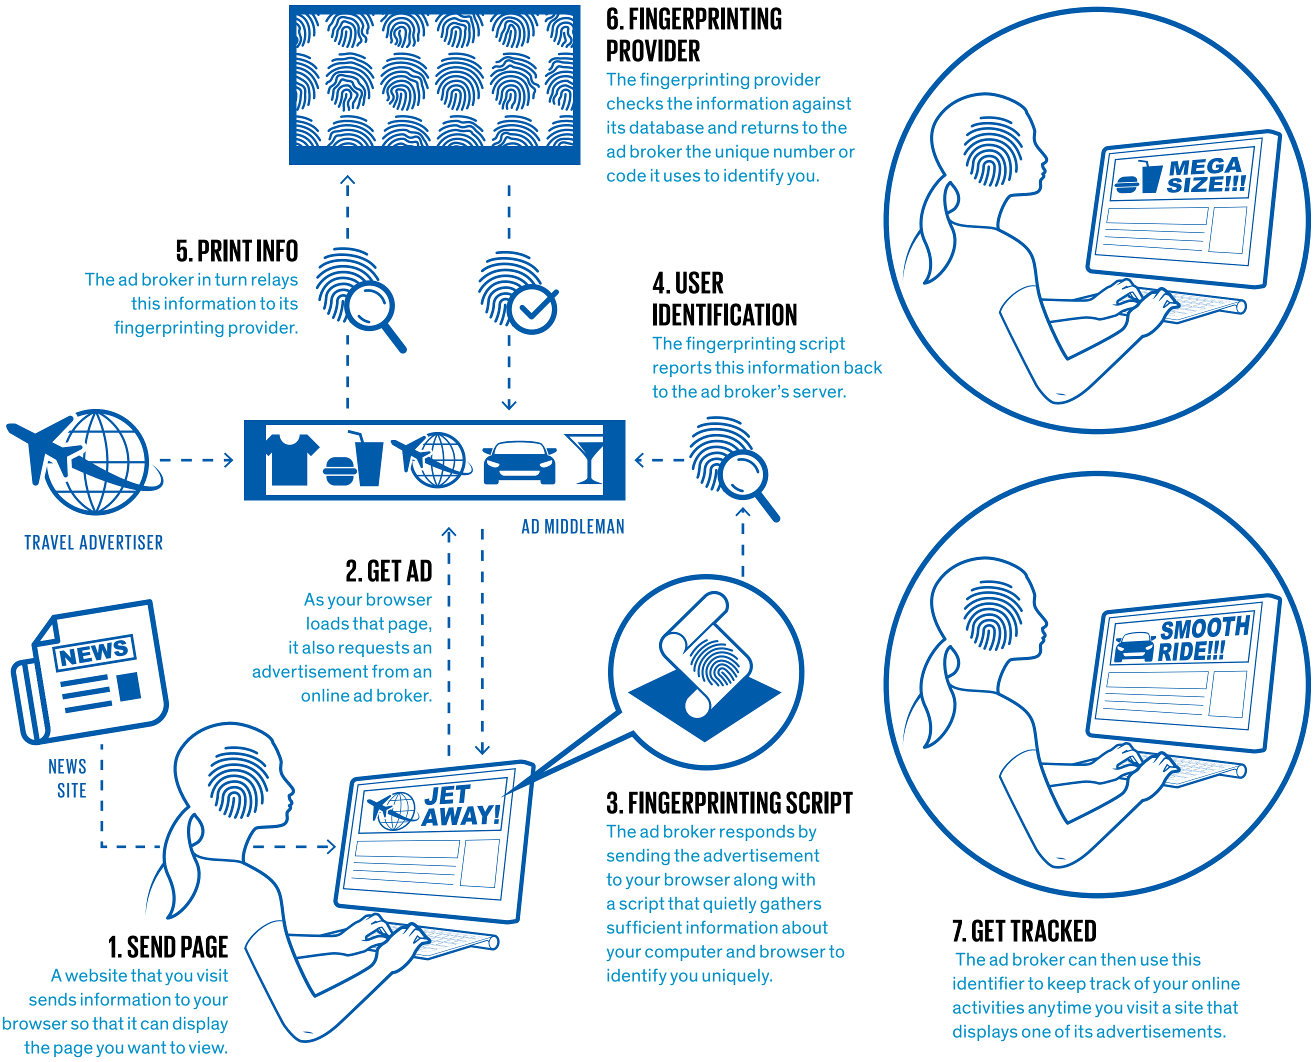

1. SEND PAGE
A website that you visit sends information to your browser so that it can display the page you want to view.

2. GET AD
As your browser loads that page, it also requests an advertisement from an online ad broker.

3. FINGERPRINTING SCRIPT
The ad broker responds by sending the advertisement to your browser along with a script that quietly gathers sufficient information about your computer and browser to identify you uniquely.

4. USER IDENTIFICATION
The fingerprinting script reports this information back to the ad broker's server.

5. PRINT INFO
The ad broker in turn relays this information to its fingerprinting provider.

6. FINGERPRINTING PROVIDER
The fingerprinting provider checks the information against its database and returns to the ad broker the unique number or code it uses to identify you.

7. GET TRACKED
The ad broker can then use this identifier to keep track of your online activities anytime you visit a site that displays one of its advertisements.

We and other researchers are indeed trying to come up with better software to thwart fingerprinting. A straightforward solution might be to stop the fingerprinting scripts from ever loading in browsers, similar to the way ad blockers work. By maintaining a blacklist of problematic scripts, an antifingerprinting extension could detect their loading and prohibit their execution.

One challenge is that the blacklist would have to be revised constantly to keep up with the changes that trackers would surely make in response. Another issue is that we don't know whether the loading of fingerprinting scripts is necessary for the functionality of certain websites. Even if it's not required now, websites could be changed to refuse loading of their pages unless the fingerprinting scripts are present and operational, which would discourage people from trying to interfere with them.

A more effective way of approaching the problem would be for many people to share the same fingerprint. To some extent that is happening now with smartphones, which can't be customized to the degree that desktop or laptop computers can. So phones currently present fewer opportunities for fingerprinters. It might be possible to make other kinds of computers all look alike if Web browsing were done through a cloud service, one that treats the browser running on the user's PC simply as a terminal. Trackers would then be able to detect only the cloud browser's fingerprint.

Companies offering cloud-based browsing already exist, but it's not clear to us whether the browsers that are exposed to potential fingerprinters actually operate in the cloud. Still, there's no reason to think that a system for preventing fingerprinting with a cloud browser couldn't be engineered. For some of us, anyway, it could be worth adopting, even if it involved monthly charges. After all, doing nothing has a price, too—perhaps one as steep as forfeiting online privacy for good. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/webfingerprinting0814

