IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 04.16

TOP TECH CARS 2016

Mercedes-Benz F 015 concept car

IEEE

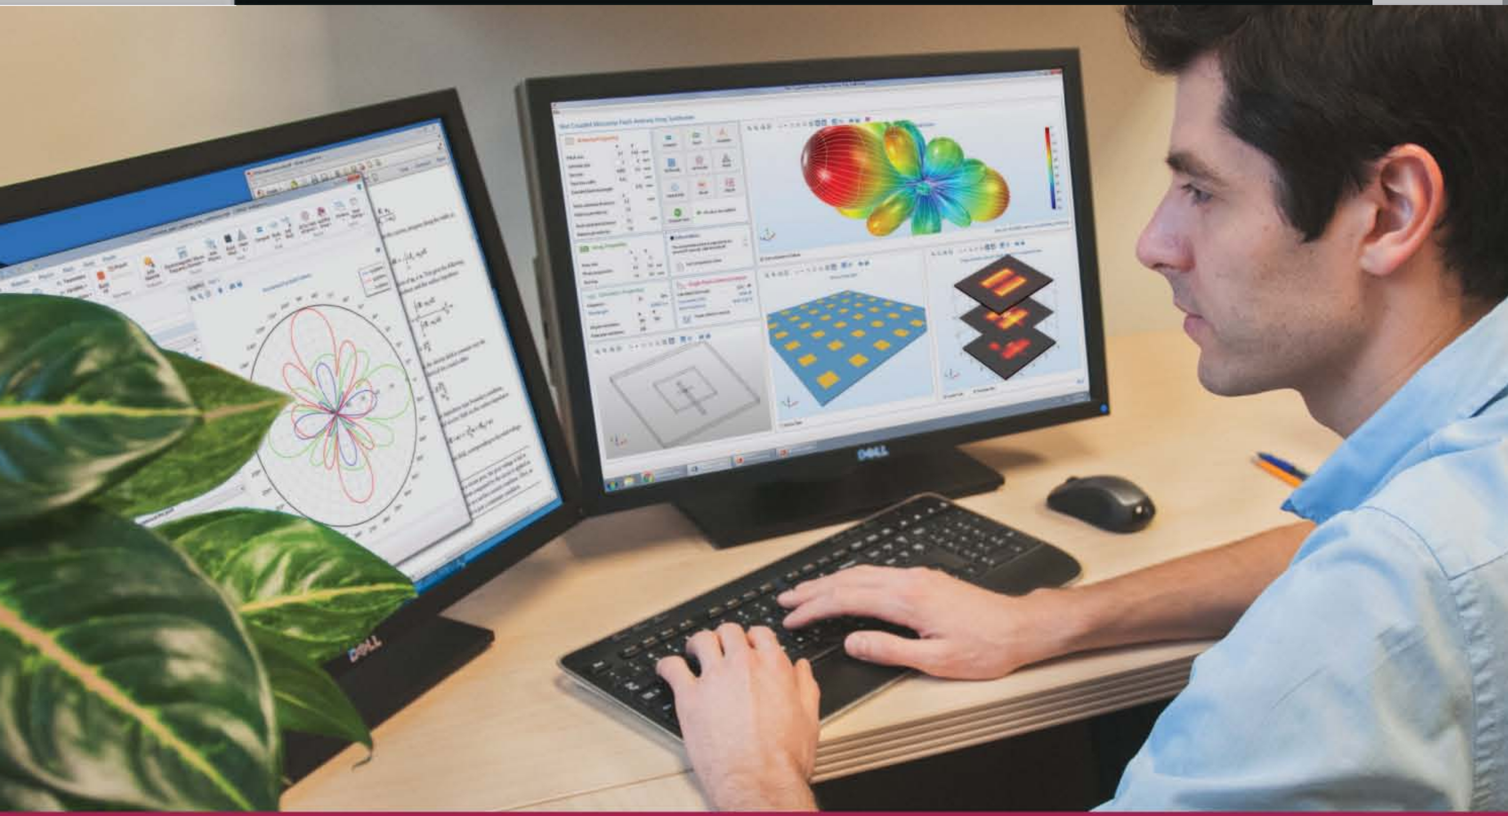

MULTIPHYSICS FOR EVERYONE

The evolution of computational tools for numerical simulation of physics-based systems has reached a major milestone.

Custom applications are now being developed by simulation specialists using the Application Builder in COMSOL Multiphysics®.

With a local installation of COMSOL Server™, applications can be deployed within an entire organization and accessed worldwide.

Make your organization truly benefit from the power of analysis.

comsol.com/application-builder

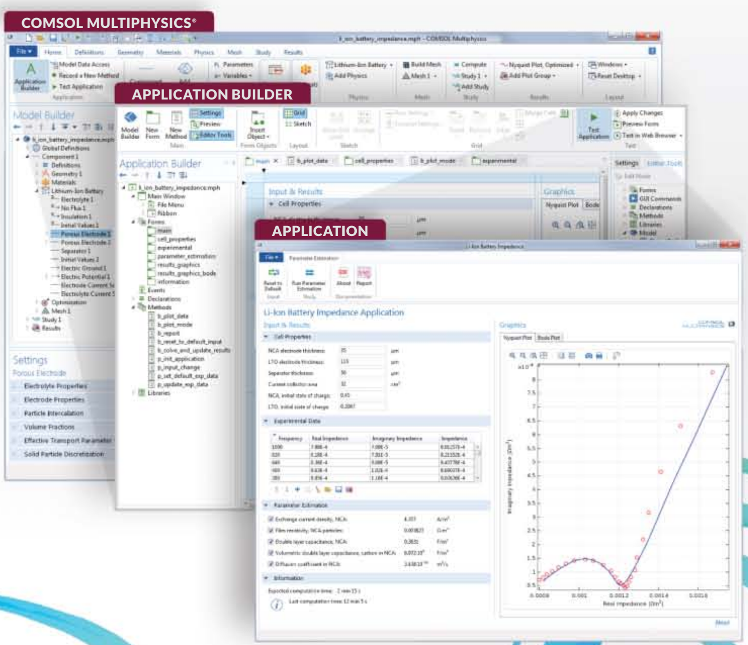

COMSOL

© Copyright 2016 COMSOL. COMSOL, COMSOL Multiphysics, Capture the Concept, COMSOL Desktop, COMSOL Server, LiveLink, and Simulation for Everyone are either registered trademarks or trademarks of COMSOL AB. All other trademarks are the property of their respective owners, and COMSOL AB and its subsidiaries and products are not affiliated with, endorsed by, sponsored by, or supported by those trademark owners. For a list of such trademark owners, see www.comsol.com/trademarks

IEEE SPECTRUM

On the Cover **Photograph by Stefan Milev/Wildfoxrunning**

PHOTOGRAPH BY The Voorhes

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

SMART DEVICES REQUIRE

SMARTER

AUTOMATED TEST SYSTEMS

The old approach to automated test isn't scaling, but you already knew that. Look at your balance sheet. To test smart devices, you need a smarter test system built on a platform of NI PXI, LabVIEW, and TestStand. More than 35,000 companies deploy NI technology to lower their cost of test—what are you waiting for?

Prepare for the future at ni.com/automatedtest

NATIONAL INSTRUMENTS™

NI PXI, LabVIEW, and TestStand

©2016 National Instruments. All rights reserved. LabVIEW, National Instruments, NI, ni.com, and NI TestStand are trademarks of National Instruments. Other product and company names listed are trademarks or trade names of their respective companies. 25202

DEPARTMENTS_04.16

IEEE SPECTRUM



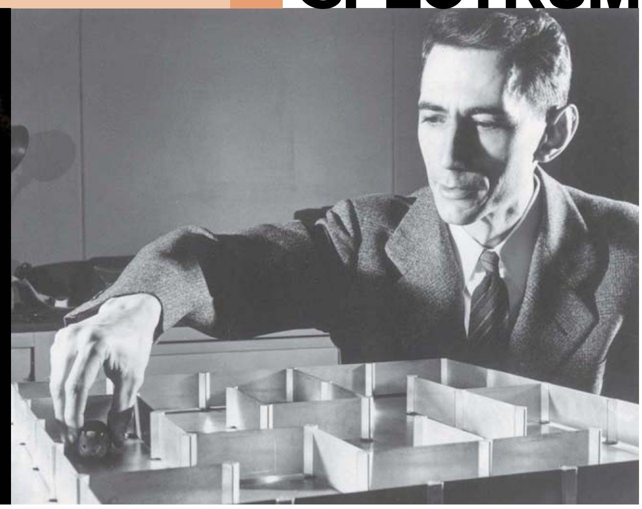

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Charged Particle Tracing Simulations—7 April
- Essential Technologies for Smart Machine Control (sMC)—Available On Demand
- Nonlinear Magnetic Materials With COMSOL Multiphysics—Available On Demand
- **NEW PRODUCT RELEASE LIBRARY** http://spectrum.ieee.org/static/new-product-release-library
- **WHITE PAPERS** http://spectrum.ieee.org/whitepapers

The Institute

Available at theinstitute.ieee.org

- **WOMEN IN ENGINEERING** This month we focus on issues that affect female engineers. IEEE Women in Engineering has revamped its program to better meet the needs of its members. We also report on what VMWare and other companies are doing to attract and retain female engineers. And several of IEEE's top female leaders give advice on getting ahead.
- **BALANCING ACT** What does it take to be a successful engineer while also juggling family responsibilities? Ericsson's global head of talent acquisition, Gunjan Aggarwal, offers words of wisdom.
- **ZERO-FUEL AIRCRAFT** Member Paige Kassalen, chair of the Pittsburgh IEEE Women in Engineering affinity group, is working on the Solar Impulse project. It aims to fly a fixed-wing airplane around the world using only solar power.

FROM LEFT: ADRIAN LEE; DAVID ALAYETO/GETTY IMAGES; KEYSTONE/GETTY IMAGES

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2016 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 53, issue No. 4, 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/about/corporate/governance/p9-26.html.

BACK STORY_

Notes From Underground

ON A DAMP, cold day in January, Mark Peplow found himself at the bottom of a massive hole in the ground. By choice. He was there to report on the construction of Crossrail, London's vast new US $21 billion underground rail network.

The hole was one of a pair of 41-meter-deep shafts in East London, on a spit of land overlooking the Thames. Back in late 2012, two tunnel-boring machines had been lowered into the shafts. The rotating heads on the 1,000-metric-ton behemoths (named Elizabeth and Victoria) chewed their way westward, finally reaching Central London in May 2015. When the Crossrail network is fully operational in 2019, it will carry some 200 million passengers every year.

In "The Digital Underground," in this issue, Peplow describes the ingenious engineering being brought to bear on Crossrail: wireless sensors and lasers that monitor construction, smart components that warn of impending failure, and a 3-D model of the whole network that can be explored from an iPad.

But on this visit, he says, he saw "the blunt end of this construction project." Peplow [above] and his hosts had reached the tunnel floor by descending a cramped staircase crafted from scaffolding and boards. Halfway down, they encountered a small shrine bolted to the wall. Site manager Peter Kelly explained that it was for Saint Barbara, patron saint of mine workers and tunnelers.

Returning to the surface, Peplow noticed workmen clearing a short, rubble-strewn stretch. Their breath clouding the air from the exertion, they were loading hunks of concrete into wheelbarrows. "There's no doubt that Crossrail is a high-tech railway," Peplow says, "but building it still involves a whole lot of mud, concrete, shovels, and sweat." ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, Past Forward is in *IEEE Spectrum*, Vol. 53, no. 4 (INT), April 2016, p. 56, or in *IEEE Spectrum*, Vol. 53, no. 4 (NA), April 2016, p. 64.

04.16

IEEE SPECTRUM

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

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

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

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

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

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

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

SENIOR ASSOCIATE EDITORS
Rachel Courtland, r.courtland@ieee.org
Eliza Strickland, e.strickland@ieee.org
ASSOCIATE EDITORS
Celia Gorman (Multimedia), celia.gorman@ieee.org
Amy Nordrum, a.nordrum@ieee.org
VIDEO PRODUCER Kristen Clark, kristen.clark@ieee.org
ASSISTANT EDITOR Willie D. Jones, w.jones@ieee.org
SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

ADMINISTRATIVE ASSISTANT
Ramona L. Foster, r.foster@ieee.org

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

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

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

EDITORIAL ADVISORY BOARD
Susan Hassler, *Chair*; David Brock, Jason Cong*, Sudhir Dixit, Limor Fried, Robert Hebner, Joseph J. Helble, Grant Jacoby, Leah Jamieson, Deepa Kundur, Norberto Lerendegui, Steve Mann, Allison March, Jacob Østergaard, Umit Ozguner, Thrasos Pappas, John Rogers, Jonathan Rothberg, Umar Saif, Takao Someya, Yu Zheng, Kun Zhou, Edward Zyszkowski
* Chinese-language edition

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

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

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.

COPYRIGHTS AND TRADEMARKS *IEEE Spectrum* is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Careers, EEs' Tools & Toys, EV Watch, Progress, Reflections, Spectral Lines, and Technically Speaking are trademarks of IEEE.

TOM LAWSON

IEEE

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

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

MANAGING DIRECTOR, PUBLICATIONS Michael B. Forster

IEEE BOARD OF DIRECTORS
PRESIDENT & CEO Barry L. Shoop, president@ieee.org
+1 732 562 3928 FAX: +1 732 465 6444

PRESIDENT-ELECT Karen Bartleson
TREASURER Jerry L. Hudgins **SECRETARY** Parviz Famouri
PAST PRESIDENT Howard E. Michel

VICE PRESIDENTS
S.K. Ramesh, Educational Activities; Sheila Hemami, Publication Services & Products; Wai-Choong "Lawrence" Wong, Member & Geographic Activities; Bruce P. Kraemer, President, Standards Association; Jose M.F. Moura, Technical Activities; Peter Alan Eckstein, President, IEEE-USA

DIVISION DIRECTORS
Maciej Ogorzalek (I); Hirofumi Akagi (II); Celia L. Desmond (III); William W. Moses (IV); Harold Javid (V); Rob Reilly (VI); Alan C. Rotz (VII); John W. Walz (VIII); Ray Liu (IX); Kazuhiro Kosuge (X)

REGION DIRECTORS
Ronald A. Tabroff (1); Timothy P. Kurzweg (2); James M.Conrad (3); Robert C. Parro (4); Francis B. Grosz Jr. (5); Thomas Coughlin (6); Witold M. Kinsner (7); Costas M. Stasopoulos (8); Antonio C. Ferreira (9); Ramakrishna Kappagantu (10)

DIRECTORS EMERITUS Eric Herz, Theodore W. Hissey

IEEE STAFF
EXECUTIVE DIRECTOR & COO James Prendergast
+1 732 502 5400, james.prendergast@ieee.org
CHIEF INFORMATION OFFICER Cherif Amirat
+1 732 562 6399, c.amirat@ieee.org
PUBLICATIONS Michael B. Forster
+1 732 562 3998, m.b.forster@ieee.org
CORPORATE ACTIVITIES Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
HUMAN RESOURCES Shannon Johnston, SPHR
+1 732 562 6343, s.johnston@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Eileen M. Lach, +1 212 705 8990, e.m.lach@ieee.org
CHIEF MARKETING OFFICER Patrick D. Mahoney
+1 732 562 5596, p.mahoney@ieee.org
EDUCATIONAL ACTIVITIES Jamie Moesch
+1 732 562 5514, j.moesch@ieee.org
CHIEF FINANCIAL OFFICER Thomas R. Siegert
+1 732 562 6843, t.siegert@ieee.org
TECHNICAL ACTIVITIES Mary Ward-Callan
+1 732 562 3850, m.ward-callan@ieee.org
MANAGING DIRECTOR, IEEE-USA Chris Brantley
+1 202 530 8349, c.brantley@ieee.org
IEEE PUBLICATION SERVICES & PRODUCTS BOARD
Sheila Hemami, *Chair*; John Baillieul, Babak D. Beheshti, Herbert S. Bennett, Jennifer T. Bernhard, Eddie Custovic, Jean-Luc Gaudiot, Ron B. Goldfarb, Glenn Gulak, Lawrence Hall, Christofer Hierold, James M. Irvine, Hulya Kirkici, Carmen S. Menoni, Paolo Montuschi, Thrasos Pappas, Michael Pecht, Michael Polis, Sorel Reisman, Tariq Samad, Fred Schindler, Gianluca Setti, Gaurav Sharma, Curtis A. Siller, David Tian, Ravi Todi, H. Joel Trussell, Leung Tsang, Stephanie M. White, Steve Yurkovich

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

CONTRIBUTORS_

Katherine Bourzac

Bourzac, a freelance reporter based in Silicon Valley, formerly worked the nanotech beat for *MIT Technology Review*. In February she covered the IEEE International Solid-State Circuits Conference for *IEEE Spectrum*. The stories she came back with included one about a silicon cochlea to let robots hear more like humans and another story, in this issue [p. 10], about trying to cram complex neural networks into GPUs. "I was fascinated by the way the neuroscience informs the design of circuits," she says.

Paul McFedries

In *Spectrum*'s bimonthly "Technically Speaking" column, McFedries tracks how language constantly evolves in response to new technologies. In this issue, he looks at the terminology that's cropping up to describe the digitization of analog information [p. 24]. McFedries's website Word Spy chronicles neologisms generally. In June, he'll be releasing a collection of his columns, also called *Technically Speaking*, through his publishing company, Word Spy Press.

Richard Stevenson

Contributing editor Stevenson has been stalking the field of gallium oxide power electronics for about five years, watching for signs that it was about to become big. Early this year, he saw those signs and pounced [p. 9]. "What grabbed my attention was making them on [inexpensive] sapphire substrates," he says. From his expertise in compound semiconductors, he knew that the devices were getting cheap enough to unseat rivals, potentially leading to smaller and more efficient power electronics.

Christopher Tozzi

How did Tozzi, a Howard University historian who specializes in revolutionary-era France, come to write "Linux Turns 25" [p. 46]? It turns out he has also worked as a Linux server administrator and has a book about open-source software in the works. He found writing the book comparatively easy: "As someone whose prior experience in history research often entailed working with archival notes scribbled by half-literate 18th-century soldiers, I thought it was great to be able to view everything on a computer screen!"

Lawrence Ulrich

Ulrich covers autos for *The New York Times* and others, and he provides the airlines plenty of business, too, with a reporting beat that keeps him circling the globe. In one four-week marathon early this year, he flew back and forth between his home in Brooklyn, N.Y., and Detroit, Los Angeles, Miami, and San Francisco, all while writing and revising our annual review of automotive tech. It's the sixth year that he's done it for *Spectrum* ["Top Ten Tech Cars," p. 34].

04.16

PHOTO-ILLUSTRATIONS BY **Gluekit**

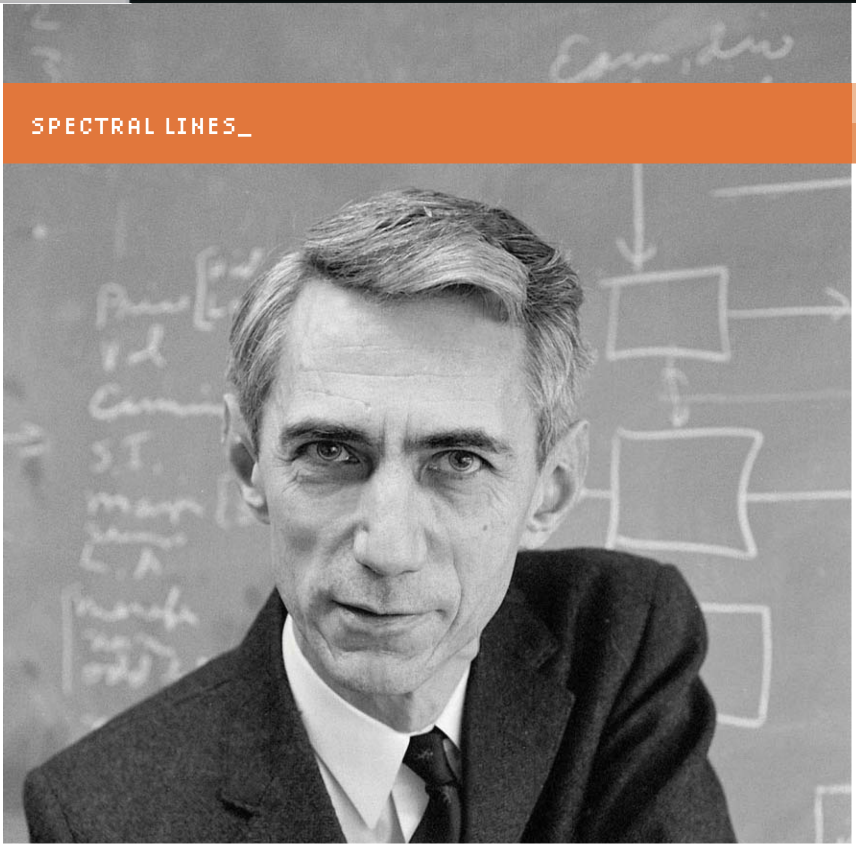

ALFRED EISENSTAEDT/THE LIFE PICTURE COLLECTION/GETTY IMAGES

Celebrating Claude Shannon

The engineer's life proves that practice can precede theory in the process of creating knowledge

While I was doing research for a biography of Claude Shannon's graduate advisor, Vannevar Bush, I learned that Bush's mentorship of Shannon in the fields of electrical engineering and applied math decisively influenced Shannon's career choices and thus the history of computing and digital media.

Early on, Bush recognized Shannon's singular talents. In a 1939 letter, Bush described his then 23-year-old student as a peculiar "genius" poised "to accomplish something striking and even useful."

Shannon did just that in the 1940s. While at Bell Labs he published a series of seminal articles in highly technical journals. The papers, which aimed to solve practical problems of telephony, launched a revolution. In "A Mathematical Theory of Communication," published in 1948, Shannon presented a unifying theory for the transmission of information that could be applied to telephones, radio, television, or any other system. By viewing all forms of information as reducible to binary digits, Shannon hypothesized that these *bits*—a term coined by a Bell Labs colleague—could be transmitted, error-free, even through a "noisy" communication channel.

Of all Vannevar Bush's legacies—he organized the Manhattan Project, after all, and built powerful analog computers in the 1930s—one of his greatest was certainly his support and encouragement of Shannon.

The centennial of Shannon's birthday will occur on 30 April, and amid the celebrations few will recall that his insights were also distorted and misunderstood. As the years passed and fulsome praise arrived, Shannon grew enigmatic and withdrawn. Prone to pranks and gambits, such as building unicycles and juggling obsessively, he matched the acclaim with modesty.

His theory of information, he wrote in a brief 1956 article titled "The Bandwagon," has "perhaps been ballooned to an importance beyond its actual accomplishments." He continued, "Authors should submit only their best efforts, and these only after careful criticism by themselves and their colleagues." He worried that his insights were being "oversold."

Historians Paul Nahin and William Aspray have convincingly shown that Shannon was correct to want to share credit. He built his theories on the shoulders of earlier giants, notably his personal hero, George Boole, the 19th-century logician. "It is because of Shannon that Boole is rightfully famous today, but it is because of Boole that Shannon first gained the attention of the scientific community," Nahin wrote in his joint biography, *The Logician and the Engineer.*

What then would I tell Shannon were I to meet him while wandering through those corridors at MIT? I would not dare risk his displeasure by joining the "bandwagon" and proclaiming him "the father of the information age." Rather, I would say that the signal lesson of his career stems from how he pursued knowledge about the world. Shannon showed that engineering creates new knowledge as readily as science does and that, moreover, engineering practice can *precede* scientific theory.

Shannon's life exposes the fiction that, in the creation of knowledge, theory must always precede practice and that engineers with "dirty hands" can never spawn the rich ideas that scientists later spend generations extending, revising, and deepening. While Shannon deserves acclaim for launching the digital revolution, his centennial is equally an occasion to celebrate all those persistent practitioners—those use-inspired problem solvers—who endure the condescension of gazers at stars and navels.

To those who insist that practice can never give birth to new pathways of inquiry and paradigms of knowledge, Shannon provides the definitive negative confirmation. And that is worth celebrating on Shannon's birthday, or any other day.
—G. PASCAL ZACHARY

G. Pascal Zachary is a professor at Arizona State University's Walter Cronkite School of Journalism and Mass Communication and is the author of *Endless Frontier: Vannevar Bush, Engineer of the American Century* (MIT Press).

