IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 11.15

IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 11.15

ANOTHER GREEN WORLD

Superadvanced space telescopes could discover a life-supporting exoplanet within five years P. 26

IEEE

IEEE SPECTRUM

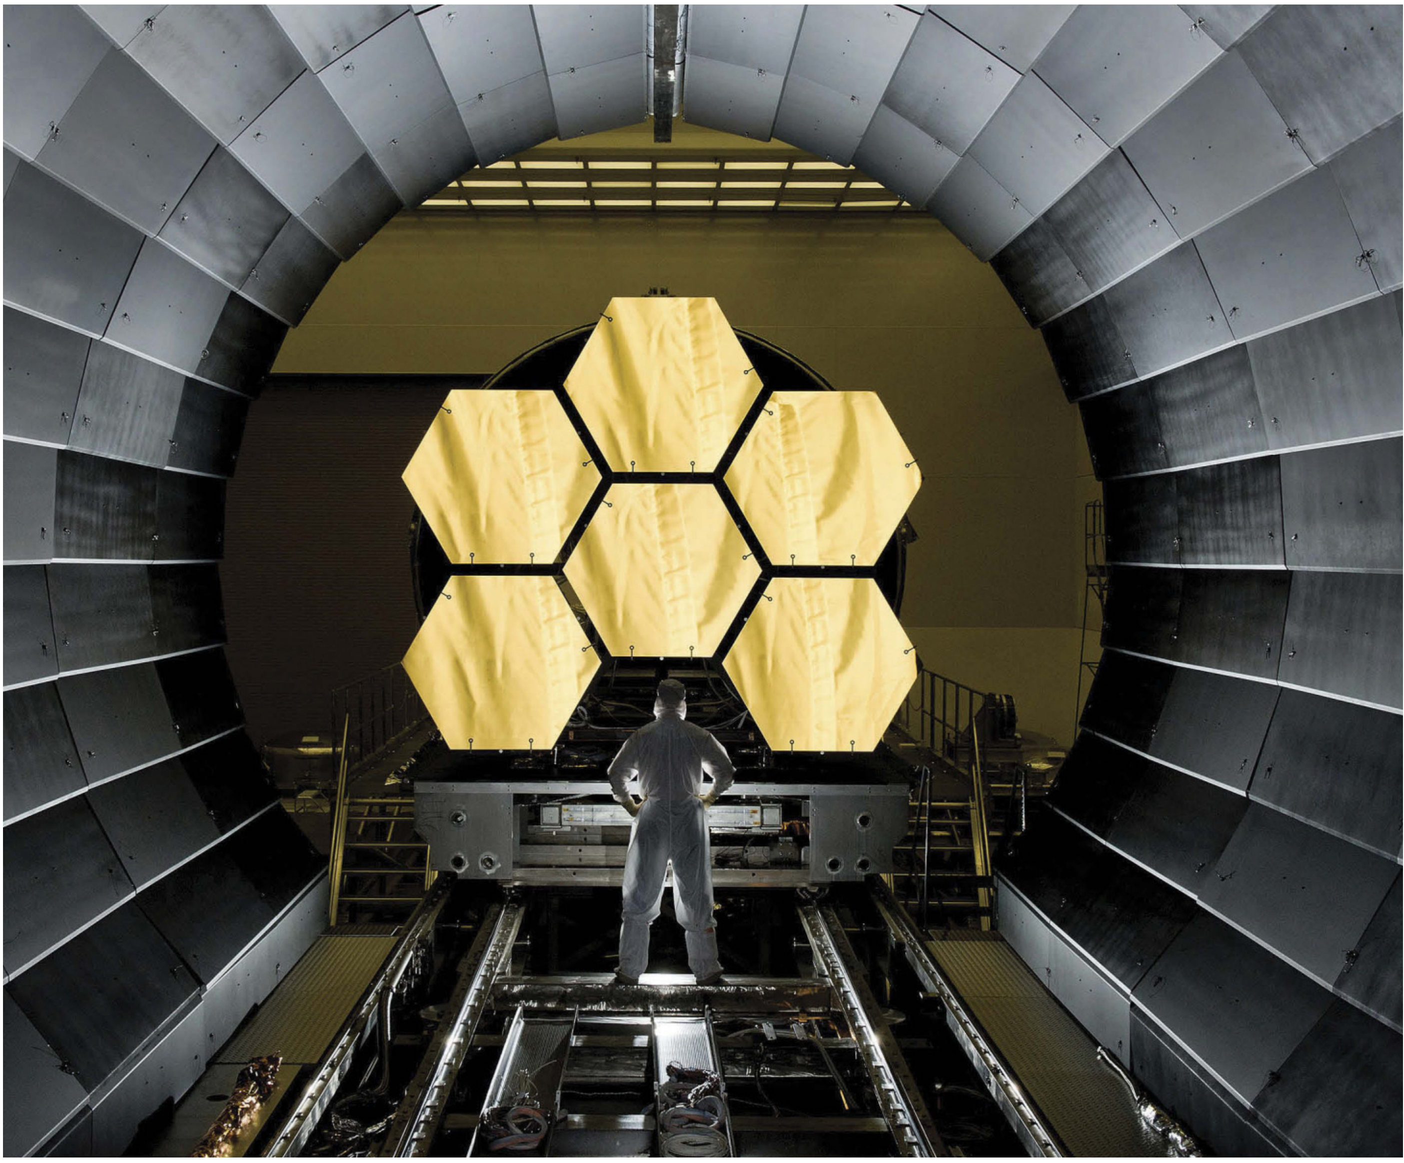

DAVID HIGGINBOTHAM/MSFC/NASA

On the Cover **Illustration for *IEEE Spectrum* by Bryan Christie Design**

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

DEPARTMENTS_11.15

IEEE SPECTRUM

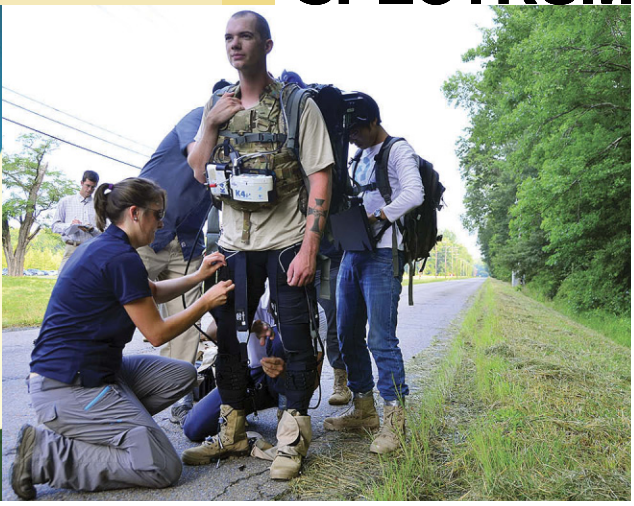

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Solve Your Industrial Networking Challenges With HMS's Anybus IP on Xilinx All Programmable Devices—3 November
- Simulating Graphene-Enhanced Devices—5 November
- Advanced PCB Rule Checking for Signal Integrity and EMC—12 November
- Multipaction & Breakdown Analysis of RF Components—19 November

- **STRATASYS 3-D PRINTER SWEEPSTAKES**
  http://spectrum.ieee.org/static/3d-printer-sweepstakes
- **MASTER BOND RESOURCE LIBRARY**
  http://spectrum.ieee.org/static/master-bond-resource-library
- **WHITE PAPERS**
  http://spectrum.ieee.org/whitepapers

The Institute

Available at theinstitute.ieee.org

- **CLOSE ENCOUNTER** After a decade-long journey through the solar system, NASA's New Horizons probe recently made its closest approach to Pluto, about 12,500 kilometers above the planet's surface. IEEE Member Ron Schulze and Senior Member Willie Theunissen helped design the probe's high-gain antenna dish, which is currently gathering scientific information.

- **PRESERVING THE INFORMATION AGE** This month the Computer History Museum, in Mountain View, Calif., will receive an IEEE Milestone Special Citation. It houses the world's largest collection of computer memorabilia, including the Cray-1 supercomputer and an Apple I.

- **THE INTERNET OF ME** The next generation of consumer connectivity will be the central topic of the IEEE International Conference on Consumer Electronics, to be held from 9 to 11 January in Las Vegas.

FROM LEFT: THOMAS TRUTSCHEL/GETTY IMAGES; TAVIS COBURN; EVAN ACKERMAN

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2015 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 52, issue No. 11, International edition. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

BACK STORY_

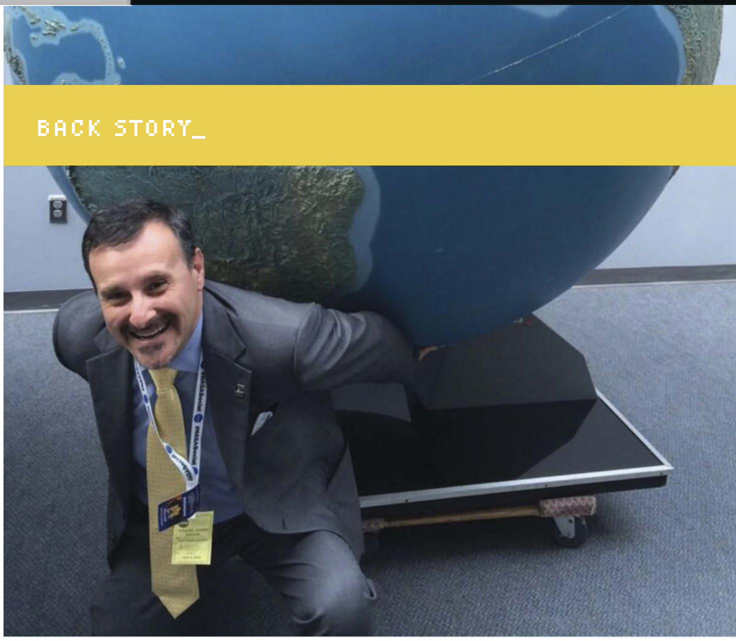

BARBIE BUCKNER

Planetary Pursuits

THE HUBBLE SPACE TELESCOPE'S long-awaited successor, NASA's James Webb Space Telescope (JWST) is set to launch in just three years. The 6.5-metric-ton scope will be able to pick up infrared emissions from some of the universe's earliest stars and galaxies. And, perhaps most alluring, it could help us study other worlds in our Milky Way that might be capable of supporting life, as astronomers Alberto Conti and Mark Clampin explain in this issue.

Conti [above], who began working on JWST at the Space Telescope Science Institute in Baltimore, now works for the telescope's main contractor, Northrop Grumman Corp. In both positions, he's made education and outreach a priority. In 2013, he was part of a team that organized three days of outreach for the South by Southwest conference in Austin, Texas. The JWST event featured more than 60 talks, a 6-meter-wide science visualization screen provided by Microsoft, and the opportunity to Skype with engineers in the telescope's clean room at NASA's Goddard Space Flight Center in Greenbelt, Md.

Rising above it all was a full-scale model of JWST, a telescope that boasts a sunshield nearly as big as a tennis court. Even though the model weighed more than 5 metric tons, Conti says, "it had to be bolted to the ground for wind." All told, some 15,000 people attended the event. "I usually work a lot, but that was a lot of work," he says. "I didn't sleep for those three days. It was crazy."

Conti, pictured here at the end of a JWST social-media event at NASA Goddard, says he's not above goofing off a bit, if it helps get the word out. "I feel like we have a responsibility to let the general public know that these [kinds of] missions have a tremendous value for humankind," he says. "I think we should do as much as we can. So if being silly helps, I'll be silly every day." ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, And Now This... is in *IEEE Spectrum*, Vol. 52, no. 11 (INT), November 2015, p. 56, or in *IEEE Spectrum*, Vol. 52, no. 11 (NA), November 2015, p. 76.