ILLUSTRATION BY **MCKIBILLO**

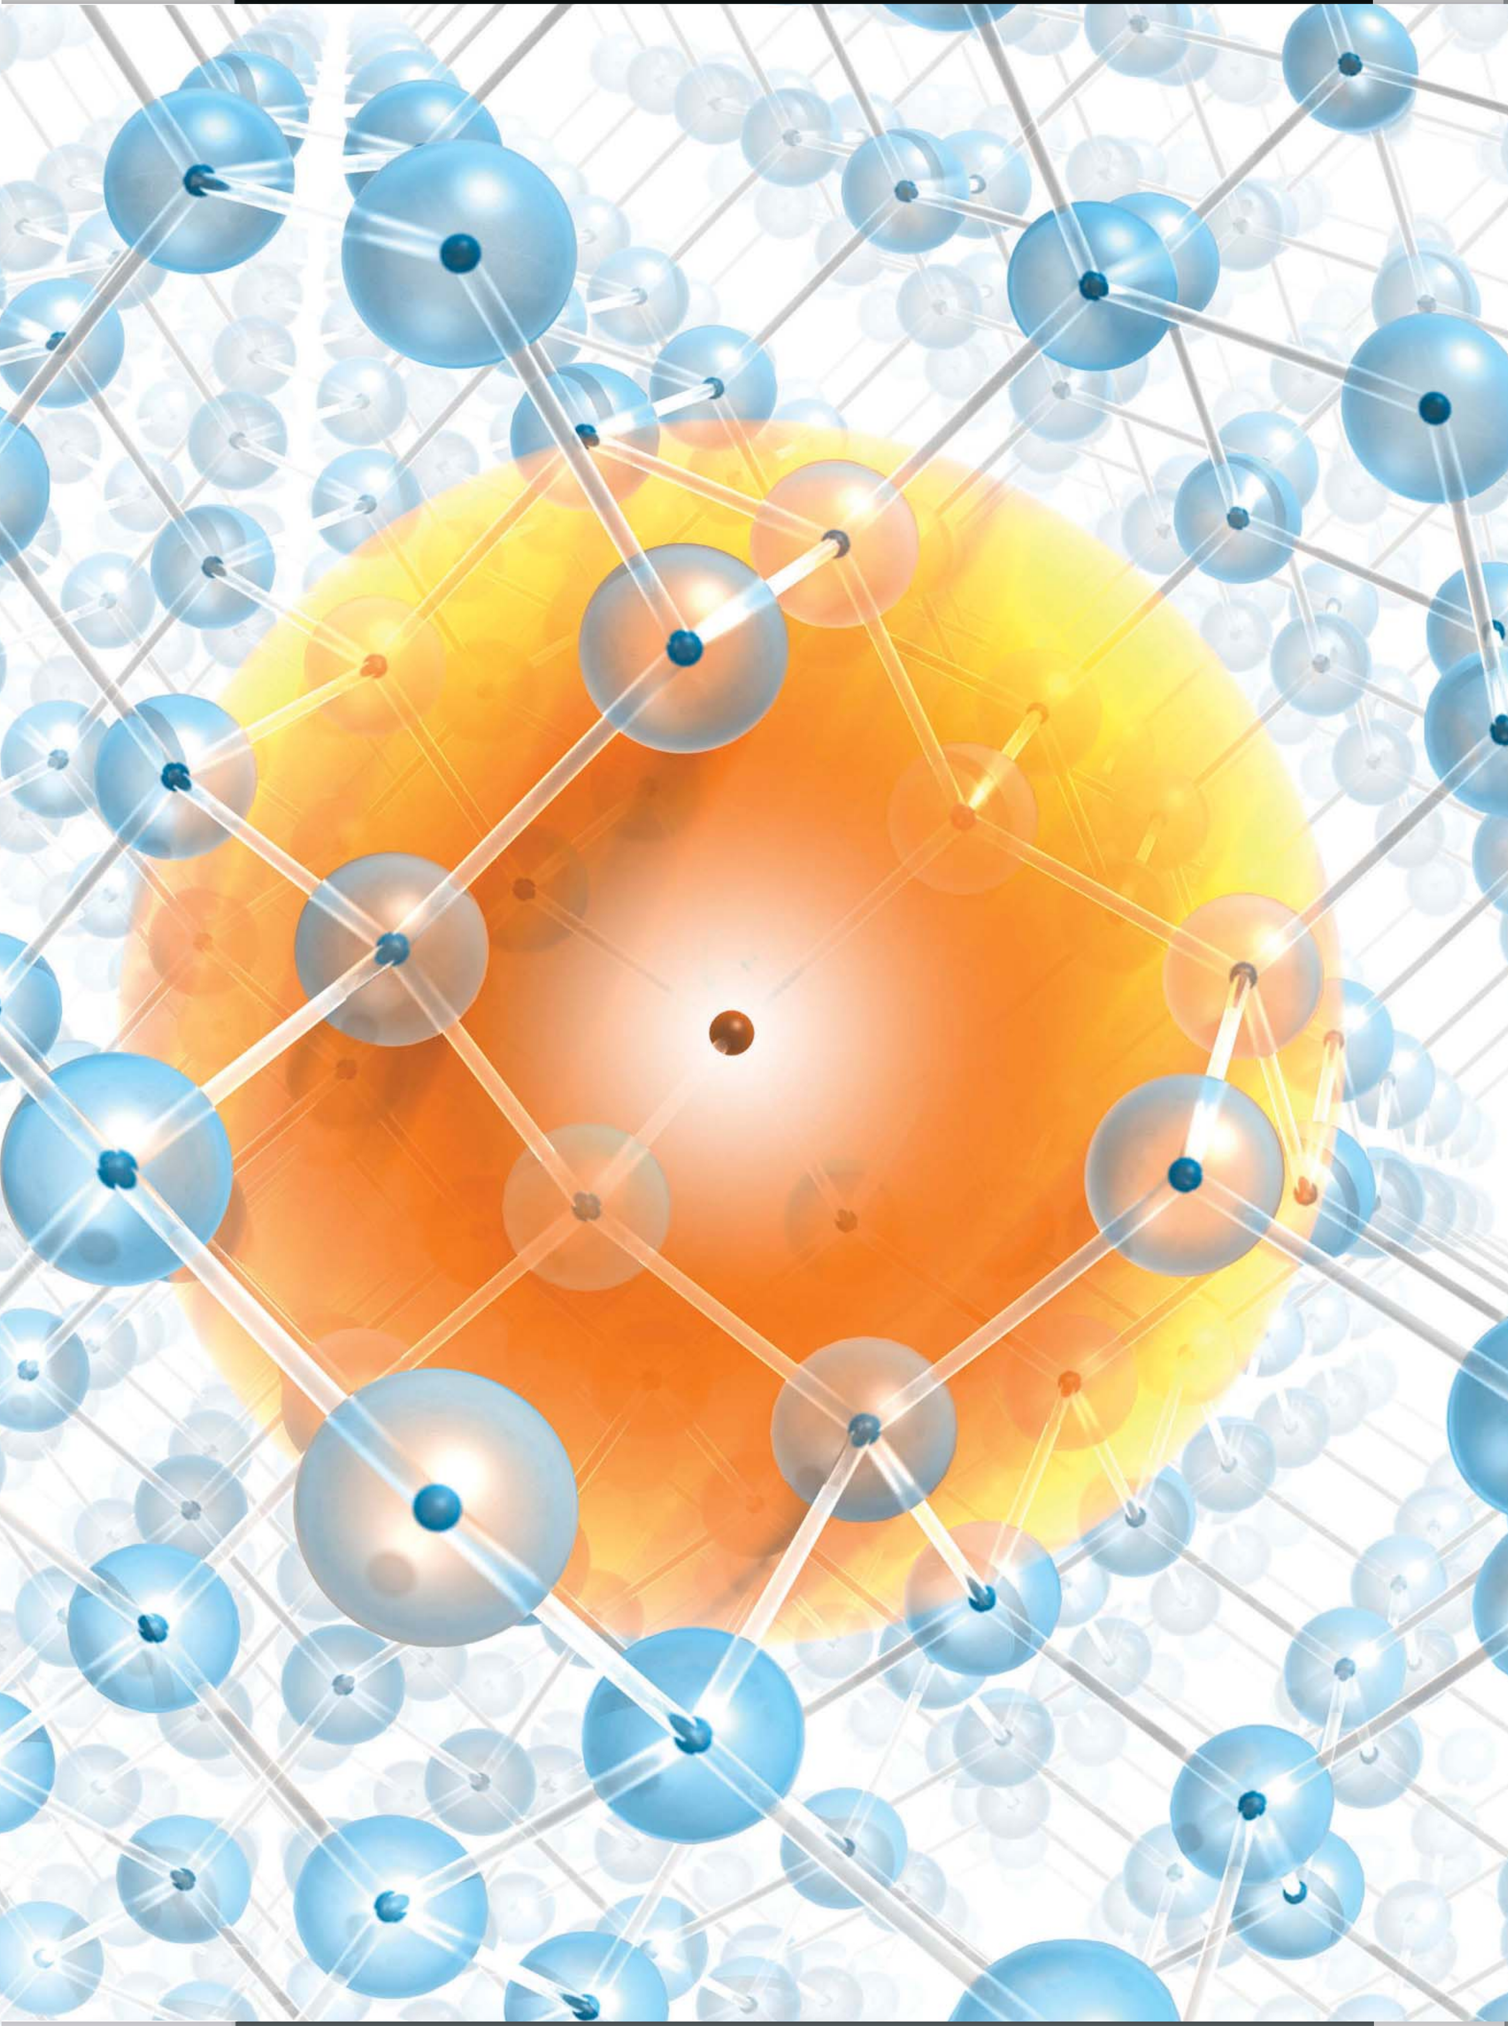

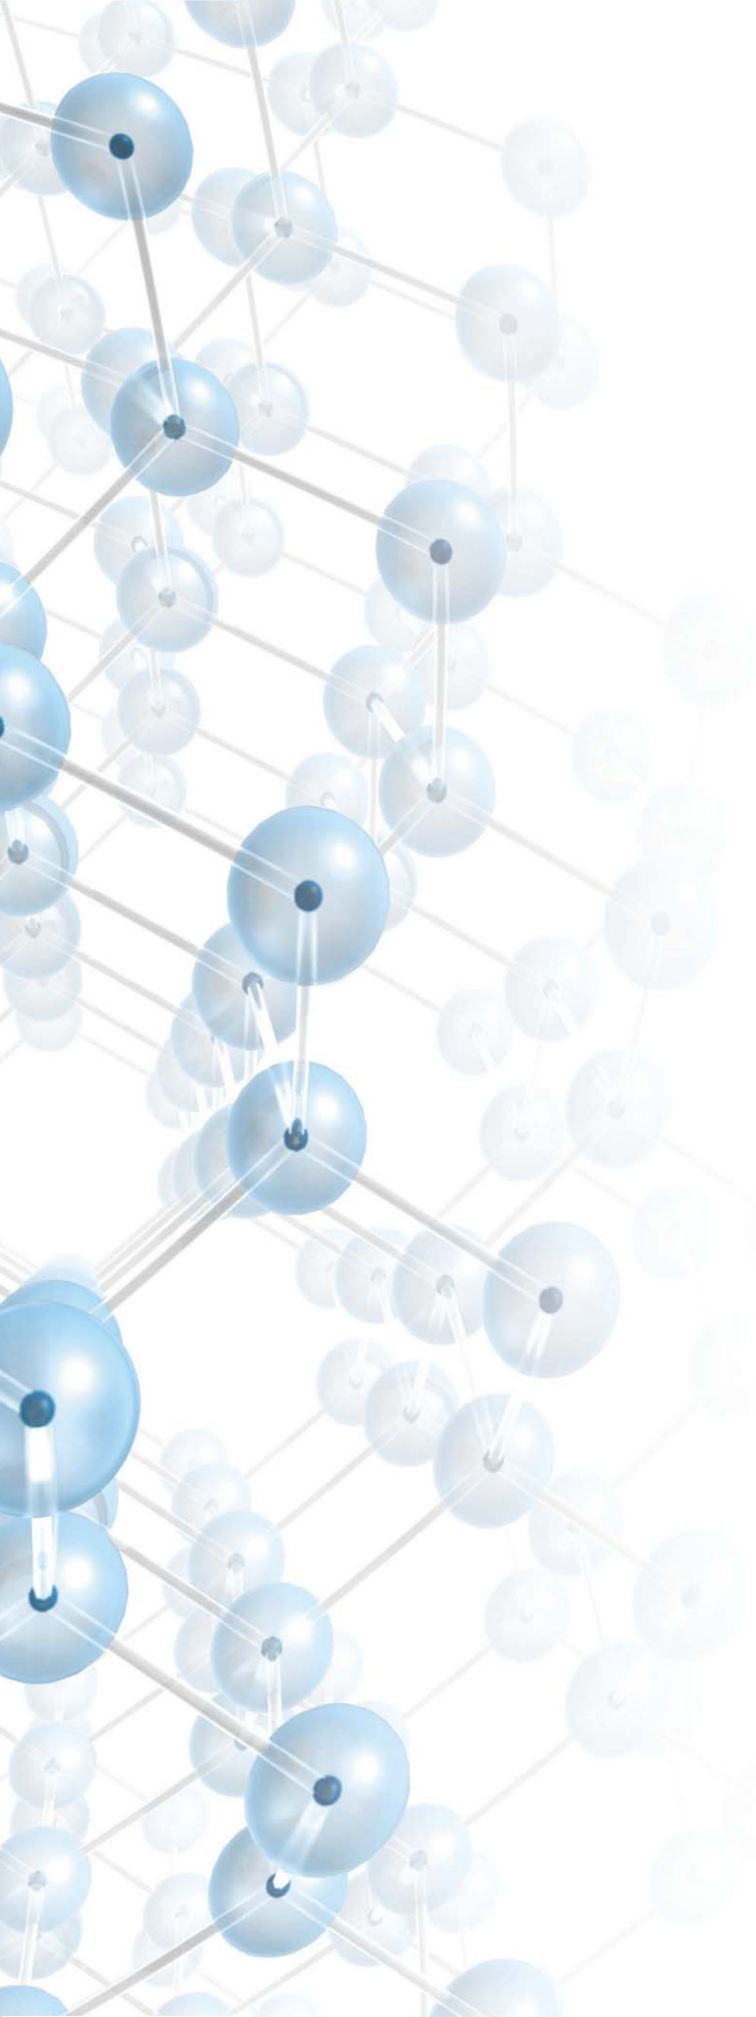

SILICON'S SECOND ACT

Can this semiconductor workhorse take computing into the quantum era?

BY CHEUK CHI LO & JOHN J.L. MORTON

ILLUSTRATION BY BRYAN CHRISTIE DESIGN

GRAND ENGINEERING CHALLENGES OFTEN REQUIRE AN EPIC LEVEL OF PATIENCE. THAT'S CERTAINLY TRUE FOR QUANTUM COMPUTING. FOR A GOOD 20 YEARS NOW, WE'VE KNOWN THAT QUANTUM COMPUTERS COULD, IN PRINCIPLE, BE STAGGERINGLY POWERFUL, TAKING JUST A FEW MINUTES TO WORK OUT PROBLEMS THAT WOULD TAKE AN ORDINARY COMPUTER LONGER THAN THE AGE OF THE UNIVERSE TO SOLVE. BUT THE EFFORT TO BUILD SUCH MACHINES HAS BARELY CROSSED THE STARTING LINE. IN FACT, WE'RE STILL TRYING TO IDENTIFY THE BEST MATERIALS FOR THE JOB.

Today, the leading contenders are all quite exotic: There are superconducting circuits printed from materials such as aluminum and cooled to one-hundredth of a degree above absolute zero, floating ions that are made to hover above chips and are interrogated with lasers, and atoms such as nitrogen trapped in diamond matrices.

These have been used to create modest demonstration systems that employ fewer than a dozen quantum bits to factor small numbers or simulate some of the behaviors of solid-state materials. But nowadays those exotic quantum-processing elements are facing competition from a decidedly mundane material: good old silicon.

Silicon had a fairly slow start as a potential quantum-computing material, but a flurry of recent results has transformed it into a leading contender. Last year, for example, a team based at Simon Fraser University in Burnaby, B.C., Canada, along with researchers in our group at University College London, showed that it's possible to maintain the state of quantum bits in silicon for a record 39 minutes at room temperature and 3 hours at low temperature. These are eternities by quantum-computing standards–the longevity of other systems is often measured in milliseconds or less–and it's exactly the kind of stability we need to begin building general-purpose quantum computers on scales large enough to outstrip the capabilities of conventional machines.

As fans of silicon, we are deeply heartened by this news. For 50 years, silicon has enabled steady, rapid progress in conventional computing. That era of steady gains may be coming to a close. But when it comes to building quantum computers, the material's prospects are only getting brighter. Silicon may prove to have a second act that is at least as dazzling as its first.

● ● ●

WHAT IS A QUANTUM COMPUTER? Simply put, it's a system that can store and process information according to the laws of quantum mechanics. In practice, that means the basic computational components–not to mention the way they operate–differ greatly from those we associate with classical forms of computing.

For example, as bizarre as it sounds, in the quantum world an object can exist in two different states simultaneously–a phenomenon known as superposition. This means that unlike an ordinary bit, a quantum bit (or qubit) can be placed in a complex state where it is both 0 and 1 at the same time. It's only when you measure the value of the qubit that it is forced to take on one of those two values.

When a quantum computer performs logical operations, it does so on all possible combinations of qubit states at the same time. This massively parallel approach is often cited as the reason that quantum computers would be very fast. The catch is that often you're interested in only a subset of those calculations. Measuring the final state of a quantum machine will give you just one answer, at random, that may or may not be the desired solution. The art of writing useful quantum algorithms lies in getting the undesired answers to cancel out so that you are left with a clear solution to your problem.

The only company selling something billed as a "quantum computing" machine is the start-up D-Wave Systems, also based in Burnaby. D-Wave's approach is a bit of a departure from what researchers typically have in mind when they talk about quantum computing, and there is active debate over the quantum-mechanical nature and the potential of its machines (more on that in a moment).

The quarry for many of us is a universal quantum computer, one capable of running any quantum or classical algorithm. Such a computer won't be faster than classical computers across the board. But there are certain applications for which it could prove exceedingly useful. One that quickly caught the eye of intelligence agencies is the ability to factor large numbers exponentially faster than the best classical algorithms can. This would make short work of cryptographic codes that are effectively uncrackable by today's machines. Another promising niche is simulating the behavior of quantum-mechanical systems, such as molecules, at high speed and with great fidelity. This capability could be a big boon for the development of new drugs and materials.

To build a universal quantum computer capable of running these and other quantum algorithms, the first thing you'd need is the basic computing element: the qubit. In principle, nearly any object

SPEEDING UP SEARCH

IN A CLASSICAL SEARCH ALGORITHM, hunting for a particular string in an unstructured database involves looking at every entry in succession until a match is found. On average, you'd have to run through half, or $N/2$, of the queries before the correct entry is located. Grover's algorithm, a quantum search algorithm named for computer scientist Lov Grover, could speed up that work by simultaneously querying all entries. The process still isn't instantaneous: Finding the correct one would take, on average, $\sqrt{N}$ queries. But it could make a difference for large databases. To search a trillion entries, the scheme would require 0.0002 percent of the number of queries needed in the classical approach. Here's how it works.

1 The input, a quantum version of the search string, is set up. It contains N different states, one of which is the index of the string you're looking for. All N states exist in superposition with one another, much like Schrödinger's cat, which can be both dead and alive at the same time. At this point, if the input is observed, it will collapse into any one of its N component states with a probability of $1/N$ (the square of the quantum state amplitude shown in the y-axis of the diagram).

2 The input is fed into the database, which has been configured to invert the phase of the correct entry. Here, phase is a quantum attribute. It can't be directly measured, but it affects how quantum states interact with one another. The correct entry is highlighted in one step, but we can't see it. The probability of observing the correct state is still the same as that of all the others.

3 To get around this observation problem, a quantum computer can be made to perform a simple operation that would invert all of the amplitudes of the states about their overall mean. Now, when the input is measured, it will be more likely to collapse into the correct answer. But if N is large, this probability will still be quite small.

4 To increase the probability of observing the correct entry, Grover's algorithm repeats steps 2 and 3 many times. Each time, the correct state will receive a boost. After $\sqrt{N}$ cycles, the probability of observing that state will be very close to 1 (or 100 percent).

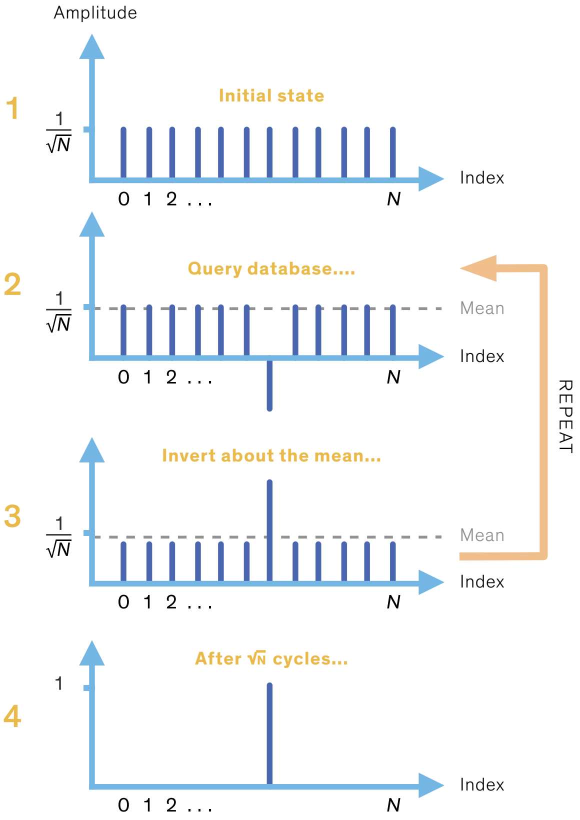

that behaves according to the laws of quantum physics and can be placed in a superposition of states could be used to make a qubit.