NEWS

3,000 K

ORIGINAL TEMPERATURE OF THE COSMIC MICROWAVE BACKGROUND. IT'S NOW JUST 2.7 K

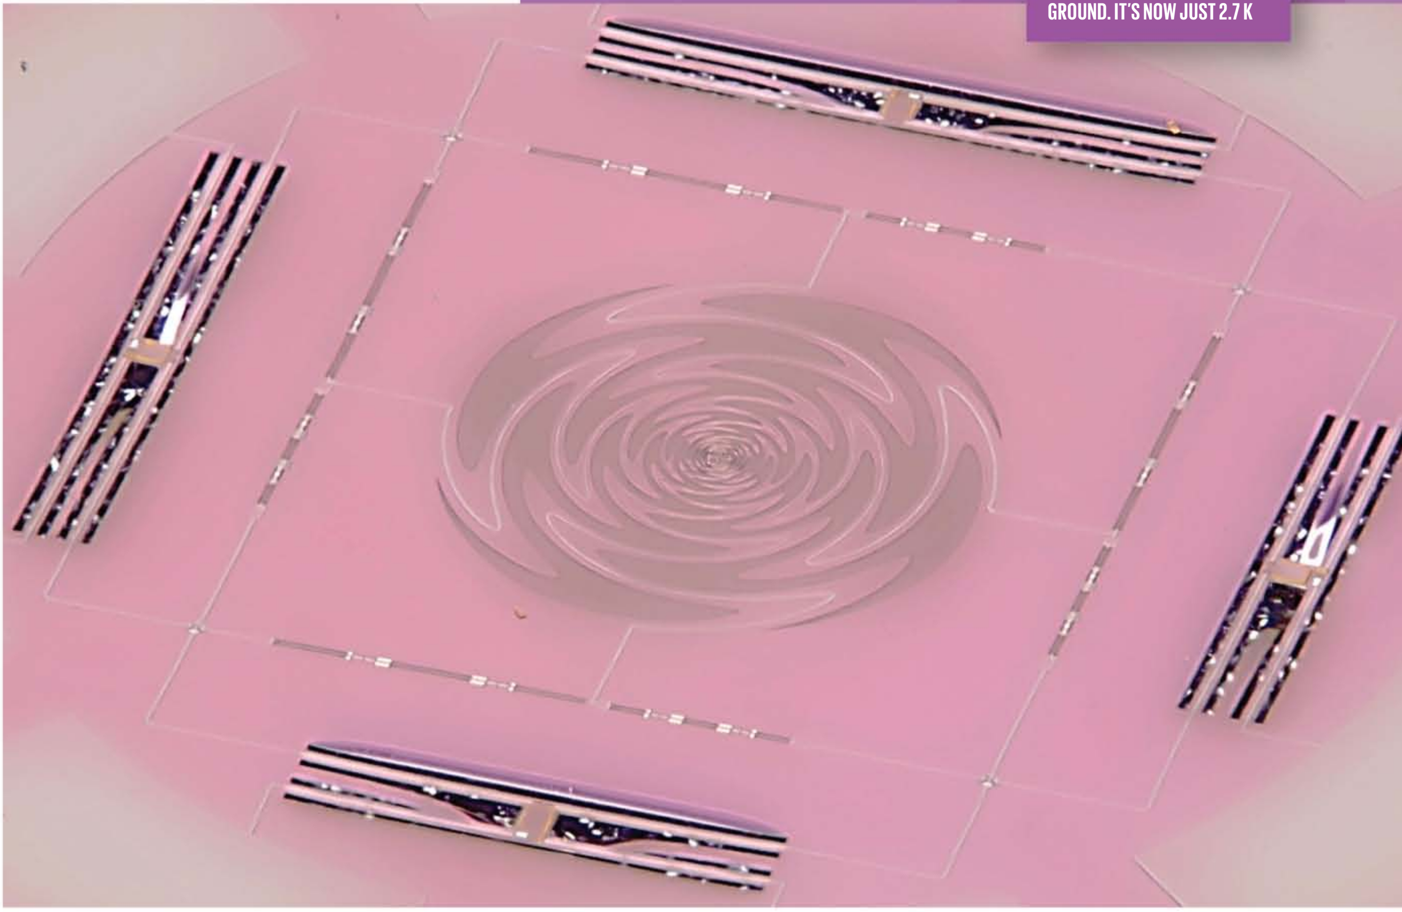

UNIVERSITY OF CALIFORNIA, BERKELEY

SWIRLY ANTENNAS SEE THE ANCIENT COSMOS

New detectors will hunt for gravitational waves from the beginning of time

SWIRLS FOR SWIRLS: This sinuous antenna picks up multiple microwave bands and sends them to surrounding detectors.

At long last, gravitational waves have been found. Now a different sort of hunt for those space-time ripples is picking up speed.

On a mountaintop in Chile and at the South Pole, a new generation of superconducting detectors is beginning the search for the imprint of gravitational waves on the cosmic microwave background (CMB), the universe's oldest observable light. While gravitational waves passing through Earth can tell us about relatively recent events involving black holes or neutron stars, the discovery of an ancient signature could provide a window into the universe a minuscule fraction of a second after the big bang.

Cosmologists have been hunting for years for evidence of such primordial gravitational waves, which should show up as a "swirl" in the CMB's polarization. And »

ON THE HUNT: Two new telescopes will join the Huan Tran in Chile [left]. All three will eventually boast sinuous antenna detector arrays [above].

for a short time it seemed the cosmic quarry had been bagged: In 2014, the South Pole-based BICEP2 experiment declared that it had found the characteristic swirls—a particular polarization pattern called the B-mode [see "The Chips That Saw the Big Bang's Fingerprints," *IEEE Spectrum*, May 2014]. But in the end, the source was identified as something much closer to home—dust from our own galaxy.

A new generation of superconducting receivers, packed with more detectors, could help physicists isolate such confounding signals. While BICEP2 was sensitive to a single frequency—150 gigahertz—these new experiments can pick up multiple frequencies simultaneously. Because the CMB's 13.8-billion-year-old light has a different spectrum than that of Milky Way dust emission, acquiring data at various frequencies can allow physicists to better identify and subtract the unwanted light.

"BICEP2 really illustrated [that] with one frequency you can only detect one type of signal—in other words, if you have one unknown and one measurement, you can only fit one thing to that one unknown," says Brian Keating of the University of California, San Diego, who worked on the experiment. "Now the name of the game is multifrequency coverage."

There are a number of ways to measure the polarization of the CMB at different frequencies. An experiment can swap multiple single-frequency receivers in and out or mount them on separate telescopes. Another option, employed by the Atacama Cosmology Telescope, in Chile, uses three-dimensional feedhorn antennas that pick up and funnel a wide band of microwaves to a multifrequency detector system.

The Polarbear team, of which Keating is a part, has settled on an approach that uses an array of silicon lenses, each over a single flat antenna. Developed by groups based at the University of California, Berkeley, and UC San Diego, the "sinuous antenna" contains four zigzagging niobium arms arranged more or less like a plus sign. Because of their fractal nature—having a structure that repeats at different scales—the antennas are capable of picking up a wide range of frequencies. The long parts of the arms pick out a particular polarization depending on their orientation. The rest of the system is similar to those used in other CMB receivers; the signals are filtered by frequency and sent to ultrasensitive superconducting devices called transition edge sensor bolometers, which register the microwave signals. The ability to pick up multiple frequencies simultaneously will not only help distinguish signals from unwanted sources of polarization swirls; it will also let physicists boost sensitivity without drastically increasing the size of the detector array.

High sensitivity per unit area is critical for CMB experiments, which demand temperatures less than a degree above absolute zero to operate effectively. That area "is more expensive than Manhattan real estate," Keating says. "What you want to do is to make each square centimeter extract as much information from the photon field as you can."

The basic antenna design was patented by engineer Raymond DuHamel in 1987. More recent work, by Gabriel Rebeiz of UC San Diego, and others, paired the

LEFT: ADRIAN LEE; RIGHT: POLARBEAR COLLABORATION structure with a silicon lens, with the aim of creating small antennas sensitive to terahertz and millimeter-wave light. The Berkeley team collaborated with Rebeiz on optimizing the design for CMB observations.

That work included shrinking the size of the antenna, says Berkeley's Aritoki Suzuki, reducing the smallest feature size to about a micrometer and making the whole antenna and its silicon lens only about 5 millimeters across.

One challenge in this miniaturization, Suzuki says, was devising a method to connect signal-carrying feed lines to the antenna in a way that would be compatible with traditional chip-printing technology. In the end, the team found they could connect the lines by snaking them inward along the arms of the antenna.

In early March, workers were getting ready to construct two new telescopes that will house some of these detector arrays, more than 5,000 meters above sea level on Chile's Cerro Toco. The first of the new receivers, dubbed Polarbear-2A, should arrive later this year. It will be sensitive to two frequencies and boast a detector count six times that of its single-frequency predecessor, which is mounted on the Polarbear team's nearby Huan Tran telescope. Before the end of 2017, the group aims to have multifrequency receivers mounted on both of the new telescopes as well as on Huan Tran to create what will be called the Simons Array.

Each approach to multifrequency detection comes with benefits and trade-offs, says Clarence Chang of the University of Chicago. Chang and his colleagues have opted to use the Berkeley team's design to upgrade the sensor array on the South Pole Telescope. That telescope's new receiver will be sensitive to three frequencies—two for CMB polarization and one for studying the evolution of galaxy clusters. The group plans to begin work installing the new array later this year, at the start of the austral summer. —RACHEL COURTLAND

NEWS

POWER ELECTRONICS' COOL NEW FLAVOR

With a bigger bandgap, gallium oxide could leapfrog other challengers to silicon

Ideally, the electronic components that route electricity through power supplies, inverters, and electric motors are cheap, efficient, and capable of handling high voltages. Judged in these terms, gallium oxide could be the best material yet, according to recent work by Flosfia, a startup in Kyoto.

That's because silicon—the incumbent material for making diodes and transistors for the power electronics market—is cheap but not very efficient. And although this weakness is addressed by devices made from silicon carbide and gallium nitride, both have had limited commercial success due to high prices. Flosfia's diodes are already performing more efficiently than those made from SiC and GaN.

The superiority of these gallium oxide devices stems from the material's approximately 5-electron-volt bandgap—way higher than that of gallium nitride (about 3.4 eV) or silicon carbide (about 3.3 eV). Bandgap is a measure of the energy required to kick an electron into a conducting state. A bigger bandgap enables a material to withstand a stronger electric field, making it possible to use a thinner device for a given voltage. That's a big deal because the thinner the device, the lower its resistance, and thus the more efficient it is.

Gallium oxide devices do not excel in all areas. Their Achilles' heel is poor thermal conductivity. "When you make a high-power device, you need to have a good thermal conductivity to extract the heat out of the device," explains Hong Lin, senior market and technology ana-

SEMICONDUCTORS IN COMPARISON: Flosfia's gallium oxide devices exploit the material's big bandgap and breakdown field while making up for its poor thermal conductivity. (Values are normalized to silicon's [blue].)

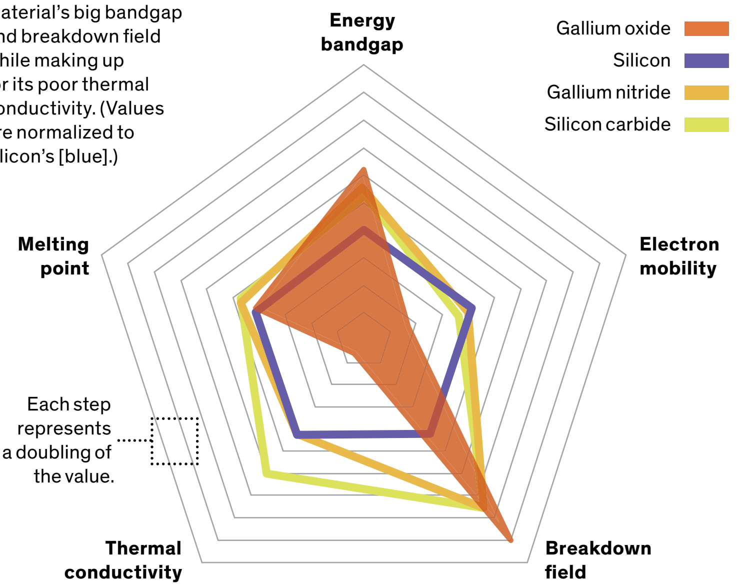

SOURCE: YOLE DÉVELOPPEMENT, LYON, FRANCE lyst at Yole Développement in Lyon, France.

Flosfia's engineers have improved the device architecture to address that very issue. In particular, they found a way to make the diode chip thinner, according to Naonori Kurokawa, a partner at University of Tokyo Edge Capital, a Flosfia investor. The key is to grow the gallium oxide crystal on a sapphire substrate.

Flosfia chief technology officer Masaya Oda found that "the gallium oxide epilayer can easily lift off the sapphire substrate," says Kurokawa. Separating the device from its foundation allows the chip to be bonded to a highly conductive, heat-sucking material, enabling operation at a lower temperature.

Using sapphire makes a lot of sense. "It's supercheap and already marketed, because of light-emitting-diode manufacture," says Kurokawa. The only downside is that the crystal quality is not as high as it would be using a gallium oxide substrate, which at the moment is quite expensive.

The process Flosfia uses for making gallium oxide devices was invented by company cofounder and Kyoto University professor Shozuo Fujita. In it, the sapphire substrate is heated and a fine mist of particles is swept into the chamber on a gust of nonreactive "carrier gas." The mist, which contains metal compounds, decomposes when it hits the hot substrate and forms a film of gallium oxide. The whole process can be cycled through rapidly because, unlike with other methods, the chamber never has to be completely evacuated. And that drives down costs.

Engineers from Flosfia detailed the results of diodes made with this growth process in the February 2016 edition of *Applied Physics Express*. One device combines a 531-volt breakdown voltage—the potential needed to reverse the flow of current—with an on-resistance of 0.1 milliohm per square centimeter, exceeding the limits of what is possible with silicon carbide.

The highest breakdown voltage Flosfia reported is 855 V. This is not that high for a wide-bandgap diode—devices made from SiC can handle 10 kilovolts or more.

Kurokawa explains that the diodes have modest breakdown voltages because they are bare chips. Introducing insulating layers into this very simple device should lead to significant improvement, he says.

Later this year Flosfia will start to provide samples of its diodes to potential customers. Further ahead, plans include a ramp-up of diode production in 2018 and the development and launch of accompanying transistors.

However, it is not yet clear whether this will ignite a gallium oxide power electronics industry. "Today we are comparing silicon with silicon carbide," says Lin, who expects gallium oxide devices to undergo a similar evaluation only once they are commercialized and considered as worthy contenders. That is still some way off.

—RICHARD STEVENSON

NEURAL NETWORKS ON THE GO

Engineers are trying to squeeze outsize AI into mobile systems

Artificial intelligence systems based on neural networks have had quite a string of recent successes: One beat human masters at the game of Go, another made up beer reviews, and another made psychedelic art. But taking these supremely complex and power-hungry systems out into the real world and installing them in portable devices is no easy feat. This February, however, at the IEEE International Solid-State Circuits Conference in San Francisco, teams from MIT, Nvidia, and the Korea Advanced Institute of Science and Technology (KAIST) brought that goal closer. They showed off prototypes of low-power chips that are designed to run artificial neural networks that could, among other things, give smartphones a bit of a clue about what they are seeing and allow self-driving cars to predict pedestrians' movements.

Until now, neural networks—learning systems that operate analogously to networks of connected brain cells—have been much too energy intensive to run on the mobile devices that would most benefit from artificial intelligence, like smartphones, small robots, and drones. The mobile AI chips could also improve the intelligence of self-driving cars without draining their batteries or compromising their fuel economy.

Smartphone processors are on the verge of running some powerful neural networks as software. Qualcomm is sending its next-generation Snapdragon smartphone processor to handset makers with a software-development kit to implement automatic image labeling using a neural network. This software-focused approach is a landmark, but it has its limitations. For one thing, the phone's application can't learn anything new by itself—it can only be trained by much more powerful computers. And neural networks experts think that more sophisticated functions will be possible if they can bake neural-net-friendly features into the circuits themselves.

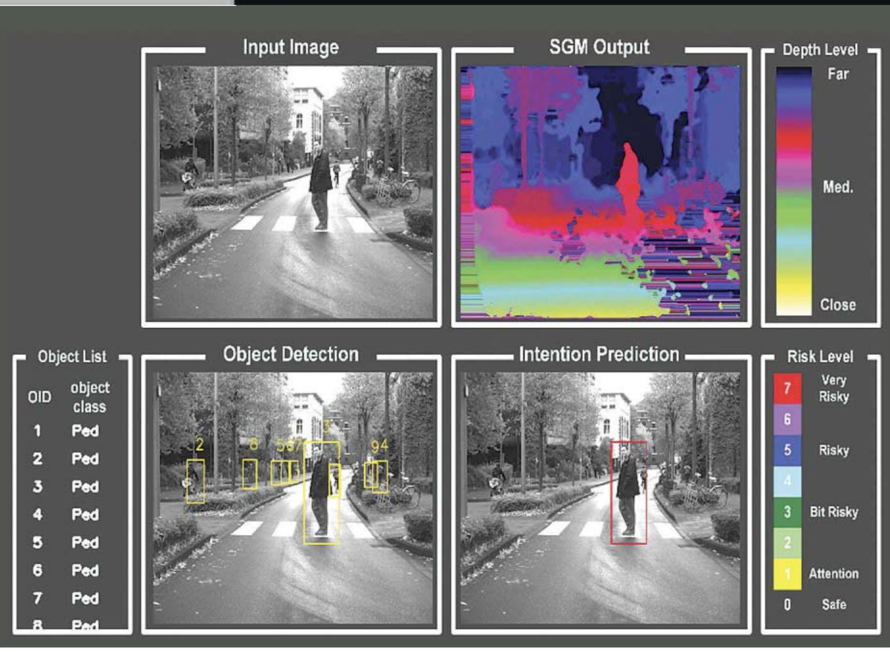

WHAT'S YOUR INTENTION? Engineers demonstrated a neural network chip for cars that predicts the risk of hitting a pedestrian.

KAIST

The larger a neural network is, the more computational layers it has, and the more energy it takes to run, says Vivienne Sze, an electrical engineering professor at MIT. No matter the application, the main drain on power is the transfer of data between processor and memory. This is a particular problem for convolutional neural networks, which are used for image analysis. (The "convolutional" in the name hints at the many steps involved.)

For the human brain, drawing on memories to make associations comes naturally. A 3-year-old child can easily tell you that a photo shows a cat lying on a bed. Convolutional neural networks can also label all the objects in an image. First, a system like the image-recognition champ AlexNet might find the edges of objects in the photo, then begin to recognize those objects one by one—cat, bed, blanket—and finally deduce that the scene is taking place indoors. Yet even doing this kind of simple labeling is very energy intensive.

Neural networks, particularly those used for image analysis, are typically run on graphics processing units (GPUs), and that's what Snapdragon will use for its Scene Detect feature. GPUs are already specialized for image processing, but much more can be done to make circuits that run neural networks efficiently, says Sze.

Sze, working with Joel Emer, also an MIT computer science professor and senior distinguished research scientist at Nvidia, developed Eyeriss, the first custom chip designed to run a state-of-the-art convolutional neural network. They showed they could run AlexNet, a particularly demanding algorithm, using less than one-tenth the energy of a typical mobile GPU: Instead of consuming 5 to 10 watts, Eyeriss used 0.3 W.

Sze and Emer's chip saves energy by placing a dedicated memory bank near each of its 168 processing engines. The chip fetches data from a larger primary memory bank as seldom as possible. Eyeriss also compresses the data it sends and uses statistical tricks to skip certain steps that a GPU would normally do.

Lee-Sup Kim, a professor at KAIST and head of its Multimedia VLSI Laboratory, says these circuits built for neural-network-driven image analysis will also be useful in airport face-recognition systems and robot navigation. At the conference, Kim's lab demonstrated a chip designed to be a general visual processor for the Internet of Things. Like Eyeriss, the KAIST design minimizes data movement by bringing memory and processing closer together. It consumes just 45 milliwatts, though to be fair it runs a less-complex network than Eyeriss does. It saves energy by both limiting data movement and by lightening the computational load. Kim's group observed that 99 percent of the numbers used in a key calculation require only 8 bits, so they were able to limit the resources devoted to it.

"It's a challenge deciding between generality versus efficiency," says Nvidia's Emer. Sze, Emer, and Kim are trying to make general-purpose neural-network chips for image analysis, a sort of NNPU. Another KAIST professor, Hoi-Jun Yoo, head of the System Design Innovation and Application Research Center, favors a more specialized, application-driven approach to making neural-network hardware.

One system Yoo described was meant for self-driving cars. It's designed to run convolutional networks that identify objects in the visual field and also to use a different type of algorithm, called a recurrent neural network. The "recurrent" refers to the system's temporal skills—they excel at analyzing video, speech, and other information that changes over time. In particular, Yoo's group wanted to make a chip that would run a recurrent neural network that tracks a moving object and predicts its intention: Is a pedestrian on the sidewalk going to enter the roadway? The system, which consumes 330 mW, can predict the intention of 20 objects at once, almost in real time—the lag is just 1.24 milliseconds.

Another difference between Yoo's system and the MIT chip is that his hardware, which he calls an Intention Prediction Processor, continues its training while on the road. Yoo's designs integrate what he calls a deep-learning core, a circuit that's designed to add to the neural nets' training. For convolutional neural networks, this deep-learning training is typically done on powerful computers. But Yoo says that our devices should adapt to us and learn on the go. "It's impossible to preprogram all events. The real world is diverse and almost impossible to predict," says Yoo. —KATHERINE BOURZAC

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

U.S. FISSION FIZZLES

With carbon markets and power prices in doubt, nuclear power is no longer affordable

U.S. nuclear power plant operators are fighting a war on two fronts: Crashing prices for natural gas and accelerating market penetration of renewable energy have both contributed to dramatic drops in wholesale power price levels–in some states, they've fallen by more than two-thirds over the past decade. This has left nuclear power, whose operating costs are pretty much fixed, with few options other than surrender.

That marks quite a reversal, says Gregory Jaczko, former chairman of the U.S. Nuclear Regulatory Commission. "It's been a widely held belief that nuclear is incredibly cheap to operate. That was the case 10 years ago, when nuclear plants were cash cows. That's not the case today, especially as the plants age," he says.

Fission is already giving ground. Two plants, in Wisconsin and Vermont, shut down in 2013 and 2014, respectively. More shutdowns are anticipated in Massachusetts, New Jersey, and New York, and at least half a dozen more plants are teetering on the brink of insolvency.

Nuclear operators had been expecting President Obama's Clean Power Plan, which would have established national carbon regulations and increased the cost of fossil-fuel-generated electricity, to offer them a reprieve. But their hope was short-lived: The U.S. Supreme Court decided in February to stay implementation of the plan.

Operators' next best hope is that state governments will be motivated to step in and save them. State governments worry that substituting gas-fired generation for nuclear would, in the words of New York governor Andrew Cuomo, "eviscerate the emission reductions" achieved through state renewable energy programs. New York's Public Service Commission projects that annual carbon emissions could rise by more than 12 million metric tons, working against a pledge to cut emissions by 40 percent of 1990 levels by 2030.

A LUCKY BREAK: The Quad Cities nuclear generating station near Cordova, Ill., is getting some needed help from a regional grid operator.

States are also voicing concerns over grid reliability and the potential loss of local jobs. The New York Independent System Operator recently projected that losing two threatened upstate reactors, operated by New Orleans-based Entergy Corp. and Chicago-based Exelon Corp., would help create a power supply deficit by 2019.

By June, state regulators plan to detail special payments for New York's nuclear plants. For companies that open their books and show that their plants are losing money, New York is vowing to make them whole.

"It's probably the stickiest energy policy question there is in the electric sector," says Julien Dumoulin-Smith, a senior power markets analyst with New York City-based UBS Investment Research.

New York is not alone. Several states, including Illinois and Ohio, are seeking to give nuclear power plants an extra boost. Exelon lists the dual-reactor Quad Cities site in Illinois among its uncompetitive plants. Dumoulin-Smith says local power prices are being "pummeled" by wind power, noting that Quad Cities sits just across the border from Iowa, which leads the United States in wind power penetration.

Quad Cities, along with Exelon's nine other reactors in Illinois, caught a break from regional grid operator PJM, which last year began offering more supplements to power plants that commit to staying in its energy market for several years, thus backstopping its power supply. Those Illinois plants raked in more than US $1.5 billion from PJM's capacity market last September, ensuring that Quad Cities will remain in operation through mid-2018.

But experts say that tougher times could lie ahead, as the extension of tax breaks for wind and solar power voted through the U.S. Congress in December fuels further growth in renewable generation. New York's program appears to anticipate this, positing that assistance to nuclear reactors is more of a temporary lifeline than a long-term guarantee. State aid is, according to a white paper from the N.Y. Public Service Commission, a means of supporting "a smooth emission-free transition from nuclear to nonnuclear resources" in the event that energy prices "are not able to support the continued financial viability of the [fission] plants." –PETER FAIRLEY

KEVIN SCHMIDT/ALAMY

If your 5G idea works here...

it will work here.