11.15

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

IEEE BOARD OF DIRECTORS
PRESIDENT & CEO Howard E. Michel, president@ieee.org
+1 732 562 3928 FAX: +1 732 465 6444

PRESIDENT-ELECT Barry L. Shoop
TREASURER Jerry L. Hudgins **SECRETARY** Parviz Famouri
PAST PRESIDENT J. Roberto B. de Marca

VICE PRESIDENTS
Saurabh Sinha, Educational Activities; Sheila Hemami, Publication Services & Products; Wai-Choong "Lawrence" Wong, Member & Geographic Activities; Bruce P. Kraemer, President, Standards Association; Vincenzo Piuri, Technical Activities; James A. Jefferies, President, IEEE-USA

DIVISION DIRECTORS
Ellen J. Yoffa (I); Hirofumi Akagi (II); Harvey A. Freeman (III); William W. Moses (IV); Susan K. Land (V); Rob Reilly (VI); Wanda K. Reder (VII); John W. Walz (VIII); Marina Ruggieri (IX); Kazuhiro Kosuge (X)

REGION DIRECTORS
Vincent P. Socci (1); Timothy P. Kurzweg (2); Mary Ellen Randall (3); Robert C. Parro (4); J. Derald Morgan (5); Thomas Coughlin (6); Amir G. Aghdam (7); Costas M. Stasopoulos (8); Norberto M. Lerendegui (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 +1 732 562 3998
CORPORATE ACTIVITIES Elena Gerstmann
+1 732 981 3452, e.gerstmann@ieee.org
EDUCATIONAL ACTIVITIES Douglas Gorham
+1 732 562 5483, d.g.gorham@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
HUMAN RESOURCES Shannon Johnston, SPHR
+1 732 562 6343, s.johnston@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Eileen M. Lach, +1 212 705 8990, e.m.lach@ieee.org
CHIEF MARKETING OFFICER Patrick D. Mahoney
+1 732 562 5596, p.mahoney@ieee.org
CHIEF FINANCIAL OFFICER Thomas R. Siegert
+1 732 562 6843, t.siegert@ieee.org
TECHNICAL ACTIVITIES Mary Ward-Callan
+1 732 562 3850, m.ward-callan@ieee.org
MANAGING DIRECTOR, IEEE-USA Chris Brantley
+1 202 530 8349, c.brantley@ieee.org
IEEE PUBLICATION SERVICES & PRODUCTS BOARD
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

PHOTO-ILLUSTRATIONS BY **Gluekit**

11.15

CONTRIBUTORS_

Tam Harbert

A contributing editor at *IEEE Spectrum*, Harbert has long been intrigued by the incessant battles over intellectual property. For our May 2013 issue, the Washington, D.C.–based journalist explored one such struggle by following a single patent through its twisted, troubled life. When Google announced plans to begin experimenting with ways to keep patents out of the hands of "trolls" [p. 12], it was another chance to look at tech IP's "constantly evolving business model," Harbert says.

Sue Ellen Haupt

Haupt leads the Weather Systems and Assessment Program at the National Center for Atmospheric Research, in Boulder, Colo. She and NCAR's William P. Mahoney describe new tools to help utilities predict how much power they can harvest from the wind [p. 46]. Haupt thinks energy traders will also find the tools useful. Early in her career, she did day-ahead energy trading for an electric utility, when wind power was a novelty. "We had one turbine in the parking lot," Haupt recalls, "but it produced very little power."

Andrew "bunnie" Huang

In our April issue, Huang wrote about how the slowing of Moore's Law could be a good thing. In this issue, he and coauthor Sean Cross give a supporting example: the Novena laptop [p. 38]. The Novena uses open-source hardware and software, which means developers can understand its workings. And because chips are advancing at a slower pace these days, Huang says, "garage innovators can take a couple of years to develop a chip-based product without finding that the completed version is obsolete."

Alexander B. Magoun

Magoun is an outreach historian at the IEEE History Center in Hoboken, N.J. In this issue, he writes about vacuum tubes that were developed at RCA in the 1930s and remain popular today [p. 56]. From 2000 to 2009, Magoun was head of the David Sarnoff Library in Princeton, N.J. The enduring appeal of these vacuum tubes doesn't surprise him: "As RCA's David Sarnoff once said, so long as an old technology has something useful and unique to offer, it will always enjoy a market."

John Ousterhout

Ousterhout is probably best known for creating the scripting language Tcl (Tool Command Language) while at the University of California, Berkeley. He then spent 14 years in industry. Returning to academia in 2008 as a Stanford professor, he realized that storage systems were not keeping up with networking. His response was a project called RAMCloud, which he describes in "The Volatile Future of Storage" [p. 32]. "I'm weird among academics," says Ousterhout. "I like to build systems that really work."

Technological Progress and the Perpetual Learning Curve

As technology evolves at warp speed, so must engineers

I TALK TO MY IPHONE 6 PLUS TO COMPOSE TEXT MESSAGES, and I'm not alone. Many people are abandoning typing, which can be hard on the hands. Why let your fingers take a pounding when you can get what you want by speaking?

Thanks to steady improvements in voice recognition, orality—vocal expression—is unexpectedly reviving. But we aren't speaking more with other people; telephony between humans is collapsing. Instead, we speak to our digital devices as never before.

The shift is part of a broader transformation in everyday life. Our artifacts are in upheaval. Compact discs are being creatively destroyed by streaming music. Tickets to sports and other entertainment events are now images on a smartphone screen. Keys and ID cards are vanishing, replaced by biometrics.

Even learning to drive an automobile is on the endangered list. Neither of my two 20-something children has a driver's license. My son says he may never get one. "Between Uber and driverless cars," he says, "I won't ever learn to drive."

Typing and driving—both courses I took in high school in the 1970s—are no longer part of the skill set that defines adulthood. For engineers who design and build the platforms upon which human life depends, rapid and radical changes demand a response. Engineers, to stay relevant and employable, must change along with the rest of us.

Consider the shifts in coding, a lucrative field that, despite high pay and status, seems addicted to change. Take the size of coding teams. When I wrote a book 20 years ago about the making of Windows NT, the 250-person team that wrote the massive program at Microsoft seemed to foretell the size and shape of software work to come. Success as a programmer in megateams depended on finding a niche, on hyperspecializing and fitting in. Today smaller, agile teams of as few as 15 coders instead value versatility.

Solo coding, meanwhile, was once the stuff of legends. Gates, Brin, and Zuckerberg all wrote initial versions of popular programs. Today, the lone-wolf coder, fueled by a me-against-the-world sensibility, is an endangered species.

Pivotal, a large software house in San Francisco, has halted solo programming, opting instead for two-person teams that enable each programmer to monitor the other. The buddy system seems to boost reliability and reduce the tendency of coding "divas" to hold projects hostage.

Programming languages also go in and out of style and shape career trajectories. Some languages rise and fall on the fortunes of particular hardware platforms, but whatever the cause of the fluctuation in popularity, it's vital to engineers to be able to see the trends. It's no surprise that *IEEE Spectrum*'s Top Programming Languages app has tens of thousands of views each month. Coders want to know what languages are hot.

Development paradigms also shift. Consider the DevOps paradigm, which seeks to overcome barriers between development and operations. Since its debut in 2009, DevOps has rapidly gained followers.

The imperative to keep pace with change unites engineers with other humans. Change isn't always a journey into the unknown; at times, the experience returns us to the familiar, to older, even dormant practices. By enabling our new digital tools to harvest best practices from the past, we all benefit.

Whatever the speed, we can still go home again—at least until the next adaptation is impossible to resist. —G. PASCAL ZACHARY

G. Pascal Zachary, a professor of practice at Arizona State University, is the author of *Showstopper!* (1994), on the making of Windows NT.

IAN SPANIER/GETTY IMAGES

NEWS

456

BILLIONS OF BITCOIN HASH FUNCTIONS PER SECOND ON 7 OCTOBER 2015

THOMAS TRUTSCHEL/GETTY IMAGES

BITCOIN NEEDS TO GET ITS ACT TOGETHER

By early next year, key players must agree on a single solution or the currency could split in two

By and large, Bitcoiners are no strangers to adversity. In fact, they rather seem to welcome it. The community is a battery of contrarianism that lives off the loathing of every institution it seeks to usurp–banks, governments, payment processors. Like the psychic slime in *Ghostbusters II*, the community gets stronger with every hateful blow.

But what happens when dissent bubbles up from within Bitcoin? This is what we are now witnessing. In September, a pair of Bitcoin programmers, Mike Hearn and Gavin Andresen, splintered off from the central group of developers and released a competing version of the Bitcoin software, called BitcoinXT. It includes a controversial rule change that the pair argue will alleviate Bitcoin's snowballing scaling problem: At up to 1 megabyte apiece, the size of the components, or blocks, that make up the complete record of Bitcoin transactions (the blockchain) is widely agreed to be too small for the currency's future. But others warn that BitcoinXT's solution will place the currency on a path toward centralized control.

The greatest threat to Bitcoin's stability, however, is not the possibility that BitcoinXT will be adopted but rather that it will be only partially adopted. The new version is programmed to go into effect on 16 January 2016 if 75 percent of the main processors of Bitcoin transactions, called miners, signal their consent. (The percentage is calculated on the basis of computing »

Who Rules Bitcoin and What Powers They Have

KEY PLAYERS	ROLE	WHAT THEY CAN DO
Developers of the core Bitcoin client Maintainer: Wladimir van der Laan Key contributors: Gavin Andresen, Matt Corallo, Corey Fields, Jeff Garzik, Luke-Jr, Gregory Maxwell, Peter Todd, Pieter Wuille	Developers contribute new code to Bitcoin Core.	**Veto proposals** before they are incorporated into the source code, forcing dissenters, such as BitcoinXT's supporters, to publish a competing version of the source code.
Large miners and mining pool operators AntPool, BitFury, BTCChina, F2Pool, KnCMiner, Slush	At minimum, miners verify transactions, bundle them into new blocks, and broadcast them to other miners and full nodes. Responsible miners also validate new transactions. Miners are rewarded for their effort with bitcoins and thus also act as the mint for the currency.	**Vote** on whether to adopt BitcoinXT. It will be activated only **if 75 percent of miners signal consent** to the upgrade. Even then, miners will still be free to decide which of the two currencies they will mine.
Users and wallet providers Armory, Bitcoin-QT, BitGo, Blockchain, Electrum, MultiBit HD	Users are the people who create new Bitcoin transaction requests. They use software called wallets to relay requests to the Bitcoin peer-to-peer network. The software also provides users with their current balances.	**Choose one version** of the currency over the other. The more wallets compatible with BitcoinXT, experts think, the greater the value of that version of the currency.
Payment processors BitPay, Circle, Coinbase, GoCoin, Stripe	Payment processors offer software that acts as surrogate wallets for merchants that want to offer Bitcoin as a payment option but do not want to be exposed to the currency's volatility. They can conduct Bitcoin transactions for their clients and then pay the clients in nondigital currencies.	**Refuse to process payments** from one or another of the competing blockchains.
Full nodes Any user or miner may also operate a full node.	Full nodes are the traffic checkpoints on the Bitcoin network. Each node carries a complete version of the Bitcoin blockchain. When new blocks are added, full nodes refer to past transactions in order to verify that no bitcoins are being spent twice and that no fundamental rules have been broken.	**Cause confusion** if they refuse to upgrade. In particular, a situation could arise where wallets that do not independently verify new blocks (and therefore blindly trust other full nodes) initially accept transactions that they will later reject.
Exchanges Bitstamp, BTCChina, Coinbase, Coinsetter, Cryptsy, Kraken	Online exchanges facilitate trades between Bitcoin and other currencies.	**Choose which of the two currencies** they will list, effectively giving them control over the availability of each.

power.) However, there are many other agents in the Bitcoin peer-to-peer network–users, verification nodes, payment processors, and exchanges–all of which must be in agreement for the network to remain healthy (see table). Together they work to update and secure an ever-growing chain of transaction records that functions as the accounting ledger we call Bitcoin. All parties must agree on the rules of the game and run compatible software; otherwise the transaction chain can bifurcate into two separate chains with part of the network working on one and another part working on the other. In Bitcoin, this is called a hard fork, and it results in two competing currencies.

To imagine the potential chaos of this scenario, consider what it would be like if your bank account suddenly underwent a digital mitosis, splitting into two nearly identical versions of itself but with slightly different names. And now imagine what would happen if the rest of the world couldn't agree on which version of your account was the official one.

"It's a really weird situation. It's not where you want to go. The thing about money is it's useful because everyone accepts the same token," explains Wladimir van der Laan, a programmer from the Netherlands. Since April, he's been operating as the maintainer for the Bitcoin core source code, which basically means he makes the last call on what changes get folded into the ever-evolving Bitcoin software.

I met Van der Laan in Montreal in September after the close of a tense, hastily organized conference, the purpose of which, essentially, was to save Bitcoin. I asked him to speculate about all the different ways that a fork could play out. And there are many (see table). Van der Laan did his best to sketch out how each player could muck up the works—causing wallet balances to go up and down sporadically and miners to waste their energy on coins that have no value—but in the end he just laughed. "We don't know," he said. "That's the problem. There's uncertainty. That's why everyone is afraid of this fork."

And that's why they all ended up in Montreal. Over the course of the weekend, developers put aside their flaring Reddit tempers and their gnawing Twitter grudges and began to carve out a strategy. They'll have to create a new protocol version that will obviate the BitcoinXT fork attempt by solving Bitcoin's scaling problem in a way that everyone can agree to.

In the past, forks have been avoided because developers like Andresen, Hearn, and Van der Laan have managed to reach a consensus before throwing the code to the public. But as the Bitcoin project has grown, so too has the clutter of voices and agendas. And agreement is no longer so easy to come by.

"We started with a world where Satoshi [Nakamoto] was God and what Satoshi said was the word of God," says Andresen, referring to the pen name of the anonymous Bitcoin architect. Now that Satoshi has gone silent and Bitcoin is in the hands of the people who use it, there are two possible ways forward. The developers could unite, salvage what remains of Satoshi's celestial aura, and continue to shepherd the Bitcoin core software as the one true protocol. Or Bitcoin could evolve by the laws of natural selection.