Since quantum behavior is typically most evident at small scales, most natural qubits are tiny objects such as electrons, single atomic nuclei, or photons. Any property that could take on two values, such as the polarization of light or the presence or absence of an electron in a certain spot, could be used to encode quantum information. One of the more practical options is spin. Spin is a rather abstruse property: It reflects a particle's angular momentum—even though no physical rotation is occurring—and it also reflects the direction of an object's intrinsic magnetism. In both electrons and atomic nuclei, spin can be made to point up or down so as to represent a 1 or a 0, or it can exist in a superposition of both states.

It's also possible to make macroscopic qubits out of artificial structures—if they can be cooled to the point where quantum behavior kicks in. One popular structure is the flux qubit, which is made of a current-carrying loop of superconducting wire. These qubits, which can measure in the micrometers, are quantum weirdness writ large: When the state of a flux qubit is in superposition, the current flows in both directions around the loop at the same time.

D-Wave uses qubits based on superconducting loops, although these qubits are wired together to make a computer that operates differently from a universal quantum computer. The company employs an approach called adiabatic quantum computing, in which qubits are set up in an initial state that then "relaxes" into an optimal configuration. Although the approach could potentially be used to speedily solve certain optimization problems, D-Wave's computers can't be used to implement an arbitrary algorithm. And the quantum-computing community is still actively debating the extent to which D-Wave's hardware behaves in a quantum-mechanical fashion and whether it will be able to offer any advantage over systems using the best classical algorithms.

Although large-scale universal quantum computers are still a long way off, we are already getting a good sense of how we'd make one. There are several approaches. The most straightforward one employs a model of computation known as the gate model. It uses a series of "universal gates" to wire up groups of qubits so that they can be made to interact on demand. Unlike conventional chips with hardwired logic circuitry, these gates can be used to configure and reconfigure the relationships between qubits to create different logic operations. Some, such as XOR and NOT, may be familiar, but many won't be, since they're performed in a complex space where a quantum state in superposition can take on any one of a continuous range of values. But the basic flow of computation is much the same: The logic gates control how information flows, and the states of the qubits change as the program runs. The result is then read out by observing the system.

Another, more exotic idea, called the cluster-state model, operates differently. Here, computation is performed by the act of observation alone. You begin by first "entangling" every qubit with its neighbors up front. Entanglement is a quantum-mechanical phenomenon in which two or more particles–electrons, for example–share a quantum state and measuring one particle will influence the behavior of an entangled partner. In the cluster-state approach, the program is actually run by measuring the qubits in a particular order, along particular directions. Some measurements carve out a network of qubits to define the computation, while other measurements drive the information forward through this network. The net result of all these measurements taken together gives the final answer.

For either approach to work, you must find a way to ensure that qubits stay stable long enough for you to perform your computation. By itself, that's a pretty tall order. Quantum-mechanical states are delicate things, and they can be easily disrupted by small fluctuations in temperature or stray electromagnetic fields. This can lead to significant errors or even quash a calculation in midstream.

On top of all this, if you are to do useful calculations, you must also find a way to scale up your system to hundreds or thousands of qubits. Such scaling wouldn't have been feasible in the mid-1990s, when the first qubits were made from trapped atoms and ions. Creating even a single qubit was a delicate operation that required elaborate methods and a roomful of equipment at high vacuum. But this has changed in the last few years; now there's a range of quantum-computing candidates that are proving easier to scale up [see "Quantum Contenders"].

Among these, silicon-based qubits are our favorites. They can be manufactured using conventional semiconductor techniques and promise to be exceptionally stable and compact.

● ● ●

IT TURNS OUT THERE ARE a couple of different ways to make qubits out of silicon. We'll start with the one that took the early lead: using atoms that have been intentionally placed within silicon.

If this approach sounds familiar, it's because the semiconductor industry already uses impurities to tune the electronic properties of silicon to make devices such as diodes and transistors. In a process called doping, an atom from a neighboring column of the periodic table is added to silicon, either lending an electron to the surrounding material (acting as a "donor") or extracting an electron from it (acting as an "acceptor").

Such dopants alter the overall electronic properties of silicon, but only at temperatures above -220 °C or so (50 degrees above absolute zero). Below that threshold, electrons from donor atoms no longer have enough thermal energy to resist the tug of the positively charged atoms they came from and so return.

QUANTUM CONTENDERS

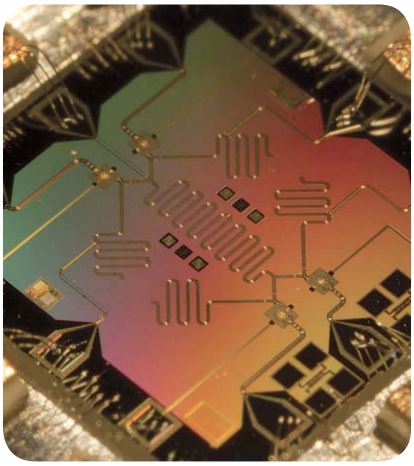

SUPERCONDUCTING: Qubits can be made from a loop of a superconducting material, such as aluminum, paired with thin insulating barriers that electrons can tunnel through. There are various ways to construct a qubit in this system. One is to use the direction of current running around the loop to make a "flux qubit." When the qubit is in a superposition of states, current flows in both directions at the same time. The start-up D-Wave Systems is making 1024-qubit systems using flux-qubit technology. But researchers have generally prioritized device development over system size; the largest systems in the laboratory have incorporated only 5 qubits. These more recent laboratory qubits are known as charge qubits and are often based on total electronic charge.

The stability of superconducting qubits has improved remarkably over the past decade, and they can be entangled with one another with good fidelity through superconducting buses. But the space required is quite large—a qubit can measure in the millimeters when the resonator needed to control it is included. Extremely low temperatures, in the tens of millikelvins, are also needed for optimal operation.

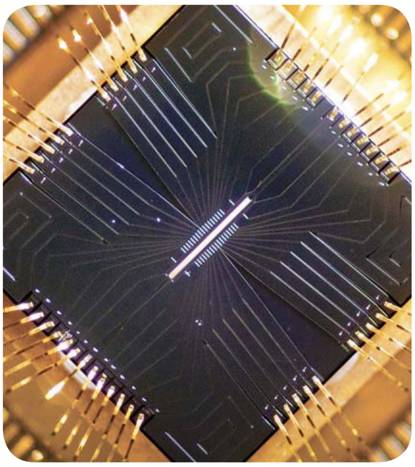

ION TRAPS: The outermost electron of an ion such as calcium can be used to create a qubit that consists of two states, which can be defined either by the electron's orbital state or its interaction with the atom's nucleus. Ion traps were among the earliest quantum-computing systems investigated, beginning in the 1990s. They have since been miniaturized and can be implemented on a chip with electrodes, which are used to suspend ions in midair and move them around. Ion traps have been made that can hold as many as 10 qubits at a time.

Since the ions are made to hover, qubits created in this fashion can be well isolated from stray fields and are thus quite stable. There are some disadvantages to this approach, however. The qubits must be constructed in an ultrahigh vacuum to prevent interactions with other atoms and molecules. And ion qubits must be pushed together to entangle them, which is difficult to do with high precision because of electrical noise.

TOP: ERIK LUCERO; BOTTOM: JOINT QUANTUM INSTITUTE

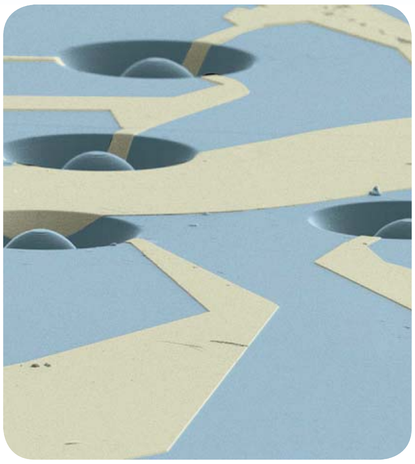

DIAMOND: Atomic defects in diamond have emerged as one of the leading methods for creating qubits in recent years. Such defects are what give diamonds their color (a nitrogen-doped diamond has a yellowish tint). One of the most promising qubits is a nitrogen atom that occupies a place near a vacant site within a diamond lattice. Just as in doped silicon, this defect can be used to make two different kinds of qubit. One can be constructed from the combined spin of two electrons that are attracted to the nitrogen atom. A qubit can also be made using the spin of the nucleus of the nitrogen atom.

Such diamond qubits are attractive because they interact readily with visible light, which should enable long-range communication and entanglement. The systems can stay stable enough for computation up to room temperature. One challenge researchers must tackle is the precise placement of the nitrogen atoms; this will present an obstacle to making the large arrays needed for full-scale general-purpose quantum computers. To date, researchers have demonstrated they can entangle two qubits. This has been done with two defects in the same diamond crystal and with two defects separated by as much as 3 meters.

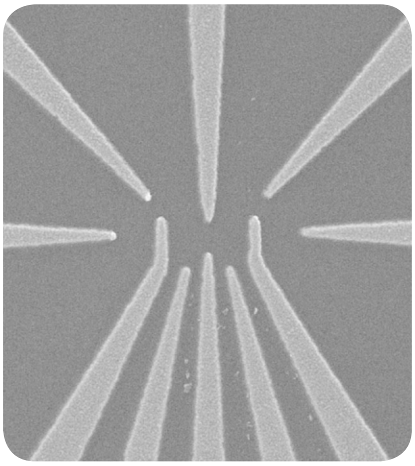

SILICON: There are a few options for constructing qubits with silicon. As with diamond, dopant atoms can be added to the crystal; phosphorus and arsenic are common choices. Either the spin of the dopant atom's nucleus or that of the electrons in orbit around it can be used to construct a qubit. Similar spin qubits can also be made artificially, by using electrode and semiconductor structures to trap electrons inside quantum dots.

Using silicon that has been purified of all but one isotope has helped boost the stability of qubit systems; the material now holds the record for the longest qubit coherence times. Silicon also has an advantage when it comes to fabrication, because systems can be constructed using the tools and infrastructure already put in place by the microelectronics industry. But the small size of quantum dots and, to a greater extent, donor systems will make large-scale integration challenging. While scalable architectures exist on paper, they have yet to be demonstrated. So far, research has largely been restricted to single-dopant systems.

TOP: HANSON LAB/TU DELFT; BOTTOM: CHRISTIE SIMMONS/UNIVERSITY OF WISCONSIN–MADISON

This phenomenon, known as carrier freeze-out, describes the point at which most conventional silicon devices stop working. But in 1998, physicist Bruce Kane, now at the University of Maryland, College Park, pointed out that freeze-out could be quite useful for quantum computing. It creates a collection of electrically neutral, relatively isolated atoms that are all fixed in place—a set of naturally stable quantum systems for storing information.

In this setup, information can be stored in two ways: It can be encoded in the spin state of the donor atom's nucleus or of its outermost electron. The state of a particle's spin is very sensitive to changing magnetic fields as well as interactions with nearby particles. Particularly problematic are the spins of other atomic nuclei in the vicinity, which can flip at random, scrambling the state of electron-spin qubits in the material.

But it turns out that these spins are not too much trouble for silicon. Only one of its isotopes—silicon-29—has a nucleus with nonzero spin, and it makes up only 5 percent of the atoms in naturally occurring silicon. As a result, nuclear spin flips are rare, and donor electron spins have a reasonably long lifetime by quantum standards. The spin state of the outer electron of a phosphorus donor, for example, can remain in superposition as long as 0.3 millisecond at 8 kelvins before it's disrupted.

That's about the bare minimum for what we'd need for a quantum computer. To compensate for the corruption of a quantum state—and to keep quantum information intact indefinitely—additional long-lived qubits dedicated to identifying and correcting errors must be incorporated for every qubit dedicated to computation. One of the most straightforward ways to do this is to add redundancy, so that each computational qubit actually consists of a group of qubits. Over time, the information in some of these will be corrupted, but the group can be periodically reset to whatever state the majority is in without disturbing this state. If there is enough redundancy and the error rate is below the threshold for "fault tolerance," the information can be maintained long enough to perform a calculation.

If a qubit lasts for 0.3 ms on average and can be manipulated in 10 nanoseconds using microwave radiation, it means that on average 30,000 gate operations can be performed on it before the qubit state decays. Fault tolerance thresholds vary, but that's not a very high number. It would mean that a quantum computer would spend nearly all its time correcting the states of qubits and their clones, leaving it little time to run meaningful computations. To reduce the overhead associated with error correction and create a more compact and efficient quantum computer, we must find a way to extend qubit lifetimes.

One way to do that is to use silicon that doesn't contain any silicon-29 at all. Such silicon is hard to come by. But about 10 years ago, the Avogadro Project, an international collaboration working on the redefinition of the kilogram, happened to be making some in order to create pristine balls of silicon-28 for their measurements.

ILLUSTRATIONS BY James Provost

IN SILICO QUBITS

THERE ARE TWO potential ways to build a quantum computer from silicon. Donor nuclei, embedded in silicon crystals, and their outermost electrons might be read using single electron transistors [left]. Two sets of gates could be used to set the states of the qubits (hyperfine gates) and control their interactions (exchange gates). Potentially easier to fabricate on large scales will be quantum dots [right], which use electrodes to trap individual electrons in thin layers of material. Illustrated here is one device concept, which employs layers of silicon and silicon germanium.

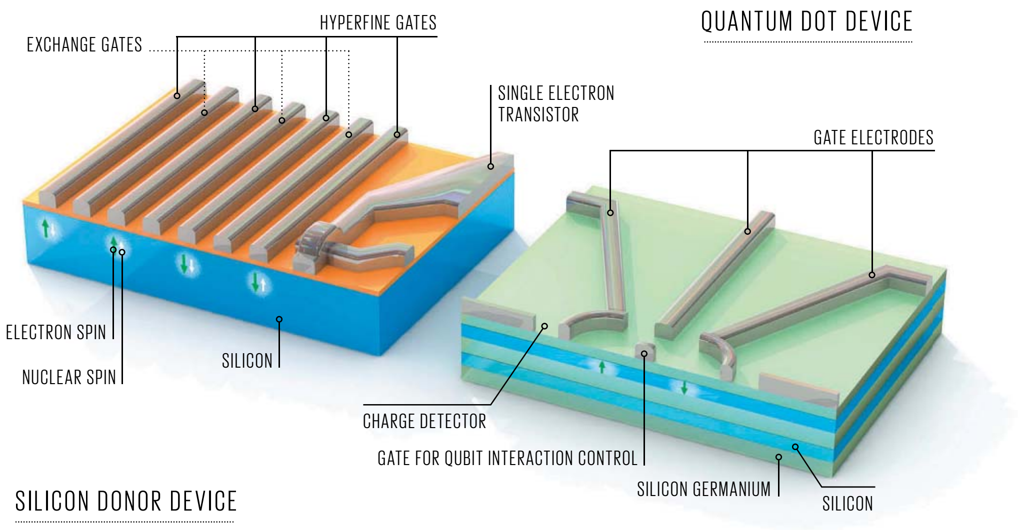

Using a series of centrifuges in Russia, the team acquired silicon that was some 99.995 percent silicon-28 by number, making it one of the purest materials ever produced. A group at Princeton University obtained some of the leftover material and, in 2012, after some careful experimental work, reported donor electron spin lifetimes of more than a second at 1.8 kelvins–a world record for an electron spin in any material. This really showed silicon's true potential and established it as a serious contender.