50,000 soccer fans streaming live-action mobile video. All at once.

Soon it will be reality. A world connected as never before. Always on. Amazingly fast. Massive traffic. Instant information. That's the promise of 5G. New spectrum. New waveforms. Millimeter-waves. Massive MIMO. More. Keysight offers the world's first 5G simulation, design and test environment able to emulate your real-world 5G wireless thesis. With deep expertise to help you tackle 5G risk and complexity. So you can go from 5G ideas to 5G reality faster.

HARDWARE + SOFTWARE + PEOPLE = 5G INSIGHTS

5G Get the latest app notes, white papers and tutorials
www.keysight.com/find/5G-Insight

USA: 800 829 4444 CAN: 877 894 4414

© Keysight Technologies, Inc. 2016

KEYSIGHT TECHNOLOGIES
Unlocking Measurement Insights

Agilent's Electronic Measurement Group is now Keysight Technologies.

PHOTOGRAPH BY NASA/Alamy

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

SPACE RESCUE

REMOTE CONTROL might work for aircraft. But for agriculture? Not so much. That's what NASA learned recently when an experiment in autonomous gardening aboard the International Space Station began to wilt. In mid-November, an astronaut activated a system designed to support the growth of this zinnia plant, with strict instructions on how to care for it thereafter. But by late December, the space team noticed that the plant wasn't thriving. So they took matters into their own hands, using their judgment to decide when to water it. By 16 January, when this photo was taken, the zinnia had rebounded and bloomed.

THE BIG PICTURE NEWS

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

Explore Every Detail

There are many factors in a successful design: meeting challenging technical specifications, behaving as intended in a realistic environment, and arriving within the budget and deadline. These are some of the challenges that CST STUDIO SUITE® can be used to overcome.

The 2016 version of the CST® industry-leading electromagnetic simulation software lets you delve deeper into your designs than ever before. Applications ranging from nanoscale graphene devices through to antenna installed performance on airliners can be simulated quickly and accurately thanks to the array of new modeling and simulation features.

Look beneath the surface and discover the potential of your ideas with CST STUDIO SUITE 2016.

CST

CST – COMPUTER SIMULATION TECHNOLOGY | www.cst.com | info@cst.com

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

RESOURCES

99.9%

DECLINE IN THE POPULATION OF THE WHITE-RUMPED VULTURE IN THE LAST 30 YEARS

DAVID ALAYETO/GETTY IMAGES

VULTURE VOYEUR
A SENSOR-PACKED EGG MONITORS NESTS FROM THE INSIDE

Vultures are nature's garbage collectors, helping the environment by consuming dead animal carcasses. In this way, they are essential in stopping the spread of diseases such as rabies (which vultures are immune to). However, recent years have seen the number of vultures decline, particularly in South Asia, where some species are close to becoming extinct due to the toxic effects of a drug used to treat cattle. To protect vulture populations, one of the things we need to know more about is their breeding behavior. ● A year ago we at Microduino were approached by the International Centre for Birds of Prey (ICBP) as part of a project to achieve this end. The ICBP needed to create an electronic egg to monitor vulture nests. The conservationists there wanted the egg to include a host of sensors that could measure both its internal temperature and the temperature gradient across its surface, as well as barometric pressure, humidity, carbon dioxide levels, light intensity, and the egg's rotation and movement. The data would then be transmitted to a relay node and uploaded to the cloud.

EXTINCTION EVENT: Vital to ecosystems and once the most common birds of prey in the world, vultures are becoming increasingly rare.

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

Initially, the ICBP had considered building a system around standard Arduino microcontrollers. The agency soon realized, however, that mother vultures would be quite displeased to see an Arduino jutting out from their eggs. The size of the board as well as the sensory and communication requirements made the standard Arduino unsuitable. So the ICBP turned to our system. Microduino makes a family of Arduino-compatible microcontrollers and modules that can be stacked on top of each other and that are just 25.4 millimeters wide by 27.9 mm long.

We decided to separate the system into three parts: a data-collection terminal that would gather and store data and perform some initial processing (the electronic egg itself), a data-relay terminal that would receive and retransmit processed data wirelessly, and a data repository in the cloud that researchers could access.

Our initial plan was to stick all the necessary modules and sensors into a 3-D-printed egg and pat ourselves on the back. However, the ICBP presented us with another logistical challenge: To avoid disturbing nesting vultures, the egg had to be capable of operating independently for 70 days!

Consequently, we reapportioned the division of labor among the three elements of our system: Data processing and storage responsibilities were shifted to the data-relay terminal to reduce the egg's power consumption.

We then set about building the egg. We built an enclosure from laser-cut wood that would fit within an artificial eggshell. We placed inside the enclosure a Microduino core, a Bluetooth Low Energy (BLE) module, and a multisensor 10DOF module, which incorporates a three-axis gyroscope, a three-axis accelerometer, a magnetic field strength sensor, and a barometer. In addition, there are fourteen DS18B20 temperature sensors that cover the entire inner-shell surface and one SHT21 humidity sensor. These all use I2C connections to communicate with the core. An 1,800-milliampere-hour lithium-ion battery provides power. We then placed the enclosure in the egg,

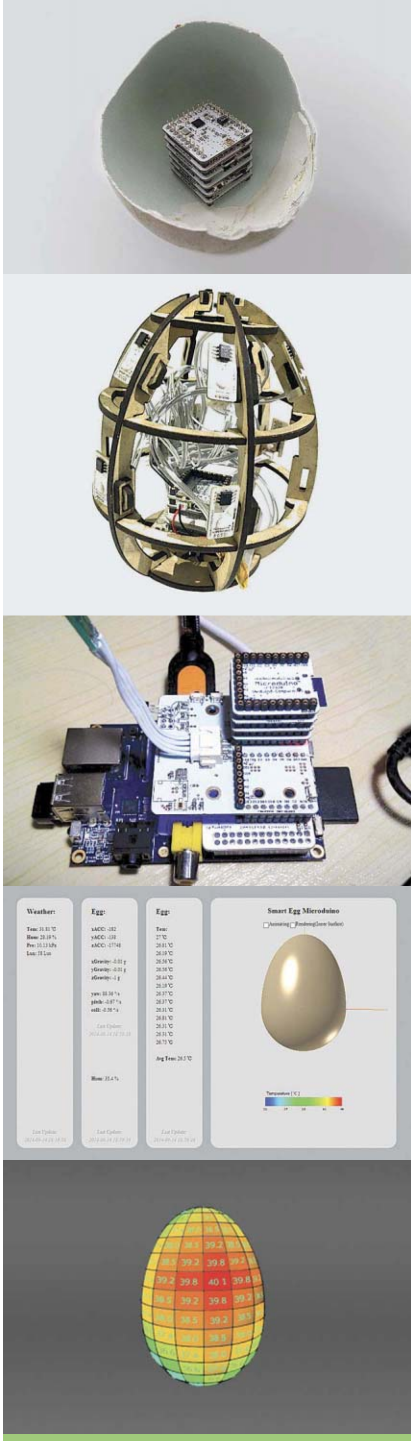

A real vulture eggshell [top] with a Microduino stack illustrates the volume we had to work with. The stack and sensors are held in place by a wooden enclosure [second from top] and data is transmitted to a relay station [third from top] outside the nest. A cloud server tracks each egg's data [bottom two images].

MICRODUINO (5)

which was fabricated with a selective laser sintering machine using PA2200 nylon. This material has qualities similar to those of vulture eggshells.

Moving on, we built the data-relay terminal by combining a Wi-Fi-enabled Raspberry Pi and a stack consisting of a Microduino Core+ (which uses a more powerful processor than the one we placed inside the egg), a Bluetooth module, a real-time clock (RTC) module, and a weather station module.

The Pi talks to the stack via a custom board that creates a serial interface between the Core+ and the Pi's general-purpose input/output connector.

The terminal is placed a short distance away from the egg and serves multiple purposes. First, it receives all the data collected from the egg wirelessly via Bluetooth. Second, it monitors the conditions outside the egg with its own light, temperature, humidity, and barometric sensors. Third, the terminal saves all the data from both the egg and its own sensors and stores it in the Pi. When connected to the Internet, the data is uploaded to a cloud server.

After building multiple eggs and data-relay terminals in this manner, we created our own cloud server to monitor the project data in real time. Each egg has a unique ID. Using the information uploaded by the data-relay terminals, we are able to construct a 3-D model of each egg's surface-temperature gradient as well as display other relevant data in real time.

We are now finalizing the project, and the ICBP will deploy the eggs for field-testing this month. A successful field test would mean we can use this method not only to benefit vulture conservation but also for a variety of other environmental efforts.

Our opportunity to work with the International Centre for Birds of Prey was amazing and eye-opening. Sometimes as we marvel at the cool new technologies that are coming out every day, we begin to lose focus on all the good we can do with our knowledge. We at Microduino would like to *egg* you on to make a positive difference in the world with all the tools at your disposal!

—BIN FENG, BRUCE LIU & KEJIA PAN

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

RESOURCES_PROFILE

REINVENTING RETAIL
THE LATEST SILICON VALLEY ADVENTURE IS IN BRICK-AND-MORTAR

RETAIL RADICAL: Vibhu Norby in his Silicon Valley store, where customers can buy bleeding-edge products and companies can get instant feedback.

Vibhu Norby isn't like most consumers; he doesn't wait until a tech product shows up at Best Buy or Costco. He's a classic early adopter who does most of his shopping on Kickstarter or Indiegogo, happy to preorder and wait months for delivery. Only after delivery does he sometimes realize that the product shipped doesn't quite fulfill the promise of the product pitched. He didn't think much about how other people bought gadgets, or brick-and-mortar retail in general.

But last February, Norby left his job at the smart thermostat company Nest (now part of Google) on a mission to fix retail. And in December he opened his first store, B8ta, in downtown Palo Alto, Calif.

At Nest, Norby discovered how tricky it is to sell hardware to the average consumer, who, unlike him, wants to see it, touch it, and ideally, try it at a store before spending a couple of hundred dollars.

"Nest," Norby says, "did a good job of getting the product onto retail shelves, but inventory wasn't being kept up, store associates weren't being trained, and boxes were ending up scattered all over the floor. It was frustrating."

Norby talked to his peers at other young consumer-electronics hardware companies, and found out that they were even more frustrated—they couldn't even get their products into stores. "One guy recently told me Apple told him it'd be six months before they could schedule a meeting; Fry's told him it would take them a year to decide," he says.

In an industry where you can build a prototype in hours and get a new product on a manufacturing line in China in months, nothing should take a year, Norby thought. "Retail," he says, "needs to be fixed. It's in a lot of trouble, and if it's not fixed, it will take hardware manufacturers down with it."

Here's how B8ta works. Companies pay a monthly subscription fee to put their product in a store. They can choose to have B8ta keep inventory on hand to sell on-site, or have store employees direct potential customers to company websites for direct sales or preorders. Every product in the store sits next to a B8ta-designed touch-screen terminal that displays product information or anything else the manufacturer wants to put there; the manufacturer can make changes in this information remotely at any time, tweaking marketing materials and pricing.

B8ta uses security cameras in the store to track user engagement—how long a visitor looks at a product and its accompanying promotional material—and reports that to the manufacturer. B8ta staff members also gather information when people purchase products, in particular about what exactly they expect to use a gadget for. "Use cases are huge for anybody making a product," Norby says.

Norby says his staff members make a real effort to familiarize themselves with the products—ideally, by getting trained directly by the product designers. B8ta's staffers are paid more than is typical for retail (Norby won't say how much), largely because they are the smartest folks he can hire. On the day I visited the store, the greeter was a person with a chemical engineering degree from Brown University.

"As a software engineer," Norby says, "I think of retail as the ultimate distributed system. Each store is a node in a network. If you operate stores like software, sell them like software as a service, and measure the success like software, not only can you open retail to small manufacturers, you can change how retail is done as a business."

According to Norby, it shouldn't matter if consumers go into a store just to browse, and then ultimately make their purchases later online. "We changed the measure of success to engagement with a product, not with sales," he says. "In the world of the Internet, why does it matter where someone swipes a credit card?" **—TEKLA S. PERRY**

Location: Palo Alto, Calif. **Founded:** May 2015 **Employees:** 12 full-time, 5 part-time **Funding:** Undisclosed

TEKLA S. PERRY

WHAT WOULD MARVIN MINSKY READ?

KEY WORKS FROM THE AI TITAN'S FAVORITE AUTHORS

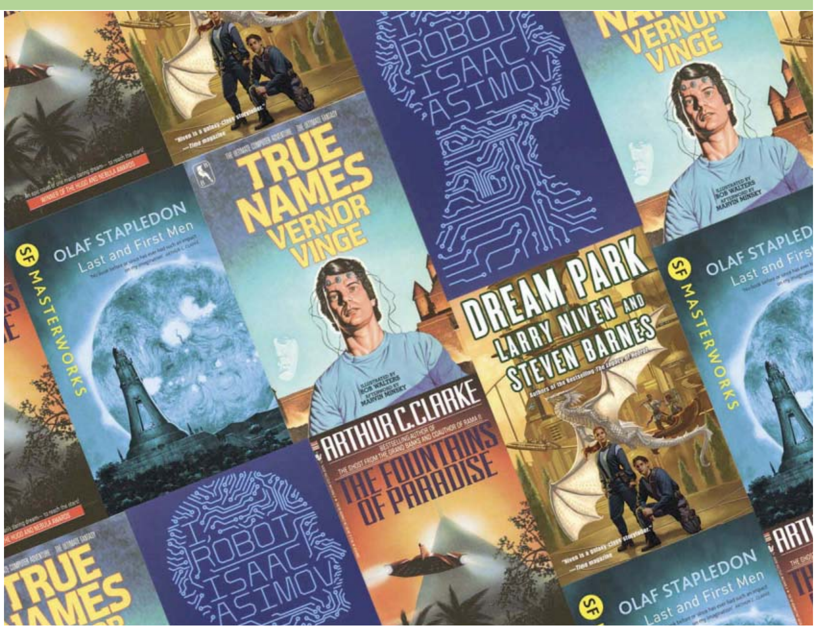

CLOCKWISE FROM TOP RIGHT: BANTAM SPECTRA; TOR BOOKS (2); BANTAM SPECTRA; GOLLANCZ

Marvin Minsky was one of the founders of artificial intelligence, and with John McCarthy, he created MIT's Computer Science and Artificial Intelligence Laboratory. Minsky, who recently passed away at the age of 88, was also a lifelong devotee of science fiction, to the almost total exclusion of all other forms of fiction. (He also penned a science fiction novel of his own, *The Turing Option*, written with Harry Harrison.)

Minsky often discussed his favorite authors but less often mentioned specific works. So, to commemorate this great visionary, the editors at *IEEE Spectrum* put together this selection of books by some of his favorites, picking titles that touch on the themes of computation and consciousness that most seem to resonate with Minsky's work—or simply share his penchant for audacious ideas.

Last and First Men by Olaf Stapledon. The book's narrative spans from the 1930s (when it was released) to about 2 billion years into the future. The book chronicles the rise and fall of *Homo sapiens* and 17 successive species that descend from it, focusing largely on their mental development and philosophies. The early chapters' reliance on ethnic stereotypes may sound a sour note for modern readers, but the book is notable for its ambition, the introduction of concepts such as genetic engineering, and its influence on later authors.

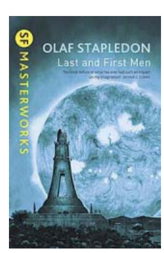

I, Robot by Isaac Asimov. This collection of short stories originally published from 1940 to 1950 lies at the core of Asimov's entire science fiction oeuvre, and explores the unintended consequences that can arise from following what appear to be the straightforward rules of behavior stated in Asimov's famous "Three Laws of Robotics." The story "Runaround" had a huge impact on the young Minsky's thinking, and he later became close friends with Asimov.

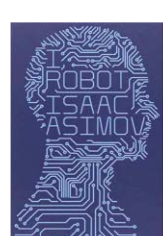

The Fountains of Paradise by Arthur C. Clarke. Published in 1979, this book brought the concept of the space elevator to the masses, featuring the construction of a superstrong cable stretching from a small island on the equator to high Earth orbit. Minsky was so taken with the idea that he worked with other researchers to develop real-world designs for the elevator and to bring a related concept called the "space fountain" to reality.

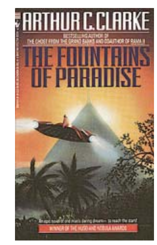

Dream Park by Larry Niven and Steven Barnes. It might be natural here to point to Niven's *Ringworld* (1970) as an example of the big-idea thinking Minsky loved, but we plumped instead for 1981's *Dream Park* as a nod to the kind of playful mind that delighted in thinking up the ultimate useless machine. The novel is a murder mystery set in the middle of a live-action role-playing game, which takes place in a computer-controlled arena so vast it makes Star Trek's holodecks look like very small potatoes indeed.

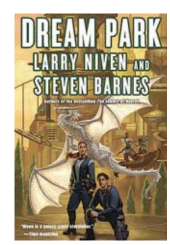

True Names by Vernor Vinge. William Gibson's 1984 *Neuromancer* became the breakout hit of early cyberpunk science fiction, but Vinge's 1981 novella is revered by aficionados. It introduces the virtual reality metaphor for cyberspace that would become de rigueur in later cyberpunk novels, and also sets the stage for Vinge's later writing on the singularity. Minsky later contributed a lengthy afterward to the novella, linking it to his own conceptions of and research on the workings of the mind. **—STEPHEN CASS**

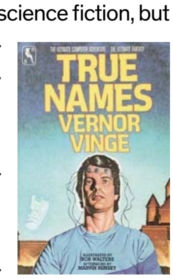

IEEE SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

INDIA'S STARTUP BOOM

AMIT AGRAWAL SAYS FOREIGN INVESTORS NEED TO AVOID THE FAMILIAR

MOHD ZAKIR/HINDUSTAN TIMES/GETTY IMAGES

ENCOURAGING ENTREPRENEURS: India's prime minister, Narendra Modi, is backing a plan to improve the business climate for startup companies.

What's the current state of the startup scene in India? "Everything is sparkling," says Amit Agrawal, director of the investment advisory firm Access Capital Advisors, in the New Delhi area. Agrawal matches startups with private investors, and he says the entrepreneurs he works with are full of optimism: "They see the sun, and nothing [else]."

The numbers back up their confidence. In 2015, investors poured US $9 billion into Indian startups, according to Your Story, a news and research website that covers India's entrepreneurial sector. And now the government wants to add its support: In January, Indian prime minister Narendra Modi announced a "Startup India: Action Plan" with much fanfare. The provisions, many timed to take effect on 1 April, include tax exemptions and regulatory easements for young companies. Modi also promised new investment funds, incubators, and research parks to help startups thrive.

The buzz is attracting foreign investors, Agrawal says, but these foreigners often get

into trouble by picking startups that are familiar. "They try to replicate a company that works in the U.S. culture or Singapore culture," he says. "But what works there won't necessarily work here." To pick winners, Agrawal says, you have to understand the quirks of Indian culture.

For example, he points to car-sharing startups that have succeeded in the United States and Europe, such as Zipcar and Car2go, which let urban drivers grab cars off the street whenever they want. But even though India's city dwellers have the same needs as their Western counterparts, startups trying to copy the Zipcar business model are struggling. In India, hiring a car with a driver is often cheaper than renting a "self-drive" car, Agrawal notes. Indian roadways are famously chaotic, and car rental companies feel more assured that their vehicles will return in one piece if they're driven by professionals: "There's a cultural logic to it, not a business logic," he says.

In addition to studying the Indian way of life, foreign investors must also recognize the challenges posed by the country's uneven development. Outside the cities, basic amenities like a steady supply of electricity and paved roads aren't guaranteed, so the logistics of manufacturing and delivering products can become an ordeal. That's why so many Indian startups are sticking to the "click model"—offering purely online services—as opposed to the "brick model," Agrawal says. As examples of such successful startups, he highlights companies that provide online booking for rail travel, as well as those offering mobile banking services.

On a final, cautionary note, Agrawal notes that Indian startups often have a difficult time if they're trying to compete with prices found in traditional local markets. A number of startups with grocery delivery apps are vying for position, he says, and currently offer steep discounts to lure customers. "But as soon as they stop providing discounts, I am not buying anything there, I promise you," he says. There are plenty of small neighborhood shops that happily deliver cheap goods to nearby doorsteps. Similarly, startups that offer on-demand cleaning services have trouble trumping local options. "If I just peek out of my balcony, I can shout, 'My maid is not in today. Can you come up and clean the utensils?' " says Agrawal, explaining that some neighborhood woman will always answer such a call.

With so many startups jumping in, there will inevitably be lots of flops. Agrawal estimates that in 2015 more than 600 startups backed by enthusiastic angel investors failed to achieve the next step—attracting venture capital firms that believe the company has a successful business model. But there's no stigma attached to failure in this bright startup scene, he says. Entrepreneurs simply turn to their next sparkling idea.

—ELIZA STRICKLAND

MIT PROFESSIONAL EDUCATION
Short Programs

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

COME TO MIT FOR A WEEK

Short Programs offer more than 50 intensive programs addressing new and evolving technologies, challenges and breakthrough solutions. Courses combine learning from MIT's groundbreaking research with real-world application perspectives from industry, government and academia. Earn CEUs and a Certificate of Completion in a variety of topic areas:

Topics include:
- Biotechnology / Pharmaceutical
- Computer Science
- Crisis Management
- Data Modeling and Analysis
- Design, Analysis, and Manufacturing
- Energy / Transportation
- Imaging
- Innovation
- Leadership / Communication
- Radar
- Real Estate
- Sustainability
- Systems Engineering

USE CODE IEEE10 AND SAVE 10%
when you register and pay fees by April 15.

NEW Professional Certificate Program in Innovation and Technology
Focused on building competitive advantage through innovation and technology, this credential signifies your depth of knowledge and skill in this high-value area. You can earn your Professional Certificate by completing four on-campus courses.

To learn more, visit shortprograms.mit.edu/IEEE

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

MANUFACTURING POWERS

MANUFACTURING HAS BECOME BOTH BIGGER AND SMALLER. During the past 10 years the worldwide value of manufactured products has grown, in inflation-adjusted terms, by more than 60 percent, surpassing US $12 trillion in 2015. ● Meanwhile, the *relative* importance of manufacturing is dropping fast, retracing the earlier retreat of agriculture (now just 4 percent of the world's economic product). Based on the United Nations' uniform national statistics, the manufacturing sector's contribution to global economic product declined from 25 percent in 1970 to about 15 percent by 2015. ● The decline has registered in the stock market, which values many service companies above the largest manufacturing firms. At the end of 2015, Facebook, that purveyor of updated selfies, had a market capitalization of nearly $300 billion, about 50 percent more than Toyota, the world's premier maker of passenger cars. And SAP, Europe's largest software provider, was worth about 75 percent more than Airbus, Europe's largest maker of jetliners. ● And yet manufacturing is still important for the health of a country's economy, because no other sector can generate nearly as many well-paying jobs. Take Facebook, which at the end of last year had 12,691 employees, versus the 344,109 that Toyota had at the end of its fiscal year, in March 2015. Making things still matters. ● The top four economies remain the top four manufacturing powers, accounting for about 55 percent of the world's manufacturing output in 2015. China is at the top of the list, followed by the United States (whose gross national product is still nominally No. 1), Japan, and Germany. But these countries differ markedly in the relative importance of manufacturing to their economies. The sector contributed about 28 percent of China's GDP in 2014, second only to South Korea, with 30 percent. In the same year, manufacturing's share came to about 23 percent in Germany, 19 percent in Japan, and only 12 percent in the United States total.

If you rank countries by per capita manufacturing value, then Germany, with more than $9,500 in 2014, came out on top among the big four, followed by Japan with nearly $7,000, the United States with about $6,500, and China with only $2,100. But the global leader here was Switzerland, with more than $15,000 per capita. Think about not only the sales of high-value products, such as the pharmaceuticals of Novartis and Roche and the watches of the Swatch Group (which includes Longines, Omega, Tissot, and other famous brands) but also machinery (ABB) and agrochemicals (Syngenta).

The share of manufactured goods in a country's total merchandise trade indicates the sector's importance in earning foreign exchange. Here, once again, China was at the top, with 94 percent of its 2014 exports attributable to manufactured items, tied with Bangladesh and Cambodia. Manufacturing accounted for 93 percent of exports in Israel, 90 percent in Japan and Switzerland, 86 percent in South Korea, 83 percent in Germany, and 62 percent in the United States.

The net balance of international trade in manufactured items is also revealing because it indicates two things: the extent to which a nation can satisfy its need for products and the demand for its products abroad. Switzerland again came out on top, with a 2014 per capita surplus of about $5,200, followed by Germany with $5,100 and South Korea with nearly $4,400. The United States, however, had a manufacturing trade *deficit* of about $1,800 per capita.