Undeniably, there is a vivid strain of anarchists in the Bitcoin community who would welcome a "survival of the fittest" approach. But if the Montreal conference is any indication, most of the developers favor a process whereby the people who know the most about Bitcoin take on the brunt of the decision making.

They have until January to prove they are fit for the job. —MORGEN E. PECK

NEWS

MOLTEN SALT TOWER REBOOTS SOLAR THERMAL POWER

For the first time, solar thermal can compete with natural gas during nighttime peak demand

Solar power projects intended to turn solar heat into steam to generate electricity have struggled to compete amid tumbling prices for solar energy from solid-state photovoltaic (PV) panels. But the first commercial-scale implementation of an innovative solar thermal design could turn the tide. Engineered from the ground up to store some of its solar energy, the 110-megawatt plant is nearing completion in the Crescent Dunes near Tonopah, Nev. It aims to simultaneously produce the cheapest solar thermal power and to dispatch that power for up to 10 hours after the setting sun has idled photovoltaics.

"When the grid wants 110 MW, we'll provide 110 MW. There will be no variability," says Kevin Smith, CEO for SolarReserve, the plant's developer, based in Santa Monica, Calif.

Crescent Dunes, due to come on line by the end of this year, uses over 17,000 mirrors to focus sunlight on a heat receiver atop a 165-meter-high tower—a layout resembling California's massive Ivanpah solar power tower. However, while Ivanpah's receiver heats steam and pipes it directly to turbine generators, SolarReserve's heats a molten mixture of nitrate salts that can be stored in insulated tanks and withdrawn on demand to run the plant's steam generators and turbine when electricity is most valuable. Smith expects that NV Energy, the Las Vegas-based utility contracted to buy Crescent Dunes' output, will want it mostly dur-

PILLAR OF SALT: More than a million square meters of mirrors focus on a tower of molten salt to generate power for the Las Vegas Strip.

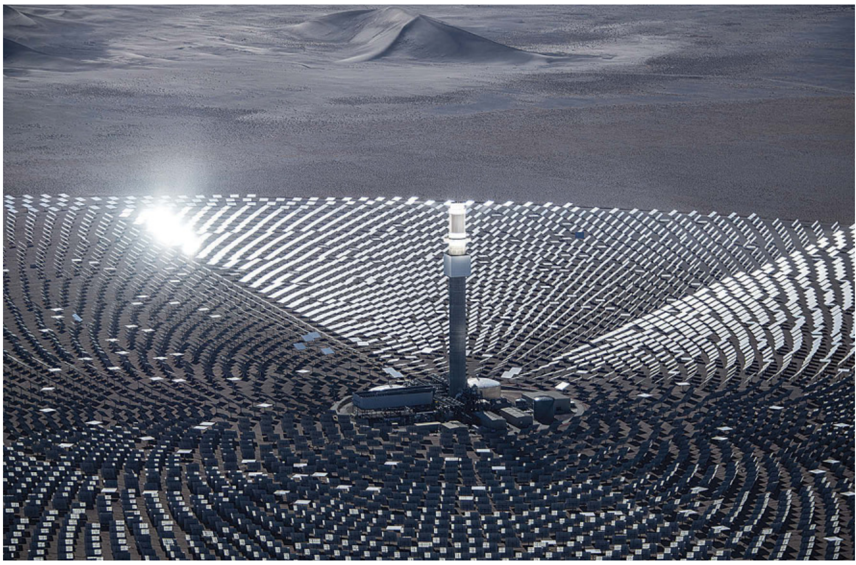

SOLARRESERVE

A HOT TIP: Power towers can burn birds. SolarReserve's killed 115 of them in an early test before the plant amended its standby procedures to eliminate midair bright spots.

ing the utility's unusually late demand peak, which the Vegas Strip's nightlife routinely stretches toward midnight.

Mark Mehos, thermal systems group manager at the National Renewable Energy Laboratory (NREL), says molten salt towers akin to SolarReserve's are "the next-generation technology" for solar thermal power. Plants without storage may never be able to compete with PV, says Mehos. And while molten salt storage is often added to trough-style plants, which use hectares of parabolic mirrors to heat synthetic oil flowing through pipes suspended above them, salt towers are cheaper and more efficient, he says.

Eliminating the heat exchange between oil and salts trims energy storage losses from about 7 percent to just 2 percent. The tower also heats its molten salt to 566 °C, whereas oil-based plants top out at 400 °C. That temperature boost squeezes 5 to 6 percent more power from the plant's steam turbines and enables a tank of salt to hold two to three times as much energy. The temperature advantage could grow: In September, SolarReserve won a US $2.4 million grant from the U.S. Department of Energy to develop a ceramic receiver that can withstand 732 °C.

Now the company just needs to make Crescent Dunes run. SolarReserve finished its construction early this year and was targeting first power generation at press time in early October. Smith says that one of the biggest complexities to master is the repositioning of the mirrors—more than 1 million square meters of reflective glass—every 60 seconds to keep their beams focused on the receiver. A 20-MW demonstration-scale plant completed in 2011 by Spanish solar thermal developer Sener Grupo de Ingeniería is running well, according to Mehos, but it must coordinate about one-sixth the number of mirrors as the new Nevada plant.

> Perhaps the biggest complexity was repositioning the mirrors—more than 1 million square meters of them—every 60 seconds to keep them focused on the receiver

Crescent Dunes promises to be the first of many big molten salt towers. Sener and fellow Spanish solar thermal developer Abengoa have large towers under construction in Morocco and Chile, respectively. SolarReserve expects to break ground this year in South Africa on its second plant. With the first plant completed, costs are coming down. Crescent Dunes' generation earns about $190 per megawatt-hour, including the value of federal subsidies, whereas Smith says the company's South Africa plant will get $125/MWh with no subsidies.

At present, solar thermal makes the most sense in markets where, unlike the United States, natural gas is pricey enough to make gas-fired plants less attractive for managing peak demand. But stronger growth is projected in the United States as well, because climate protections there threaten to restrict the use of natural gas to balance power generated by increasing levels of intermittent renewable power. NREL has predicted that solar thermal plants could be doing as much grid balancing in California as gas-fired and hydro plants combined by 2030, when that state's utilities are mandated to use 50 percent renewable electricity.

That is, unless declining battery costs make storing PV power more economical. James Nelson, a California-based energy modeler with the Union of Concerned Scientists, says this intrasolar race will be good for the environment: "I'm glad we have multiple technologies that can help the transition from gas power plants."

—PETER FAIRLEY

RODNEY BROOKS'S ONE-ARMED GAMBIT

The famed roboticist's new industrial robot can feel its way through precision motions

"Everything is a mess," Rodney Brooks says, as we make our way through a brick-walled room filled with robots in various states of assembly here at Rethink Robotics' headquarters in Boston. The mess, we should point out, makes Brooks, the company's founder and CTO, very happy. It means things are busy. The past few weeks, he says, have been particularly hectic. That's because today Rethink is making its collaborative factory robot Sawyer available for purchase.

Rethink introduced Sawyer, the company's second robot, earlier this year, but it didn't go on sale until mid-September. Like Baxter, its predecessor, Sawyer is specifically designed to safely operate alongside human workers, who can teach the robot to perform tasks simply by moving its arm through a series of motions (see "Rise of the Robot Worker," *IEEE Spectrum*, October 2012). But whereas Baxter has two arms and can help with packaging and material-handling tasks, the single-arm Sawyer is designed for tasks that require more precision, including machine tending and circuit-board testing.

To achieve that, Rethink redesigned its series elastic actuators, adding a new titanium spring as well as compact, lightweight servos with a high-gear ratio made by Harmonic Drive. This means Sawyer can measure forces on its joints with greater precision, allowing it to perform

SOLARRESERVE

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

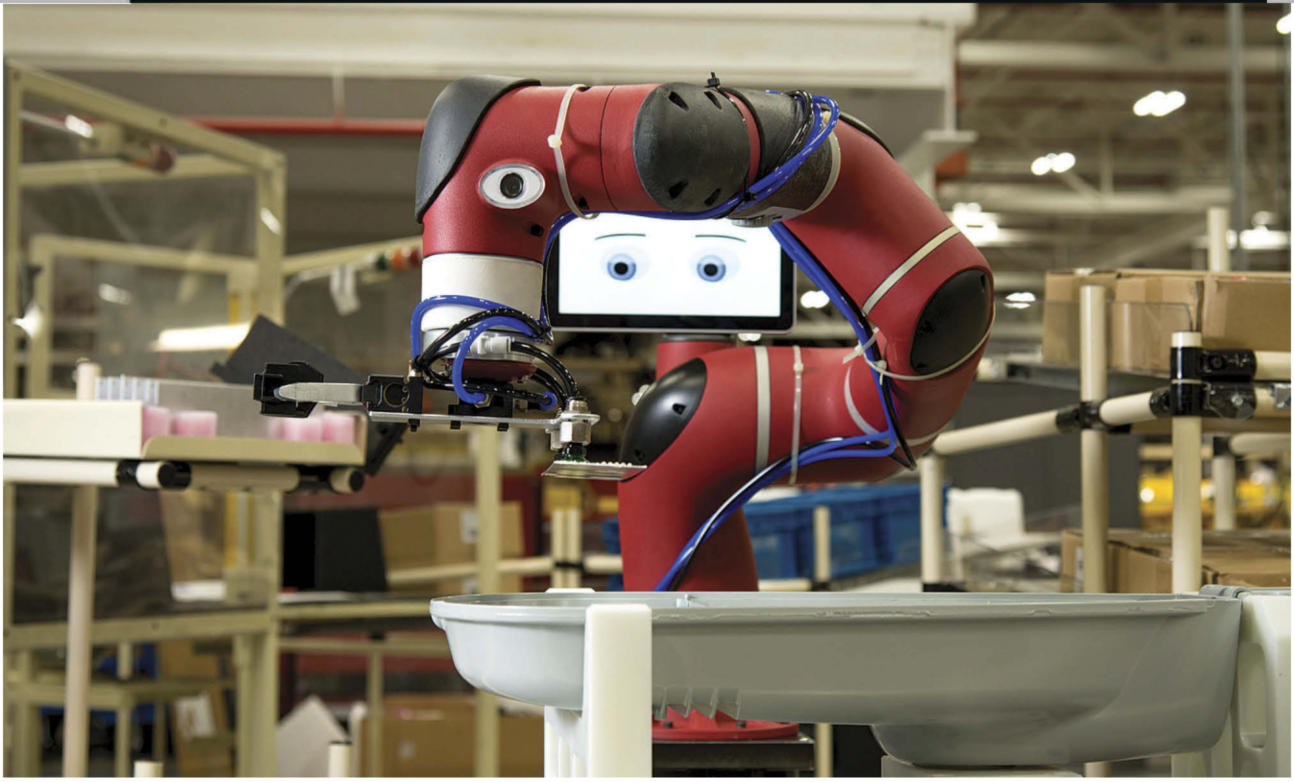

RETHINK ROBOTICS

WORKING, NOT VOGUING: Sawyer has one arm, but it's pretty strong and incredibly precise. This means that the robot may have more applications than its two-armed older brother, Baxter.

tasks like placing a part into a machine by "feeling" when the part is on the right spot. It also means Sawyer is more expensive, at US $29,000 (Baxter is $25,000).

"There are many important improvements over Baxter that make Sawyer suitable for many tasks that were not suited for Baxter," says Satyandra K. Gupta, a professor of mechanical engineering and director of the Maryland Robotics Center at the University of Maryland, in College Park.

Like Baxter's arms, Sawyer's appendage has seven degrees of freedom, but the new robot's arm has a greater reach–1.26 meters–and it's stronger, capable of carrying 4 kilograms. Sawyer is also more precise: It can perform tasks that require 0.1 millimeters of tolerance, much better than what Baxter can do. The precision is helped along by a camera in the bot's wrist, making recognition and inspection of parts easier. The camera also allows Sawyer to reregister its coordinate system, in case the robot or objects around it are moved around.

Rethink hopes that these features will help Sawyer stand out from the competition–most notably the UR3, UR5, and UR10 robotic arms from the Danish firm Universal Robots. That firm has sold about 4,000 robots all over the world and was acquired by automatic test equipment company Teradyne for $350 million in May.

Rethink doesn't say how many robots it has shipped, and it's unclear whether Baxter is selling "like hotcakes," as Brooks had hoped. With the development of a second robot, the stakes are higher for the 7-year-old company, which has raised $113.5 million from investors that include Bezos Expeditions, Charles River Ventures, Draper Fisher Jurvetson, GE Ventures, and Goldman Sachs.

Perhaps that's one reason why Rethink wants to offer Sawyer to customers worldwide. With Baxter, the company put a lot of emphasis on the North American market. For Sawyer, Brooks says that in addition to North America, "our big targets to start with are China, Japan, Western Europe, and Mexico."

Dan Kara, a robotics analyst at ABI Research, in Oyster Bay, N.Y., says Sawyer should be "very well received" by manufacturers, particularly in China, that are striving to become more agile and flexible as labor costs increase and their clients demand smaller volume runs and higher levels of customization. But he notes that Sawyer "is very much a new platform," so Rethink still has a lot of convincing to do if it wants to ramp up its sales.