Our group has since shown that the spins of some donor atoms–bismuth in particular–can be tuned with an external magnetic field to certain "sweet spots" that are inherently insensitive to magnetic fluctuations. With bismuth, we found that the electron spin states can last for as long as 3 seconds in enriched silicon-28 at even higher temperatures. Crucially, we found lifetimes as high as 0.1 second in natural silicon, which means we should be able to achieve relatively long qubit lifetimes without having to seek out special batches of isotopically pure material.

These sorts of lifetimes are great for electrons, but they pale in comparison to what can be achieved with atomic nuclei. Recent measurements led by a team at Simon Fraser University have shown that the nuclear spin of phosphorus donor atoms can last as long as 3 minutes in silicon at low temperature. Because the nuclear spin interacts with the environment primarily through its electrons, this lifetime increases to 3 hours if the phosphorus's outermost electron is removed.

Nuclear spins tend to keep their quantum states longer than electron spins because they are magnetically weaker, and thus their interaction with the environment is not as strong. But this stability comes at a price, because it also makes them harder to manipulate. As a result, we expect that quantum computers built from donor atoms might use both nuclei and electrons. Easier-to-manipulate electron spins could be used for computation, and more stable nuclear spins could be deployed as memory elements, to store information in a quantum state between calculations.

The record spin lifetimes mentioned so far were based on measuring ensembles of donors all at once. But a major challenge remained: How do you manipulate and measure the state of just one donor qubit at a time, especially in the presence of thousands or millions of others in a small space? Up until just a few years ago, it wasn't clear how this could be done. But in 2010, after a decade of intense research and development, a team led by Andrea Morello and Andrew Dzurak at the University of New South Wales, in Sydney, showed it's possible to control and read out the spin state of a single donor atom's electron. To do this, they placed a phosphorus donor in close proximity to a device called a metal-oxide-semiconductor single-electron transistor (SET), applied a moderate magnetic field, and lowered the temperature. An electron with spin aligned against the magnetic field has more energy than one whose spin aligns with the field, and this extra energy is enough to eject the electron from the donor atom. Because SETs are extremely sensitive to the charge state of the surrounding environment, this ionization of a dopant atom alters the current of the SET. Since then, the work has been extended to the control and readout of single nuclear spin states as well.

SETs could be one of the key building blocks we need to make functional qubits. But there are still some major obstacles to building a practical quantum computer with this approach. At the moment, an SET must operate at very low temperatures–a fraction of a degree above absolute zero–to be sensitive enough to read a qubit. And while we can use a single device to read out one qubit, we don't yet have a detailed blueprint for scaling up to large arrays that integrate many such devices on a chip.

ILLUSTRATION BY Bryan Christie Design

THERE IS ANOTHER APPROACH to making silicon-based qubits that could prove easier to scale. This idea, which emerged from work by physicists David DiVincenzo and Daniel Loss, would make qubits from single electrons trapped inside quantum dots.

In a quantum dot, electrons can be confined so tightly that they're forced to occupy discrete energy levels, just as they would around an atom. As in a frozen-out donor atom, the spin state of a confined electron can be used as the basis for a qubit.

The basic recipe for building such "artificial atoms" calls for creating an abrupt interface between two different materials. With the right choice of materials, electrons can be made to accumulate in the plane of the interface, where there is lower potential energy. To further restrict an electron from wandering around in the plane, metal gates placed on the surface can repel it so it's driven to a particular spot where it doesn't have enough energy to escape.

Large uniform arrays of silicon quantum dots should be easier to fabricate than arrays of donor qubits, because the qubits and any devices needed to connect them or read their states could be made using today's chipmaking processes.

But this approach to building qubits isn't quite as far along as the silicon donor work. That's largely because when the idea for quantum-dot qubits was proposed in 1998, gallium arsenide/gallium aluminum arsenide (GaAs/GaAlAs) heterostructures were the material of choice. The electronic structure of GaAs makes it easy to confine an electron: It can be done in a device that's about 200 nanometers wide, as opposed to 20 nm in silicon. But although GaAs qubits are easier to make, they're far from ideal. As it happens, all isotopes of gallium and arsenic possess a nuclear spin. As a result, an electron trapped in a GaAs quantum dot must interact with hundreds of thousands of Ga and As nuclear spins. These interactions cause the spin state of the electron to quickly become scrambled.

Silicon, with only one isotope that carries nuclear spin, promises quantum-dot qubit lifetimes that are more than a hundred times as long as in GaAs, ultimately approaching seconds. But the material faces challenges of its own. If you model a silicon quantum dot on existing MOS transistor technology, you must trap an electron at the interface between silicon and oxide, and those interfaces have a fairly high number of flaws. These create shallow potential wells that electrons can tunnel between, adding noise to the device and trapping electrons where you don't want them to be trapped. Even with the decades of experience gained from MOS technology development, building MOS-like quantum dots that trap precisely one electron inside has proven to be a difficult task, a feat that was demonstrated only a few years ago.

As a result, much recent success has been achieved with quantum dots that mix silicon with other materials. Silicon-germanium heterostructures, which create quantum wells by sandwiching silicon between alloys of silicon and germanium and have much lower defect densities at the interface than MOS structures, have been among the front-runners. Earlier this year, for example, a team based at the Kavli Institute of Nanoscience Delft, in the Netherlands, reported that they had made silicon-germanium dots capable of retaining quantum states for 40 microseconds. But MOS isn't out of the running. Just a few months ago, Andrew Dzurak's group at the University of New South Wales reported preliminary results suggesting that it had overcome issues of defects at the oxide interfaces. This allowed the group to make MOS quantum dots in isotopically pure silicon-28 with qubit lifetimes of more than a millisecond, which should be long enough for error correction to take up the slack.

● ● ●

AS QUANTUM-COMPUTING RESEARCHERS working with silicon, we are in a unique position. We have two possible systems–donors and quantum dots–that could potentially be used to make quantum computers.

Which one will win out? Silicon donor systems–both electron and nuclear spins–have the advantage when it comes to spin lifetime. But embedded as they are in a matrix of silicon, donor atoms will be hard to connect, or entangle, in a well-controlled way, which is one of the key capabilities needed to carry out quantum computations. We might be able to place qubits fairly close together, so that the donor electrons overlap or the donor nuclei can interact magnetically. Or we could envision building a "bus" that allows microwave photons to act as couriers. It will be hard to place donor atoms precisely enough for either of these approaches to work well on large scales, although recent work by Michelle Simmons at the University of New South Wales has shown it is possible to use scanning tunneling microscope tips to place dopants on silicon surfaces with atomic precision.

Silicon quantum dots, which are built with small electrodes that span 20 to 40 nm, should be much easier to build uniformly into large arrays. We can take advantage of the same lithographic techniques used in the chip industry to fabricate the devices as well as the electrodes and other components that would be responsible for shuttling electrons around so they can interact with other qubits.

Given these different strengths, it's not hard to envision a quantum computer that would use both types of qubits. Quantum dots, which would be easier to fabricate and connect, could be used to make the logic side of the machine. Once a part of the computation is completed, the electron could be nudged toward a donor electron sitting nearby to transfer the result to memory in the donor nucleus.

Of course, silicon must also compete with a range of other exciting potential quantum-computing systems. Just as today's computers use a mix of silicon, magnetic materials, and optical fibers to compute, store, and communicate, it's quite possible that tomorrow's quantum computers will use a mix of very different materials.

We still have a long way to go before silicon can be considered to be on an equal footing with other quantum-computing systems. But this isn't the first time silicon has played catch-up. After all, lead sulfide and germanium were used to make semiconducting devices before high-purity silicon and CMOS technology came along. So far, we have every reason to think that silicon will survive the next big computational leap, from the classical to the quantum age. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/silicon0814

Making Bertha

Radar is the key to Mercedes-Benz's robotic car

By JÜRGEN DICKMANN,
NILS APPENRODT
& CARSTEN BRENK

Illustrations by
ANDREW ZBIHLYJ

It is August 2013, and we are sitting in what looks like a standard S-Class Mercedes, nosing through traffic in a small town in southern Germany. The streets are narrow and jam-packed with cars, and pedestrians are everywhere. Yet nobody has a hand on the wheel, and nobody has a foot anywhere near the pedals. Still, you can't fault the driving: This car is in charge of itself.

Or herself. We and our colleagues at Daimler call her Bertha, after the wife of Mercedes-Benz founder Karl Benz, who exactly 125 years earlier became the first joyrider in history when she took her two sons for a 100-kilometer jaunt in her husband's car, from Mannheim to Pforzheim. When the leather brake pads wore out, she found a shoemaker. When the fuel ran out, she bought more from a pharmacist (who marketed it as a cleaning fluid).

Her point was to prove that her husband's internal-combustion engine was ready for general service. Our point is to retrace her famous route and thus prove that autonomous driving is also a reality in the making. You can see bits and pieces of that future already. Without input from the driver, cars can now park, space themselves out on the highway, hold the center of the lane, and even stop when a crash is imminent. These building blocks of autonomous driving have already saved lives. Putting them together to make the perfect robot chauffeur, however, is still a work in progress.

On its way to Pforzheim, our car had to deal autonomously with a number of highly complex situations, including encounters with roundabouts, crossings, traffic lights, pedestrians, cyclists, and trams. As if steered by an invisible hand, Bertha negotiated heavy traffic and narrow streets; it knew just where to turn, when to change lanes, when to stop, and when to start driving again. How could it see enough to perform these feats?

Look at this S-Class from the outside and you will notice nothing out of the ordinary. Get inside, though, and the first secret appears: Behind the windshield hangs a pair of cameras two hands' breadth apart. Like your eyes, they provide depth perception. On either side of the windshield there are more cameras that work independently and across a very wide swath of territory. Their job is to recognize traffic signs. Add to that eight state-of-the-art radar sensors, invisible from the outside, which provide close to 360-degree coverage around the vehicle, sensing objects from a few centimeters to as much as 200 meters away.

Most present-day automotive radars represent cars, pedestrians, and other moving targets as points on a plane, each with an arrow indicating the target's speed and direction of motion. That's not enough information to make Bertha see, though. We had to get the car's radars to provide all the information a human driver would want.

That was tough. But in the end, we taught the radar to track pedestrians, cyclists, and other vehicles moving through junctions and roundabouts, for example. We coaxed it to provide adequate coverage for making lane changes. And we enabled it to determine the boundaries of the lane the car was in up to 140 meters ahead. We also introduced the first algorithms capable of deducing the dimensions of other vehicles or stationary objects. Even more important, those algorithms can tell the difference between a pedestrian and a fence post.

Radar must handle a thousand distractions: Every manhole, every tree, every patch of grass produces reflections

The main challenge in autonomous driving has always been how to teach the vehicle to know where it is, recognize what it sees, and react appropriately. Just as people recognize objects by taking into account their movement, color, shape, and size, autonomous vehicles are at their best when using many different types of sensors. Those so far include ultrasonic and infrared sensors, optical video systems, laser scanners, and radar.

Google's celebrated autonomous car employs an elaborate set of sensors. For detecting objects, it uses a laser scanner mounted on the roof and long-range radars affixed to the front of the car. For recognizing traffic signs and signals, it uses a high-resolution video camera. The data from all this equipment are then superimposed on a stored digital map.

MERCEDES-BENZ

This multilayered approach allows the car to manage inner-city traffic all by itself. It's a great technical accomplishment, one that has energized the entire auto industry. However, until the cost of the equipment comes down, this strategy is perhaps not so practical or economical.

Bertha's design is based on a very different approach, one that relies on compact optical cameras–and radar.

To most drivers, the word *radar* conjures police radar guns, which detect the speed of a targeted car. Unlike optical systems, radar operates well no matter what the weather, working as it does with microwaves. It measures the speed of one object relative to that of another by means of the Doppler effect, most commonly heard in the changing frequency of a train's whistle as it approaches and then retreats from you.

You might think that radar would be far easier to use in cars than in airplanes, which after all have been using radar for generations. Cars are slower; they monitor two dimensions, not three; and they look ahead just 200 meters, not kilometers. But don't forget, the sky is largely empty, reflections are few, and those objects that do appear on a radar screen will certainly be of interest. Here on Earth, though, radar must see through a thousand distractions: Every manhole, every tree, every patch of grass produces reflections.

Solutions that make sense for planes and ships often do not work in cars. The radar sets at airports and on ships typically provide the necessary 360 degrees of coverage by rotating the antenna. But that's just not practical on a car–the rotating dish would be big and conspicuous, and its moving parts might not last for very long. Then there are the radars that fly high above Earth, in satellites or in airplanes, exploiting synthetic-aperture techniques to provide imagelike representations of stationary elements down below. But that's no good for a car either, because it doesn't work up close, and it doesn't produce images fast enough.

Radar manufacturers thus have had to be very inventive to meet automotive requirements. What has aided them most is the development of compound semiconductors, such as indium gallium arsenide and silicon germanium, which can reach frequencies of 76 gigahertz or higher, making possible sensors that are small enough to fit behind the bumper and yet can distinguish a pedestrian from a car from 100 meters away. What's more, these frequencies see through rain and snow, provide good resolution over a wide field, and can be updated every 40 to 60 milliseconds, fast enough to keep a close eye on a changing traffic situation.