The United States enjoyed generations of manufacturing trade surpluses until 1982; China had chronic deficits until 1989. What are the chances of the United States redressing its massive manufactured trade imbalance with China, or of India (with only a tiny surplus in 2014) replicating China's manufacturing success? ■

ILLUSTRATION BY Chad Hagen

CURATING THE DIGITAL AGE

Out-of-copyright materials in NYPL digital collections are now available as high-resolution downloads. No permission required, no hoops to jump through: just go forth and reuse! *—New York Public Library*

AS I MENTIONED IN MY FEBRUARY 2013 COLUMN, "Balancing Act," the belief that our life offline is separate from our life online has been denounced as digital dualism. But there's less of a debate when it comes to differentiating between analog objects and digital data. Yes, the print and electronic copies of the same book contain the same words, but it's obvious to most people (and, increasingly, to researchers) that the two reading experiences are quite different. • We need to understand such differences because the world is going to see a lot more digital data in the near future. This includes **born-digital** data, which is originally created in an electronic format, as well as **born-analog** data, which starts life as a physical object and then is **reborn digital**. A great example of this digitization came earlier this year when the New York Public Library announced that it was making more than 180,000 digitized items available to anyone with an Internet connection, no questions asked. • That librarians would turn themselves into **digital curators** is no surprise, since as **analog curators** for the past few centuries they have been constantly bumping into the physical constraints of storage space and material decay. One approach is to get rid of stuff, and librarians and archivists employ a pleasing variety of terms related to the removal of unwanted or duplicate material from their collections: *Weeding* and *culling* generally refer to the removal of individual items, while *purging*, *screening*, and *stripping* are most often used for the removal of multiple related items. But the main problem with physical materials is that they possess what archivists call, poetically, *inherent vice*: the tendency for something to deteriorate over time because of some fault in the material itself (for example, the presence of lignin in cheap paper, which causes the paper to yellow) or the way the material reacts with its surroundings (for instance, the fact that bugs eat some books because they're attracted to the mold that grows in damp paper).

The digitization of analog materials can solve these problems, and engineers are constantly trying to find faster ways to turn atoms into bits. For now, though, we mostly have to rely on the skills of **scanops** (scanner operators) to generate those bits, although on their less skilled days those operators end up scanning their own body parts, such as fingers and hands, a phenomenon known as **Google hands**. Some companies are applying the principles of **crowdsourcing** and **gamification** to the digitizing realm, creating leisure activities that let users contribute to the process. (I would be remiss if I didn't mention the opposite process: turning digital Web documents and data into books and zines, a genre called the **printed Web**.)

Ideally, digitized data is online (readily available), but it might end up either offline (not available) or **nearline** (only indirectly available). It can also end up in **dark archives** (which are inaccessible to the public), **dim archives** (which are usually inaccessible but can be made accessible), or **light archives** (another term for those that are fully accessible).

Having digitized some data, the archivist now faces a new problem: the eventual obsolescence of the data structures or media used to store the data, necessitating a **format migration** (or a **media migration**) to something newer. Copying the data without changing the format or media type is called **refreshing**.

There is a large cottage industry of life coaches and self-appointed gurus who recommend, with varying degrees of urgency and stridency, that we become digital dualists and spend less time online. Fulminations against digitization are harder to find, and that's just as well, since, with enlightened institutions such as the New York Public Library leading the way, having digital access to books, photos, and other analog materials can only be a good thing. Try to ignore the fingers. ■

ILLUSTRATION BY Mark Allen Miller

Here's the answer: the egg came first.

5 DAYS. 50 EVENTS.
GO TO KNOW.

JUNE 1–5, 2016 / NEW YORK CITY / WORLDSCIENCEFESTIVAL.COM

Founding Benefactors: SIMONS FOUNDATION · Alfred P. Sloan Foundation · John Templeton Foundation

Major Supporters: ANN ZIFF · bezos family foundation · conEdison everything matters · Della Pietra Family Foundation · THE KAVLI PRIZE

Full Year Media Partners: abc7 · THE WEEK · mental_floss · POPULAR SCIENCE

Festival Media Partners: abc NEWS · ScienceNews · IEEE SPECTRUM · WNYC

University Partners: NEW YORK UNIVERSITY · CUNY The City University of New York · COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK · The Rockefeller University · THE NEW SCHOOL LANG · THE COOPER UNION

World Science Festival

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

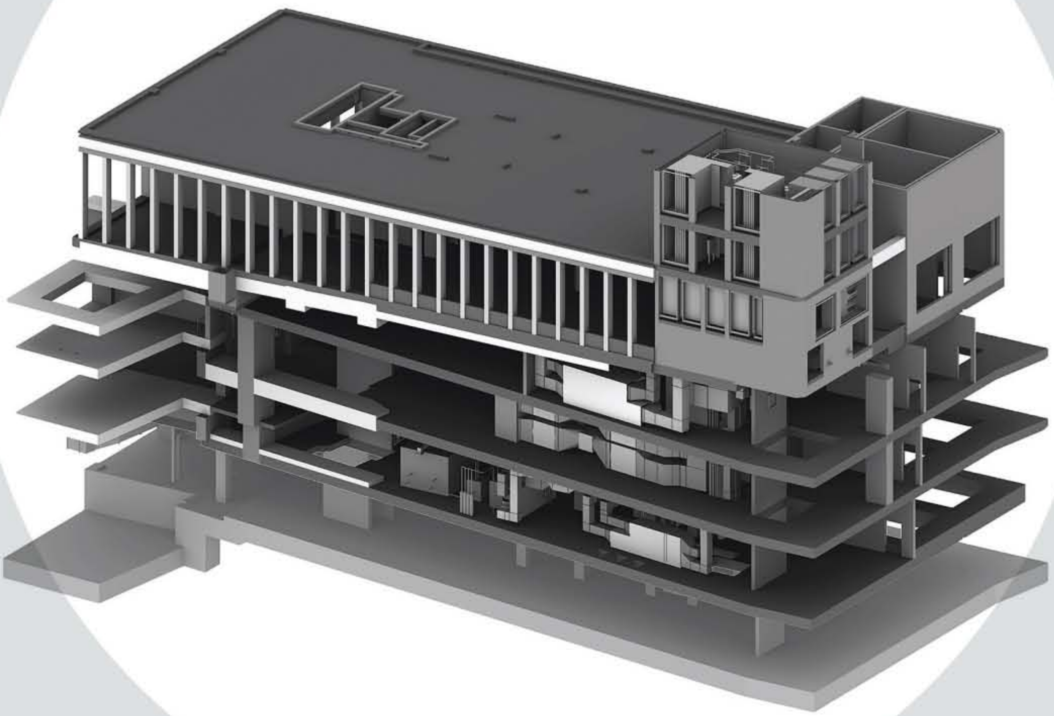

THE **DIGITAL** UNDERGROUND

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

by **MARK PEPLOW**

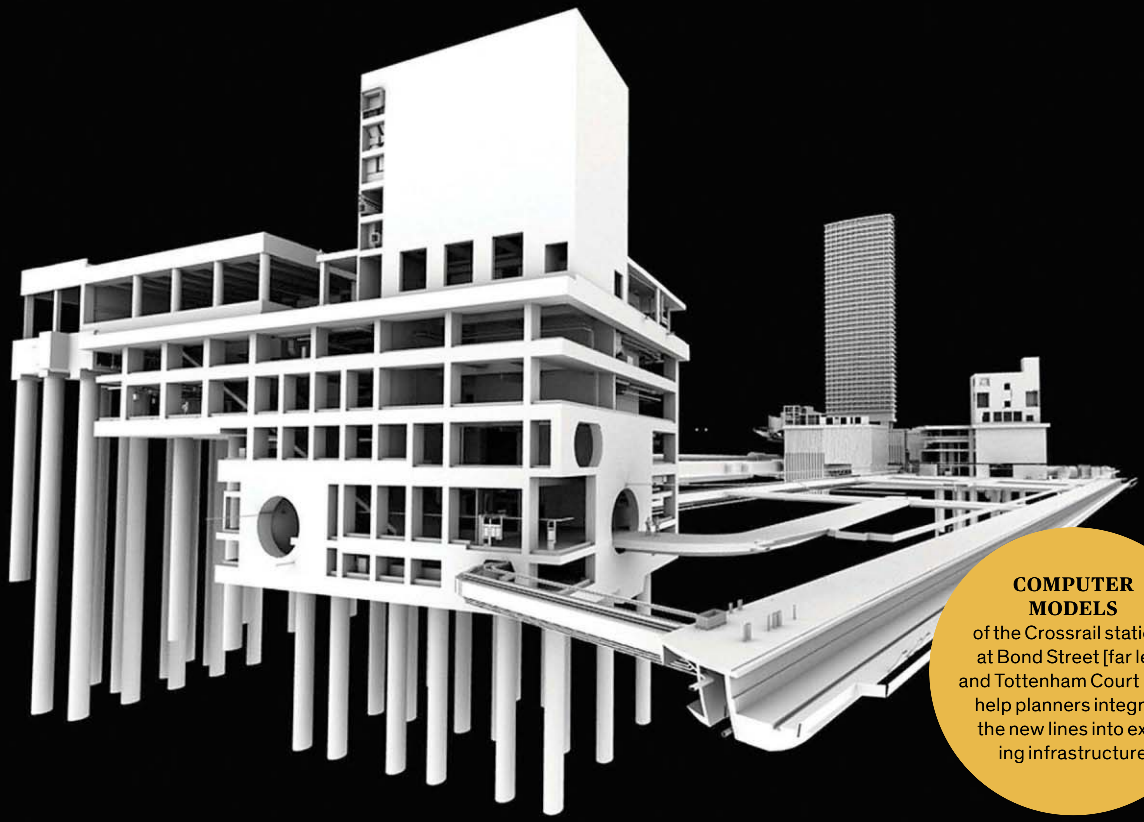

COMPUTER MODELS of the Crossrail stations at Bond Street [far left] and Tottenham Court Road help planners integrate the new lines into existing infrastructure.

London's Crossrail train network is a $21 billion test of virtual modeling

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

DIGITAL FIRST: Computer modeling of every aspect of the Crossrail network (such as the Paddington ticket hall and platforms shown here) has saved time and money by reducing costly construction errors.

AS THE RIVER THAMES meanders eastward through London, it horseshoes around a lobe of land called the Isle of Dogs. In the 19th century, the area boasted one of the world's busiest dock complexes, but by 1980 it had deteriorated into an industrial wasteland. More recently, thanks to massive redevelopment, the area has blossomed again to become Canary Wharf, an enclave of glittering glass skyscrapers that is now a global financial center. And here, in the murky waters of the North Dock, sits one of the largest and sleekest stations in London's newest railway: Crossrail.

With a £14.8 billion (about US $21 billion) budget, Crossrail is currently the biggest engineering project in Europe. Since construction began in 2009, it has bored 42 kilometers of tunnels beneath Central London, creating a subsurface rail network that will speed passengers between the eastern and western fringes of the city. Its 40 stations, 10 of them completely new, will connect with existing National Rail and London Underground lines. The first Crossrail trains will start running in 2017, and when the network is fully open in 2019, it will shuttle an estimated 200 million passengers every year.

The tunneling was completed in May 2015, to much fanfare. But Crossrail is not just about digging—it has been a hotbed of innovation. The entire network was designed in a 3-D virtual environment; once the railway is up and running, a version of this 3-D model will help managers monitor, from a tablet, countless electrical components and systems across the network. "We've built two railways—one real, one virtual reality," says Rhys Williams, the head of mechanical, electrical, and public health for Crossrail. He says one of his goals is to show how transport networks can become safer, smoother to run, and cheaper to maintain by using smart design coupled with the latest technology. And so everything about Crossrail—the escalators and elevators, lighting, ventilation, communication, the railcars themselves—is

ALL IMAGES: CROSSRAIL

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

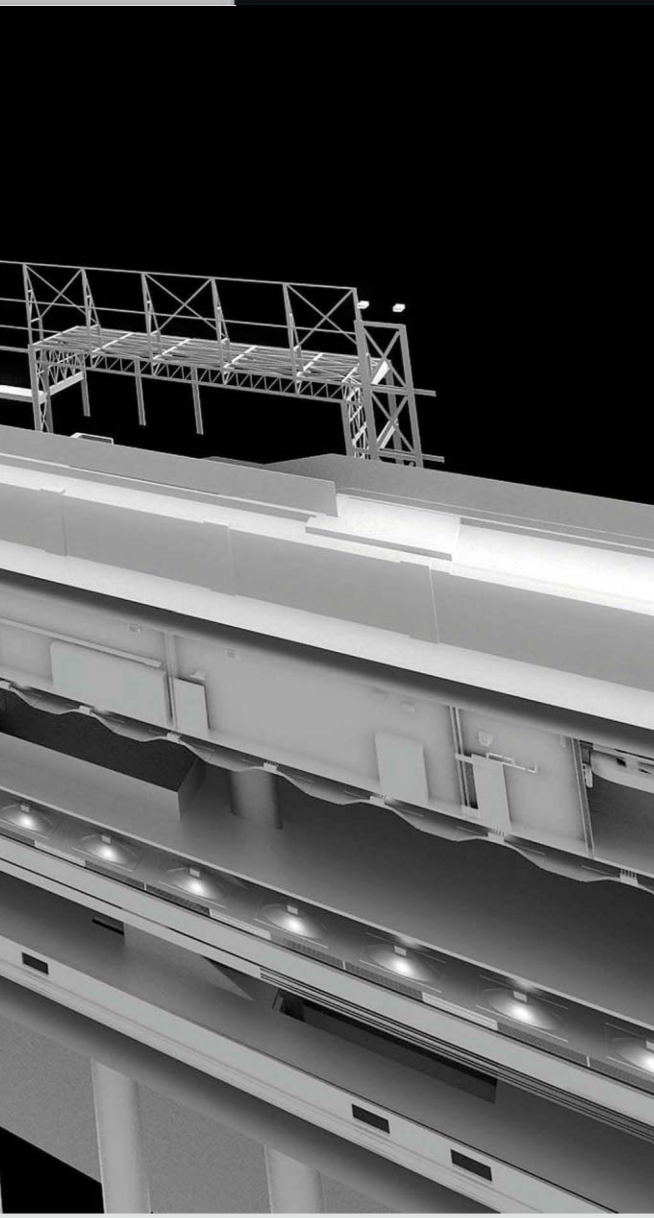

VIRTUAL MEETS REALITY: A point-cloud survey of Paddington station [above] uses laser scanning to ensure that construction conforms to building plans.

being engineered to reduce energy use, improve safety, and streamline operations.

Already, the project is setting standards that are being adopted by other transport networks in the United Kingdom and providing a test bed for digital technologies that could revolutionize megaconstruction. "The whole approach to making it smarter—a lot of that has happened for the first time on Crossrail," says Robert Mair, head of civil and environmental engineering at the University of Cambridge, in England, and a consultant to the project. "Crossrail really is a digital railway."

I FIRST VISITED THE CANARY WHARF station in October. Major construction had concluded a month earlier, and the station is a cavernous and pristine shell. From the outside, the gracefully curved timber lattice roof juts out like the prow of a ship. Inside, it's all gleaming marble floors and polished metal, nearly devoid of people. The next step will be to install ticket barriers, platform doors, and all the other furniture and fittings of a major transport hub.

Today, though, the ticket hall looks more like a dance floor. "Are you ever tempted to have a party down here?" I ask Williams. In fact, I'm told, just 10 days after construction was officially completed, London's *Evening Standard* newspaper held a splashy event here to launch its annual "1,000 power list," celebrating the capital's movers and shakers. ("There's not one single category for engineering," says Williams, with a sigh.) The paper brought in a bar and a choir and hung the ceiling with more than 1,000 glowing globes.

For history buffs, the party harks back to an earlier tunnel-based fete held nearly two centuries earlier and just 2 km away. In November 1827, master engineer Marc Brunel and his son, Isambard, threw a banquet to boost public confidence in their partly completed Thames Tunnel, which was to become the first transport tunnel beneath a major river. First used by pedestrians, it later was opened to trains, which continue to trundle through it today. Like Crossrail, the Thames Tunnel was an engineering marvel, built with the aid of the elder Brunel's tunneling shield—a cast-iron, three-story framework that kept the ground from collapsing while workmen hacked away at the bare earth. On a good day, they might make 30 centimeters of progress.

In contrast, Crossrail's eight massive tunnel-boring machines, each equipped with a 7-meter-diameter cutting head, averaged 38 meters per day. Collectively, the machines gouged out 6 million metric tons of earth, digging as much as 40 meters below ground in places and securing the fresh tunnels with concrete rings as they went.

All those new tunnels and stations will provide much-needed relief to London's overburdened transportation system. Throughout the last century and a half, the expansion of the London Underground subway network, known as the Tube, spurred the growth of London's suburbs and helped

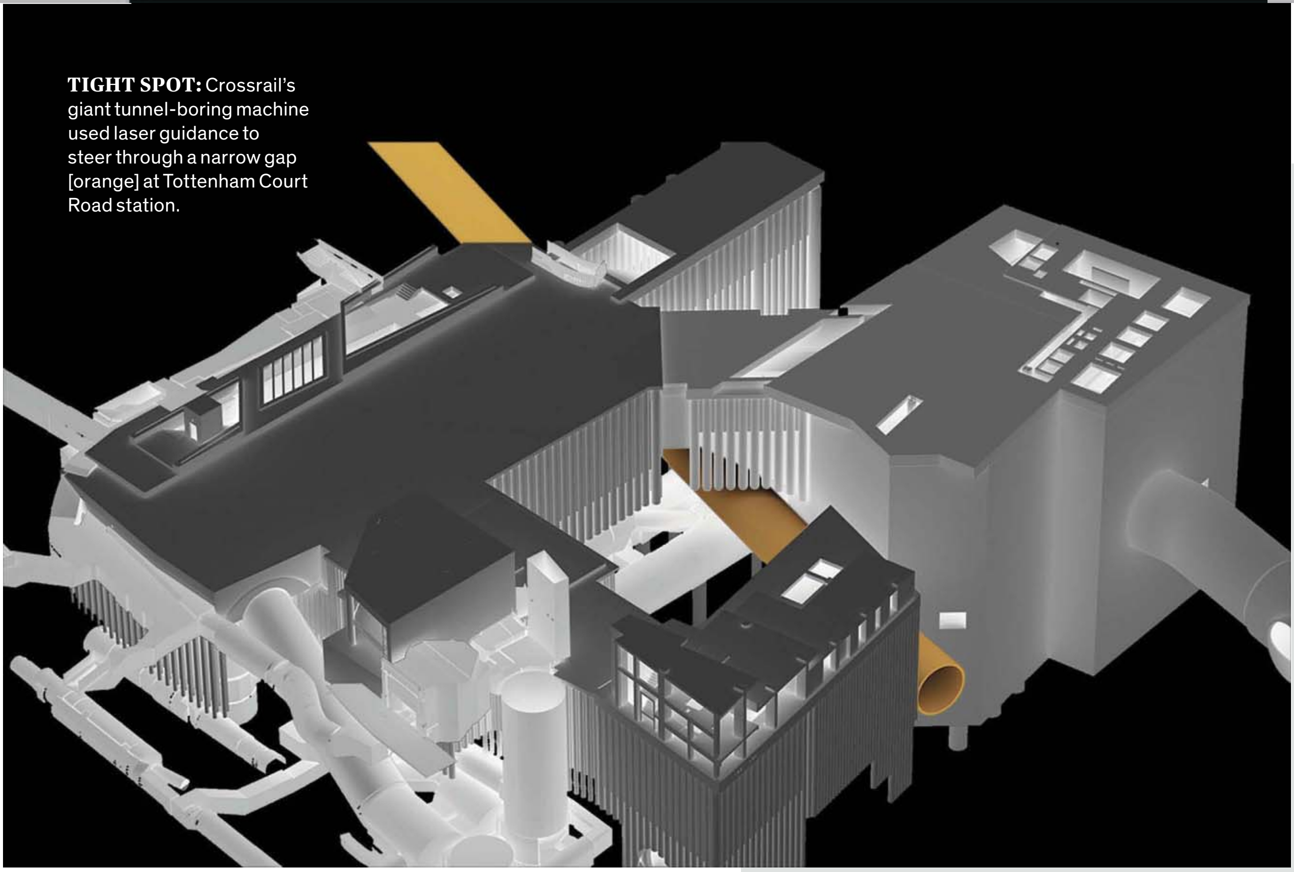

TIGHT SPOT: Crossrail's giant tunnel-boring machine used laser guidance to steer through a narrow gap [orange] at Tottenham Court Road station.

create a sprawling commuter belt. Today the city's population stands at 8.6 million, and it adds 100,000 newcomers each year. Meanwhile, the skyline is spiked with cranes and new office buildings that draw ever more people to inner London.

Roads are choked with traffic. The Tube's spaghetti of 402 km of tunnels and surface lines and 270 stations sees more than 1.3 billion passenger journeys each year. Commuters often stand for their entire journey in hot and cramped conditions.

"Suburbia was created by the railways," says Mark Lucas, who until recently worked for the borough council in Redbridge, in northeast London, where four Crossrail stations are being built. "But suburbia only works if you have good connectivity to Central London and beyond."

Though the need for Crossrail has long been clear, the will to construct it was much slower in coming. Building this massive underground infrastructure beneath a 2,000-year-old city that already has plenty of underground infrastructure has been a herculean feat. Its designers, Williams declares, knew there was little room for error. Unlike, say, an office building, Crossrail's tightly bored tunnels and underground stations won't be able to accommodate lots of extra equipment later should any system be found wanting. "We can only do this once, and we have to get it right the first time," says Williams.

TO ENSURE THEY GET IT RIGHT the first time, Crossrail's designers and engineers have relied heavily on 3-D computer modeling. Widely used in the oil and gas industry and in automotive and aerospace manufacturing, the technique has also been catching on in civilian construction.

The Crossrail model actually consists of more than 250,000 little models joined together in a database and linked to another database containing all the data and documentation about all of the railway's assets–from 1-watt LED lightbulbs to the giant fans that extract smoke in the event of a fire–as well as detailed descriptions of all the work that's going on. Overall, the spatial data takes up a few terabytes, with the associated documentation adding a further 5 terabytes so far. "I don't think building information modeling has been done on this scale anywhere in the world before," says Williams.

The model saves time and money because it offers a common set of data for the thousands of contractors and consultants working on Crossrail. This arrangement helps reduce the disparities that can often arise between the work of two contractors. Contractors can also make changes to their part of the model, which, after review, get incorporated into the master model. Malcolm Taylor, head of technical information at Crossrail, says that previous projects he worked on used multiple methods to track the same information. "We got rid of all this software we didn't need. All our information is now in one place," Taylor says. "It's saved millions of pounds per year."

Once the construction of a space is completed, teams do what's called a "point-cloud survey" using a rotating laser scanner that sends out coherent beams of light and catches the reflections from millions of surrounding points, building up an extremely precise 3-D image of the shape of what's been built. This ground truth is then compared with the virtual model. As long as the match is within a centimeter or so, the model is updated with the laser survey data. Rare mismatches may indicate a problem with the construction that needs to be fixed, says Taylor.

The model will also become a crucial tool for tracking and managing Crossrail's many systems once the railway is running. For example, each station has computer-controlled lighting that dims automatically when nobody is in the area and sends alerts when a fixture needs a new bulb. Another system oversees emergency equipment. Meanwhile, cameras will monitor the movement of people through corridors—if the pace slows to a crawl during rush hour, for instance, the system is designed to automatically redirect commuters along different corridors.

All this data and more will flow wirelessly to base stations and then on high-speed fiber-optic links to the Crossrail control center in northeast London. Managers viewing this information within the virtual model will be able to zoom in on an area of interest and check the operation and maintenance logs for a specific component. They'll also be able to access the model through a tablet computer. "We're giving them something they've never had before—a maintenance manual on an iPad," says Williams.

The model even features an augmented-reality interface. Workers are able to hold an iPad up to a station wall, for example, and see a virtual view of the underlying electrical wiring and ductwork overlaid on the wall. The image of the wall is simply what the iPad's camera is seeing; the image of what's behind the wall is generated from the model based on the iPad's coordinates as determined by GPS (when it's available) and Bluetooth beacons. It's a handy feature to avoid accidentally drilling holes through cables and to make sure things fit before they get to the site.

A geographical information system was used to create Crossrail Maps, an application that will allow Crossrail employees to query a particular location of London on a map and then dive into the Crossrail assets there. "We can click on a station and then on a defined space, such as a room," Taylor explains. "Because every asset we create has a set of *x*, *y*, and *z* coordinates, we can call up all the assets in that room from another linked database and get all the information about each one." Previously, locating a component meant digging through architectural drawings and archived manuals. "What would normally have taken many minutes—and sometimes hours—can now be done in seconds," Taylor says.