Rethink says the companies already using Sawyer include GE Lighting–where the robot works on a production line positioning parts into a light fixture–and office furniture maker Steelcase, which taught Sawyer how to place parts into a welding machine.

So why didn't Rethink build Sawyer first, instead of Baxter, if that's what more customers seem to want? Like a parent who doesn't want to pick a favorite child, Brooks says there is a need for both. "And in the future there may be more robots," he says. But when asked what kind of robot could be next for the company, he declines to elaborate. "We're not going to come up with a 50-kilogram-payload robot. I'm pretty sure of that. But there will be different robots for different niches." –ERICO GUIZZO

GOOGLE VERSUS THE TROLLS

Internet giant snaps up patents directly from inventors

Google's Patent Purchase Promotion, which the company says received "thousands" of submissions during a three-week window, may prompt similar experiments in keeping patents out of the hands of what it considers the bad guys of intellectual property.

The experimental program was an attempt to intercept patents that individual inventors, operating companies, and others may have otherwise sold to organizations that don't make products but rather use the patents to extract license fees from operating companies, which do. Such organizations are commonly called nonpracticing entities, patent assertion entities, or (less politely) patent trolls. The program offered a chance for anybody to sell patents to Google at a price set by the patent holder. Google wound up buying 28 percent of the offered patents that it deemed relevant to its business, according to Kurt Brasch, the company's senior product licensing manager.

Even if the number of patents was only a few thousand, it's an impressive result given that applicants could submit only one patent at a time rather than entire portfolios, says Matt Moyers, senior director at Black Stone IP, an investment bank. The prices Google disclosed also seem in line with today's market values, he notes. The median price–excluding those offered at US $1 billion or more–of all submissions was $150,000, and Google paid prices ranging from $3,000 to $250,000.

That's probably low, however, compared with what the company might have paid just a few years ago. Prices for software patents have fallen dramatically since the U.S. Supreme Court's decision last year in *Alice Corp. v. CLS Bank International*, which raised the bar on software patentability.

"Since then courts have been killing off software patents at a pretty high rate, which has been driving down the value of software patents in general," says Matthew Ellsworth, partner and patent attorney at Sheridan Ross, based in Broomfield, Colo. He estimates that a patent portfolio worth $2 million to $3 million three years ago might be worth less than $100,000 today.

Prices for patents in Google's experiment

US $150,000: Median price of all patents submitted

47 percent: Submissions offered under $100,000

$3.5 billion: Highest submission price

$3,000: Lowest price Google paid

$250,000: Highest price Google paid

If Google had allowed more time, it would have likely received many more submissions. Some potential sellers either didn't see the notice or couldn't react in time. Ellsworth says some of his clients were interested in submitting patents, but "they couldn't wrap their heads around it fast enough and decided to let the opportunity pass rather than make an uninformed decision."

Even though Google received "a number of questions asking if we could extend the program," it stuck with its original schedule in the interests of fairness, says Brasch. He says the company has not yet decided if it will run the program again.

Google was surprised by the participation of individual inventors–who made up 25 percent of submissions–and by brokers, which submitted about half of the patents owned by operating companies. Those results showed that "the friction in the market is a barrier to both buyers and sellers," says Brasch.

Individual inventors especially appreciated the speed of the transaction, Brasch says. Google took just a few weeks to decide whether to buy the patents, but most buyers usually take their time evaluating a patent, which enables trolls to swoop in and make a tempting offer to inventors who are cooling their heels.

Still, the majority of patents sold to trolls come from operating companies, Google's prime target with this program. Operating companies made 75 percent of the submissions, and Google received many inquiries from them about the progress and success of the program, Brasch says. "Clearly, there was interest in what we learned," he says. "This is something that the industry could learn from. It opened a dialogue with other operating companies that I think will provide some benefits down the road."

The submissions from brokers were not necessarily all from operating companies. "I wouldn't be surprised if some of those were actually nonpracticing entities or licensing agents," says Moyers. In fact, a Google spokesperson confirmed that the company had indeed considered patents submitted through brokers by trolls, although Google would not specify whether it had purchased any.

The experiment nevertheless seemed to pique interest in ways to improve the patent marketplace. "If Google prompts a trend where many operating companies are going to do something similar, it starts to create a real and true marketplace for patents," says Moyers. "And that is what the industry desperately needs right now." —TAM HARBERT

PHOTOGRAPH BY ChinaFotoPress/Getty Images

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

DEEP-DISH PEEPER

ENGINEERS WORKING on the Five-hundred-meter Aperture Spherical radio Telescope, or FAST, being built into a valley in Guizhou, China, hope to finish installing the radio telescope's 4,450 triangular panels by next September. The panels, whose attitudes can be adjusted to form parabolas that point in different directions, will let scientists pick up faint signals originating more than 7 billion light-years away. To accommodate the data collected by the giant sensor—which will overtake the Arecibo Observatory in Puerto Rico as the world's largest and most sensitive radio telescope—a supercomputer capable of executing 10^{15} floating-point operations per second is being built to accompany it.

THE BIG PICTURE NEWS

RESOURCES

1915

THE YEAR THE "KETTERING BUG," THE FIRST DRONE AIRCRAFT, FLEW

IS U.S. DRONE RACING LEGAL? MAAAAYBE

Star Wars fans will remember the pod-racing scene from *Episode I: The Phantom Menace*, where the young Anakin Skywalker speeds through the landscape in a levitating scooter. A similar adrenaline rush is available through the new sport of drone racing, which typically involves quadcopters zooming around a course low to the ground. The quadcopters are piloted by people wearing video goggles or some other means of obtaining a first-person view, a form of radio-control flight that goes, naturally enough, by the acronym FPV. ▶

ILLUSTRATION BY Tavis Coburn

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

Multiple drone-racing leagues have sprung up in the United States, and this July saw the first Drone Nationals (officially, the 2015 Fat Shark U.S. National Drone Racing Championships), in Sacramento, Calif.

The attraction of drone racing is easy enough to understand. What's puzzling is how an organized sport could emerge in the face of what appears to be a legal prohibition on the whole shebang.

You see, in June 2014 the U.S. Federal Aviation Administration issued its "interpretation" of current regulations on model aircraft, an interpretation that bars FPV flight, or at least the use of video goggles for FPV. Whether you can use a video monitor is open to interpretation of the interpretation.

The FAA's justification is that, according to law, a model aircraft must be flown "within visual line of sight of the person operating the aircraft." With its interpretation, the FAA defined this phrase to mean that the pilot needs to keep the model in sight at all times, and it specifically prohibited the use of video goggles.

You'd think that after that, the FAA would frown on the Drone Nationals. But the FAA fully supported the competition and even sent two representatives to attend.

So does the FAA allow model aircraft to be flown by FPV or not? After the agency issued its 2014 interpretation, I concluded it didn't and ceased pursuing a hobby I had enjoyed for several years. But discovering the FAA's embrace of the Drone Nationals spurred me to revisit the question.

Be warned: The answer is still muddy. The first source I approached was Brendan Schulman, who recently took up a post as vice president for policy and legal affairs for the Chinese drone maker DJI. He points to signs that the FAA may be reversing course on FPV in situations where there is no threat to full-scale aircraft, such as for a drone race conducted, say, within a forest or at low altitude over a stadium, as was the case at the Drone Nationals. In particular, he notes that the Notice of Proposed Rulemaking that the FAA released this past February allows FPV in certain circumstances.

But for now: to FPV or not to FPV? "The interpretive rule is not regulation," says Rich Hanson, who works on government and regulatory affairs for the Academy of Model Aeronautics, a national membership organization for aircraft modelers. "I don't envision that [the FAA] would take action" against an FPV enthusiast, he says. Still, just because the practical likelihood of getting into trouble with the FAA is low, that doesn't mean you should flout the FAA's wishes. But what does the FAA want?

John Goldfluss, who was one of the FAA representatives at the Drone Nationals, shed some light on the question, although he stresses that his views should not be taken to represent the FAA's official position. He regards the Drone Nationals as a test, in which the FAA participated to see whether loosening its stance of FPV might be warranted, at least in some very well-controlled situations.

But that still doesn't make much sense to me, for the following reason: Imagine that the FAA concludes that FPV should be allowed at events like the Drone Nationals. That's all fine. But the people who fly at the Drone Nationals also need to practice somewhere. And the vast majority of aspiring drone racers will never get the opportunity to fly at the nationals anyway. They'll just compete at local events. How can they do that if the FAA still frowns on flying radio-controlled model aircraft by FPV in general?

I warned you that the answer was muddy.

—DAVID SCHNEIDER

THE RACING FORM: FPV competitors have judges "ride" along to ensure the course is flown correctly.

THE ULTIMATE DIGITAL PICTURE FRAME

ADD GESTURE CONTROL TO AN HD DISPLAY TO SHOWCASE ARTWORK

SOME 20 YEARS AGO, BILL GATES WAS THE **king of computing, and not above boasting about his new high-tech house. In his book *The Road Ahead* he described how large monitors in the house would display great works of art, changing every day. These digital art frames could even react as you walked past, no button pushing required. Now *that*, I remember thinking, is something I would like. ● Gates predicted that the day would come when almost any middle-class American family could enjoy this kind of technology. Recently I decided to find out whether that day has arrived. ● The impetus to build my own digital art frame came from my mother, Sylvia Gibbs, a professional artist. Over the years I have framed and hung quite a few of her paintings. But as I was straightening up a closet this spring I found a portfolio full of prints and drawings she had given me years ago. They are terrific; I'd love to put them up. But it would cost thousands of dollars to have them all professionally framed. And I don't have enough wall space to hang them.**

LEFT: ROSWELL FLIGHT TEST CREW; RIGHT: W. WAYT GIBBS; ARTWORK: SYLVIA GIBBS

But then I remembered Gates: Why not display them all in a single "frame" and switch among them as I liked? Of course I could simply buy a digital photo frame designed for displaying family snapshots. But I wanted something large enough to show the art at near-original size; consumer digital picture frames have displays that top out at around 18 inches. I also wanted a frame that could look nice on a wall without cables dangling from it. I wanted to control it without buttons or a mouse, and to load art onto it easily through a shared network folder. It shouldn't cost more than $300.

The first step was obvious. I needed an energy-efficient HD monitor of at least 24 inches with good color fidelity, high brightness, wall-mounting holes, and a thin, unobtrusive bezel. I chose a Viewsonic VA2451M LED display which, for US $170, had all these features, plus downward-facing interface ports from its rear power supply and logic board. Hanging on a $10 VideoSecu ultraslim wall mount, the back face of the flat panel is only 5 centimeters away from the wall—more than enough space for the computer that I chose to be the brains of the frame, a Raspberry Pi 2 Model B ($73, including a 32-gigabyte microSD card, power supply, and Wi-Fi dongle).

I mounted the Pi inside an enclosure, hung it on two small bolts that I had screwed into the back of the monitor, and connected it to the display. It took a little careful surgery with a Dremel tool on the monitor case to make room for the Pi's power cable, but everything then fit nice and snug.

The image-handling software was trickier. Of the many picture-display programs available for the Pi, none seemed to have all the capabilities I needed: the ability to loop through all the image files in a nested folder structure; transition smoothly from one to the next after a customizable interval, scaling each image to fill the screen; and importantly, the ability to accept input from a sensor connected to the Pi's general-purpose input/output (GPIO) interface. In the end I settled on OSMC, an open-source media center recently ported to the Pi.

Once OSMC was up and running, I tackled the problem of controlling the flow of images without a wireless mouse or keyboard.

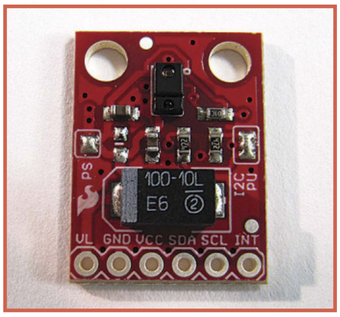

UNLIKE 20 YEARS AGO, today good-quality, high-definition, flat-panel displays, such as this 24-inch Viewsonic [top], are available at a reasonable price. Images were stored on and displayed by a Raspberry Pi [middle], which translated commands sent by a gesture sensor [left].

W. WAYT GIBBS (2)

The Pi has a dedicated header on its logic board for a camera, and my first thought was to use that. But then I found an even simpler and less expensive option at SparkFun: a thumbnail-size gesture sensor ($15) built around Avago's APDS-9960 chip.

The code library that SparkFun provides for the sensor is written for Arduinos, but I found an open-source port for the Pi online and, with a bit of help from the wiringPi website and the OSMC forum, got the sensor to communicate with the Pi via the sensor's I2C interface. I wrote a simple shell script to translate the sensor output into keystrokes that get forwarded to OSMC as commands.