Research on automotive radar dates back many decades. One of the first examples was the Eureka Prometheus project, which began in 1987 and ran until 1995. It was a collaboration between the University of Munich and a number of car companies, including our own. In 1994, the project's vehicles traveled around 1,000 km in normal traffic, mainly autonomously, on a multilane motorway near Paris. Then, in the project's finale, they drove from Munich

Radar Eyes on **All Sides**

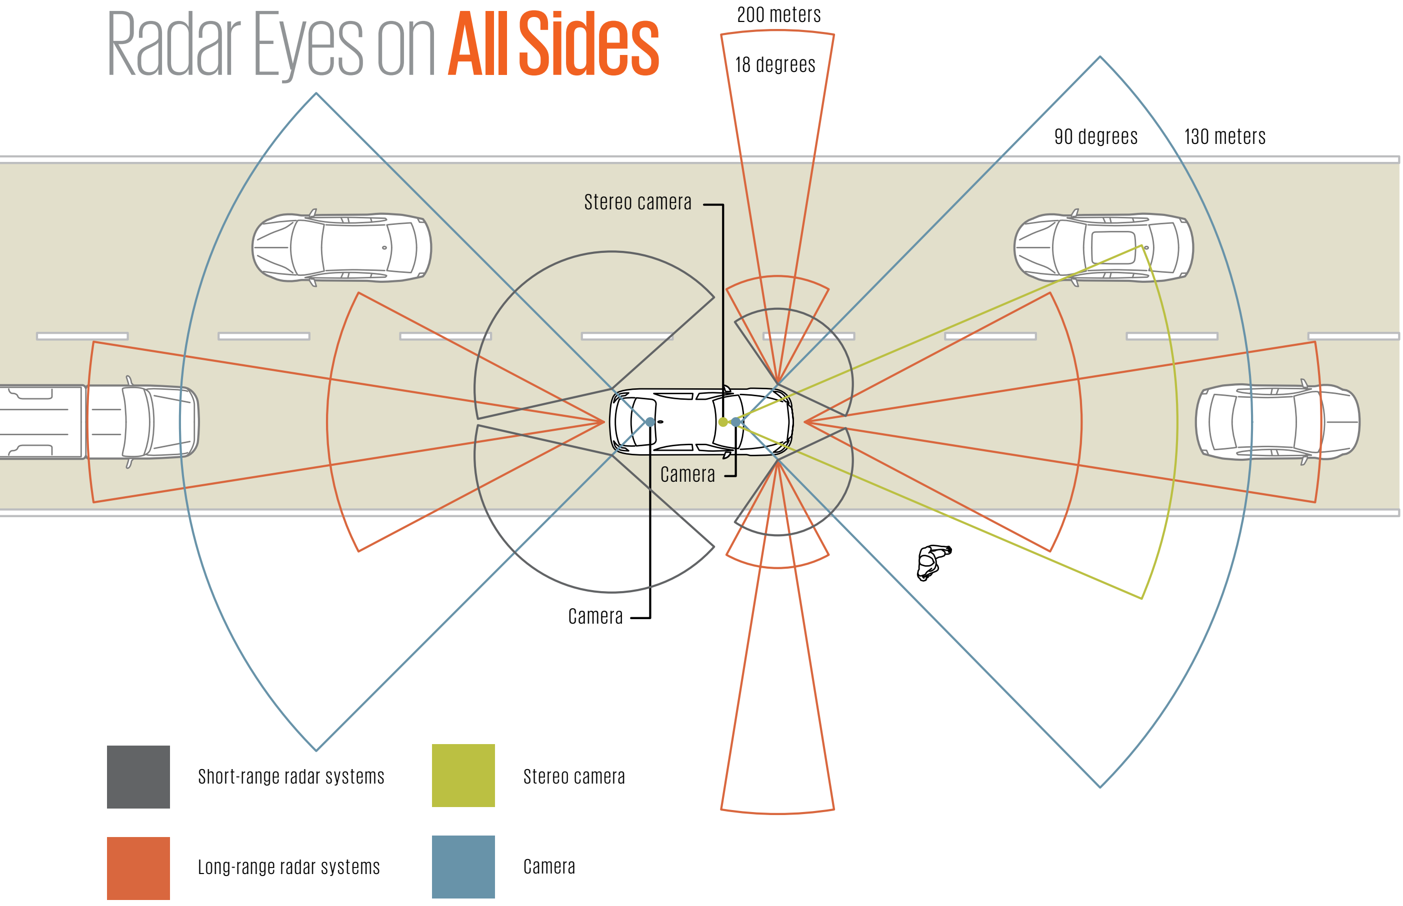

to Copenhagen. It was nearly a decade before autonomous-vehicle research got its next big boost, with the U.S. Defense Advanced Research Projects Agency's Grand Challenge competitions in 2004 and 2005, during which autonomous cars faced off in the desert. Then, in 2007, came the DARPA Urban Challenge competition, followed in 2012 by the introduction of Google's autonomous car.

The first commercial application to come out of the Prometheus project was Daimler's Distronic adaptive cruise control, which went into production in the Mercedes S-Class in 1998. (Toyota had introduced the first commercial adaptive cruise control system the year before.) The system used one long-range radar and two shorter-range units, all of them mounted in the front of the vehicle. Daimler then developed a succession of driver-assistance systems capable of detecting hazardous situations, issuing an alert, and more recently, automatically intervening to avoid an accident.

For instance, a system Daimler calls Speed Limit Assist, which went into production in 2005, warns the driver about going too fast. Another, dubbed Pre-Safe Brake, introduced the following year, automatically applies the brakes if it determines that there's a risk of colliding with the vehicle in front.

The next step was to extend such protection from the system's initial sphere of application, the highway, to urban environments. That's where the two short-range radars came in, in some Mercedes-Benz cars, in 2009. If a collision threatens, the radars of the Pre-Safe Brake system prime the brakes for immediate use. If the short-range radar determines that a crash simply cannot be avoided, the system applies the brakes some 100 ms before impact, substantially reducing damage to the car and its occupants.

Finally, in 2013, the new S-Class boasted an electronic safety "cocoon" spanning nearly 360 degrees, with both short-range radars and a stereo camera. It's enough to protect city drivers from just about any threat, even that of a rear collision. In all these systems, drivers are always in the loop: They can overrule the system, at least up to the last fraction of a second before a crash. To make the next step to full autonomy, we have tried to do as much as possible with the kinds of radars, video cameras, and other sensors already on our S-Class cars. Such sensors, by the way, are also carried in standard production vehicles from Audi, BMW, Ford, Lexus, and Volvo.

We increased the number of sensors and improved their arrangement to achieve 360-degree coverage and to see objects in greater detail. We installed two slightly modified long-range radars from Continental Automotive Group at the sides of the front bumpers to provide early detection of vehicles coming from the left or right at intersections. One additional long-range radar monitors the traffic to the rear. Finally, we mounted four short-range radars at the corners of the vehicle. These units are entirely new, having been developed in collaboration with Delphi Automotive to provide improved coverage of the car's immediate surroundings in crowded settings.

Compound semiconductors can reach frequencies of 76 gigahertz, enough to distinguish a pedestrian from a car from 100 meters away

ILLUSTRATION BY John MacNeill

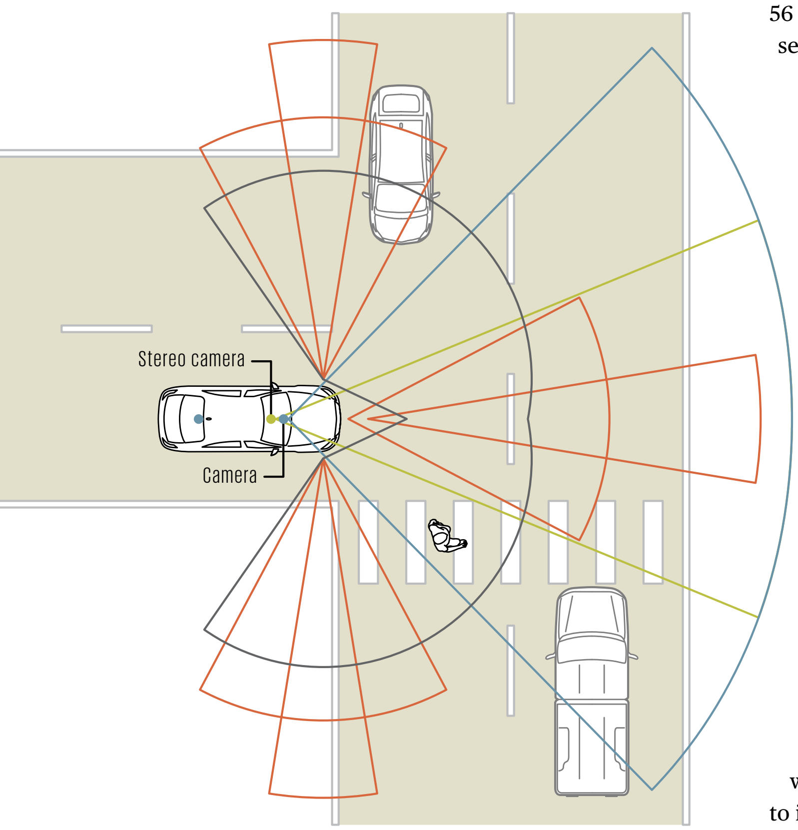

The car must get a precise fix on the location and direction of every object that might collide with it, particularly in dense traffic, in narrow streets, and when facing oncoming traffic. Without such information, Bertha would have hung back timidly at the entrances of roundabouts, waiting perhaps for hours to get a chance to enter.

The beams from these automotive radars are each steered electronically, so they require no rotating antennas or any other moving parts. This way, the system can point the radar in different directions and focus the beam accordingly on objects of interest. We also took advantage of today's higher radar frequencies, which allow for finer resolution of both range and speed. We were able to put it all together to map the environment with the help of algorithms originally devised for use in laser scanners or for image processing.

To help the radar system distinguish people from lampposts, we took two steps. First, we increased the Doppler sensitivity of the device to the point where our system can determine which of a pedestrian's two feet is advancing and which is stationary. Second, we improved the Kalman filtering, a method often used to interpret noisy data collected over a period of time. These two refinements are what let Bertha confidently conclude, yes, that's a pedestrian, moving in this direction and at that speed.

Of course, optical systems can usefully complement the radar. But radar alone works under all weather conditions, provides a full 360 degrees of coverage, and sees up to 200 meters ahead. The stereoscopic cameras in the front, by contrast, span only 56 degrees and can see only 40 meters ahead. Thus, radar must serve as the ultimate backup.

One case where optics work better is in keeping track of the edges of a traffic lane, so that the system can predict the car's trajectory and keep the car within its lane. Video systems are the first choice, but when the light is glaring or the snow is blinding, the radar system must step in. Radars are color-blind, but they can get reflections from objects at the side of the road, like guard rails or even just loose pieces of gravel. To exploit these reflections and to predict the course of winding roads, we developed algorithms that match clothoids, the special curves used in many roads to avoid sudden lateral accelerations.

The most valuable lessons are those that road testing reveal. Of course, such trials have helped us to solve a lot of everyday problems, but more important, they have shown us what still needs to be solved.

One example: Before Bertha's test drive, all our interest was focused on tracking moving objects, but now we know we must also deal with things that stay put. Parked cars often hide the approach of a pedestrian at the side of the road, which means the sensors don't have enough time to predict the pedestrian's movement and adjust the car's driving accordingly. Another example of where we can improve is to extend the car's zone of awareness to include any other autonomous car that may be in front of it. When many cars drive themselves, each of them stands to benefit from what the others are detecting (and planning). Data from all sensors will be shared. Our group has already begun working out how to do that for radar.

Radar, unlike optical systems, provides 360 degrees of coverage and sees 200 meters ahead in any weather

While the engineering of fully autonomous vehicles remains out of reach for the moment for commercial car manufacturers, some benefits accrue immediately after incremental improvements in the various enabling technologies.

The beauty of this project is that although full autonomy may lie well in the future, many of the steps toward that goal produce immediate payoffs. For instance, electronic stability control systems, which Mercedes-Benz introduced back in the 1990s, cut U.S. accident rates by 27 percent for cars and by 67 percent for sport-utility vehicles, according to an analysis by the U.S. National Highway Traffic Safety Administration. Sensors that help drivers react to danger are useful now, not just years from now, when they will be integrated into fully robotic cars.

For the time being, we're not trying to supplant drivers so much as to relieve them of tedium and protect them against their all-too-human blind spots. Should the car one day become a faithful chauffeur, so much the better. Some of us, however, will always love driving and strive to do it ourselves—perhaps with a little help from time to time. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/mercedes0814