CROSSRAIL'S TRAINS are not the short, skinny rattletraps of the ancient Tube network. They are full-size, state-of-the-art trains, each one able to carry up to 1,500 people at a time. Accordingly, Crossrail has much longer station platforms and wider tunnels than does the Tube.

Those tunnels, 7.1 meters in diameter, posed a particular challenge. The settling of the earth increases with the square of the tunnel's diameter, so engineers had to figure out how to prevent buildings above them from sinking or tilting. To protect the buildings, Crossrail's engineers used a technique known as compensation grouting, developed in the 1990s by the Geotechnical Consulting Group, a London-based company cofounded by Robert Mair, of the University of Cambridge. It involves drilling a horizontal fan of narrow tubes into the earth directly above where the tunnel will be bored and then fitting the tubes with pipes that can dispense liquid cement. On the surface, workers affix optical prisms to surrounding buildings, and then a surveying instrument known as a theodolite bounces laser beams between the prisms, monitoring their precise locations as the tunnel is bored. If a building moves, the laser beams fall out of alignment, and workers promptly pump cement into the fan of tubes to strengthen the surrounding ground. Remarkably, no buildings were damaged during the digging. "These large tunnels have been constructed under literally thousands of buildings, and they have all been protected," says Mair.

Construction engineers also used laser guidance to steer a tunnel-boring machine through the tightest spot in the network, at Tottenham Court Road station, where two Tube lines already intersect amid a tangle of pipes, cables, and sewers. Here, the giant digger managed to thread through a narrow gap that left just 85 centimeters' clearance above one platform tunnel, and only 35 cm below an escalator.

During construction, engineers have been able to monitor the status of tunnels, shafts, and stations thanks to inexpensive wireless sensors and fiber-optic cable. Mair and his colleagues at the Cambridge Centre for Smart Infrastructure and Construction (CSIC) developed the real-time monitoring technology, and it's helped them assess factors like the exter-

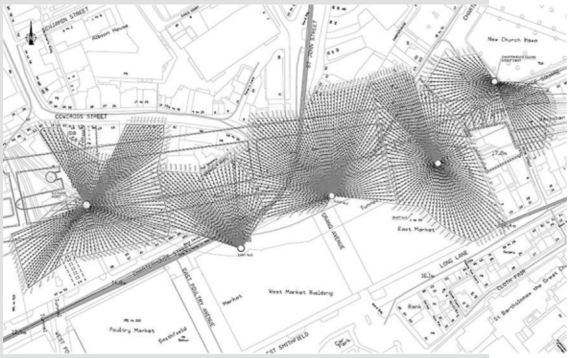

DID THE EARTH MOVE? To tunnel beneath a 2,000-year-old city without disturbing existing structures, engineers used a technique known as compensation grouting. If laser monitoring aboveground [top] revealed unwanted movement, workers injected liquid concrete into narrow tubes to stabilize the ground.

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

nal stress on a tunnel. Mair's team embedded cheap optical fibers into the concrete walls; the scattering of light along the fibers reveals how much strain they're under. Mair says this monitoring will help to determine if the tunnels were built stronger than they really need to be. Such information could help make future designs less expensive.

Meanwhile, a spinoff of the CSIC called UtterBerry has been using Crossrail to test low-power wireless sensors that can monitor not only strain but also temperature, humidity, and many other aspects of construction work. Wireless sensors have been used in construction for more than 20 years, but they have typically been heavy, weighing 1 kilogram or more, with battery lives measured in weeks. An UtterBerry weighs just 15 grams and can operate for more than a year on one charge. In principle, an engineer could carry hundreds of them at once and mount them in a tunnel over the course of a day.

Heba Bevan, a Ph.D. student at CSIC, designed the small sensor as a classroom project. Unlike conventional sensors, she says, UtterBerries collect readings only when they need to, crunching the raw data with an ARM Cortex processor. By tracking changes in temperature and humidity, for example, the sensors can calculate how the weather might be causing a concrete panel to expand or contract; any expansion or contraction that exceeds such expected seasonal changes would require further investigation. And if one UtterBerry detects unusual vibrations or some other anomaly, it can tell other units in the area to up their sampling rate to boost the overall quality of data being gathered.

UtterBerry's first project for Crossrail was in late 2012, and Bevan founded the company the following year. Her full-time staff of five just won a contract to provide yearlong sensor monitoring for another major U.K. construction project.

"Usually it's hard for smaller companies to get into construction," says Bevan. "But Crossrail has really encouraged this." In the future, she believes that sensors like these won't be added as an afterthought but rather designed into a bridge, building, or tunnel itself, to provide lifetime monitoring from construction to demolition.

DOWN ON THE PLATFORM at Canary Wharf, I can hear the relentless noise of hammering and drilling emanating from further down the tunnel. Later, a mobile gantry will lay rails on top of the ties, followed by a 465-meter-long "concreting" train that will pour its load around the ties to fix everything in place. Then another train will install high-voltage cable, to provide power in and around the tunnels. The trains themselves will draw electricity from 25-kilovolt overhead lines, so that they can travel seamlessly between Crossrail and the aboveground train system.

Canary Wharf's escalators sweep upward from the platform, and although they look fairly conventional, they're not. Unlike the elderly, power-hungry machines still found in some Tube stations, which break down frequently and can take months to refit, the Crossrail escalators rely more on off-the-shelf replacement parts, which will speed up maintenance. Thanks to machining and casting techniques borrowed from the aircraft industry, the parts are much lighter, too.

About 85,000 people will stream up and down these escalators every day once Canary Wharf opens in late 2018. At the Liverpool Street and Farringdon stations, there will even be inclined, glass-walled elevators that travel at a 30-degree angle right next to the escalators. "Kids going up the escala-

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

tors can be waving to mum, who's in the inclined lift with the twins in a buggy," says Williams.

A handful of inclined elevators are operational in Europe and the United States (at the Huntington Metro station in Alexandria, Va., for example), he adds, but Crossrail is introducing them to the U.K. and setting the U.K. standards for how these elevators should be fitted. His team has also convinced Tube engineers to take up the technology. Last October, Greenford station became the first Tube station to have an inclined elevator.

"A few years ago, the prevailing attitude was, 'My granddad didn't do it like that,' " says Williams. "But now they are championing these works."

BACK IN WILLIAMS'S OFFICE, I ask what's still up in the air as this megaproject enters its final phases. Rather a lot, it turns out. Despite the huge amount of construction work that has already been completed, many aspects of Crossrail do not yet have an agreed-upon final design. That's because the complexity and uncertainty involved in building large underground structures means that detailed "fit out" designs are begun only after it's fairly clear how the space will look.

In answer to my question, Williams waves at some tall stacks of paper. "All of this," he says with a smile. Patting one stack, he adds, "This one is [the design for] Woolwich station. It's only just arrived, so we haven't started the design review yet." He hopes the remaining design reviews will be done by the end of July, so that the fit-out work can start in earnest on these stations. "The only part of Crossrail that we've actually finished is the academy," he adds.

He's referring to the Tunneling and Underground Construction Academy in Ilford, a facility that opened in 2011 to train more than 10,000 people to work in Crossrail's tunnels. The tunnels and stations themselves have been a valuable proving ground for hundreds of young engineers. Many of them could move on to Crossrail 2, a north-south twin that (if approved) would be built from about 2020 to 2030. Williams believes that the lessons learned on Crossrail could make that project run much more smoothly.

Though it hasn't even opened yet, the new railway system has already won over some key people. Christian Wolmar, a prominent transport journalist who last year campaigned to be mayor of London, is as supportive of Crossrail as he is opposed to some other big infrastructure projects. He has skewered HS2, the U.K.'s proposed high-speed railway from London to Birmingham, Manchester, and Leeds, as "the wrong scheme in the wrong place." But ask him about Crossrail and his enthusiasm is unbridled.

"It's a game changer," Wolmar says. "The impact is going to be massive. People will flock to it."

Meanwhile, Crossrail's opening date looms, and Williams and others continue to push ahead on all fronts. Their efforts, they hope, will be an enduring engineering legacy. "We're not building this just for the full Crossrail opening in December 2019," says Williams. "We're looking at the next 120 years." ■

DIGGING DEEP: Transforming Crossrail's 42 kilometers of tunnels into working rail lines has involved, among other things [from left], the actual tunneling, completed in May 2015; spraying concrete embedded with optical fibers, which reveal any strain in the tunnel walls; connecting Crossrail tunnels with existing passages; and installing rails and power lines.

POST YOUR COMMENTS at http://spectrum.ieee.org/crossrail0416

LEAVE THE DRIVING TO… *THEM*

+

TOP TEN TECH CARS 2016

BY LAWRENCE ULRICH

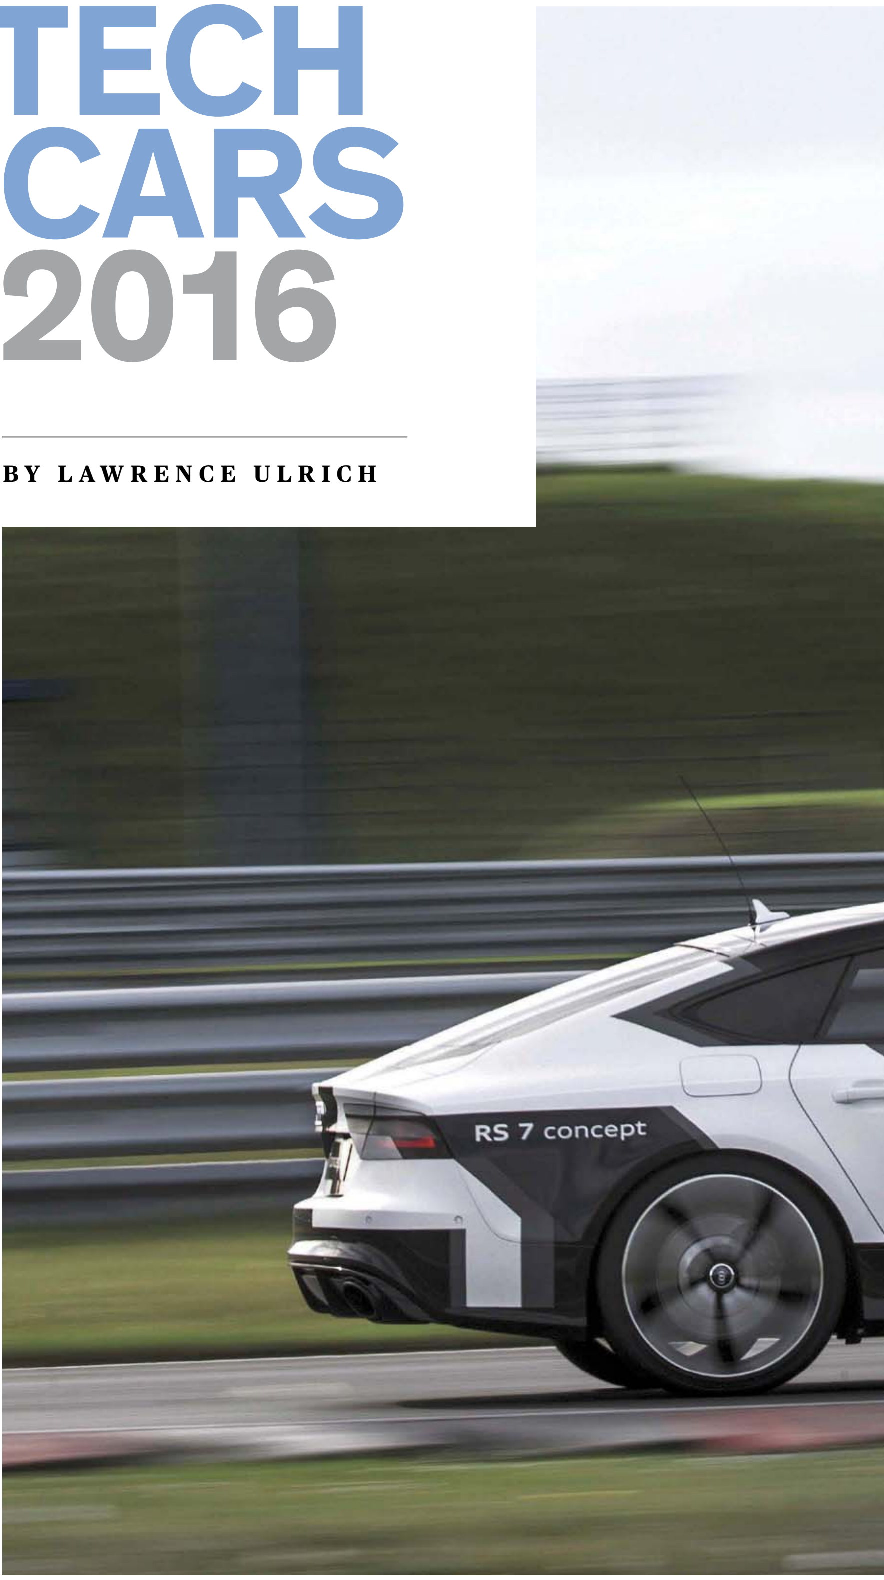

AUDI

"We wanted flying cars; instead we got 140 characters," is venture capitalist Peter Thiel's famous credo.

But though a freeway in the sky seems as fantastical as ever, we are going to get something even better: a self-driving car.

Such a robot, fully aware of its environment, with 360-degree vision and peerless driving skills, is a matter of when, not if. Humans' fascination with these machines seems limitless, even though autonomous cars could turn us into mere cargo. And unlike airborne cars, self-drivers could prevent the 1.2 million deaths caused by traffic accidents every year.

It's no utopian fantasy. Among our Top Ten Tech Cars this year are a robotic Audi that tears around racetracks like a professional driver and an electric Tesla whose impressive autopilot skills are as close as the nearest showroom.

Booming sales should also help accelerate the technological pace. Americans parked 17.5 million new cars in their driveways in 2015, more than any year in history, and the Chinese bought even more. That left the industry awash in profits and able to spend heavily on R&D to bring pioneering cars and technologies to market.

So, carbon-based life form, the message is clear. If you enjoy driving, get your fill while you can. In our list of the 10 cars that are rocking our technological world, we've included more than a few choices to help maximize your motoring pleasure. After all, we're only human.

Audi Autonomous RS7
Fast lap to the future

Audi's autonomous cars are becoming quite the world travelers: Recall the much-ballyhooed first robotic drive from San Francisco to New York City, about a year ago. Impressive stuff, though honestly, humans can hold their own at pulling into a rest stop.

Here in Spain, the man-vs.-machine competition will be at higher speed and for higher stakes. I'm about to take on Robby, the autonomous RS7 sport sedan that's designed to rock a racetrack at speeds that would blister Google's cartoonish bubble car. If a human driver can't keep up, it occurs to me, then our obsolescence draws that much closer. It's only a matter of time before governments and automakers pry the ignition keys out of our fallible, accident-prone hands for good.

Robby is looking cool and confident in the pits at Parcmotor Castelloli, near Barcelona. And for good reason: The Audi weighs 400 kilograms (882 pounds) less than Bobby, the RS7 that holds the world speed record for autonomous cars, at 240 kilometers per hour (149 miles per hour).

I take to calling the newer car Robby the Robot, after the glass-skulled automaton from the 1956 sci-fi movie *Forbidden Planet*. Miklos Kiss, Audi's head of predevelopment for driver assistance systems, introduces us to his two autonomous brainchildren. Popping

PRICE	Not for sale; prototype
POWER PLANT	4.0-L V-8 with twin turbochargers; 418 kW (560 hp)
OVERALL FUEL ECONOMY	12.4 L/100 kilometers (19 mpg)

Bobby's hatch, we find it full to bursting with computer gear. Robby's, in contrast, has plenty of room left over for luggage. There's a single MicroAutoBox brain and power supply and two other small computers.

Incredibly, Audi's latest differential GPS unit can fix Robby's position to within 2 centimeters, vastly better than today's GPS standard of roughly 1.5 meters. This price-no-object system also uses redundant cameras to triangulate and thus to confirm the car's location. Keeping a fully autonomous vehicle safely within lanes will require zeroing in to 50 cm (around 20 inches).

Controllers adjust the engine, electric steering, transmission, and brakes, with redundant fail-safes: There's a spare power supply and brake controllers if the first ones conk out. A 4.0-liter biturbo V-8 produces a villainous exhaust note that echoes off the dusty canyon walls.

I slide into Robby's shotgun seat with some trepidation, thinking that the driver's side will remain spookily unoccupied. But surprise! There's an Audi engineer in the seat, along for the passive ride but holding a plunger connected to a cord. If something goes wrong and he lets go of the plunger's button, the Audi will slow and halt on the track. Theoretically.

The checkered flag waves. The RS7 launches itself down the front straight and charges into the first corner, the steering wheel twirling, the ghosts fully in charge of this machine. The brake and throttle pedals aren't moving at all because the computer commands are bypassing the old analog connections.

Before I arrived, Audi engineers had manually driven Robby around this Spanish track to measure its barriers and "geo-fence" a safe zone beyond which Robby will not go. Like a real-life slot car, the Audi locks onto its programmed track line, its path varying by only a few centimeters. Yet the RS7 also reacts in real time to conditions such as a slippery track or wearing tires, dialing back power or correcting the steering if it begins to slide off its satellite-guided path.

Robby turns out to be a smoothy, a race-instructor type who puts up great, flowing lap times yet keeps the car utterly balanced and composed. Its dead-consistent laps vary by less than 0.3 second, including a best this day of 2:09.2.

Now, it's my turn, and I jump aboard a production Audi RS7. Suddenly, I'm John Connor taking on Skynet, a human fighting for his increasingly pitiful and superfluous species. After one reconnaissance lap on this unfamiliar track, I turn the Audi loose. Please, Lord, don't let me lose to a stinking machine.

I drive back into the pits and dash over to the timer. A few Audi engineers applaud, a bit grudgingly, I'm thinking. But my lap of 2:05.4 is nearly 4 seconds quicker than Robby's. Even after hundreds of laps in recent weeks, Robby's best is still 2 seconds behind my first trip around Castelloli. Take that, you remorseless Terminator, German accent and all.

But as my adrenaline subsides, I'm forced to concede that adrenaline is among my biological advantages. Robby may know speed, but the words "race" or "win" simply aren't part of his vocabulary. Yet.

Where I punished the tires and pushed the limits, Robby stayed emotionless, programmed to run safe, endlessly repeatable laps. The last thing Audi needs is for a self-driving car to disintegrate against a wall or injure its passengers, setting back autonomous driving by a few years. A few tweaks of the algorithms, a few more generations, and Robby's offspring will be chip-enabled Michael Schumachers. (Please, Audi, name your next car Ricky Bobby.) They'll scan and pick out us humans up ahead and make us eat their digital dust, if they choose. Or they'll chauffeur our miserable hides straight to the police if we act up, as did Tom Cruise's 2054 Lexus in *Minority Report*. What's to stop them?

Yes, the rise of the machines seems inevitable. But as my race with Robby showed, some humans will still put up a fight before going to the scrap heap. ■

BRAIN IN A BOOT: The Audi RS7's computing hardware fills the trunk space.

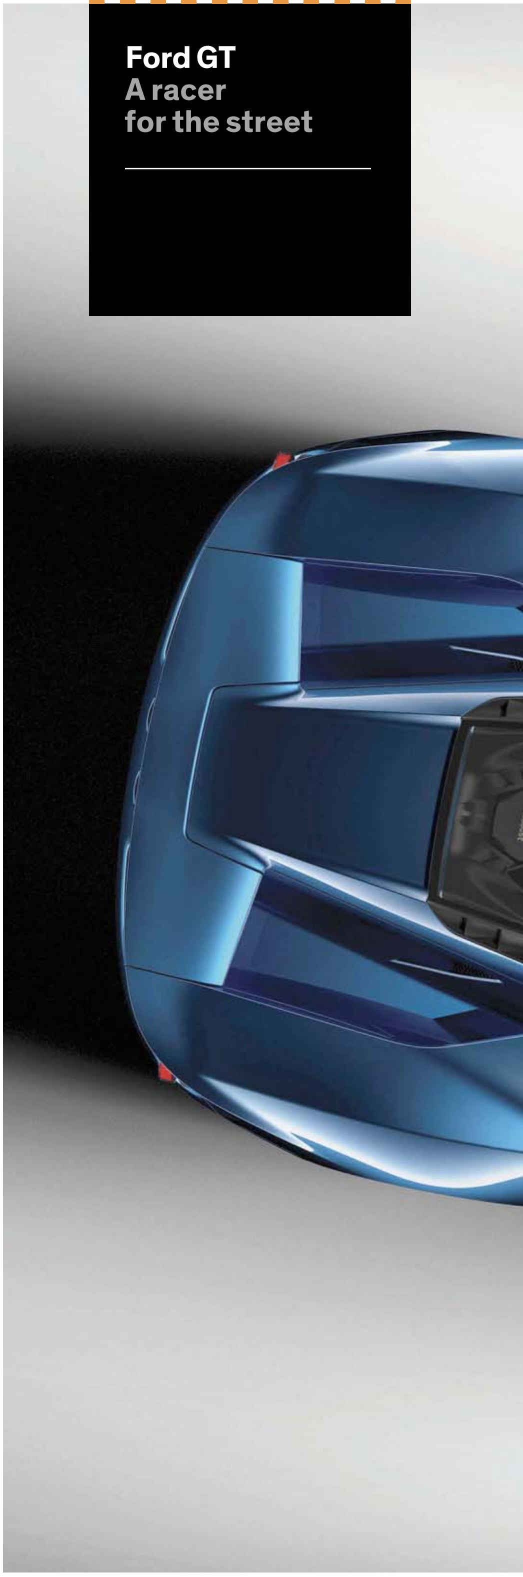

Ford GT
A racer for the street

LEFT: AUDI; RIGHT: FORD

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

Nearly a half century ago, the Ford GT40 went to the 24 Hours of LeMans and crushed mighty Ferrari, sparking an enduring legend. Now Ford looks for a LeMans déjà vu this summer with a reborn racing GT, followed by 250 annual copies of a roughly US $400,000 scissor-doored wonder car for the street.

The GT eschews a V-8 for a downsized twin-turbo V-6 based on its Daytona-winning LMP2 race engine. Ford is promising the best power-to-weight ratio of any production car in the world, with a hand-laid carbon-fiber tub and body for this mid-engine monster.

An active suspension lets the Ford hunker down at triple-digit speeds to reduce drag, while an active air brake at the rear rises and angles as needed to boost aero downforce or slow the car into corners. The gorgeous fuselage is billionaire bait, but the bravura style is wedded to pure function. A curved pair of flying buttresses perform dual tricks: The winged roof channels direct air to the rear spoiler, and their hollow sections contain piping for the turbo intercoolers: Engine intake air is hoovered from beneath the car, compressed into the turbos, then snaked through the winglets and down again to hyperventilate the V-6. Heated air from the intercoolers flows rearward and exits through tubes in the center of the rear taillights.