When it was all wired up and debugged, the effect was extraordinary. I waved my hand to the right over the sensor, and the display advanced to the next picture. A wave left brought back the previous image. The sensor also detects upward, downward, inward, and outward gestures, all of which I connected to different OSMC functions—one of them being to shut the system down for the night. I connected a momentary switch to the reset header on the Pi and hid it on the back of the monitor to enable rebooting in the morning.

Rummaging around my workshop, I found a nice piece of mahogany about the right size to cover the power cords and sensor. I carved out a channel for the cables and made a niche with a peephole in it for the sensor. This screws to the monitor using the same hole that normally secures it to a stand.

All that remained was to upload my mother's portfolio. The prints are too fragile to run through a document feeder, and some are too large for my flatbed scanner. But I own an impressive overhead scanner, the Fujitsu SV600, which is the perfect tool for the job. It doesn't make contact with the art, so there is no risk of damage, and it can scan items (including bound books and 3-D objects) up to 43 centimeters wide at high resolution and with accurate color. In less than 30 minutes, I'd scanned a thick stack of prints and drawings to the Pi's pictures folder.

There they were, works by Mom. And, in a nice bit of turnabout, through its sensor, this digital art gallery was staring intently at me.
—W. WAYT GIBBS

YELOHA
BRINGING SOLAR INTO THE SHARING ECONOMY

A SUNNY TRIO: Gathered in their Boston offices are Yeloha's founders: [from left] Idan Ofrat, Amit Rosner, and Paolo Tedone.

Sleek, shiny solar panels decorate only the most elite roofs. While the solar industry has been experiencing an upswing in installations, the qualifications required to get a solar system on your roof are surprisingly capricious. Most suppliers offering subsidized systems require that you have owned your home for the past two years, the house faces south, and the roof gets ample sun exposure, among other things.

For those locked out of solar, say hello to Yeloha, a Boston startup that wants to make solar sharable. Company cofounder Amit Rosner went through the solar system checklist, but found out his Boston home is part of the 49 percent of solar-ineligible households in the United States. His cofounders, Idan Ofrat and Paolo Tedone, can't host solar systems either. "This is crazy," Rosner says. "None of us can go solar. Most of our friends or family can't go solar. No matter how much we want it or need it, we cannot go solar."

The company's business model was inspired by successful sharing-economy platforms, such as Uber and Airbnb. Yeloha matches "sun partners" with "sun hosts." Hosts offer up their prime sunlit roofs for free panel installations. They share the energy the panels create with sun partners, and get credits off their bills for 20 to 30 percent of the energy produced with no need to switch from their usual utility providers. All the solar electricity the sun hosts generate is fed to the electric grid. Yeloha works with utilities so that sun hosts and sun partners are credited with their share. To evaluate a sun host, Yeloha's in-house software first analyzes the host's building location to make sure there is adequate space available and sun exposure. If the software gives a green light, one of Yeloha's installation partners does an on-site check.

The sun partners—those who don't have solar-friendly roofs—pay a subscription fee and save about 5 to 10 percent on their electricity bills. However, partners can save more depending on how many panels they subscribe to. Yeloha's subscriptions start at one panel per year for US $65, promising to generate 336 kilowatt-hours annually.

Yeloha began in 2012 after Rosner worked at a photovoltaic development company called SolarEdge Technologies, headquartered in Hod Hasharon, Israel. He studied data analytics as a computer science major at the Hebrew University of Jerusalem. These skills came in handy when the Yeloha team created algorithms to simulate solar on roofs.

The computer system they built processed live solar-power performance data. For two years, they collected data from thousands of actual solar arrays of homeowners who signed up on Yeloha's first platform. The analysis of the data helped forecast solar potential on nearby roofs and identify candidates that could generate enough electricity to offset their bills and have extra to share.

Since its official launch in April, Yeloha has connected thousands of people in hundreds of cities with solar. "It sounds crazy, but we really want to have solar panels everywhere," he says. "There is no excuse today for not going solar."

David Feldman, a senior financial analyst at the National Renewable Energy Laboratory's Washington, D.C., office, says that while shared solar is a small part of the total solar market, the business model is becoming more popular. "A lot more programs are being set up across the United States," Feldman says. "States are passing laws that support shared solar, and utilities are becoming more interested in this business model."

Feldman's 2015 report estimates that shared solar could make up 32 to 49 percent of the solar market in 2020, producing up to 11 gigawatts of electricity in the United States.

Yeloha continues to expand beyond sunny household roofs to all kinds of properties, such as theaters and places of worship. Rosner says shared solar can give everyone an opportunity to be a part of a community while doing the right thing for the environment. "We all live under the same sun," he says. —**LAUREN J. YOUNG**

Name: Yeloha **Location:** Boston **Founded:** April 2015 **Employees:** 15 **Funding:** US $1 million seed round from the company's early platform, plus $3.5 million series A round

YELOHA

BLASTOFF!
FOUR GAMES FOR ARMCHAIR ASTRONAUTS

A ROCKY MOMENT: In *Elite: Dangerous*, players have the option to mine the rings of gas giants for resources—or pirate from those who do.

FRONTIER DEVELOPMENTS

In the early days of video games, players often found themselves flying spaceships.

Titles like *Asteroids*, *Defender*, and *Galaxian* dominated arcades, while home systems saw big sales for games such as *Cosmic Ark* and *Elite*. But in the 1990s, the rise of first-person shooters and strategy games pushed spaceships into the background, although entries like the massively multiplayer game *Eve Online* kept the genre alive. Now, fortunately for would-be pilots, there are a number of new spaceship-based games available that have the added bonus of being particularly rewarding for those with a scientific or engineering bent.

First up, there's *Elite: Dangerous*. Released last December, this is a reboot of the Elite game series by one of its original creators and is available for PCs, Macs, and the Xbox for US $45. The original *Elite* was released in 1984, and pioneered vast procedurally generated environments and the kind of open-ended "sandbox" game play that's now common in many modern titles. Players could fly from star system to star system, trading goods, mining asteroids, or engaging in space piracy as they accumulated cash to upgrade their spacecraft. *Elite: Dangerous* successfully retains most of the elements that made the original great, but now, rather than an arbitrary collection of fictional stars, players fly around a beautifully detailed representation of the Milky Way, based on real stellar data. (You also have the option of multiplayer action.) Want to visit the Orion Nebula, or gaze on the supermassive black hole at the center of the galaxy? No problem—assuming you survive the journey.

Next is *Kerbal Space Program* (*KSP*), a game that finally emerged from beta earlier this year after a lengthy gestation period and is available for PCs and Macs for $40. In *KSP*, you run the space program of Kerbin, a small planet populated by amiable green Kerbals. Starting with very simple rocket parts that can barely make orbit, your job is to design and assemble spacecraft to accomplish ever more elaborate missions. The Kerbals' solar system has been set up to stack the odds in the player's favor: For example, Kerbin's small size means that escape velocity from the surface is only 3.4 kilometers per second, compared to 11 km/s for Earth. But all the orbital dynamics are played straight—players must learn how to come to grips with fuel-to-mass ratios and spacecraft power budgets, perform orbital plane changes, raise and lower their apoapsis and periapsis, and so on. It's even possible to perform slingshot and aerobraking maneuvers.

StarMade is the least developed of the games mentioned here. Currently free to play, it's at the alpha stage of development, and so has rough edges. Still, it has potential, and allows a huge range of spacecraft designs. Anyone familiar with the insanely popular *Minecraft* will recognize *StarMade*'s blocky look and feel and its basic crafting concept. However, rather than mining minerals on a completely flat world to make medieval weapons or magical items, in *StarMade* the goal is to create bigger and better spacecraft in which to explore the universe. Raw materials mined from asteroids can be converted into components such as flight computers, thrusters, and reactor cores, and then assembled into huge vehicles.

The last title I'll mention here isn't really a game, but rather crosses the line into pure simulation. *Orbiter*, which is free to play and available for PCs, tries to replicate as closely as possible the dynamics and engineering systems of real (and fictional) spacecraft amid a detailed representation of our solar system. Unlike in *KSP*, creating your own spacecraft isn't for the casual player, but while *KSP* is played using a third-person perspective, *Orbiter* puts you right into the action in the pilot's seat. (Longtime readers of *IEEE Spectrum* may remember that in 2013 I wrote a Hands On article about building a custom USB controller to aid with the faithfully duplicated difficulties of docking to the game's simulated International Space Station.) It's been about five years since the last official release of *Orbiter*, but development continues and eager players can grab betas, the latest of which was released in September. **—STEPHEN CASS**

AMERICAN EXCEPTIONALISM

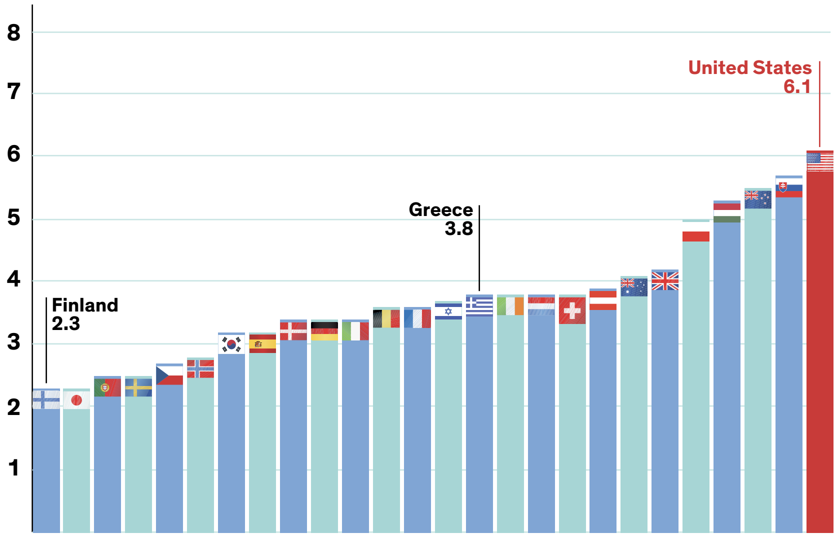

INFANT MORTALITY RATE, 2010: This revealing measure is a good proxy for overall quality of life. It is therefore telling that the United States comes in 26th, well below Greece, whose financial troubles recently have put it in the headlines.

SOURCES: CDC/NCHS, LINKED BIRTH/INFANT DEATH DATA SET (U.S. DATA); OECD 2014 (ALL OTHER DATA)

BELIEF IN "AMERICAN EXCEPTIONALISM"—THAT UNIQUE blend of ideals, ideas, and love of liberty made so powerful by great technical and economic accomplishments–is alive and well. Even President Obama, a reluctant endorser to begin with, has come around. Early in his presidency (in April 2009), he affirmed his belief by essentially denying it: "I believe in American exceptionalism, just as I suspect that the Brits believe in British exceptionalism and the Greeks believe in Greek exceptionalism." By May of 2014, he had relented: "I believe in American exceptionalism with every fiber of my being." ● But such proclamations mean nothing if they cannot stand up to the facts. And here what really matters is not the size of a country's gross domestic product or the number of warheads or patents it may possess but the variables that truly capture its physical well-being and educational standard. These variables are simply life, death, and knowledge. ● Infant mortality is an excellent proxy for a wide range of conditions including income, quality of housing, nutrition, education, and investment in health care. Very few babies die in those affluent countries where people live in good housing and well-educated parents (themselves well nourished) feed them properly and have access to medical care. How does the United States rank among the world's roughly 200 nations? The latest available comparison (for 2010) shows that with 6.1 of every 1,000 live-born babies dying in the first year of life, the United States does not figure among the top 25 nations. Its infant mortality is far higher than in France (3.6), Germany (3.4), and Japan (2.3). And the U.S. rate was 60 percent higher than in Greece, a country portrayed in the press as an utter basket case.

Excusing that very poor rating by saying that the European countries have homogeneous populations does not work: Modern France and Germany are full of recent immigrants (just spend some time in Marseille or Düsseldorf). What matters more is parental knowledge, good nutrition, the extent of economic inequality, and access to universal health care, the United States being (notoriously) the only modern affluent country without the latter.

And looking at the journey's end gives an almost identically poor result: In 2013, U.S. life expectancy ranked 34th worldwide, an average of 79 years for both sexes, which is–again–behind Greece (81), as well as Portugal (81) and South Korea (82). Canadians live three years longer on average, Italian men four, and Japanese women (at 87) six years longer compared to their U.S. counterparts.

Educational achievements of U.S. students (or a lack thereof) are scrutinized with every new edition of the Organization for Economic Co-operation and Development's Program for International Student Assessment, or PISA. The latest results for 15-year-olds show that in math, the United States ranks just below Russia, Slovakia, and Spain, but far lower than Canada, Germany, and Japan. In science, U.S. schoolchildren place just below the mean PISA score (497 versus 501); in reading, they are barely above it (498 versus 496)–and they are far behind all the populous, affluent Western nations. PISA, like any such study, has its weaknesses, but large differences in relative rankings are clear: There is not even a remote indication of any exceptional U.S. educational achievements.