AMERICA'S FORGOTTEN ROCKETEER

HOW FRANK MALINA AND HIS UNLIKELY TEAM LAID THE FOUNDATION FOR THE U.S. SPACE PROGRAM

BY JAMES L. JOHNSON

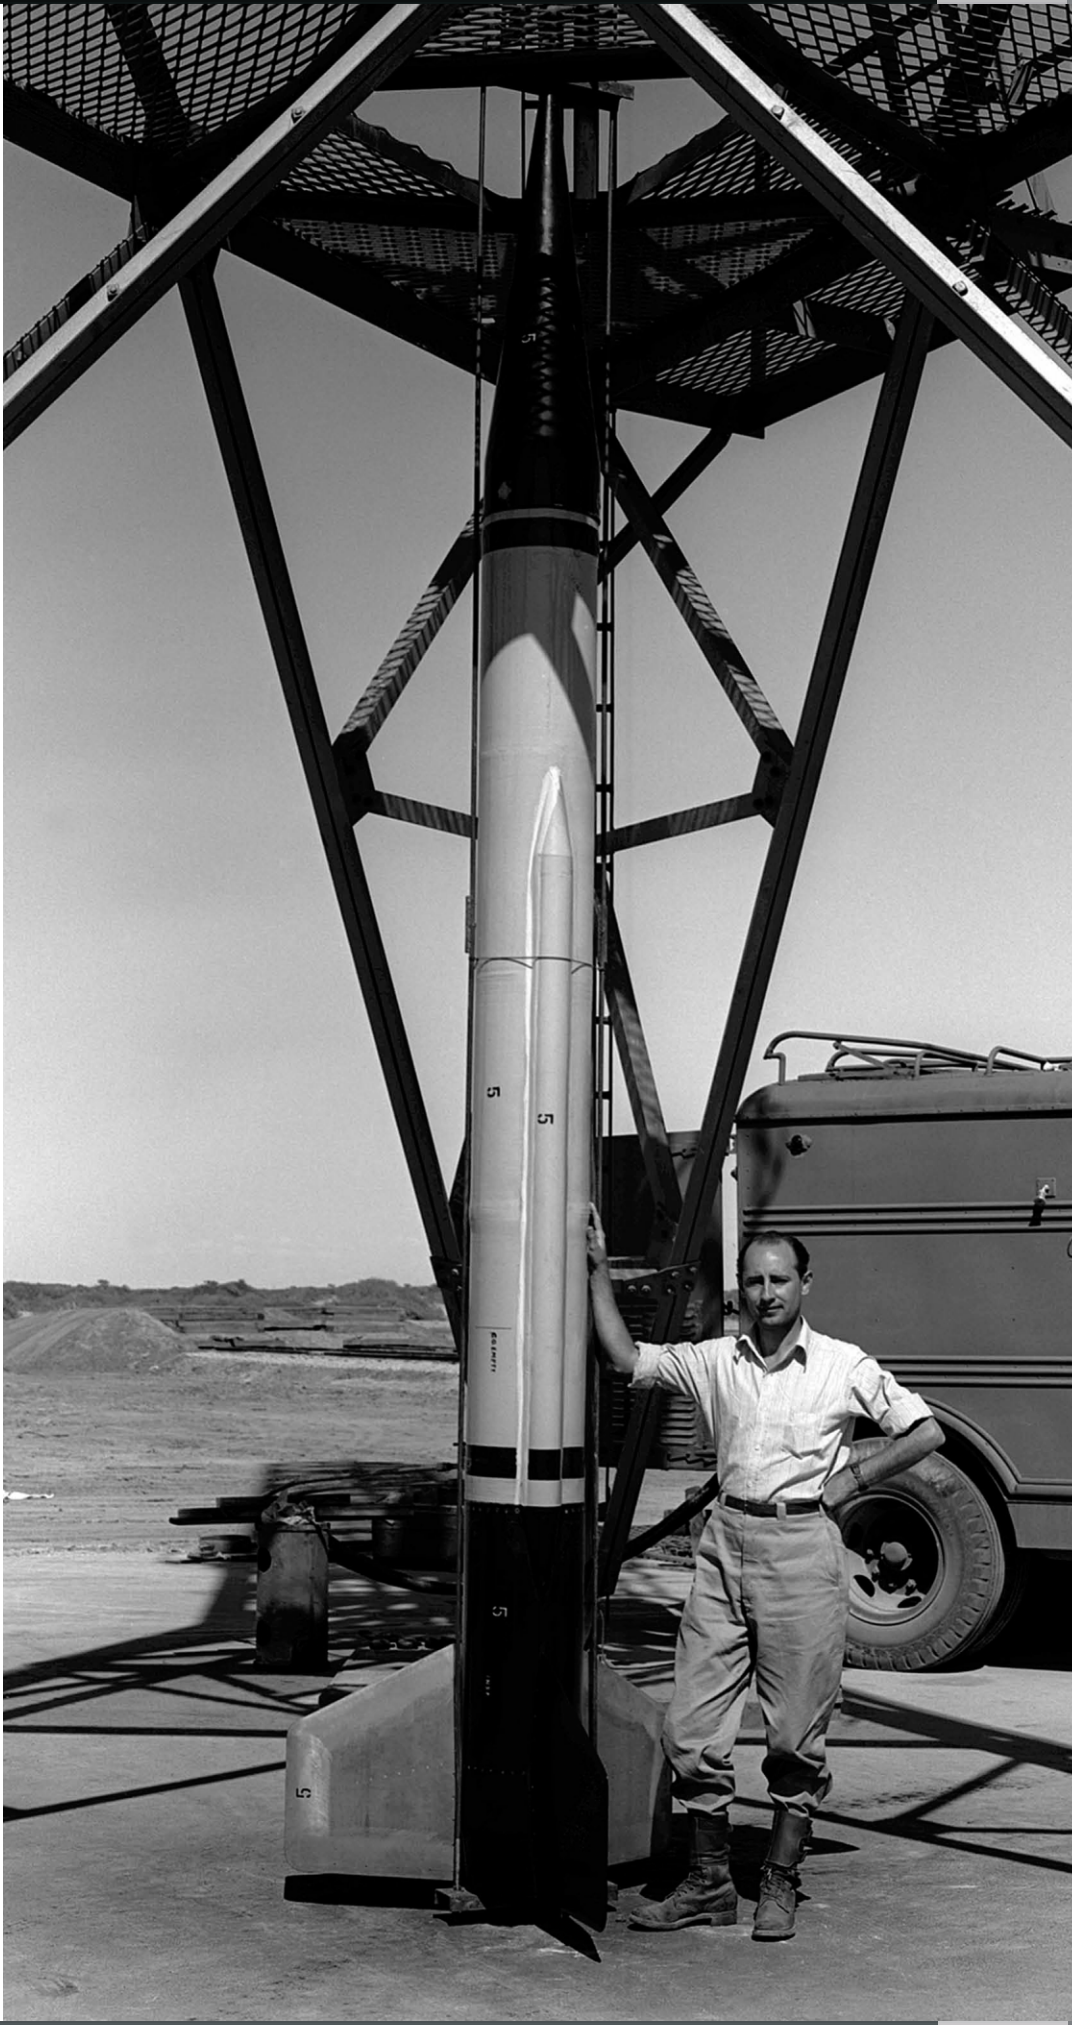

In the early decades of the 20th century, rocket science wasn't considered the brainy endeavor it is now. Far from it: Simply expressing an interest in the field was enough to provoke ridicule. Becoming a rocket scientist was enough to get you ostracized from whatever field you were in before. ¶ Frank Malina didn't care. Overcoming incredible institutional resistance and rather daunting technical and financial odds, the engineer, while still a grad student at Caltech in the mid-1930s, started up a research program that would lay the foundations for U.S. rocket and missile development. During the run-up to World War II, that work took on new significance. By the war's end, Malina had become the top American rocket expert and had cofounded the Jet Propulsion Laboratory, which today is one of the world's premier space research organizations. ¶ And yet, you've probably never heard of him. Most histories of the U.S. space program treat Malina and his group as a footnote. They say the real work started only after the war, with the arrival of Wernher von Braun, Hitler's chief rocket scientist. Without the German's genius, the story goes, U.S. extraterrestrial explorations would never have gone so far so fast.

That version of events, though, overlooks the key contributions made by Malina and his team of engineers, scientists, and technicians, who not only advanced the state of rocketry but did so on a fraction of the funding that their German counterparts enjoyed. Between 1936 and 1946, Malina's team pioneered the use of solid propellants, which in the decades following World War II became crucial in both missiles and launch vehicles, and they also did fundamental work with liquid propellants. Equally important was Malina's institutional legacy, in cofounding both JPL and the Aerojet Engineering Corp. (now Aerojet Rocketdyne), a major aerospace player to this day.

What makes Malina's story all the more compelling is that he was a man of great contradictions: A professed pacifist, he nevertheless designed powerful rockets to further the war effort. A communist sympathizer, he made a fortune through his stake in Aerojet. A consummate engineer, he opted to abandon his research career while still in his 30s and would eventually dedicate himself full-time to artistic pursuits. And yet, this sometimes deeply conflicted individual did more than anyone to legitimize the pursuit of rocket propulsion and to pave the way for others to pursue their paths to the stars.

LIKE MOST OF THE EARLY ROCKETEERS, Malina was drawn to the subject because rockets meant space travel. Born in 1912 in the tiny town of Brenham, Texas, Malina as a boy devoured Jules Verne's classic *From the Earth to the Moon,* which vividly imagined an extraterrestrial trip. Even as an adolescent, Malina had an engineer's mind-set. In a college essay on interplanetary travel, he enumerated the great difficulties that would need to be overcome, including the vast distances to traverse, the hostile atmosphere upon arrival, and the lack of any means of communication between that distant point and Earth.

In 1934, after getting a bachelor's in mechanical engineering from Texas A&M University, he headed to Caltech. There, he had the good fortune to begin working for the renowned aerodynamicist Theodore von Kármán, who became his thesis advisor. Von Kármán led one of the world's foremost centers of aeronautical research, the Guggenheim Aeronautical Laboratory, California Institute of Technology (GALCIT). As the limits of propeller propulsion for high-speed flight became obvious, von Kármán and others eagerly sought alternatives.

Late in 1935, after hearing a fellow student's presentation on rocket-powered aircraft, Malina found his old interest in space-

MEMENTO MEI: Frank Malina, shown at left in 1945, helped legitimize rocketry as an engineering discipline by leading a major research program and by cofounding the Jet Propulsion Laboratory. Yet his significant contributions have been all but forgotten.

JPL (2)

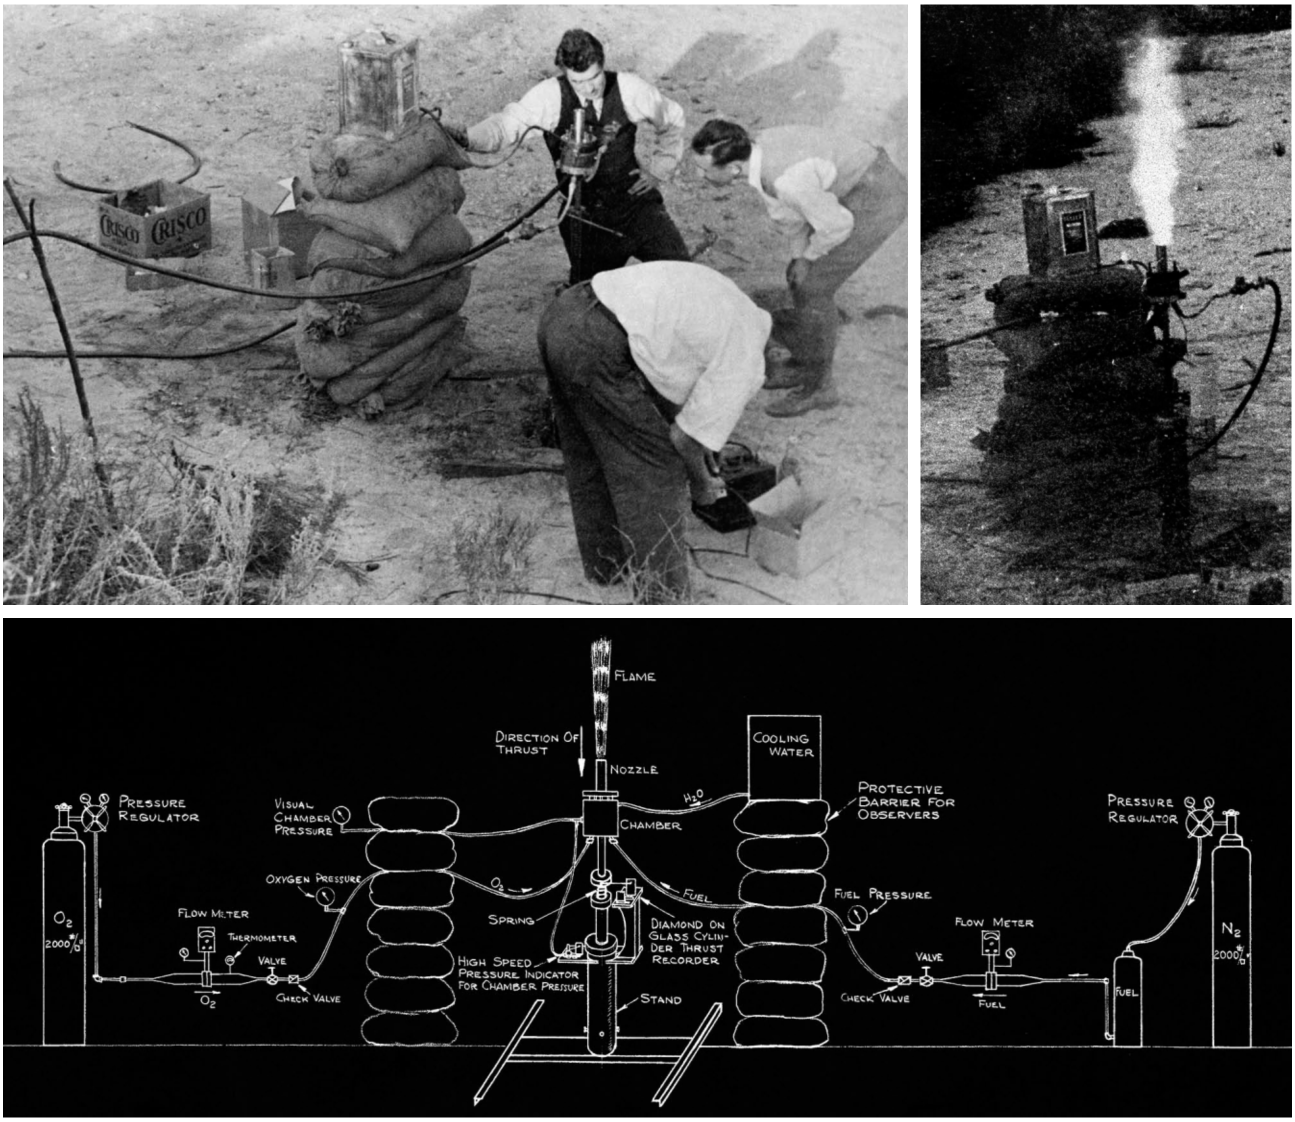

RISKY BUSINESS: Starting in 1935, Malina, then a Caltech grad student, led a colorful and dedicated team of fellow students and rocket enthusiasts, advised by aerodynamics pioneer Theodore von Kármán. The team, which became known around campus as the Suicide Squad for their reckless experiments, scored their first major success in the fall of 1936. Hauling their makeshift equipment out to the Arroyo Seco [top left], a dry river basin on the western edge of Pasadena, Calif., they successfully fired a small methyl-alcohol-fueled rocket motor [top right] on their third day of tests. At bottom is Malina's sketch of the experimental apparatus; note the "protective barrier for observers," otherwise known as a pile of sandbags.

JPL (3)

flight rekindled. Coincidentally, a newspaper article about the presentation drew the attention of Pasadena resident John "Jack" Whitesides Parsons, a self-taught chemist who had been experimenting with powder rockets for some time. Parsons and a mechanic friend visited Caltech, seeking advice on building a liquid-propellant rocket. The student who'd given the presentation directed them to Malina, and Malina, recognizing what he would later describe as Parsons's "uninhibited fruitful imagination," agreed to work with Parsons.

Malina and Parsons were an odd pair. Methodical and reserved, the 23-year-old Malina was very much the academic. Parsons, two years younger, had enough ingenuity, boldness, and exuberance to more than compensate for the fact that he had no degree beyond high school. He also dabbled in magic and the occult. Somehow, this unlikely duo advanced rocket science further than either of them could have possibly done alone.