It's all executed so beautifully that we were pleased to gawp at the thing at the recent Detroit Auto Show. But I'll be happier when Ford finally lets us drive it. ■

The gorgeous fuselage is billionaire bait, but the bravura style is wedded to pure function

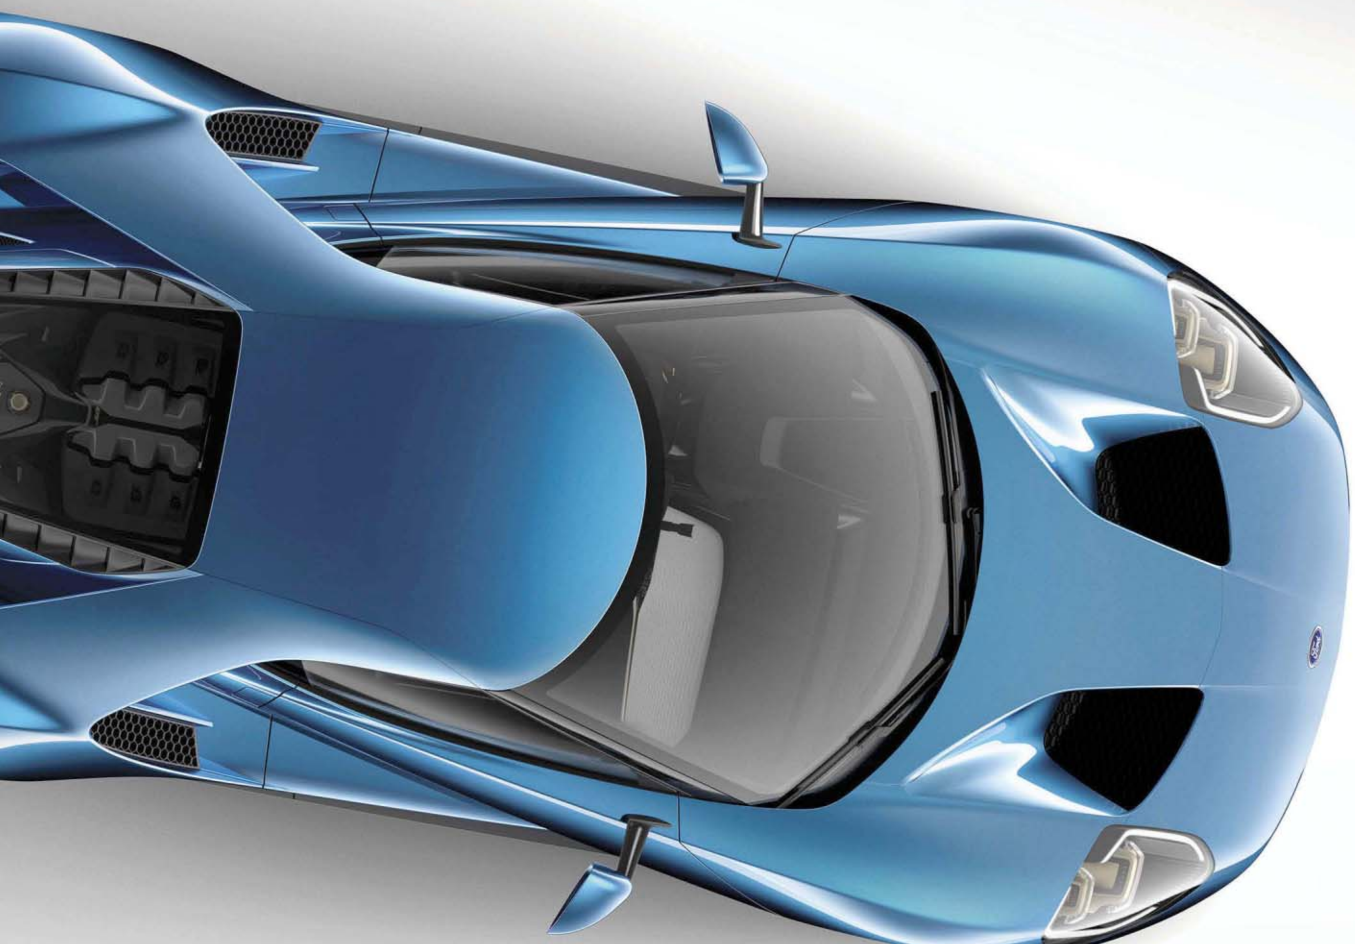

PRICE	US $400,000
POWER PLANT	3.5-L V-6 with dual turbochargers; 485 kW (650 hp)
OVERALL FUEL ECONOMY	N/A

Ferrari 488 Spider
No more natural breathing

Hustling Ferrari's latest fantasy through Italy's Emilia-Romagna region, we take all of 3 seconds to salute the end of an era and to hold on tight as another begins. That's how much time it takes the glorious 488 Spider to reach 100 kilometers per hour (62 miles per hour). Closing a book on seven decades of howling, superhigh-revving, naturally aspirated engines, every new Ferrari will now be turbocharged, a hybrid, or both. And while Ferrari's new turbos can't hit the operatic, 9,000-rpm tenor notes of its predecessors (not yet, anyway), it's game over in every other regard.

Although the midmounted V-8 displaces 0.6 liters less than the departed 458 model, it pumps out a shocking 20 percent more power. And the 760 newton meters (561 foot-pounds) of torque surpasses the old 458 by fully 40 percent. The result shatters the record for power per liter for any production-car V-8 even as fuel consumption drops by 14 percent.

That unbeatable one-two punch of power and efficiency is why such stalwarts as Ford, Mercedes, and Porsche have gotten the memo: Join the turbo revolution, or die. In 8.7 seconds, or less time than a Toyota Prius takes to reach 100 km/h, the Ferrari is doing 200 km/h (124 mph).

With warp-speed acceleration achieved, Ferrari looked to tackle another, more elusive supercar goal: making the Ferrari easy to drive. The Side Slip Control 2 system assesses a driver's skill level in real time, applying its Formula One–bred stability and traction systems to maximize speed in any situation. Twitch a finger on the carbon-fiber steering wheel and the Ferrari reacts in 0.06 second. Flick the column-mounted paddles and the seven-speed F1 automated gearbox downshifts

PRICE	US $275,000
POWER PLANT	3.9-L V-8 with dual turbochargers; 493 kW (611 hp)
OVERALL FUEL ECONOMY	13.8 L/100 km (17 mpg)

40 percent faster than before and upshifts 30 percent more rapidly.

The style is *bellissimo*, naturally, but the beauty springs from pure aero function. The signature is the new "blown spoiler," a discreet cove atop the rear deck that funnels air to pin the Ferrari to the pavement, with no need for a rear spoiler that would add aero drag and spoil the lines.

After slicing up the countryside like a haunch of prosciutto, we find final affirmation on a looping ascent to Forte di San Leo, a soaring promontory and medieval fortress. The sash-wearing mayor and townspeople pour out to greet our Ferraris, snapping enough photos to fill a family album. Yes, driving a 488 Spider in Italy is almost like cruising in the Popemobile. But while the pope can't make you infallible, this car just might. ■

With warp-speed acceleration achieved, Ferrari looked to tackle another, more elusive supercar goal: making the Ferrari easy to drive

Mercedes-Benz F 015 Concept
The mobile lounge

PRICE Not for sale

POWER PLANT Hydrogen fuel cell

OVERALL FUEL ECONOMY N/A

For years, automakers have rhapsodized about how our cars would become mobile offices and living spaces. And then they botched even the simple stuff, like letting you dial up the Backstreet Boys from your iPod on the car's sound system. But fear not. The Mercedes-Benz F 015 will be the rolling roost of your dreams. You'll just have to wait at least until 2030, when Mercedes thinks this kind of hydrogen-powered, fully autonomous vehicle will become viable.

Bigger than an S-Class, the Benz concept looks like a *Clockwork Orange* hipster lounge, with its walnut-veneered floor and wall-wrapping touch and gesture displays. Mercedes sees it as a retreat that will maximize privacy or productivity in the hectic urban zones of the future. The mod theme continues with two front white-leather-clad lounge seats that swivel rearward (after all, the "driver" usually won't need to attend to the road). The steering wheel telescopes into the dash during autonomous mode. And any passenger can take charge of vehicle functions, such as speed and which of the 360-degree views to project inside the car.

The concept car makes clever use of optical technologies to communicate with cars and pedestrians. A forward-looking laser, for example, can beam messages onto the pavement, including a whimsical image of a zebra crossing or the words "please go ahead." You get into the car using a smartphone app, which opens enormous clamshell-style portals for easy access to the lounge space inside. A hydrogen-fueled F (for "fuel") cell plug-in hybrid drive system could deliver a 1,100-kilometer driving range, Mercedes says, including 200 km on battery power.

If you care to trust Mercedes's crystal ball, by 2030 hydrogen cars will be a common sight. Unless gasoline still costs $2 a gallon. ■

LEFT: FERRARI; RIGHT: STEFAN MILEV/WILDFOXRUNNING

Hyundai Tucson Fuel Cell
Fill 'er up— with hydrogen

PRICE
US $499 per month on lease

POWER PLANT
Hydrogen fuel cell with 99 kW (134 hp) electric motor

OVERALL FUEL ECONOMY
4.7 L/100 km (50 mpge)

Hydrogen cars have been such a tease. For decades, carmakers have held out the hope of clean energy, free of fossil fuels and tailpipe emissions, without telling us where the hydrogen will come from and what it will cost.

The Tucson answers the second question, at least: Lease the Tucson in California—where a US $5,000 state subsidy reduces the payment to around $440 a month—and Hyundai will pick up the tab for three years' worth of hydrogen fuel.

The Hyundai is basically an electric car with a fuel cell fed by compressed hydrogen in place of a lithium-ion battery. A pair of tanks hold 5.6 kilograms of hydrogen at 69,000 kilopascals (10,000 psi). It flows into a fuel cell, where electrons are stripped from hydrogen atoms, creating ions that flow one way as electricity flows the other way, powering an electric motor. Then the ions combine with oxygen in the air to form water vapor—the only emission.

The Tucson drives like any electric car, only slower than most. Figure more than 10 seconds to reach 100 kilometers per hour (62 miles per hour). That's what you get when a 99-kilowatt (134-horsepower) AC motor must move about 2 metric tons of SUV.

The advantages of hydrogen include good, old-fashioned fill-ups, taking not much longer than a gasoline stop, and also a reasonable driving range of 426 km (265 miles), according to Hyundai.

As for the "where," there are fewer than 10 public hydrogen stations in California. However, the state plans to invest $47 million to open another 28 stations and to bring that number up to 100 over a decade. ■

LEFT: HYUNDAI; RIGHT: GM

Chevrolet Volt
Lighter, faster, and more frugal

Today's US $2-a-gallon gasoline doesn't bode well for the plug-in Volt's showroom sales in America. But when gas prices spike again, the Volt hybrid may end up looking like one of the smartest fuel savers on the planet—and for one-third the price of a Tesla Model S.

This second-generation Volt is a green sweetheart to drive, and improved in nearly every way over the original: sleeker, lighter, faster, quieter, and more efficient.

Let's cut to the chase. On a light-footed test drive north of San Francisco, I managed precisely 60 all-electric miles (97 kilometers) before the gasoline engine kicked in. That beat the official 53-mile estimate, itself a 40 percent improvement over the first-generation Volt. At that point, the Volt's new direct-injection engine smoothly blended combustion with electric power to offer 675 km of total range—no range anxiety in this car. As we crossed the 106-mile mark on this trip, we'd burned exactly one gallon of gasoline—3.8 liters—along with $1.50 worth of wall electricity.

In other words, we spent $3.50 to cover 106 miles. Go ahead and try to do that in a Mini Cooper. There's not a gasoline, diesel, or conventional hybrid on the road that can match that efficiency, which equated to 2.0 L/100 km (120 mpge) over the electric portion and better than 70 mpg overall. As advertised, owners who commute fewer than 53 miles round-trip can punch into work every day without ever using a drop of gasoline. Purely coincidentally, the Volt's battery supplies a half-gallon's worth of gasoline energy. So if you manage 53 miles (85 km) on a charge—which we exceeded without even glancing at the helpful gauges that coach you toward efficient driving—that works out exactly to the official U.S. estimate of 106 mpge.

And you won't sacrifice performance. The new Volt's curb weight drops by more than 90 kilograms (200 pounds), to about 1,607 kg, and the propulsion unit alone is 45 kg lighter. The new four-cylinder engine is more powerful, runs on regular rather than premium fuel, and operates at lower rpm to reduce the drone that plagued the original.

The dual electric motors are markedly revised, sharing no common parts with the first-gen Volt. Total horsepower

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

remains at 149, but it now has a very hefty 401 newton meters (296 foot-pounds) of torque—21 more than before. The upshot is an 8.4-second surge to 60 mph (97 km/h). Perhaps more important, it gets to 30 mph in just 2.6 seconds, 0.7 second quicker than before—and remarkably, just two-tenths of a second slower than the 220-horsepower Volkswagen GTI that served as my "rabbit" during the driving

We spent $3.50 to cover 106 miles. Go ahead and try to do that in a Mini Cooper

test. One of the Volt's coolest new features is a steering-wheel paddle that triggers regenerative braking. Grab the paddle as you're entering turns or rolling up to stoplights and it feels like downshifting in a sports car, even as you're saving energy.

Battery mass drops by 10 kg, and there are 192 lithium-ion cells in the Volt's T-shaped, under-floor battery pack, down from 288 before. Yet thanks to a tweak in battery chemistry, capacity is up 8 percent, to 18.4 kilowatt-hours—the secret to the Volt's newly extended electric range.

And at $26,495 after a $7,500 federal tax giveaway (er, credit) in the United States, the Volt is a bargain. The Chevy is not only two-fifths the cost of a Tesla and one-fifth the cost of a BMW i8, it's actually less than the $33,500 price of the average new car in America—and $1,200 less than its lower-performing predecessor.

The Volt isn't the only Chevy that will test American appetites for electrified cars: Late this year, the Chevrolet Bolt goes into production, a pure EV hatchback that promises 200 all-electric miles, for an estimated price of around $30,000 after federal and local tax breaks.

And what will they call the Bolt's electric successor? Well, it will have to be "Jolt," because "Colt" has already been taken. ■

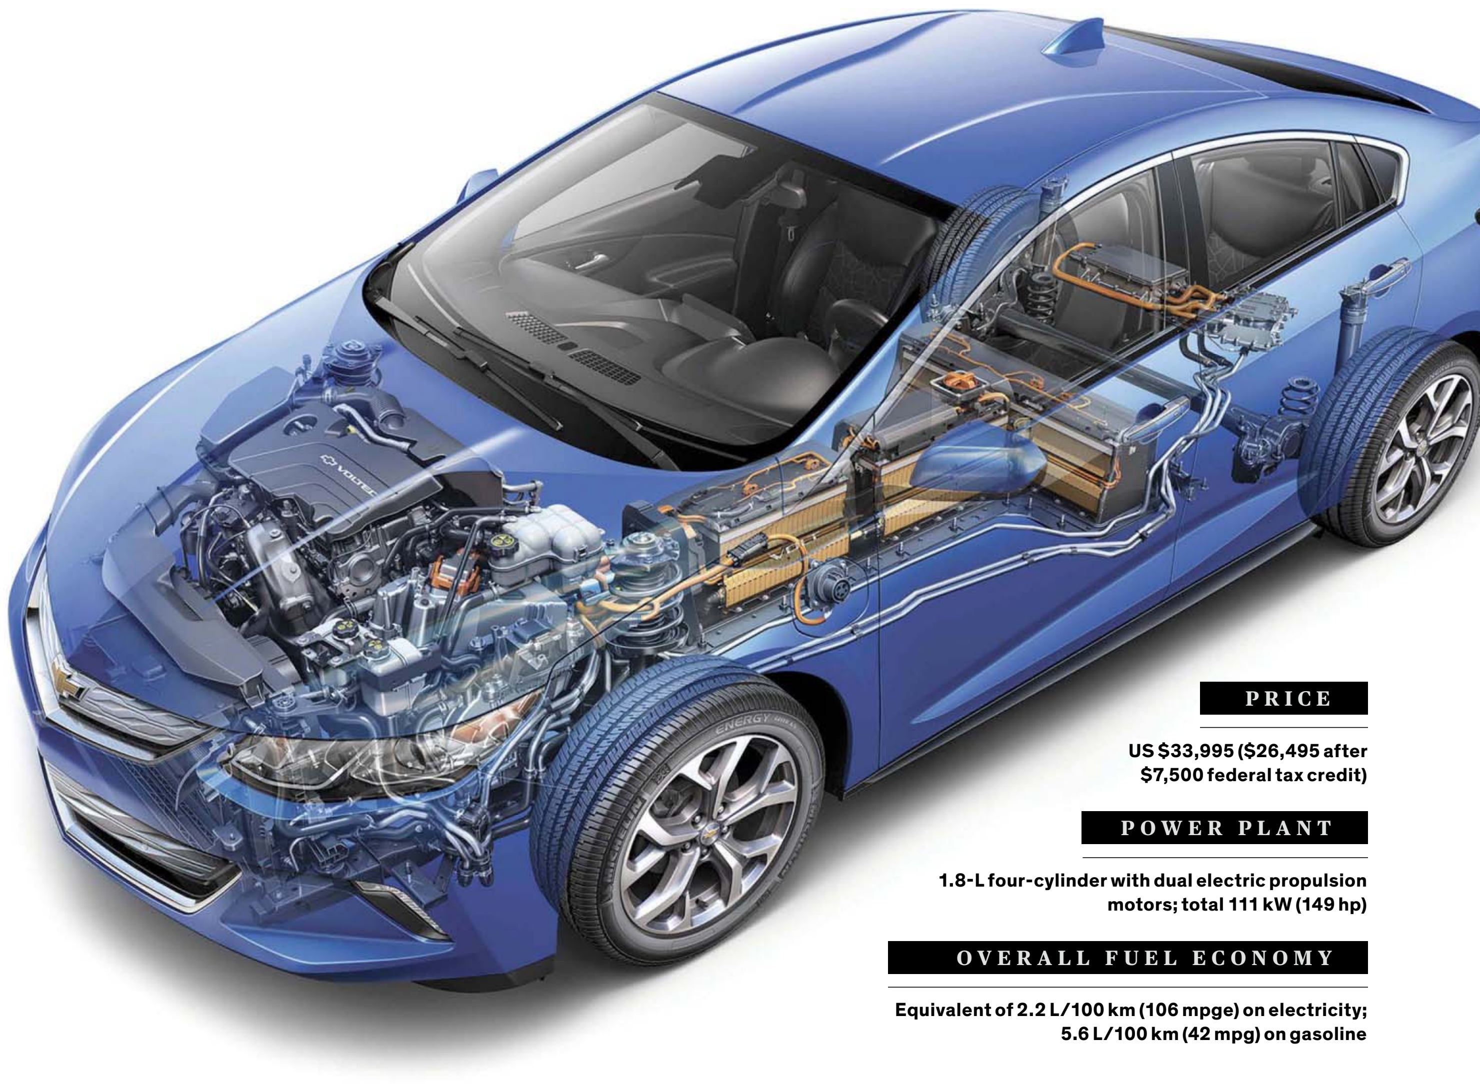

PRICE

US $33,995 ($26,495 after $7,500 federal tax credit)

POWER PLANT

1.8-L four-cylinder with dual electric propulsion motors; total 111 kW (149 hp)

OVERALL FUEL ECONOMY

Equivalent of 2.2 L/100 km (106 mpge) on electricity; 5.6 L/100 km (42 mpg) on gasoline

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

BMW 7 Series
Bells and semiautonomous whistles

The BMW 7 Series is the world's first production car that can operate with no driver aboard, if only to dazzle the neighbors when it eases into your home garage.

It's more than just a gimmick. Press a button on the remote, with its palm-size LCD readout, and watch the BMW drive itself into a garage or back its way out. With its camera and ultrasonic sensors, the BMW can fit into small or stuffed garages that are too tight to allow opening the car's doors. This roboparking feature is offered only in Europe for now, but BMW is pressing U.S. regulators for approval.

The car can thrill when you're behind the wheel, as I discovered on a track test at Monticello Motor Club in New York's Catskills region. As a riposte to the S-Class, Mercedes's technical tour de force, BMW's flagship sedan is faster and sharper handling. It also brings several technical firsts of its own.

One of them is an infrared camera that converts hand signals into action: Twirl a finger in the air to adjust audio volume. Point at the central touch screen to accept an incoming phone call, or wave a haughty hand to dismiss it. Pinch fingers to pan a 360-degree exterior camera view for parking.

And while controlling the many gadgets in some luxury cars can be as enjoyable as a tax audit, the BMW offers a clever back-seat solution. It's a pop-out Samsung tablet, nicely integrated into an aluminum-framed center console. The Wi-Fi–connected tablet manages just about everything, including a Bowers & Wilkins audio system with 1,400 watts and 16 speakers, and a panoramic roof with 15,000 LEDs forming a starlight pattern. While you're star-gazing, choose one of eight onboard perfume scents, including the water-fresh fragrances of the "Blue Suite" and the woodsy notes in the "Green Suite." Or you can ionize the atmosphere—especially nice if you are cruising around in polluted China. On the

PRICE	US $99,395 (750i xDrive version)
POWER PLANT	4.4-L V-8 with dual turbochargers; total 332 kW (445 hp)
OVERALL FUEL ECONOMY	12.4 L/100 km (19 mpg)

road, as opposed to parking, the driving can be semiautonomous. The BMW will steer itself along highway curves at up to 210 kilometers per hour (130 miles per hour), with hands-free operation for up to 15 seconds—at which point the car flashes an alert to get you back into the loop.

BMW, whose electrified i3 and i8 pioneered the carbon-fiber chassis in mass-produced cars, follows up on the 7 Series with what it calls carbon-core construction. The company says it's the first production car to bond carbon fiber with both steel and aluminum. The woven fiber forms the 7 Series' central tunnel, roof pillars, and other components, stiffening the structure and melting away some 135 kilograms (300 pounds) versus the previous edition, getting down to a svelte 1,915 kg (4,225 pounds) for the 740i version.

Crank up to the 750i version and you get a new 4.4-liter V-8 with dual turbochargers that nestle cleverly between the cylinder banks for an especially compact engine and maximum thermal efficiency.

An active version of BMW's famous twin-kidney front grille opens only when cooling air is needed, reducing drag. Active electric four-wheel steering turns the rear wheels opposite the fronts for a remarkably tight low-speed turning radius; then, it steers four wheels in tandem to boost stability and confidence at autobahn speeds. And both robust cornering and comfy cruising are aided by a satellite-based navigation system that adjusts the air suspension, including blissful bump absorption whenever it knows there are no curves ahead.

For those who'd rather save fuel than burn it, a forthcoming plug-in hybrid version will charge up on the home BMW i Wallbox; BMW says it will cover 37 km (23 miles) on electricity alone at up to 120 km/h (75 mph). Slated to become the first four-cylinder luxury flagship sedan ever sold in America, the 740e xDrive gets its urge from an efficient 2.0-liter turbo with a boost from an electric motor integrated into the eight-speed transmission.

All that tech doesn't come cheap. But this 7 Series impresses whether you're sitting in the front seat or in back. Or even watching it park itself. ■

Porsche Mission E Concept
Take that, Tesla

PRICE **Not for sale**

POWER PLANT **Twin electric motors; 447 kW (600 hp)**

OVERALL FUEL ECONOMY **N/A**

When Porsche floats a big idea in sports cars, it doesn't fool around. The 911, for example, has been a design and performance benchmark since 1963. So now comes the Mission E, a knee-weakening electric beauty with 600 all-wheel-drive horsepower, much of it coming from a pair of permanently excited electric motors adapted from its LeMans-winning 919 Hybrid race car. It's just a concept, but Porsche has pledged to get a version in showrooms by 2018. When that happens, figure on a dash to 60 miles per hour in less than 3.5 seconds, a 500-km (311-mile) driving range and a quick 15-minute charge to 80 percent battery capacity via a world's-first 800-volt architecture.

The four-seater is swathed in carbon fiber, aluminum, and steel, and Porsche is targeting a curb weight below 4,400 pounds (1,995 kilograms—about 10 percent lighter than its obvious rival, the Tesla Model S). A five-dial instrument panel nods to the seminal 911. An eye-scanning camera recognizes which screen menus a driver is looking at, then magnifies them via steering-wheel buttons. Those gauges move in parallax to accommodate the driver's seating position. Side mirrors are ditched for cameras that beam the exterior view onto corners of the windshield, perfect for watching lesser cars disappear as the Mission E charges ahead. ■

LEFT: BMW; RIGHT: PORSCHE

Volvo XC90
Safety still comes first

PRICE
US $69,095 with tax credit up to $4,600 (T8 plug-in version)

POWER PLANT
2.0-L four-cylinder with supercharger, turbocharger, and dual electric motors; total 298 kW (400 hp)

OVERALL FUEL ECONOMY
9.8 L/100 km (24 mpg)

Sweden's Volvo had been in virtual hibernation in recent years, cut loose from Ford, sold to a Chinese owner, and resting on past innovations. That changes with the new XC90, a luxury SUV that's hogging such awards as North American Truck of the Year and reaffirming Volvo's reputation for safety.

The XC90 wraps three rows and seven passengers in a Scandinavian body that looks as solid as a glacier. It's the kind of SUV that you might expect to have a V-8, such as the one in the old XC90. But the new one cuts the cylinder count in half: the Volvo Twin Engine combines a turbocharger and a supercharger in a 2.0-liter four to make a very respectable 236 kilowatts (316 horsepower)—and a downright burly 298 kW in the T8 plug-in hybrid version, thanks to a 65-kW electric-motor assist. That T8 hustles to 100 kilometers per hour (62 miles per hour) in just 5.6 seconds, and its E-mode lets it cover up to 40 km on electricity alone.

Volvo's painfully antiquated infotainment systems are replaced by a Tesla-like vertical touch screen that responds to swipes, pinches, and other smartphone-style functions. And the XC90 antes up with two world's firsts in self-driving safety: In addition to staying in its lane, as some other cars do, the Volvo takes evasive steering and braking action if it senses an imminent departure from the road.

And if the Volvo leaves the road anyway, the system electrically cinches front seat belts to secure occupants as much as possible. Energy-absorbing seat material deforms to mitigate spinal injuries from vertical impacts, such as from going airborne and then bottoming out in a ditch. The XC90 also tries to avoid collisions by hitting the brakes if the driver should turn into an oncoming vehicle's path with too little speed or room to make it.