Some readers might find these facts discomforting, but there is nothing arguable about them. In the United States, babies are more likely to die and high schoolers are less likely to learn than their counterparts in other affluent countries. Politicians may look far and wide for evidence of American exceptionalism, but they won't find it in the numbers, where it matters. ■

INFOGRAPHIC BY Erik Vrielink

UNDERSTANDING OPENNESS

The choice between closed versus open innovation is not always simple

MANY, MANY YEARS AGO I GAVE A TALK AT THE U.S. Federal Communications Commission. I remember nothing of that day except for a passing comment made by the speaker before me, a software pioneer. He said that the greatest achievement of the computer industry was not the personal computer but the concept of an open platform on which third parties could innovate, and that this open platform was what the telecom industry lacked. • I thought this was a silly thing to say. It had been easy enough for the computer industry to standardize the communications bus in a little steel box, but it would be quite another thing to create an open platform from a worldwide network made up of a myriad of different technologies. This would never happen. • A few years later I realized how wrong I had been. The speaker's prescient comment had been made before the evolution of the Internet. A simple protocol–TCP/IP–had turned a vast network of networks into the virtual equivalent of a little steel box. Of course, regulatory and legal revolutions were also involved, but the end result was that users were enabled–and even incentivized–to innovate on the telecom network. And they did. • Still later I returned to Bell Labs as an alumnus to give a talk. "Once we controlled the network...," I began. I looked out on the audience with nostalgia. The room was full of engineers and scientists, but this august assembly was a mere drop in the bucket compared to the millions upon millions of users who were now guiding the evolution of network services. Once we had thought this was our exclusive charter.

I was recalling this chain of events in the context of a discussion about the automotive industry, which is edging toward an open platform as cars become more and more like computers with wheels. Driving an open platform down the highway sounds a little scary, but maybe something similar could have been said about the telecom network, too.

Through the years I have attended and participated in many meetings and conferences devoted to innovation. It seems to me as if the subject is stuck in a do-while loop. Each year, speakers begin their presentations by observing that in the past we had not known how to cultivate innovation. Now, fortunately we do, each speaker will say, and then go on to describe a new organizational paradigm that is sure to increase innovation. Fast forward a year and repeat.

I think that innovation occurs quite naturally and would even be difficult to stop. I imagine an organization putting out an edict that innovation will no longer be tolerated; I can envision the engineers coming in at night and working on ideas quietly by candlelight.

Given that innovation occurs naturally, a sure way to increase innovation is to have more people involved and empowered to innovate, and this is the power of an open platform. However, there are both technical and economic risks involved. Chaos, lack of control, and market loss loom as possibilities. Sometimes it's a choice between open and innovative and closed and profitable. In his book *The Triumph of Ethernet* (Stanford University Press, 2002), Urs von Burg also argues that when improvements to a technology can't be made on an independent, modular basis but require interdependent, systematic changes, closed and proprietary approaches have the edge.

I sometimes wonder if today there are college courses on technology business strategy. When and how, and at what interface or levels, is it advisable to be open? There are many examples that can be studied, and it's an interesting and important question. ■

ILLUSTRATION BY Harry Campbell

Innovate the world, come to Carnegie Mellon ECE

Electrical and Computer Engineering at Carnegie Mellon University has a long history of excellence. Our vision is to be a worldwide creative driving force of highest scholarly and entrepreneurial quality. Our mission is to inspire, educate and produce electrical and computer engineers capable of tackling fundamental scientific problems and important societal challenges.

We are engaged in global collaborative research and teaching efforts through extensions of our academic programs and research in Silicon Valley, Portugal, Rwanda, China and Singapore.

www.ece.cmu.edu

- B.S. in Electrical and Computer Engineering
- B.S./M.S. in Electrical and Computer Engineering
- M.S. in Electrical and Computer Engineering
- M.S. in Software Engineering
- M.S./MBA in Electrical and Computer Engineering
- Ph.D. in Electrical and Computer Engineering

Electrical & Computer ENGINEERING

Carnegie Mellon University
College of Engineering

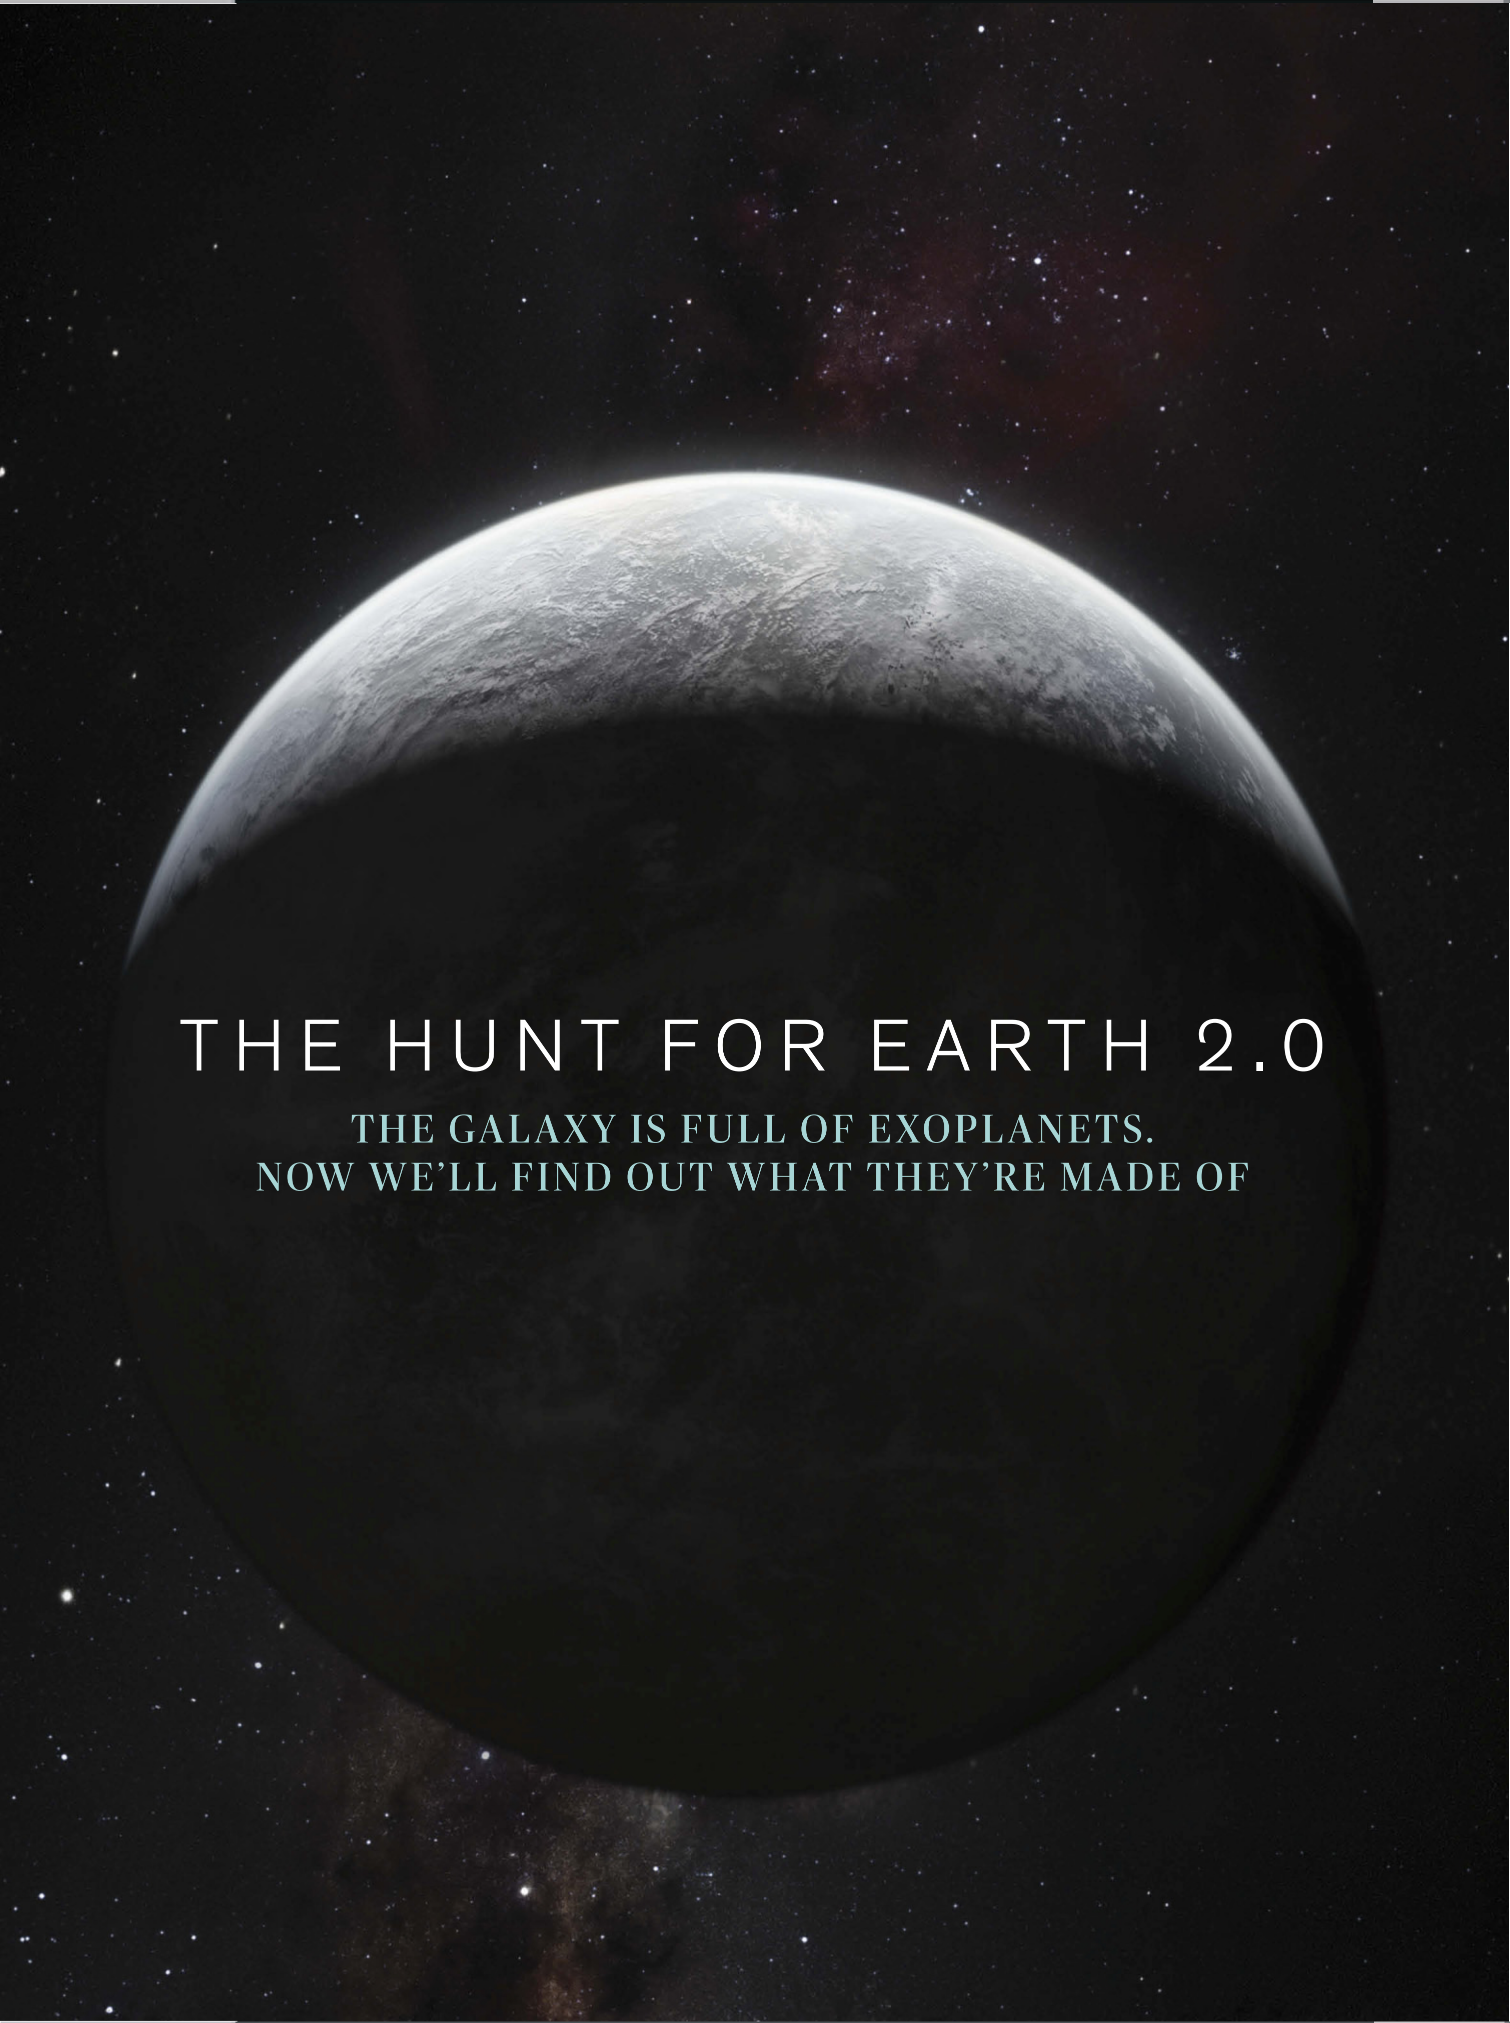

THE HUNT FOR EARTH 2.0

THE GALAXY IS FULL OF EXOPLANETS. NOW WE'LL FIND OUT WHAT THEY'RE MADE OF

BY ALBERTO CONTI & MARK CLAMPIN

Earlier this year, astronomers announced the discovery of a planet not much bigger than Earth. Dubbed Kepler-452b, the planet orbits a star like our own sun. Even more exciting, it orbits its star at just about the same distance that Earth orbits the sun, with a year that lasts just 20 days longer than our own. For the first time, astronomers had found a world that could be called—if not an Earth twin—at least a close cousin.