Malina proposed to von Kármán that he, Parsons, and Parsons's mechanic friend, Ed Forman, design a sounding rocket that would carry scientific instruments into the upper atmosphere, to an altitude of about 40 kilometers. Despite the fact that Parsons and Forman had no Caltech affiliation, von Kármán agreed to support the trio, although he could provide only advice and use of the facilities but no actual funding. Over the next several years, Malina and his crew would pursue their investigations by taking part-time jobs, scrounging spare parts and materials where they could, and "borrowing" from university labs as needed. Malina and Parsons soon attracted additional graduate students to join them, including A.M.O. Smith, who would go on to become chief aerodynamicist at Douglas Aircraft, and Hsue-Shen Tsien, who would later return to China to found its missile and space program.

ROCKETRY TAKES OFF: In 1939, Malina's team received a small grant to develop jet-assisted takeoff rockets for the U.S. Army Air Corps. During flight tests in 1941 [bottom left], the JATO rockets cut the takeoff distance of a small civilian airplane by half; above, von Kármán is shown writing notes while Malina, to his left, looks on. Eventually, the group produced rockets with enough thrust to lift bombers like the Douglas A-20 Havoc [top left].

JPL (3)

Rocketry was still in its infancy. Although rocket clubs in Europe and the United States gave amateurs an outlet for their interests, no serious university programs existed, and thus not much had been done to put theory into practice. Still, Malina's group wasn't quite starting from scratch. Back in the 1890s, Russian mathematician Konstantin Tsiolkovsky's calculations showed that extraterrestrial travel via rocket propulsion was theoretically feasible. And in 1926, the secretive U.S. engineer Robert Goddard launched the world's first liquid-fueled rocket. Its flight lasted just 2.5 seconds, rising to 12.5 meters. But his results inspired a new generation of rocket enthusiasts, including Malina and Parsons and also von Braun.

Consulting with von Kármán, Malina decided to focus first on the rocket engine. Up to then, nobody had built an engine suitable for a sounding rocket–that is, with enough thrust to reach an altitude of 40 km or so. Goddard's rockets, for example, never got higher than 2.6 km. "Until one could design a workable engine with a reasonable specific impulse," Malina later recalled, "there was no point in devoting effort to the design of the rocket shell, propellant supply, stabilizer, launching method, payload parachute, etc." Parsons, who loved nothing more than the thrill of launching rockets, argued against Malina's methodical approach. Fortunately for the sake of science, Malina won the argument.

> The rocketeers were ecstatic. There was just something awesome about setting things ablaze

The scientific rigor imposed by Malina didn't stop the team from taking extraordinary risks. In one memorable experiment, a small rocket motor they were testing misfired, releasing a corrosive cloud of dinitrogen tetroxide that rusted equipment throughout the building. "We were told to move our apparatus outside the building at once," Malina wrote. "We also were thereafter known at Caltech as the 'Suicide Squad.' "

IN LATE OCTOBER 1936, the GALCIT rocketeers were ready to test their rocket motor. Given the likely noise and possibility of an explosion, they chose a spot off campus: the Arroyo Seco, a dry river basin on the western edge of Pasadena. Hauling their makeshift equipment out to the sand, they attached the motor to a 1-meter-tall stand, positioning it so that the exhaust flame would shoot straight up into the air. A spring would measure the rocket's thrust. Hoses connected the motor to a tank of oxygen and another of methyl alcohol fuel. The tests would provide crucial data to back up their calculations, including fuel consumption, thrust, and temperatures and pressures inside the motor.

A cord fuse was supposed to ignite the fuel, but on the first three attempts it detached before it could be lit. On the fourth and final try, the fuse lit but then detached, managing to ignite some fuel that had spilled onto the equipment and also the oxygen hose. "The oxygen hose for some reason ignited and swung around on the ground, 40 feet from us," Malina wrote the next day in a letter to his parents. "We all tore out across the country wondering if our check valves would work." The check valves, designed to prevent the fuel and oxygen from backing up, did work, and although the resulting fire badly damaged their apparatus, the rocketeers were ecstatic. As Parsons and Forman already knew, there was just something awesome about setting things ablaze.

The group duly replaced the fuse with a spark plug, and on 15 November the motor burned for 5 seconds. Two weeks later, it fired for 20 seconds, this time with a deafening roar that indicated complete combustion. On their final test, on 16 January 1937, it burned for a full 44 seconds.

Over the next year and a half, the team continued their experiments as time and money allowed, and Malina published two landmark papers on their work in the *Journal of the Aeronautical Sciences*. Despite the rocketeers' steady progress, however, their results drew little outside interest.

Then came the war. As hostilities in Europe and the Pacific deepened, U.S. military leaders began casting about for any new technologies that might assist in the war effort–including rockets.

FLIGHT PATHS: After World War II, Wernher von Braun [right], father of Nazi Germany's V-2, took Malina's place in the public eye as the leading U.S. rocketeer. In 1950, a two-stage rocket [opposite page] consisting of a sounding rocket (the WAC Corporal, designed by Malina's team) atop a V-2 became the first rocket to be launched from Cape Canaveral, in Florida.


NASA

In December 1938, at von Kármán's behest, Malina presented a report on "jet propulsion" to a group of government and military advisors in Washington, D.C. (He intentionally avoided the word "rockets," which still had a poor reputation in scientific circles.) The report impressed the U.S. Army Air Corps's commanding general, Henry "Hap" Arnold, who in January 1939 gave Malina's group US $1,000 to develop rockets for jet-assisted takeoff (JATO). Six months later, Arnold gave them an additional $10,000. He hoped that JATO rockets, mounted on an airplane's wings, might help a heavily laden aircraft take off from short island runways in the Pacific.

The group experimented with solid fuels, starting with the oldest: gunpowder. After a series of unintended explosions, Malina and von Kármán felt compelled to examine the theoretical stability of burning a solid fuel under pressure. Their conclusion, now known as the von Kármán-Malina theory of constant-thrust long-duration engines, showed the process could be made stable if the pressure inside the chamber remained constant. By the summer of 1941 the group's gunpowder-based propellant was performing well enough to warrant flight tests on an actual aircraft. At March Field, in California, JATO rockets cut the takeoff distance of an Ercoupe, a fighter-size civilian aircraft, by half.

The engineers still had a problem. When the JATO rockets were stored for more than a few days, they would explode upon ignition. For the better part of a year, Parsons and Malina searched in vain for a solution. The bombing of Pearl Harbor in December 1941 only heightened the urgency.

Then, in June 1942, the self-taught Parsons had a brilliant insight. He was watching a construction crew mixing molten asphalt when it occurred to him that he could cast asphalt with an oxidizer, such as potassium perchlorate, to create a solid propellant. The combustible asphalt would act as both fuel and binder. The concept proved to be a fundamental technological breakthrough for all solid propellants, and Parsons's idea lives on in both missiles and launch vehicles, including the Polaris, Minuteman, and Titan. That same year Malina, Parsons, von Kármán, and two others formed the Aerojet Engineering Corp. to manufacture JATO rockets for both the U.S. Army and Navy.

OF COURSE, MALINA AND HIS COLLEAGUES weren't the only rocketeers. In 1943 came British intelligence reports that the Germans were constructing an extraordinarily large rocket at Peenemünde, on the Baltic Sea coast. Over the preceding decade, Wernher von Braun had been leading a well-funded, top-secret effort that was about to show the world the rocket's destructive power.

Like Malina, von Braun was born in 1912 and had been drawn to rocketry as a youth. While still an undergraduate, he joined the Verein für Raumschiffahrt (Spaceflight Society), whose 500 amateur members followed the latest developments in the field and also experimented with rockets. By 1932, their work had attracted the attention of the German army, which recruited von Braun and supported his classified doctoral thesis, "Construction, Theoretical, and Experimental Solution to the Problem of the Liquid Propellant Rocket."

In stark contrast to the U.S. military, the Germans were serious about rockets. By the end of the war, Germany would spend more on rockets than the United States spent on the Manhattan Project—$3 billion versus $1.9 billion. With that kind of largesse, von Braun's impressive facilities included not just the Peenemünde research and production center but also an additional manufacturing center near Nordhausen.

Although von Braun's team developed a reusable JATO rocket, their design was apparently never used. That work would have been eclipsed in any event by the far bigger and more sophisticated V-2 rocket. The world's first production ballistic missile, it burned liquid oxygen and ethyl alcohol, carrying a metric ton of explosives over distances of 320 km. Its guidance system relied on a pair of gyroscopes to steer the fins and vanes; it was not very accurate.

Von Braun's enterprise produced more than 6,000 V-2s, which were used primarily against London and Antwerp, Belgium, starting in September 1944. The actual manufacturing was done by prisoners from the concentration camp Mittelbau-Dora. As the historian Michael J. Neufeld has documented, von Braun went so far as to handpick detainees with technical qualifications for this work. (The prisoners were worked literally to death. In all, about

JPL

12,000 died producing von Braun's rockets; for comparison, the rockets themselves would kill an estimated 9,000 people, many of them civilians.)

In 1943, the U.S. Army shared the British intelligence reports with von Kármán and Malina and asked if they too could develop a long-range guided missile. They could, they said, and the Army provided the newly renamed Jet Propulsion Laboratory with $3 million for the first year of operation. Von Kármán was named JPL's director, and Malina was its chief engineer. Construction of a new facility began in the Arroyo Seco, just west of where Malina, Parsons, and Forman had conducted their crude experiments in 1936. Today, JPL still occupies a 72-hectare campus at that site.

BY THE END OF 1944, VON KÁRMÁN was spending much of his time in Washington, so he resigned his JPL post, and Malina, who had been overseeing daily operations anyway, was soon named acting director.

The lab's crash program to build a smaller, lighter version of the V-2 called for a small solid-fueled rocket and then a larger liquid-fueled guided missile. The former, dubbed the Private, stood 2 meters tall with a planned range of 18 km. It used asphalt as a fuel with potassium perchlorate as an oxidizer.

The liquid-fueled missile, called the Corporal, was 11 meters high and had a range of 120 km. It burned hydrazine with red fuming nitric acid as the oxidizer–a combination the JPL team had developed and patented. The innovative fuel remained liquid at room temperature, so it needed no cooling system, and it was hypergolic, meaning it ignited spontaneously when its constituents were brought into contact, so no ignition system was required. The same type of fuel would later be used in the Apollo program, to propel the command and service module and the lunar module.

As hostilities deepened, U.S. military leaders began casting about for any new technologies that might assist in the war effort—including rockets

Before his team built the guided version of the Corporal, Malina decided to start with an unguided rocket. It was dubbed the WAC, which stood for either "without attitude control" or "Women's Army Corps," because it was the Corporal's "little sister." The first was launched at the White Sands Missile Range in New Mexico on 16 September 1945, two weeks after the official end of World War II. A subsequent launch on 11 October reached an altitude of 70 km, nearly to the boundary of space. A more advanced version of the Corporal would later be deployed in Europe as the United States' first nuclear missile.

Two months after the WAC Corporal launch, Malina's team returned to White Sands to test the Private. In all, they launched 24 rockets without a single failure. Although never itself used as a weapon, the Private was the direct predecessor of all solid-fueled missiles that came after it.

FOR MALINA, THE SUCCESS of the Private and WAC Corporal was bittersweet. He finally had a high-altitude sounding rocket, but it was already clear that the same machine was also a stepping-stone to a nuclear-armed ballistic missile. Malina was at heart a pacifist who had worked on military technology only because he believed the fascists needed to be defeated. In a 1978 interview, he recalled being "caught up in the wave at the end of the war of hate for war, and fear of the development of the atom bomb, and seeing the things we had been developing for space exploration being used for military purposes."

Malina tried after the war to convince Caltech's board of directors to support unclassified high-altitude research based on the WAC Corporal. The board rejected the proposal. In 1946, he and another original member of the Suicide Squad, Martin Summerfield, wrote a paper for the Army describing how technology then available could be used to launch a satellite into orbit. The Army showed no interest.

It seemed the military was interested in Malina's ideas only as they applied to weapons. Malina decided instead to leave rocketry altogether. In 1947 he resigned from JPL and accepted a position with UNESCO in Paris, a position in which he felt he could work toward peace rather than war.

The ascendancy of the German engineers within the U.S. program rankled Malina

That wasn't the only reason for Malina's departure, however. The FBI had been investigating him since 1942, suspecting him of being a member of the Communist Party and, worse, a communist spy. In 1946, bureau agents raided his house while he was out of town. His abrupt transition from war hero to potential enemy of the state was surely galling.

Was Malina in fact a communist? In 2009 I studied Malina's considerable FBI file, and I also went through his papers at the Library of Congress. The records show clearly that Malina was likely a member of a Los Angeles branch of the Communist Party in the late 1930s. His FBI file, for instance, contains a copy of a 1939 application to the Communist Party, in what appears to be Malina's handwriting. He was also no fan of capitalism. In a 1936 letter to his parents, he wrote, "Events in Europe are certainly leading to another war. There seems to be only one hope, overthrowing of the capitalist system in all countries and an economic union of all nations." Of course, at the height of the Great Depression, countless academics, artists, professionals, and others held such views. And, as JPL historian Erik M. Conway has written, the Communist Party branch to which Malina belonged "dissolved after the shocking announcement of the Soviet Union's nonaggression pact with Germany in 1939."

As for espionage, there was perhaps reason to at least suspect Malina. Several security breaches occurred during his tenure at JPL, the most significant involving classified lab documents that turned up in the hands of a Russian courier. According to a 1942 FBI report, at least five unnamed informants identified Malina as a possible spy; the report concluded that "the loyalty of the subject would be questionable if he had to decide between our form of government and that of Russia." J. Edgar Hoover himself repeatedly prodded the U.S. attorneys to indict Malina, which they finally did in December 1952, for failing to disclose his Communist Party status to the government. His U.S. passport was also revoked.

And yet, despite numerous investigations from 1942 until 1960, the bureau never found any evidence of spying or of more than a passing interest in communism. More likely, Malina was just one of the thousands of wrongly accused Americans ensnared by the Red Scare of the early 1950s. The indictment against him was dismissed in 1954, and his passport was restored four years later.

By then, Malina had resigned from UNESCO, and, now wealthy from his stock in Aerojet, he "cut loose from everything and became an artist," as he later told an interviewer. He enjoyed some success as a kinetic sculptor, often invoking themes of science and engineering in his work. He never returned to research, although in 1960 he helped found the International Academy of Astronautics with von Kármán and others, to foster international cooperation in space exploration.