At speeds below 50 km/h, the Pilot Assist function links the camera and steering to keep the XC90 centered in its lane, actively moving the steering wheel to do the job. It's yet another vehicle poised on the brink of autonomous driving, and its priorities are clear: Safety first. ■

LEFT: VOLVO; RIGHT: TESLA MOTORS

With typical hyperbole, Tesla bills its electric Model X as the fastest, safest, and most capable SUV in history. But it's the falcon doors that take this seven-passenger electric car literally over the top. Those motorized portals rise and fold on complex hinges to make the second and third rows effortlessly accessible, even in a packed parking lot. And they incorporate ultrasonic and capacitive sensors to avoid dinging other cars or obstacles.

The crowd-pleasing doors suggest other technical wonders within. The Model X may weigh 2,470 kilograms (5,441 pounds). But with supercar-like power [see specs,

Tesla Model X
Just don't call it a DeLorean

below]—of which 35 percent comes from a front electric motor and the rest from another at the rear—the Tesla will get to 60 miles per hour (97 kilometers per hour) in a claimed 3.8 seconds, or 3.2 seconds in its infamous Ludicrous mode. That's quicker than a Porsche Cayenne Turbo. Yet the Model X can also cover 414 km—from Frankfurt to Brussels, or Detroit to Toronto—on a single charge. And you can get that charge in 30 minutes using Tesla's proprietary Supercharger network.

Tesla figures the Model X's rollover risk is half that of typical large SUVs, because its under-floor battery delivers a limbo-low center of gravity. With no gasoline engine up front, the nose can provide two things: a "frunk" for storage and a giant, energy-absorbing crumple zone.

Like mod diner stools, "monoposto" second-row seats are perched on sturdy motorized pedestals that slide fore and aft and open up bonus room for gear or passengers' feet. Those falcon doors slide open or closed at the touch of a button. A curving panoramic windshield, the largest on any production car, wraps across part of the roof. A medical-grade HEPA filter sucks up pollen, bacteria, viruses, and pollution.

And in the race toward autonomy, Tesla's Auto Pilot one-ups the others as the closest thing yet to a self-driving production car. Wrapped in camera, radar, and sonar sensors, the Tesla is the first car to add lane changes to the semi-autonomous repertoire: Flick the turn signal and the Model X checks for an opening and eases left or right. Officially, drivers are supposed to keep mitts on the wheel, but go ahead and try to resist a hands-free, show-offy demo for passengers. "Look, Ma, no hands!" ■

PRICE	US $132,000 (P90D version)
POWER PLANT	Dual AC electric motors and lithium-ion battery pack; total 568 kW (762 hp)
OVERALL FUEL ECONOMY	89 mpge, the electric equivalent of one gallon of gasoline (or 2.6 L/100 km)

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

Linux at 25

Why this open-source marvel triumphed over its rivals

By Christopher Tozzi

Revolution was brewing on the European periphery in the summer of 1991. Years of struggle by outspoken rebels against the status quo were coming to a boil, and a way of life that once seemed unassailable neared its collapse. ¶ Ask most historians to tell you about that revolution and they'll describe the events that preceded the dissolution of the Soviet Union. As an attempted coup d'état by reactionary hardliners failed and Boris Yeltsin outlawed the Communist Party, it became clear to the world that the radical fervor that began sweeping across Eastern Europe in the late 1980s would soon undo the once-mighty Soviet empire. ¶ Yet, in another corner of Europe, a revolution of a different sort was stirring. No one—not even its chief instigator—recognized its significance. Nonetheless, the code that an irreverent Finnish college student named Linus Torvalds quietly unveiled in August 1991 has ended up touching at least as many lives as did the political upheavals of the late 20th century. I'm talking, of course, about Linux. »»

PHOTOGRAPH BY The Voorhes

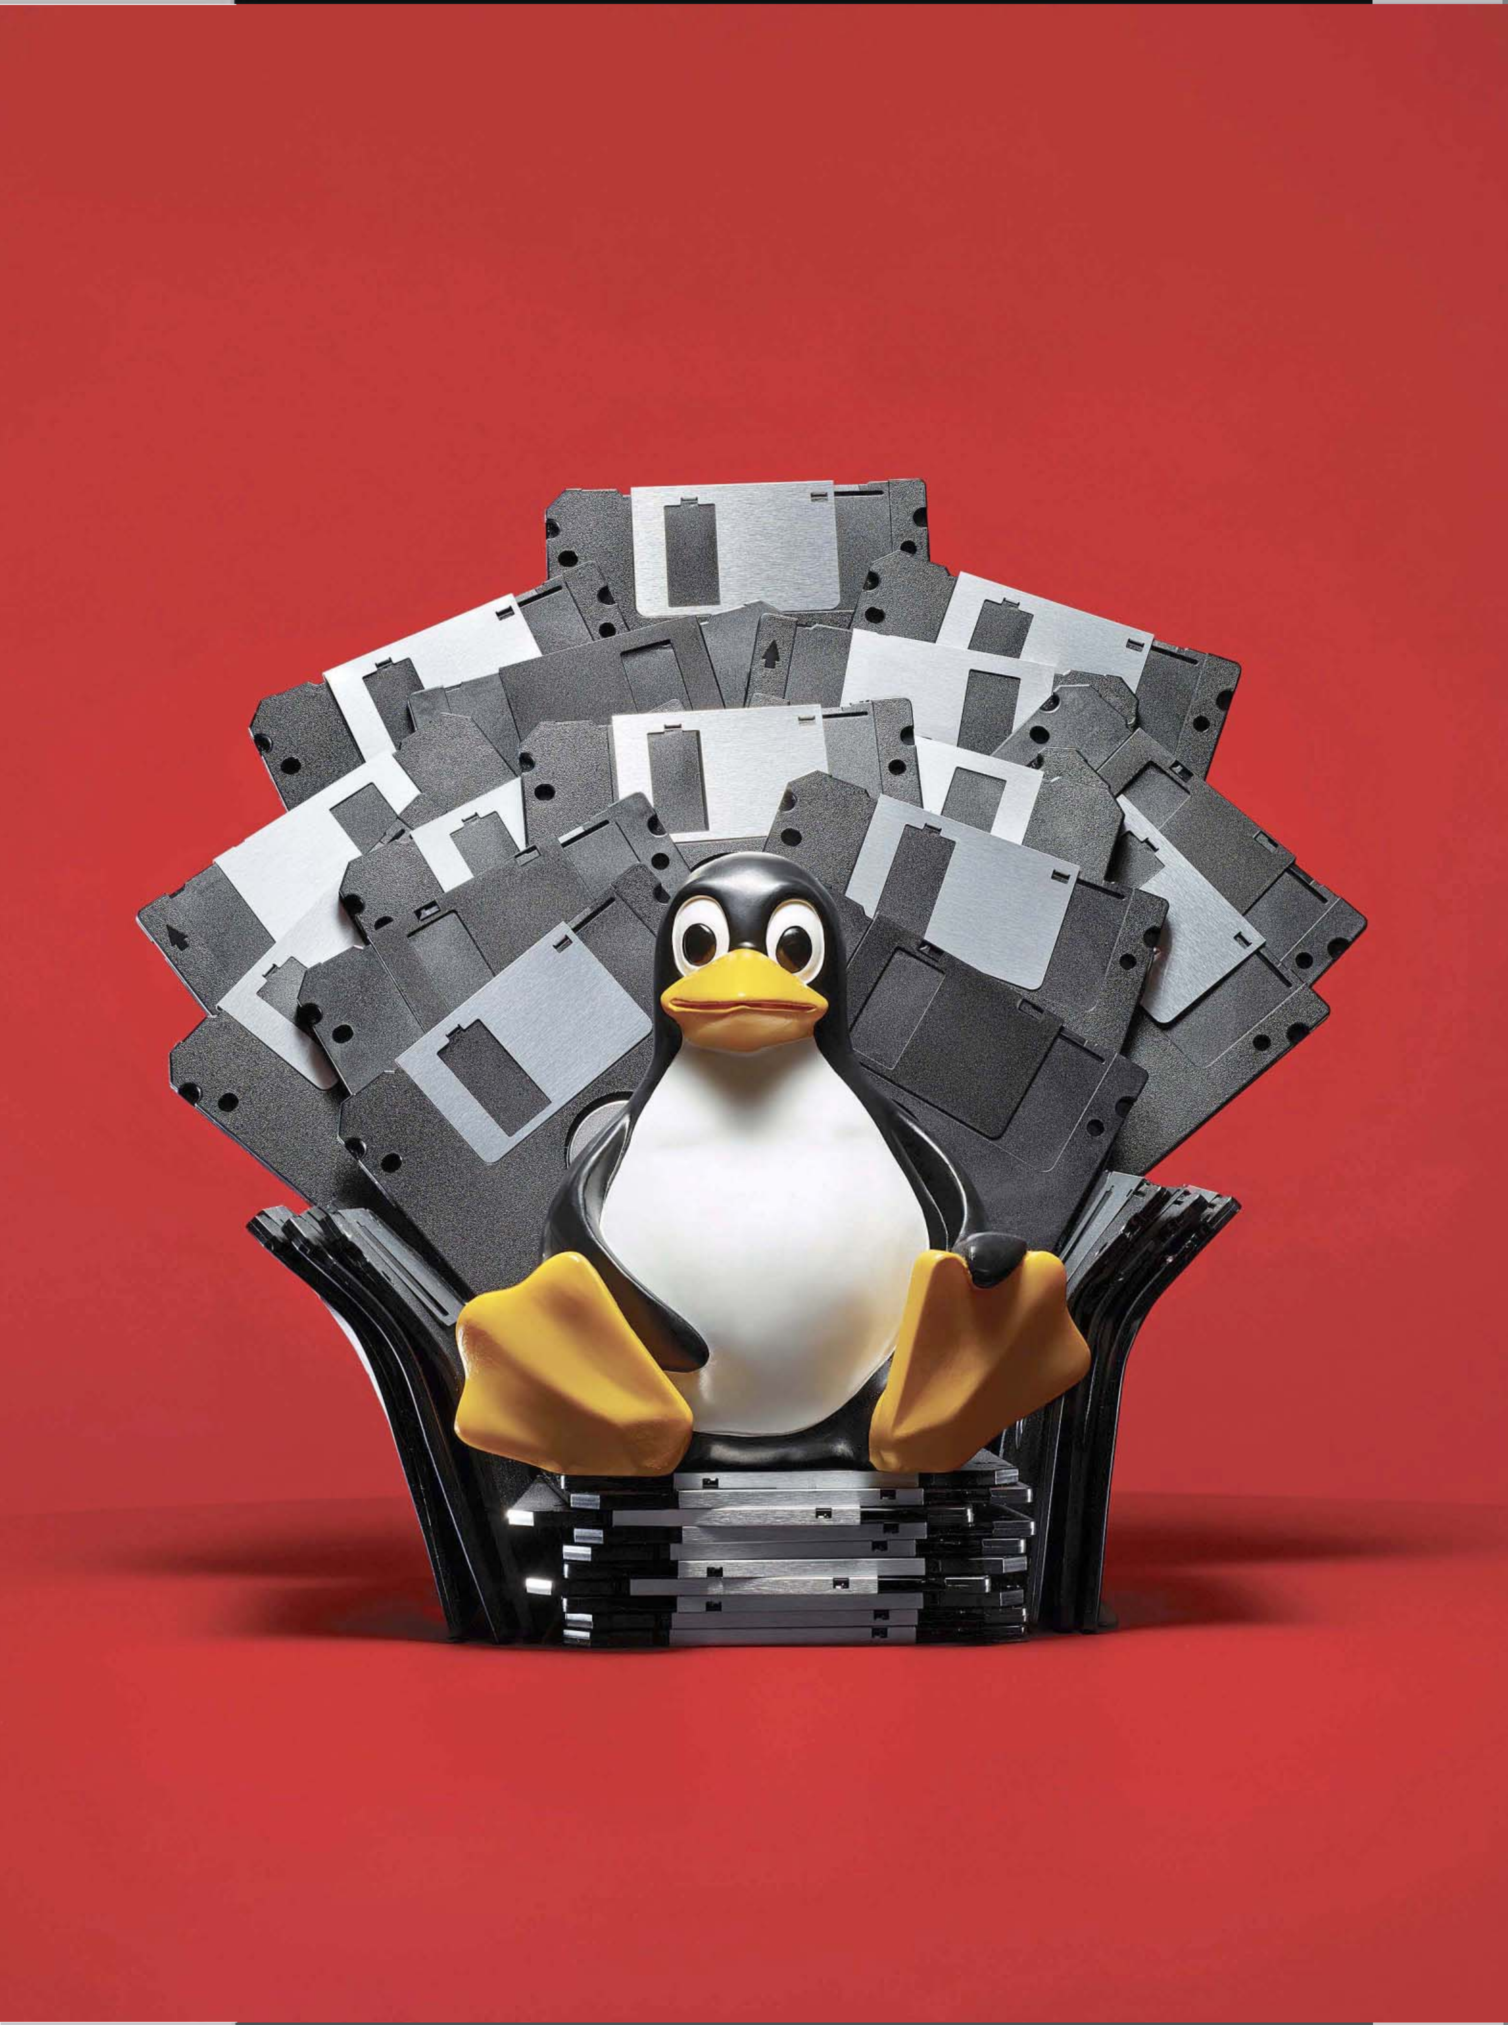

IEEE SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

Torvalds did not plan any of this. He was merely an "accidental revolutionary," as he described himself in his autobiography, *Just for Fun* (2001, HarperBusiness). Almost unwittingly, he kick-started the free-software revolution—a movement that much more prominent programmers had been trying to get off the ground for years.

It's all the more remarkable, then, that Linux, which celebrates its 25th birthday later this year, has so profoundly challenged the norms of software development. It showed programmers everywhere that a different world was possible—a world where they could share code openly, collaborate informally, and make a decent living, even if they gave away the chief product of their labor for free. The advantages of working this way have since become obvious to even the most hard-headed of business leaders, with most large software-development companies now sharing at least some of the fruits of their programmers' efforts openly.

How did Linux end up producing such radical change? And why did other free-software activists' attempts to build bigger and seemingly better systems than Linux fail to achieve as much momentum? With the insight that comes from retrospection, it's now possible to answer those questions.

Of course, in 1991 it would have been ludicrous to suggest that Linux would end up as anything notable. Torvalds had less than two years of college-level work to his name when he started writing the Linux kernel—the code that manages an operating system's core functions—early that year. He worked out of a sparse apartment in Helsinki, where his only computer was an off-the-shelf PC with an Intel 80386 processor, known colloquially as a "386."

Meanwhile, two major teams of professional programmers at some of the world's most elite computer-science research labs were developing other free-software kernels. (The term *open-source software*, more commonly used to describe Linux and similar efforts today, did not come into vogue until 1998.) These teams had each spent years laboring to create a kernel that could do what Linux alone ultimately did.

One team was working on an operating system known as Berkeley Software Distribution, or BSD, which had been in development since the late 1970s at the University of California, Berkeley. BSD was initially conceived as an enhanced version of the Unix operating system, whose code was owned by AT&T. But the BSD project evolved into an endeavor to create a Unix clone that was completely free of AT&T code.

ORIGIN OF SPECIES: Linux, along with many competitors, had its conceptual roots in the Unix operating system, which the late Dennis Ritchie, Keith Thompson, and others developed at AT&T's Bell Labs.

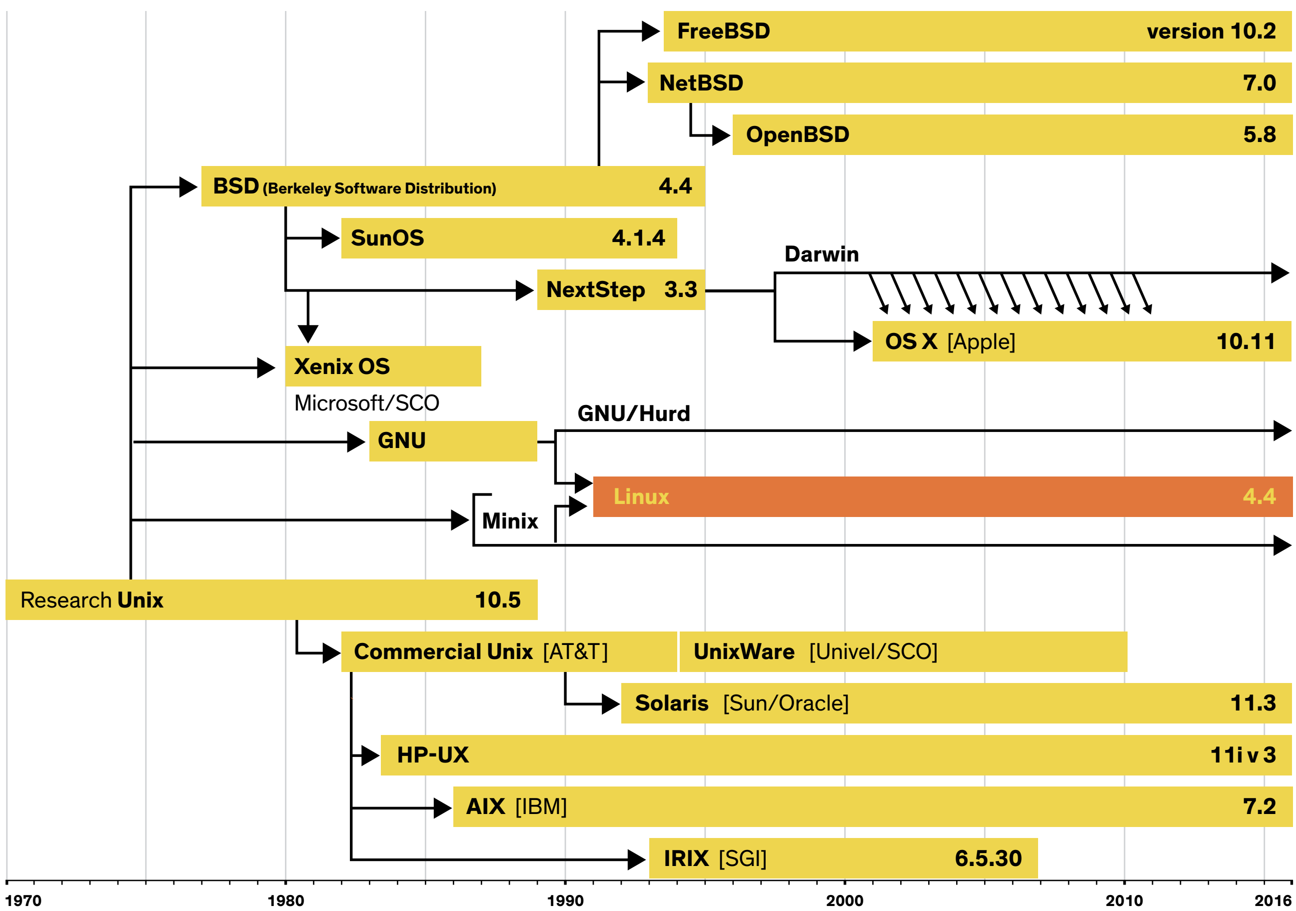

SOURCE: WIKIMEDIA COMMONS

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

LONG LIVE HACKERDOM: Richard Stallman [above, in 1998] spearheaded another effort to develop a Unix-like operating system: GNU/Hurd.

In June 1989, the BSD team rolled out Networking Release 1, or Net/1, its first package of Unix-like software that users could legally deploy without purchasing a Unix license from AT&T. But because Net/1 consisted primarily of networking code, it was hardly a replacement for Unix. Two years later, the developers released a complete operating system. The platform, called Net/2, provided what most people considered the first fully functional free-software operating system.

"Free software" is in the eye of the beholder, however, and not everyone was satisfied with the code from Berkeley. On the other side of the United States, in Cambridge, Mass., another team of programmers led by Richard Stallman was building its own freely redistributable Unix clone, called GNU, a recursive acronym for "GNU's Not Unix!"

Stallman, who viewed the BSD license as problematic because it did not require the source code of derivative works to remain available, started the GNU project in January 1984. Like the BSD developers, Stallman and his team were professional computer scientists. They had access to computer labs at MIT, where Stallman—who in 1990 received a MacArthur "genius" grant for his work on GNU—had been employed prior to launching the GNU project.

By 1991, the GNU team had produced functional components of all the important parts of the operating system—except the one that mattered most: the kernel. They had hoped they might adapt

NATHANIEL WELCH/CORBIS

Linus Torvalds Reflects on 25 Years of Linux

The creator of the open-source operating system discusses its past, present, and future

By Stephen Cass

In this abridgment of an e-mail interview with *IEEE Spectrum*, Linux's creator reflects on the last quarter century and what's to come. The complete interview is available online.

Stephen Cass: What's one thing you know now that you wish your younger self knew?

Linus Torvalds: Actually, I credit the fact that I didn't know what the hell I was setting myself up for for a lot of the success of Linux. If I had known what I know today when I started, I would never have had the chutzpah to start writing my own operating system: You need a certain amount of naïveté to think that you can do it.

I was [also] perhaps more open to outside suggestions and influence than I would have been if I had a very good idea of what I wanted to accomplish. [Developers] could join with their own vision of where things should go. I think that helped motivate lots of people.

S.C.: Is there one early technical decision made during Linux's development that you now wish had gone a different way?

L.T.: The thing about bad technical decisions is that you can always undo them. It's obviously better to always make the right decision every time, [but] I'd rather make a decision that turns out to be wrong later than waffle about possible alternatives for too long.

S.C.: Why do you think Linux never became a significant presence on mainstream desktops?

L.T.: Hey, still working on it. There are multiple reasons, but one of the big ones is simply user inertia. The desktop is...both very personal—you interact with it rather intimately every day if you work with computers—but also complicated in ways many other computing environments aren't.

Look at your smartphone. That's also a fairly intimate piece of computing technology, and one that people get pretty attached to (and one where Linux, thanks to Android, is doing fairly well). The desktop is in many ways more complex, with much more legacy baggage.

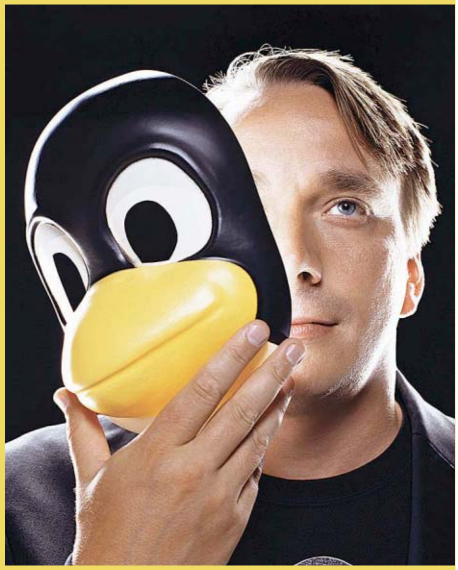

S.C.: What's the biggest challenge currently facing Linux?

L.T.: The kernel is actually doing very well. It can be very challenging to get big and invasive changes accepted, so I wouldn't call it one big happy place, but I think kernel development is working. Many other open-source projects would kill to have the kinds of resources we have.

That said, one continual challenge we've always had in the kernel is the plethora of hardware out there. The good news is that a lot of hardware manufacturers are helping. That didn't used to be true.

S.C.: What current technical trends are you enthusiastic about? Are there any that dismay you?

L.T.: I have always been interested in new core hardware, particularly CPUs. In a bigger picture, it's very interesting to see how AI is finally starting to really happen. I'm not dismayed by the fact that true AI may finally start to happen, like clearly some people are. Not at all.

S.C.: Do you think Linux will still be under active development on its 50th anniversary? What is the dream for what that operating system would look like?

L.T.: I'm not a big visionary. I'm a very plodding pedestrian engineer, and I try to keep my eyes firmly on the ground. I suspect that we've seen many more changes in how computers work in the last 50 years than we're necessarily going to see in the future. [We've] simply learned what works and what does not.

Of course, neural networks, et cetera, will change the world, but part of the point with them is that you don't "program" them. They learn. They are fuzzy. People will want smarter machines, but people will also want machines that do exactly what they're told. So our current style of "old-fashioned" computing won't be going away; it'll just get augmented.

code developed elsewhere to build the GNU kernel, but that strategy didn't pan out. It wasn't until mid-1991—close to the time when Torvalds started writing Linux—that GNU developers finally began the tedious work of developing from scratch their own kernel, which they named "the Hurd."

GNU's Hurd project was well publicized, as was the release of BSD's operating system. Even in far-off Helsinki, Torvalds was well aware of both. And yet he wrote the Linux kernel anyway. Why?

The answer has several parts. The simplest is that Torvalds did it "just for fun," as he explained in his autobiography. To him, Linux was a learning exercise. He was especially keen on using it to familiarize himself with the peculiarities of the 386 PC he had recently acquired.

Torvalds's second motivation stemmed from the deficiencies that he found in the operating system he was running at the time—Minix. He used it because BSD's Net/2 did not run on PC hardware and because GNU was not yet complete. Minix was yet another Unix-like operating system, released in 1987 by Andrew Tanenbaum, a professor of computer science at VU Amsterdam.