But what do we know about this world? Like most of the exoplanets found to date, Kepler-452b remains a mystery. It's not clear, for example, whether it boasts a rocky surface like Earth's and, if it does, whether it has oceans or a breathable atmosphere.

Today, our exoplanet catalogs are filled with many hundreds of entries but only the most basic of details: fundamental features, such as mass or radius–or in certain cases, both–and a few numbers pertaining to each planet's orbit. There is little to tell us what these planets are made of, what their climates are like, and whether they could support life.

But this state of affairs is beginning to change. More than 20 years after the first exoplanets were found, the focus is shifting to planetary characterization. Astronomers are studying the spectra of exoplanet atmospheres, filling in our picture of these distant worlds. This burgeoning field will get a big boost with NASA's James Webb Space Telescope. Slated to launch in 2018, JWST will be able to scrutinize the atmospheres of exoplanets that can't currently be studied. With any luck, it will give us our first detailed look at other worlds that might be capable of supporting life.

Twenty-five years ago, we knew of only one solar system in the Milky Way–our own. We now know of well over 1,000 stars that harbor one or more planets.

Some of the first exoplanets were discovered by looking at the spectral lines of stars–specifically at the shift in wavelength that occurs when the gravitational pull of a planet makes its star wobble back and forth with respect to Earth. Today, the majority of known exoplanets have been discovered by looking for transits. A transit occurs when a planet passes across the face of its star, creating a small dip in the amount of starlight that reaches our telescopes.

The European COROT telescope, launched in 2006, was the first telescope dedicated to hunting for exoplanet transits from space, where you don't have to contend with the fluctuations in light levels caused by Earth's atmosphere. Less than three years later, NASA launched the survey telescope Kepler. By staring for years at the same patch of sky, Kepler has identified more than 4,600 possible planets. Of those, more than 1,000 have been confirmed by ground-based observations.

The planetary systems Kepler discovered are quite diverse. Many are like ours, with multiple planets orbiting one star. Other systems are far

MARTIN KORNMESSER/ESO

more exotic, with planets that orbit two stars that also orbit one another.

Among the most significant results from Kepler is the abundance of smaller exoplanets. By extrapolating from Kepler's discoveries, astronomers now think there are billions of worlds about the size of Earth in the Milky Way. A good fraction of exoplanets are larger than Earth but smaller than Neptune. Those on the smaller end of the spectrum may have developed a rocky surface and an atmosphere like Earth instead of becoming an ice giant with a thick atmosphere and high surface pressure like Neptune.

If a smaller exoplanet orbits a star at the right distance—not too close and not too far—it could have temperatures in the right range for liquid water to exist on the surface. The location of this "habitable zone" varies with the type of star. The numbers are still quite uncertain, but Kepler's survey suggests that as many as about 20 percent of the G-type stars in the galaxy—a class of stars that includes our own sun—have roughly Earth-size planets in their habitable zones, for a total of a billion or so worlds. Small, temperate planets also seem to be quite common around smaller, cooler M-type stars, also called red dwarfs.

How do we find out what these worlds are like? The transit of a planet can yield important, basic information. The amount of light the transit blocks out, for example, gives the planet's diameter: The bigger the planet, the more light it will block. And the time interval between transits will yield the planet's orbital period.

But if we want to learn more, we need to pick apart the planet's light. Studying a planet's spectrum can reveal what its atmosphere is made of and how thick and hot it is. Studying the molecules present in the atmosphere could give insight into the environment in which the planet formed—and also whether the planet might prove hospitable to life.

Fortunately, transits give us a relatively straightforward way of studying planetary spectra. When an exoplanet passes between its host star and Earth, astronomers can measure the starlight being absorbed by the planet's atmosphere. By comparing the star's spectrum during a transit with its spectrum when the planet is not in front of it, scientists can deduce what elements are present in the planet's atmosphere. A similar exercise can be performed to look at thermal signals from the planet's atmosphere, radiated as infrared light. These signals can be separated from starlight because they will disappear while the planet passes behind its host star.

The first such atmospheric studies were done in the early 2000s, when astronomers used a spectrograph aboard the Hubble Space Telescope to measure the levels of sodium in the atmosphere of the exoplanet HD 209458b, a giant planet that sits about 150 light-years from Earth. Since then, Hubble and NASA's Spitzer Space Telescope, along with ground-based detectors, have helped add to the spectroscopic menagerie. Hubble data on GJ1214b, for instance, suggest that it might be a cloudy planet or have a thick haze of molecules made through interaction with sunlight. Last year, Hubble observations of WASP-43b were used to construct a detailed temperature profile of that planet's atmosphere, which is hot enough to melt iron on its star-facing side. Spitzer observations of 55 Cancri e, which orbits its host star so closely that its year lasts only 18 Earth hours, suggest the planet is quite dark in appearance.

These and other observations have started to give us deeper insight into what exoplanets are actually like. But they've been limited to a small group of planets that meet the right set of observing conditions: The star must be bright enough and the planet big enough for its atmosphere to generate a strong signal. To learn about a wider set of planets—in particular, ones that are small enough to be rocky like Earth—we need a bigger telescope. We also need a new list of spectroscopic candidates; the transiting planets that Kepler finds typically orbit stars that are too distant and therefore too dim to create enough of a signal.

The process of identifying such candidates is already under way. One effort, a Chile-based array of telescopes called the Next-Generation Transit Survey, is focused on worlds around the size of Neptune and recently saw first light. And in 2017, NASA's Transiting Exoplanet Survey Satellite (TESS) will launch. The spacecraft is designed to survey the whole sky and is expected to find a number of smaller exoplanets around

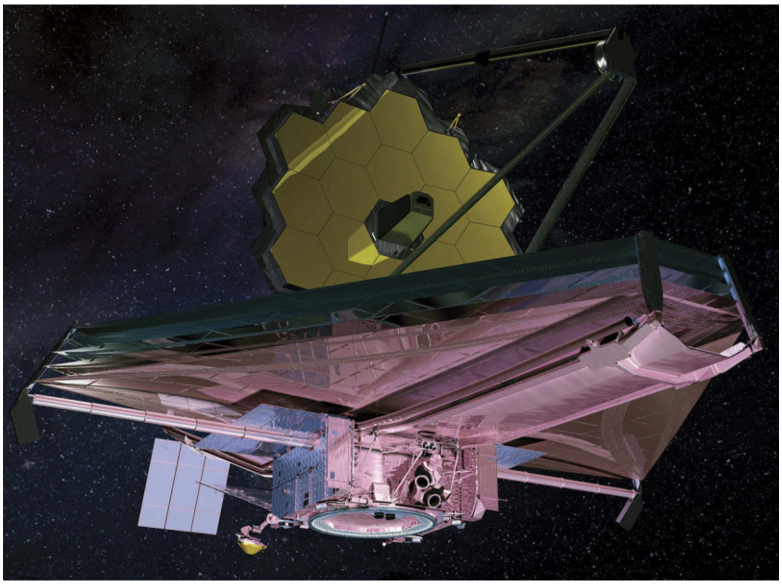

INFRARED EYE: The James Webb Space Telescope (JWST) boasts a 6.5-meter-wide mirror [yellow] that will let it pick up infrared light from a variety of sources. As a general-purpose observatory, JWST will have many targets of study, including exoplanets. A multilayer sunshield will help the telescope reach its operating temperature, by protecting it from the heat of the sun.

NORTHROP GRUMMAN

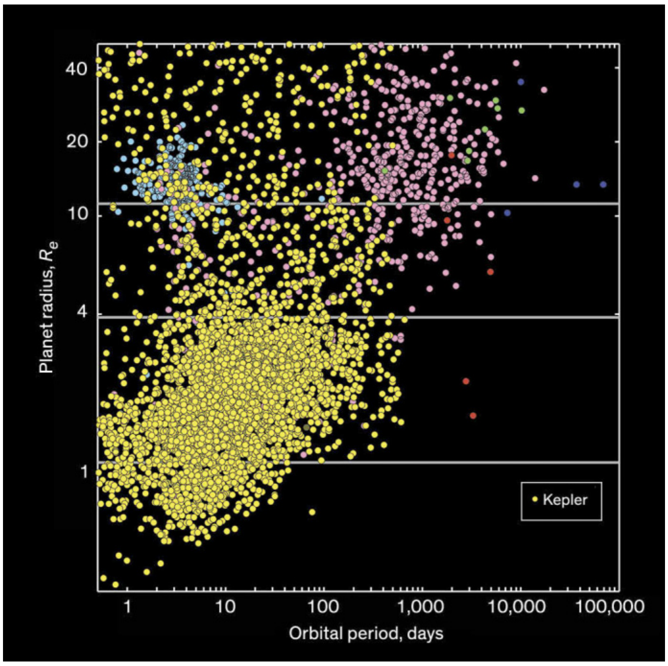

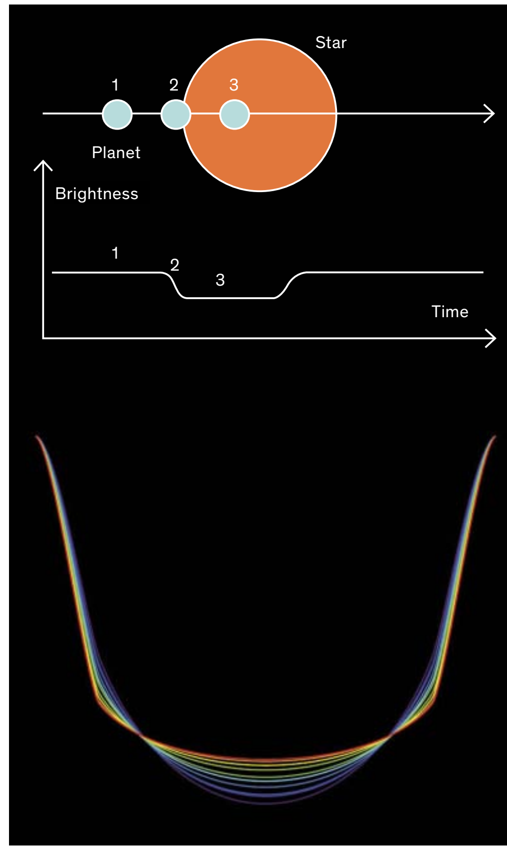

DIPS AND DOTS: The exoplanets found by Kepler [above, yellow dots] span a large range in size. Some are smaller than Mars, while others are bigger than Jupiter (non-Kepler discoveries are shown in other colors). Kepler found these planets by looking for transits [right, top], dips in a star's light when a planet passes across its face. A planet's atmosphere will absorb different amounts of starlight at different wavelengths [right, bottom]. The differences in the light that reaches Earth can be used to determine physical properties and composition.

bright stars located within about 200 light-years of Earth.

As TESS begins its work, our teams at Northrop Grumman in Redondo Beach, Calif., and at NASA Goddard Space Flight Center, in Greenbelt, Md., will be putting the final touches on the James Webb Space Telescope.

The telescope's 6.5-meter-wide mirror will be the largest ever flown in space for astronomy, giving JWST about seven times the light-collecting area of Hubble. While Hubble is most sensitive to the optical and ultraviolet part of the electromagnetic spectrum, JWST will conduct the bulk of its observations in the infrared, which will afford it an unprecedented view of the birth and evolution of galaxies and the formation of stars.

The telescope wasn't originally designed with exoplanets in mind. Even so, it will provide an impressive capacity to measure transit spectra. The near-infrared happens to be a region of the spectrum where water, oxygen, carbon dioxide, and a number of other interesting molecules have prominent spectral features.

In general it is quite tricky to find these spectral lines in transit spectra. The measurement involves detecting light from the very thin ring of a planet's atmosphere against the bright backdrop of its star. You then have to subtract one very large number from another: starlight with planet from starlight with no planet.