EVEN AS MALINA'S PLACE among rocketeers faded, von Braun's grew ever brighter. As the war in Europe transitioned to the Cold War, the U.S. government brought more than 1,500 German scientists, engineers, and technicians to the United States, including von Braun and much of his staff. Nazi party affiliations and war crimes were conveniently overlooked because the Germans' expertise was now considered crucial in the race against the Soviets. In October 1945, von Braun and part of his team arrived at Fort Bliss, in El Paso, Texas, just a short drive from where Malina's WAC Corporal was being tested. Von Braun would spend five years there, before being transferred to Huntsville, Ala., where he led the Army's rocket program. In 1960, he became the first director of NASA's new Marshall Space Flight Center, in Huntsville, overseeing the development of the heavy-lift Saturn rockets that would carry astronauts to the moon.

The ascendancy of the German engineers within the U.S. program rankled Malina. According to Malina's son Roger, "He was philosophically bitter that the Nazi engineers had become U.S. space heroes but the founders of U.S. rocketry who had dedicated the war years to working for the Allies had been dispersed." And in a 1967 article, Malina himself wrote: "Popular opinion, even the opinion of some who should know better, has been that American rocket developments lagged far behind that of Nazi Germany. This belief is false, but myths die hard."

One last series of experiments is worth noting in this regard. Starting in late 1946, JPL researchers at White Sands assembled and launched two-stage rockets consisting of a WAC Corporal atop a German V-2. The second-stage WAC Corporal would take off just as the bigger missile reached its maximum velocity. On 24 February 1949, one such "Bumper" rocket broke the altitude record by climbing to 393 km, approaching the orbit of today's International Space Station. The following year, another Bumper became the first rocket to be launched from the newly constructed Cape Canaveral. In so doing, this odd American-German hybrid ushered in the space age. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/malina0814

Joint Institute of Engineering

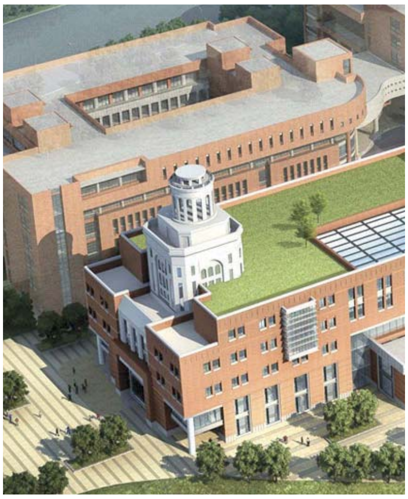

FACULTY POSITIONS AVAILABLE IN ELECTRICAL/COMPUTER ENGINEERING

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

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

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

PLEASE VISIT: jie.cmu.edu for details

SHUNDE INTERNATIONAL

Joint Research Institute

RESEARCH STAFF POSITIONS AVAILABLE IN ELECTRICAL/COMPUTER ENGINEERING

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

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

Candidates with industrial experiences are preferred.

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

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

Application review will continue until the position is filled.

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

SUN YAT-SEN UNIVERSITY Carnegie Mellon University


»IRGEND LERNT MAN«
»COGITO, ERGO SUM«


IEEE


Proceedings of the IEEE
OPTICS & PHOTONICS FOR SECURITY & DEFENSE
Proceedings of the IEEE
ADVANCES IN MAGNETIC DATA STORAGE TECHNOLOGIES

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 September 15th, 2014. For more information, visit **http://sst.nu.edu.kz.**

TU WIEN TECHNISCHE UNIVERSITÄT WIEN Vienna University of Technology

FAKULTÄT FÜR ELEKTROTECHNIK UND INFORMATIONSTECHNIK
Faculty of Electrical Engineering and Information Technology

Technische Universität Wien, Fakultät für Elektrotechnik&Informationstechnik, invites applications for the faculty position of a

FULL PROFESSOR for MICRO-and NANOSENSORS

at the Institut für Sensor- und Aktuatorsysteme (http://www.isas.tuwien.ac.at) which provides state-of-the-art silicon micro-technology and participates in the Zentrum für Mikro- und Nanostrukturen (http://zmns.tuwien.ac.at). Candidates should have an excellent scientific background and strong expertise in at least one of the following research areas: novel integrated sensor concepts, novel sensing principles for physical and chemical quantities, lab-on chip, technology of micro- and nanofabricated sensors, micro- und nanofabricated energy systems. More information concerning this position is available at **http://etit.tuwien.ac.at/fakultaet/offene-stellen-open-positions/**

Applications are expected until October 20th 2014 and should be sent along with a scientific CV and list of publication to the Dean of the Fakultät für Elektrotechnik und Informationstechnik, Technische Universität Wien, Dekanatszentrum Erzherzog Johann -Platz, Gusshausstrasse 30/4, A-1040 Wien.

Carleton UNIVERSITY
Canada's Capital University

The Department of Systems and Computer Engineering invites applications from qualified candidates for an Industrial Research Chair appointment in the field of "Sensor Technology for the Internet of Everything (IoE Sensors)" at the rank of Assistant, Associate or Full Professor. Applications will be considered until the position is filled. The Chair is sponsored by Cisco Systems Canada Co. ("Cisco"). The appointment of the Chair holder will normally be for a five (5) year term, renewable after appropriate review by the University and any federal funding partner such as NSERC.

The successful candidate must have a Ph.D. degree in electrical and computer engineering or the equivalent. At the time of appointment or within two years of the appointment, the candidate must have a membership in a Canadian professional engineering association. The successful candidate will be expected to demonstrate research excellence and to effectively teach electrical and computer engineering courses at both the undergraduate and graduate levels.

The primary purpose of the Industrial Research Chair will be to promote, support, and lead research, development and innovation through industry-linked projects in advanced IoE sensor technologies and its use in transforming business architectures to reduce cost and increase productivity. The Chair will also connect to the University's strategic goals and long-range vision to position the University as a one of the world's premier institutes of higher learning.

The successful candidate must satisfy the requirements for a Natural Sciences and Engineering Research Council of Canada (NSERC) Industrial Research Chair, be a distinguished researcher with an international reputation in innovation and an acknowledged expertise in technologies relevant to the specific scope of this position. The Cisco Chair in IoE Sensors will focus on projects and research activities that utilize state-of-the-art technologies to develop new innovative sensor solutions for the IoE for transforming advances in these technologies into innovative solutions that improve efficiency, reliability and reduce cost. The Chair is expected to attract significant funding from government and industrial partners as well as have a strong ability to supervise graduate students at the Doctorate and Masters level.

The Department of Systems and Computer Engineering is a research-intensive department and hosts a large and active community of some 45 research faculty and adjunct professors and over 230 graduate students. In addition, the department participates in a broad range of undergraduate programs covering the ECE spectrum.

Applicants are expected to submit an application electronically (to cisco-chair@sce.carleton.ca) that includes five PDF documents as follows: 1) curriculum vitae; 2) a detailed research plan; 3) a description of teaching philosophy and preference; 4) a summary of research publications; 5) and the names of five referees.

Located in Ottawa, Ontario, Carleton University is a dynamic and innovative research and teaching institution committed to developing solutions to real world problems by pushing the boundaries of knowledge and understanding. Its internationally recognized faculty provide for more than 27,000 students from every province and more than 100 countries around the world with academic opportunities in more than 65 programs of study.

All qualified candidates are encouraged to apply. Canadians and permanent residents will be given priority. All positions are subject to budgetary approval.

A more detailed description of the position is available at **www.sce.carleton.ca/dept/dept-files/Cisco-chair-position.pdf**

VODAFONE REVEALS PHONE MONITORING NUMBERS BUT MANY NATIONS BAN STATISTICS ON SURVEILLANCE

Earlier this year, the international mobile phone carrier Vodafone Group released figures showing how often it had intercepted a call or supplied call metadata at the request of a government. The most interesting aspect was not the raw numbers of requests (which can yield very misleading figures due to differences in market share and what a request covers in various jurisdictions), but how willing governments were to allow Vodafone to disclose the information. —STEPHEN CASS

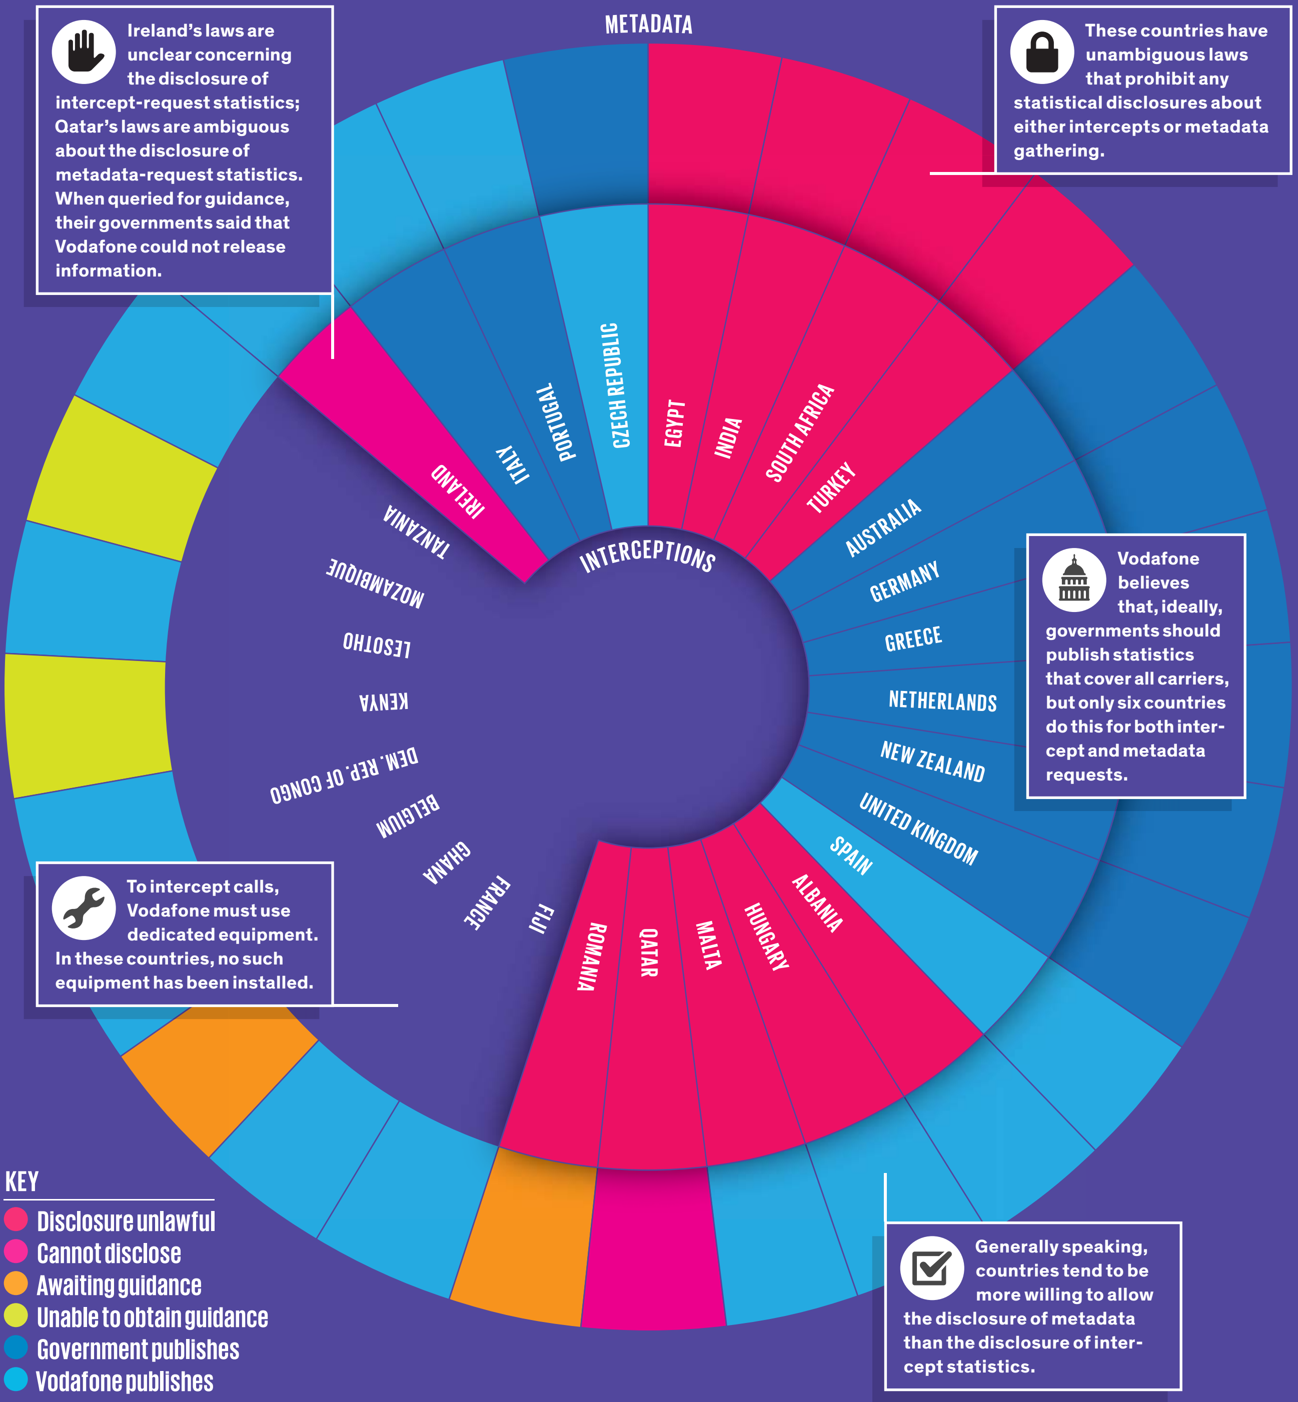

Ireland's laws are unclear concerning the disclosure of intercept-request statistics; Qatar's laws are ambiguous about the disclosure of metadata-request statistics. When queried for guidance, their governments said that Vodafone could not release information.

These countries have unambiguous laws that prohibit any statistical disclosures about either intercepts or metadata gathering.

Vodafone believes that, ideally, governments should publish statistics that cover all carriers, but only six countries do this for both intercept and metadata requests.

To intercept calls, Vodafone must use dedicated equipment. In these countries, no such equipment has been installed.

Generally speaking, countries tend to be more willing to allow the disclosure of metadata than the disclosure of intercept statistics.

SOURCE: VODAFONE GROUP

INFOGRAPHIC BY Brandon Palacio


Agilent's Electronic Measurement Group
Keysight Technologies