Tanenbaum had designed Minix as a teaching tool rather than as a full-blown operating system. It worked well enough for Torvalds, but he felt it lacked many important features. For one thing, it wasn't well suited for PCs with 386 processors. It also offered no terminal emulator, which Torvalds needed to log into his university's Unix system from his home PC.

And so at some point early in 1991—he doesn't remember exactly when it was—Torvalds began writing a terminal emulator for Minix. After completing it, he added disk and file-system drivers so that he could upload and download files from remote computers through the emulator. Those steps brought him closer to producing a complete operating-system kernel. By July, he was asking fellow Minix users on a Usenet bulletin board about documentation for POSIX, the Portable Operating System Interface. This is a technical standard for Unix-like operating systems, one that the IEEE Computer Society had

IAN WHITE/CORBIS

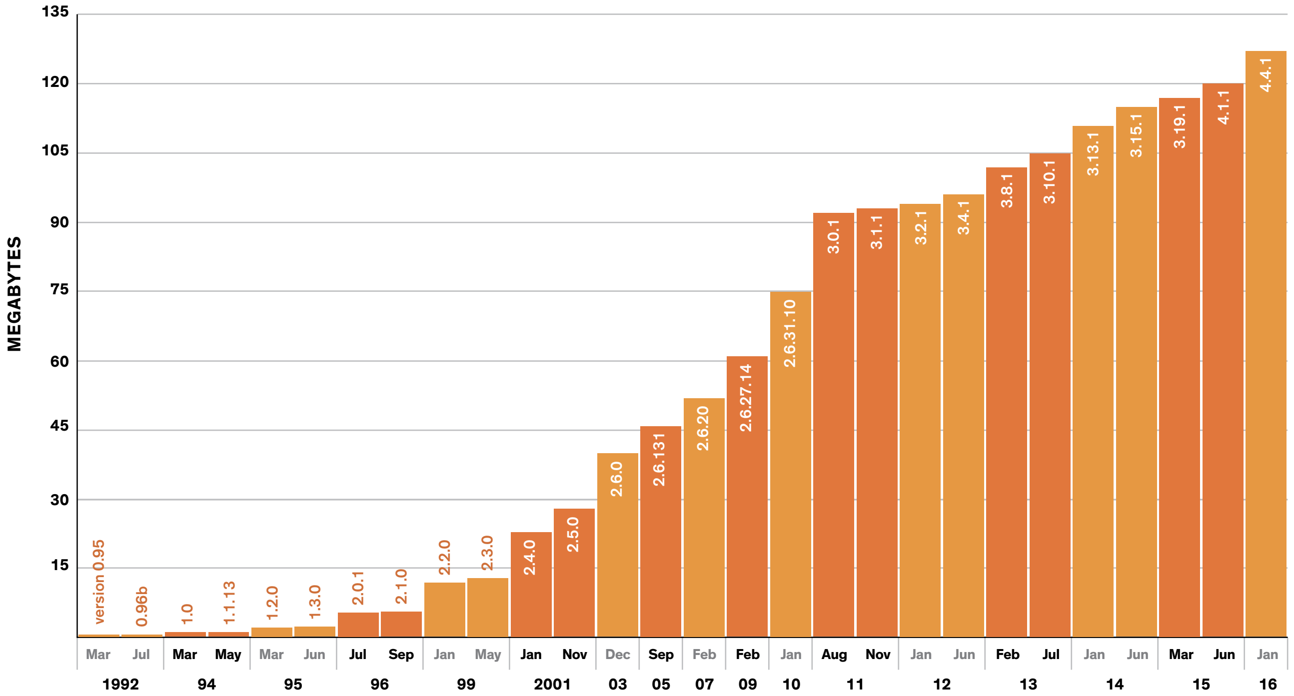

BIGGER AND BETTER: The size of the Linux kernel (shown here in GZIP-compressed file size) has grown enormously over the past quarter century.

SOURCE: HTTP://WWW.KERNEL.ORG/PUB/KERNEL/

formulated several years before. His interest in POSIX was a sure sign that he intended to write a Unix-like kernel.

At the end of August, in the same Usenet group, Torvalds announced his kernel to the world. He warned that it was "just a hobby" and "won't be big and professional like gnu." The following month, a friend posted the first version of the code on an FTP server, taking it upon himself to name it Linux even though Torvalds had dismissed that name for public use, preferring instead to call it Freax. From that server, anyone could download the code, modify it, and send the changes back to Torvalds. And many began doing so.

Of course, Linux took time to mature. Torvalds didn't release Linux 1.0, the first version he deemed of production quality, until 1994. And only gradually did others–some for fun, some in pursuit of profit–begin combining the Linux kernel with other software programs, notably the utilities that the GNU team had produced, to build complete Linux distributions. (Starting in 1994, Stallman and other GNU developers advocated for the name "GNU/Linux" to describe these distributions, but such terminology has enjoyed little following beyond some free-software purists.)

By the mid-1990s, it was clear that Linux was here to stay–and that the prospects for widespread adoption of competing free-software kernels were growing dim.

It would've made more sense if the Hurd, BSD, or even Minix had caught on. But Linux proved far more enduring. Why?

Most observers of the open-source software movement have attributed Linux's success to fortunate timing. Lars Wirzenius, a Finnish programmer who shared an office with Torvalds at the University of Helsinki when they were students, noted at the 1998 Linux Expo that "if the Hurd had been finished a few years ago, Linux probably wouldn't exist today. Or the BSD systems might have taken over the free operating system marketplace."

Tanenbaum made a similar point when he wrote in the early 2000s that legal troubles surrounding BSD in the 1990s "gave Linux the breathing space it needed to catch on." Those troubles began with a lawsuit in 1992 by Unix System Laboratories (USL), which at the time owned the Unix trademark. The company alleged that distributors of BSD-based software had improperly incorporated Unix code into their products. The case was settled out of court the next year, but then the Regents of the University of California countersued, claiming that USL had not properly credited the university for BSD code that was present in Unix.

In June 1993, Novell acquired USL, and the legal drama subsided. By February 1994, the parties had reached a settlement, which required a handful of changes to the BSD code. So in the end, the lawsuits did not limit the ability of BSD developers to distribute their operating system or of users to run it. And indeed, operating systems derived from BSD, such as FreeBSD and OpenBSD, continue to enjoy a healthy following. But BSD's legal troubles did create uncertainty, which slowed adoption at a crucial moment. Unsure whether they could legally use BSD software without paying steep licensing fees to USL, coders shied away just as Linux was evolving into a full-featured alternative.

Meanwhile, the Hurd remained far from complete: It couldn't even boot a computer until 1994, and the first alpha release, version 0.0, appeared only in 1996. By the end of the decade, GNU deemed the Hurd basically functional, and it remains under active development. But a production-quality version has yet to be released. In its absence, Linux became the basis for operating systems that otherwise consist mostly of GNU programs.

But timing alone doesn't fully account for the Linux kernel's explosive popularity. Even before the BSD lawsuits began in 1992, Linux had gained a small following. Moreover, when the legal machinations concluded in early 1994, Linux was no more mature, in technical terms, than BSD, which at any rate could boast support for a broader range of hardware platforms.

Timing was therefore only one factor. Another involved licensing. The Linux code was available under the GNU General Public License, or GPL, which is termed "copyleft"–that is, it requires programmers to make the source code for all derivative works publicly available. In contrast, the BSD licenses allowed developers essentially to do whatever they wished with derivative code, including making it closed-source.

In that sense, BSD came with fewer licensing requirements. But many programmers who worked with Unix-like operating systems didn't see that as a virtue. For them, protecting the core tenets of the hacker culture that had given rise to Unix itself–a culture that valued transparency, openness, and the sharing of code–was most important. The BSD licensing terms, more than the lawsuits, were the main reason these programmers focused their attention on Linux.

It mattered, too, that the Linux code never cost a penny. From the outset, Torvalds was profoundly opposed to the idea that anyone should have to pay for his kernel. Indeed, prior to switching to the GPL in early 1992, he had released Linux under a crude license he wrote himself, which required that users "not profit from the distribution. In fact even 'handling costs' are not acceptable."

Linux's radical lack of a price tag set it apart from most of the other freely licensed Unix-like kernels of the early 1990s. Minix cost US $169 dollars, an "outrageous" sum, in Torvalds's estimation. "Look at who makes money off minix, and who gives linux out for free," Torvalds grumbled in a Usenet post at the time of Linux's birth. "Make minix freely available, and one of my biggest gripes with it will disappear." BSD's Net/2 could be legally copied by users without paying, but an official version on disk from Berkeley cost $1,000. Some of the other BSD derivatives were similarly priced. Even GNU's software had a price–as much as $5,000 dollars for the "deluxe distribution" of GNU software and manuals–if ordered directly from those running the project.

Torvalds's adamant opposition to charging for Linux made his kernel a true outlier. It probably also added to

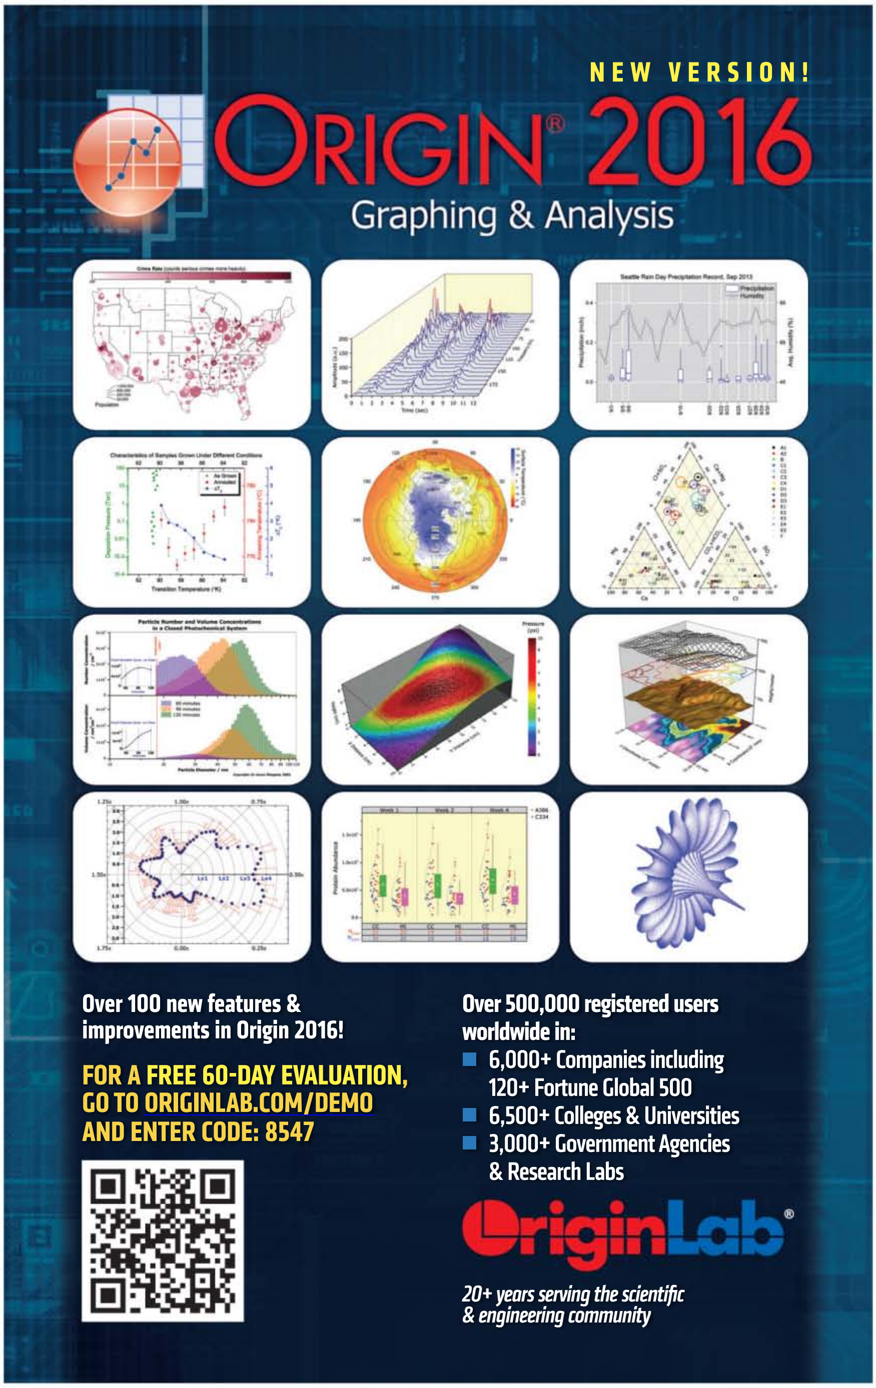

IEEE JobSite
The Right Candidate - Right Now!

Take the next steps to finding your ideal job.

The IEEE Job Site can help you explore the engineering career opportunities that may be right for you.

Take advantage of this unique member benefit.

Create a personal, confidential job profile, post your resumé and start your job search now!

IEEE

Visit the IEEE Job Site at www.ieee.org/jobs

IEEE JobSite
The Right Candidate - Right Now!

How much better would your company be if there were more like you?

Refer your Human Resource colleague to the IEEE Job Site for a free job posting and they'll find the best potential employees in the business – a fellow IEEE member.

Visit the IEEE Job Site today at www.ieee.org/jobs/hrpromo

Find the right candidate - right now!

IEEE

MIT PROFESSIONAL EDUCATION
Advanced Study Program

Your goals, your way

Attend the Advanced Study Program at MIT.

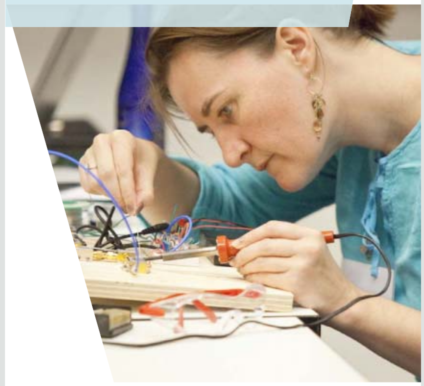

Flexible by design, the non-degree program allows you to **enroll part-time or full-time, choose from more than 2,000 MIT courses** and connect to exceptional MIT students and faculty — enabling you **to meet your educational aspirations.**

Want to learn more? Visit asp.mit.edu/ieee

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

THE INDUSTRY STANDARD SOFTWARE PACKAGES IN GROUNDING AND ELECTROMAGNETIC INTERFERENCE

Are now more powerful with the addition of new standard components such as arbitrary transformers and cables.

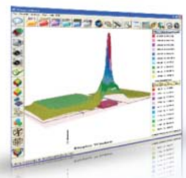

Classical Equipotential Grounding

Multiple grounding systems having any shape, in simple or complex soils, including any number of layers or heterogeneous soil volumes.

- AutoGroundDesign
- AutoGrid Pro
- AutoGround & MultiGround

Complex Frequency Domain Grounding, Interference Analysis & Environmental Impact Assessment

Perform fast, yet complex and accurate, Interference analyses on pipelines, railways, etc. Examine electromagnetic impacts to the environment.

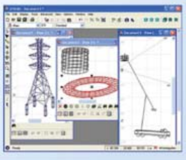

- SESTLC
- Right-Of-Way
- MultiLines & SES-Enviro
- MultiGroundZ

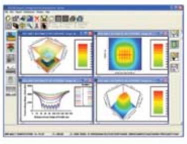

Frequency & Time Domain Arbitrary Network Analysis

Upgrade to the full power of MultiFields or CDEGS to tackle any electromagnetic problem involving aboveground and buried conductor networks including lightning or switching surge analysis.

SES

World Leader in Grounding & EMI
www.sestech.com
800-668-3737 450-622-5000

Toyota Technological Institute has an opening for an assistanat professor position in the Department of Advanced Science and Technology. Applications are encouraged from all relevant areas. For more inforamation, please refer to the website **http://www.toyota-ti.ac.jp/english/employment/assistantprofessor.html**

Position: Assistant Professor (full-time, 5-year contract, non-tenured, non-renewable)

Research field: Materials processing (Physical properties of metals, Metal-base materials, Structural/ Functional materials, Microstructural control engineering, etc.)

Qualifications: A Ph.D. in a relevant field. The successful candidate is expected to demonstrate potential to develop strong and outstanding programs in the above research field. It is also necessary for him or her to supervise students, and to teach advanced and basic courses both at the undergraduate and graduate levels.

Start date: October 2016 or at the earliest convenience

Documents: (1) A curriculum vitae
(2) A list of research activities
(3) Copies of 3 representative papers
(4) A brief summary and future plan of your research and educational statement (within one or two pages each)
(5) Names of two references including phone numbers and e-mail addresses
(6) An application form (available on our website)

Deadline: **June 15, 2016**

Inquiry: Search Committee Chair Professor Yasutake Ohishi
Phone: +81-52-809-1860 E-mail: y-ohishi1@toyota-ti.ac.jp

The above documentation should be sent to:
Mr. Takashi Hirato
Administration Division
Toyota Technological Institute
2-12-1, Hisakata, Tempaku-ku
Nagoya, 468-8511 Japan

Please write "Application for Assistant Professor of Materials Processing Laboratory" in red on the envelope.

Linux's appeal with hackers, who were wary of any activities by software distributors that smacked even remotely of commercialism.

All of these factors fostered the rich developer community that sprang up around Linux. As early as 10 October 1991, only weeks after the Linux code had become publicly available, Torvalds acknowledged that his nascent kernel "never would have seen the light of day or would have been much worse without the help of some others." He went on to name collaborators who were helping him develop Linux via the Internet.

To be sure, Torvalds was not the first programmer to embrace a decentralized, Internet-based community of developers. Keith Bostic, a lead BSD developer, did something similar in 1990, when he enticed hundreds of volunteer programmers from across the Internet to help rewrite Unix utilities in preparation for the release of Net/2. And according to Tanenbaum, Linux followed "essentially the same development model as Minix."

Still, the size and efficiency of Torvalds's following greatly exceeded anything that had come before. Within a few years, Torvalds had built a burgeoning, loosely organized community in which releasing code early and often, and relying on others to spot and fix bugs, became the means to quick improvement. By early 1994, insiders were commenting on the phenomenon. For example, Robert Young, a founder of the open-source software company Red Hat, wrote in the magazine *Linux Journal*, "The number and frequency of new releases of Linux, and drivers and utilities, are amazing to anyone familiar with traditional Unix development cycles." Such rapid innovation far outpaced what most other software projects could manage with their centralized control. It wasn't until 1997, when Eric S. Raymond wrote an essay (later turned into a book) titled "The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary" describing the way Linux developers operated, that a wider group of people began to appreciate the novelty and power of this new approach.

The momentum that Linux established in the early 1990s on the foundation of its fortunate timing, copyleft licensing, and lack of commercial ambitions among its core developers has sustained the kernel for 25 years. As of June 2015, Linux totaled 19.5 million lines of code–up considerably from just over 10,000 in 1991 and about 250,000 at the start of 1995. It is the work of more than 12,000 individual authors, some of whom have contributed just a few lines of code, others vast amounts. On average, 7.71 updates make their way into the kernel code every hour.

Linux is now used not just in many of the machines you'd recognize as a computer but also for embedded applications, meaning you'll find Linux in things like your wireless router, your e-reader, and your smart thermostat. A 2008 study estimated the kernel's total worth in monetary terms–difficult as that is to quantify for something often available for free–to be $1.4 billion. By now it must be several times this figure.

That's not bad for something a nerdy 22-year-old cooked up for kicks at the end of the Cold War. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/linux0416

IEEE.tv gets a mobile makeover

Bring an award-winning network of technology programs with you.

Available on the App Store

Available in Android Market

Go mobile or get the app.
www.ieee.tv

IEEE

香港科技大學
THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

25 INNOVATING TODAY IMAGINING TOMORROW 啟・創・未來 HKUST

Joint Faculty position
Department of Electronic and Computer Engineering (ECE) and Department of Mechanical and Aerospace Engineering (MAE)

The Department of Electronic and Computer Engineering (ECE) and **the Department of Mechanical and Aerospace Engineering** (MAE) at the Hong Kong University of Science and Technology (HKUST) invite applications for a joint faculty position at the rank of Assistant Professor in the interdisciplinary area of aerospace engineering. Applicants should have a PhD with demonstrated strength in research and a strong commitment to teaching. Successful candidates are expected to lead an active research program and promote synergy between ECE and MAE in the aerospace program. Areas of research and interests may include: Reliable Systems Design, Sensors System Design, Data Integration and Analysis, Green Aviation and UAV with focus in Aerospace applications. Other areas which can promote synergy between ECE and MAE may also be considered.

HKUST is a world renowned, international research university in Asia's most vibrant city, Hong Kong. Its Engineering School has been consistently ranked among the world's top 25 since 2004. The high quality of our faculty, students and facilities provides for outstanding opportunities for faculty to pursue highly visible research programs. All formal instruction is given in English and all faculty members are expected to conduct research and teach both undergraduate and graduate courses. The Departments of ECE and MAE have excellent computing resources, and state-of-the-art teaching and research laboratories. Currently, the Department of ECE has about 40 faculty members, 600 undergraduate students and 500 postgraduate students. The Department of MAE has 30 faculty members, about 400 undergraduate students and 350 postgraduate students. The University is committed to increasing the diversity of its faculty and has a range of family-friendly policies in place.

Starting salary will depend on qualifications and experience. Fringe benefits including medical and dental benefits, annual leave, and housing will be provided where applicable. Initial appointment will normally be on a three-year contract. A gratuity will be payable upon successful completion of contract. Re-appointment will be subject to mutual agreement.

Applications including full curriculum vitae, list of publications, names of five referees addressed to Professor Vincent Lau, Chair of the Search Committee, should be sent by email to eesearch@ust.hk. Applications will be considered until the position is filled.

More information about the departments is available on the websites http://www.ece.ust.hk and http://www.mae.ust.hk.

(Information provided by applicants will only be used for recruitment and other employment-related purpose.)

上海科技大学
ShanghaiTech University

TENURE-TRACK AND TENURED POSITIONS IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

The newly launched ShanghaiTech University is built as a world-class research university, which locates in Zhangjiang High-Tech Park. We invite highly qualified candidates to fill tenure-track/tenured faculty positions as its core team in the School of Information Science and Technology (SIST). Candidates should have exceptional academic records or demonstrate strong potential in cutting-edge research areas of information science and technology. They must be fluent in English. Overseas academic connection or background is highly desired.

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

Compensation and Benefits:
Salary and startup funds are highly competitive, commensurate with experience and academic accomplishment. We also offer a comprehensive benefit package to employees and eligible dependents, including housing benefits. All regular ShanghaiTech faculty members will be within its new tenure-track system commensurate with international practice for performance evaluation and promotion.

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

Applications:
Submit (in English, PDF version) a cover letter, a 2-page research plan, a CV plus copies of 3 most significant publications, and names of three referees to: sist@shanghaitech.edu.cn (until positions are filled). For more information, please visit Job Opportunities on **http://sist.shanghaitech.edu.cn/**

Deadline: April 30th. 2016

Faculty Positions in Robotics and Mechatronics

NAZARBAYEV UNIVERSITY

Nazarbayev University is seeking highly qualified full-time faculty members at the Assistant, Associate and Full Professor levels 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 expertise and research interests in the areas of biorobotics, unmanned aerial vehicles, motion planning, and mechanical systems design. Exceptional candidates with research interests in the broader field of 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 May 1st, 2016. Successful appointments are expected to begin on August 1st, 2016. For more information please visit **http://sst.nu.edu.kz**.

PAST FORWARD_BY EVAN ACKERMAN

THE BALLISTIC BROILER

In the 1950s, in a top-secret underground bunker somewhere in Virginia, aerospace researchers developed the MegaVeggie Roast-O-Matic 9000 to flash-broil enormous turnips. After a successful (and surprisingly tasty) test program, they transitioned from oversized root vegetables to missile nose cones, using the apparatus to simulate the intense temperature of atmospheric reentry from space. This research proved much more practical, albeit less nutritious. ■

DON CRAVENS/THE LIFE IMAGES COLLECTION/GETTY IMAGES

"Delightful" - Wired "Robot heaven" - Mashable

Welcome to the world of

ROBOTS

For iPad

Get the app now:

robotsapp.net

Download on the App Store


MathWorks®
Accelerating the pace of engineering and science
Satellites that tango,
using automatically generated code.
That's Model-Based Design.
To achieve a breakthrough in two satellites performing autonomous formation flying and rendezvous, engineers at OHB Sweden automatically generated verification tests and the final flight code from system models. Discover Model-Based Design with MATLAB and Simulink at mathworks.com/mbd
MATLAB® & SIMULINK®
©2012 The MathWorks, Inc.
Image of satellite Tango transmitted by Mango. © OHB Sweden