It's complicated in part because the signal becomes diluted as it's divided up by wavelength, which you need to do to hunt for spectral signatures. So to boost the precision of these observations, you need to maximize the number of photons that are collected. The smaller the planet and the thinner its atmosphere, the bigger a telescope's collecting area must be in order to gather a sufficient signal. At the same time, a telescope needs to be able to look at bright stars without saturating its detector. Some of JWST's instruments do this with special optics, which spread out and so diffuse the starlight.

With its bigger collecting area, JWST's observations of transiting exoplanets are expected to achieve a signal-to-noise ratio of 100,000 to 1, about five times as high as Hubble's. That precision should allow the new telescope to discern the composition and thermal structure of giant-planet atmospheres with just one or two transits. It will also, for the first time, allow us to study the atmospheres of smaller exoplanets, including potentially habitable worlds; TESS should turn up at least a few rocky worlds, orbiting in the habitable zones of nearby stars, that JWST might examine.

Most of these candidates will be orbiting M stars. Finding and studying planets that orbit those smaller stars is generally easier because a planet will block out a larger fraction of the starlight, creating a deeper, clearer signal. M stars are also thousands of degrees cooler than our own sun, so the range of distances at which liquid water might exist lies much closer in. That makes transits by planets in the habitable zone of an M star much more frequent than those around sunlike

CLOCKWISE FROM LEFT: NATALIE BATALHA/NASA AMES RESEARCH CENTER; HANS DEEG; HEATHER KNUTSON/CALTECH

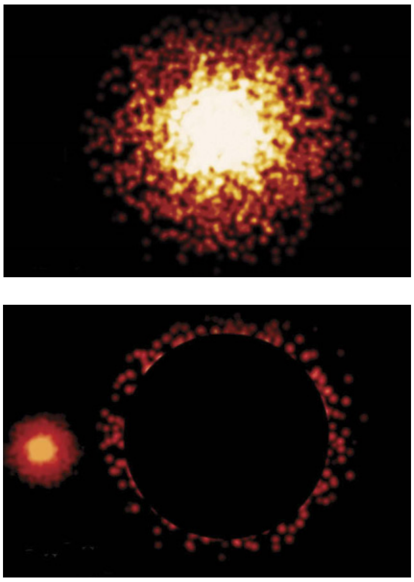

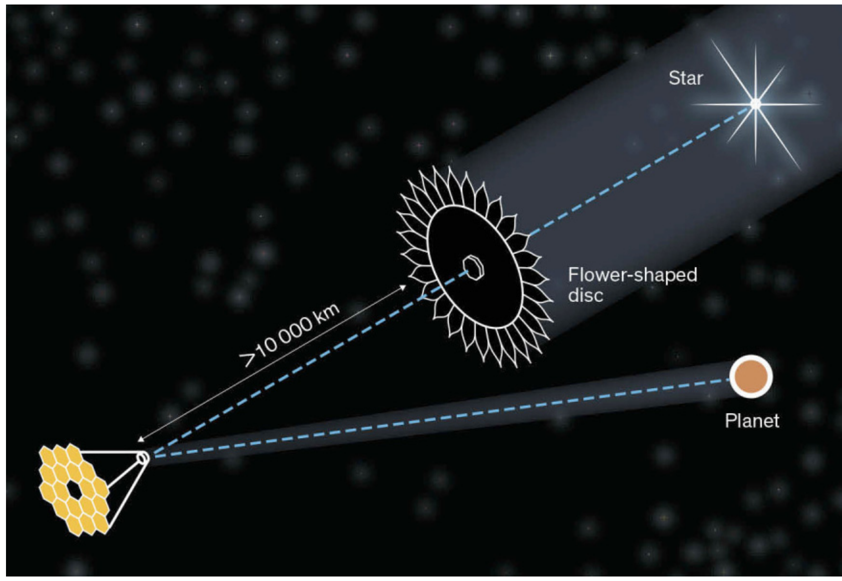

BLOCKING THE LIGHT: The view of a close-in exoplanet is easily overwhelmed by the light of its host star [top left]. A telescope can use an internal coronagraph to help reduce this light and reveal close-orbiting planets [bottom left].

FLOWER POWER: A future space telescope [above] might employ a specially designed starshade [model at near left] positioned tens of thousands of kilometers away to block a star's light so that a planet can be seen more easily.

stars. But it's still unclear how habitable M-star planets might be. The closer-in planets are likely to be bombarded early on with high levels of stellar radiation that could strip away their water. And tidal tugs from the star could lock the planet's rotation so that it always presents the same face to its star, as our moon does with Earth. JWST will give us a better sense of what such M-star planets are like.

Exploring rocky nearby worlds will be an evolving process. JWST will be able to measure the spectra of key atmospheric molecules, such as water, carbon dioxide, and methane, thereby giving us a clearer view of the atmospheres of those worlds.

Finding signs of life is another matter. What would constitute clear evidence? Water, for example, seems to be necessary for life as we know it, but the presence of water does not mean that there is life. Methane can be produced by microbes, but it can also be created by geological processes. High levels of oxygen might indicate the presence of plants, but it doesn't have to.

The more planets we characterize, the better we'll be able to understand the overall planetary population. Astronomers are already planning increasingly capable telescopes. Six years after JWST is sent to space, it will be joined in its orbit 1.5 million kilometers from Earth by PLATO, a European Space Agency spacecraft containing an array of 34 telescopes and cameras. PLATO is expected to find more transiting planets like Kepler-452b–Earth-size worlds that orbit in the habitable zone of sunlike stars.

But transiting planets represent just a small fraction of the planets in our galaxy–those that happen to orbit in such a way that they pass in front of their stars as seen from Earth. Many other worlds no doubt orbit their stars in ways that don't produce strong transit signals; the most extreme example would be one with an orbital plane perpendicular to our line of sight with the star, so that to us its orbit resembles a bull's-eye.

To widen the candidate pool and get a better, more direct look at each exoplanet, we'll need to take pictures of them. A sufficiently sensitive telescope will let us discern changes in brightness as the planet rotates, which could be due to the presence of continents and oceans. It might also let us see a "red edge," a spectral signature caused by plant life, which reflects strongly in the near infrared.

To date, direct imaging has been used to find only large, bright planets orbiting far from their host stars. Detecting the light reflected from Earth-size planets that orbit closer to their stars will be a challenge: Just one photon will be reflected for every 10 billion emitted by the planet's sunlike host. To image such a planet, a telescope will have to be sensitive enough to detect this small

TOP LEFT: JPL/NASA (2); BOTTOM: ALEX EVERS/NORTHROP GRUMMAN

ILLUSTRATION BY James Provost

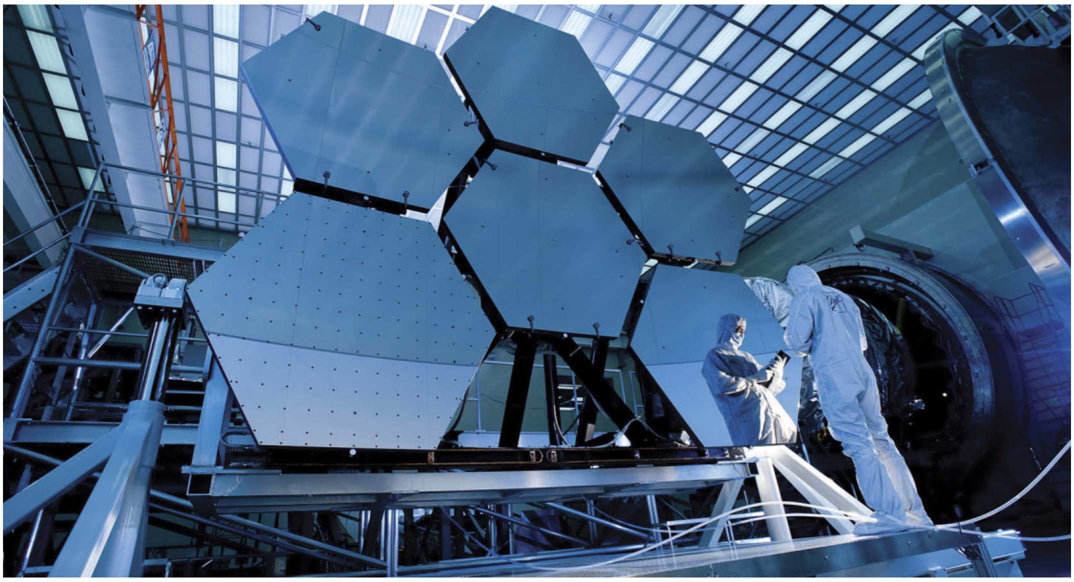

COLD TESTS: JWST's primary mirror consists of a foldable mosaic of 18 beryllium hexagons [six are shown here]. Each segment is coated with a thin layer of gold to optimize the mirror's ability to reflect infrared light. To keep their intrinsic glow to a minimum, the telescope's mirrors are designed to operate below -200 °C.

signal—and have enough resolution to separate star from planet.

Today's ground-based telescopes can easily resolve the angular separation between a small planet in the habitable zone and its star, but contrast is challenging. Internal devices called coronagraphs can block out starlight. But the turbulence in Earth's atmosphere makes it hard to create a perfect shadow, one that would block out all of a star's light but still let you see objects that are very close to it. Adaptive optics can help correct for that turbulence, but residual starlight will still overwhelm any light from small, close-in planets.

To see those planets, we'll have to get above Earth's atmosphere. A recent study, which extrapolates from our current knowledge of exoplanet populations, suggests we'd need a space telescope with a diameter of at least 12 meters in order to find and directly image about 30 Earth-size planets around sunlike stars.

Launching such a large telescope might seem daunting. But the lessons we've learned designing JWST could help us scale up. JWST's primary mirror is a mosaic of 18 lightweight hexagonal segments that fold up so that the telescope fits inside its launch vehicle. A larger telescope would just have more segments.

A large mirror will give the telescope the requisite resolution and sensitivity. To discern the planet, it will need to block out the bright light of the star. One way to do this is with a coronagraph, as mentioned. An alternative approach is to use a starshade, a large starlight-blocking structure that would fly in formation with the telescope but be positioned tens of thousands of kilometers in front of it.

Astronomers and engineers are weighing these two options. In order for a coronagraph to block out the light from a close-in exoplanet's star effectively, the telescope must be so stable that the position of an optical wave front drifts by no more than a few picometers—less than the size of a hydrogen atom—over the course of minutes. This level of stability is about 1,000 times that of JWST. Achieving it will require careful structural engineering as well as adaptive optics to counteract normal expansion and contraction due to the heat of the sun. Thanks to geometry, a distant starshade would need to stay aligned with the telescope to within only a matter of meters. Between observations, though, the starshade would have to be moved and realigned every time the telescope looked at a new star. Because of this, an internal coronagraph is likely to be the better option for an initial survey of a large number of stars.

Plans for a direct-imaging mission are still on the drawing board. Astronomers are talking about a 10-meter-class space telescope, one that might launch some time in the 2030s. Such a telescope could see changes as an exoplanet spins. In 2008, NASA's Deep Impact spacecraft looked back at Earth to test this approach. The spacecraft's camera could pick up variations in light reflected off land and water as the planet rotated.

But to truly image the surface of a distant world and identify what may be oceans and continents, we'll likely need an array of more than a dozen space telescopes. As is often done with radio antennas, these would be connected to form an interferometer that would combine data to create a telescope that is effectively as big as the largest distance between the components.

Such an array is likely many decades away. But there is still ample reason to be excited. The field of exoplanet studies is now poised to answer some of the most compelling questions humankind has ever asked: Are we alone in the universe? Are there other worlds out there that are capable of supporting life? JWST will give the first glimpses into what rocky alien worlds are like. Future missions will seek signatures of life on these worlds. One day, we could find our planet's true twin, ending the isolation that has held for billions of years, ever since life first appeared on Earth. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/exoplanet1115

THE VOLATILE FUTURE OF STORAGE

Speedy DRAM could replace hard disks for cloud computing

WHEN IT COMES TO COMPUTER STORAGE, the magnetic disk has been top dog for almost half a century. The first commercial disks appeared in 1956, and by the early 1970s their cost and capacity had improved to the point where they began to replace magnetic tape as the primary storage medium for computers. By the end of that decade, tapes had been relegated mostly to a backup role. Since then, disk technology has improved at an exponential rate, just like integrated circuits. Nowadays, a typical drive holds 20,000 times as much data as it did in 1985, and on a per-byte basis, disks cost one-millionth of what they did then.

No wonder hard disks are so pervasive. This is also why today's popular forms of computer storage, such as file systems and relational databases, were designed with disks in mind. Indeed, until recently any information kept on a computer for more than a few seconds probably ended up on disk. ● But the hard disk's reign may be coming to an end. The most obvious challenger is flash memory, which is faster, more compact, and more resistant to shock. Virtually all mobile devices, such as tablets, smartphones, and watches, already use flash instead of disk. Flash memory is also displacing hard drives in laptops and, increasingly, in large-scale applications running in data centers, where its speed is a significant advantage.

by JOHN OUSTERHOUT

ILLUSTRATION BY Greg Mably

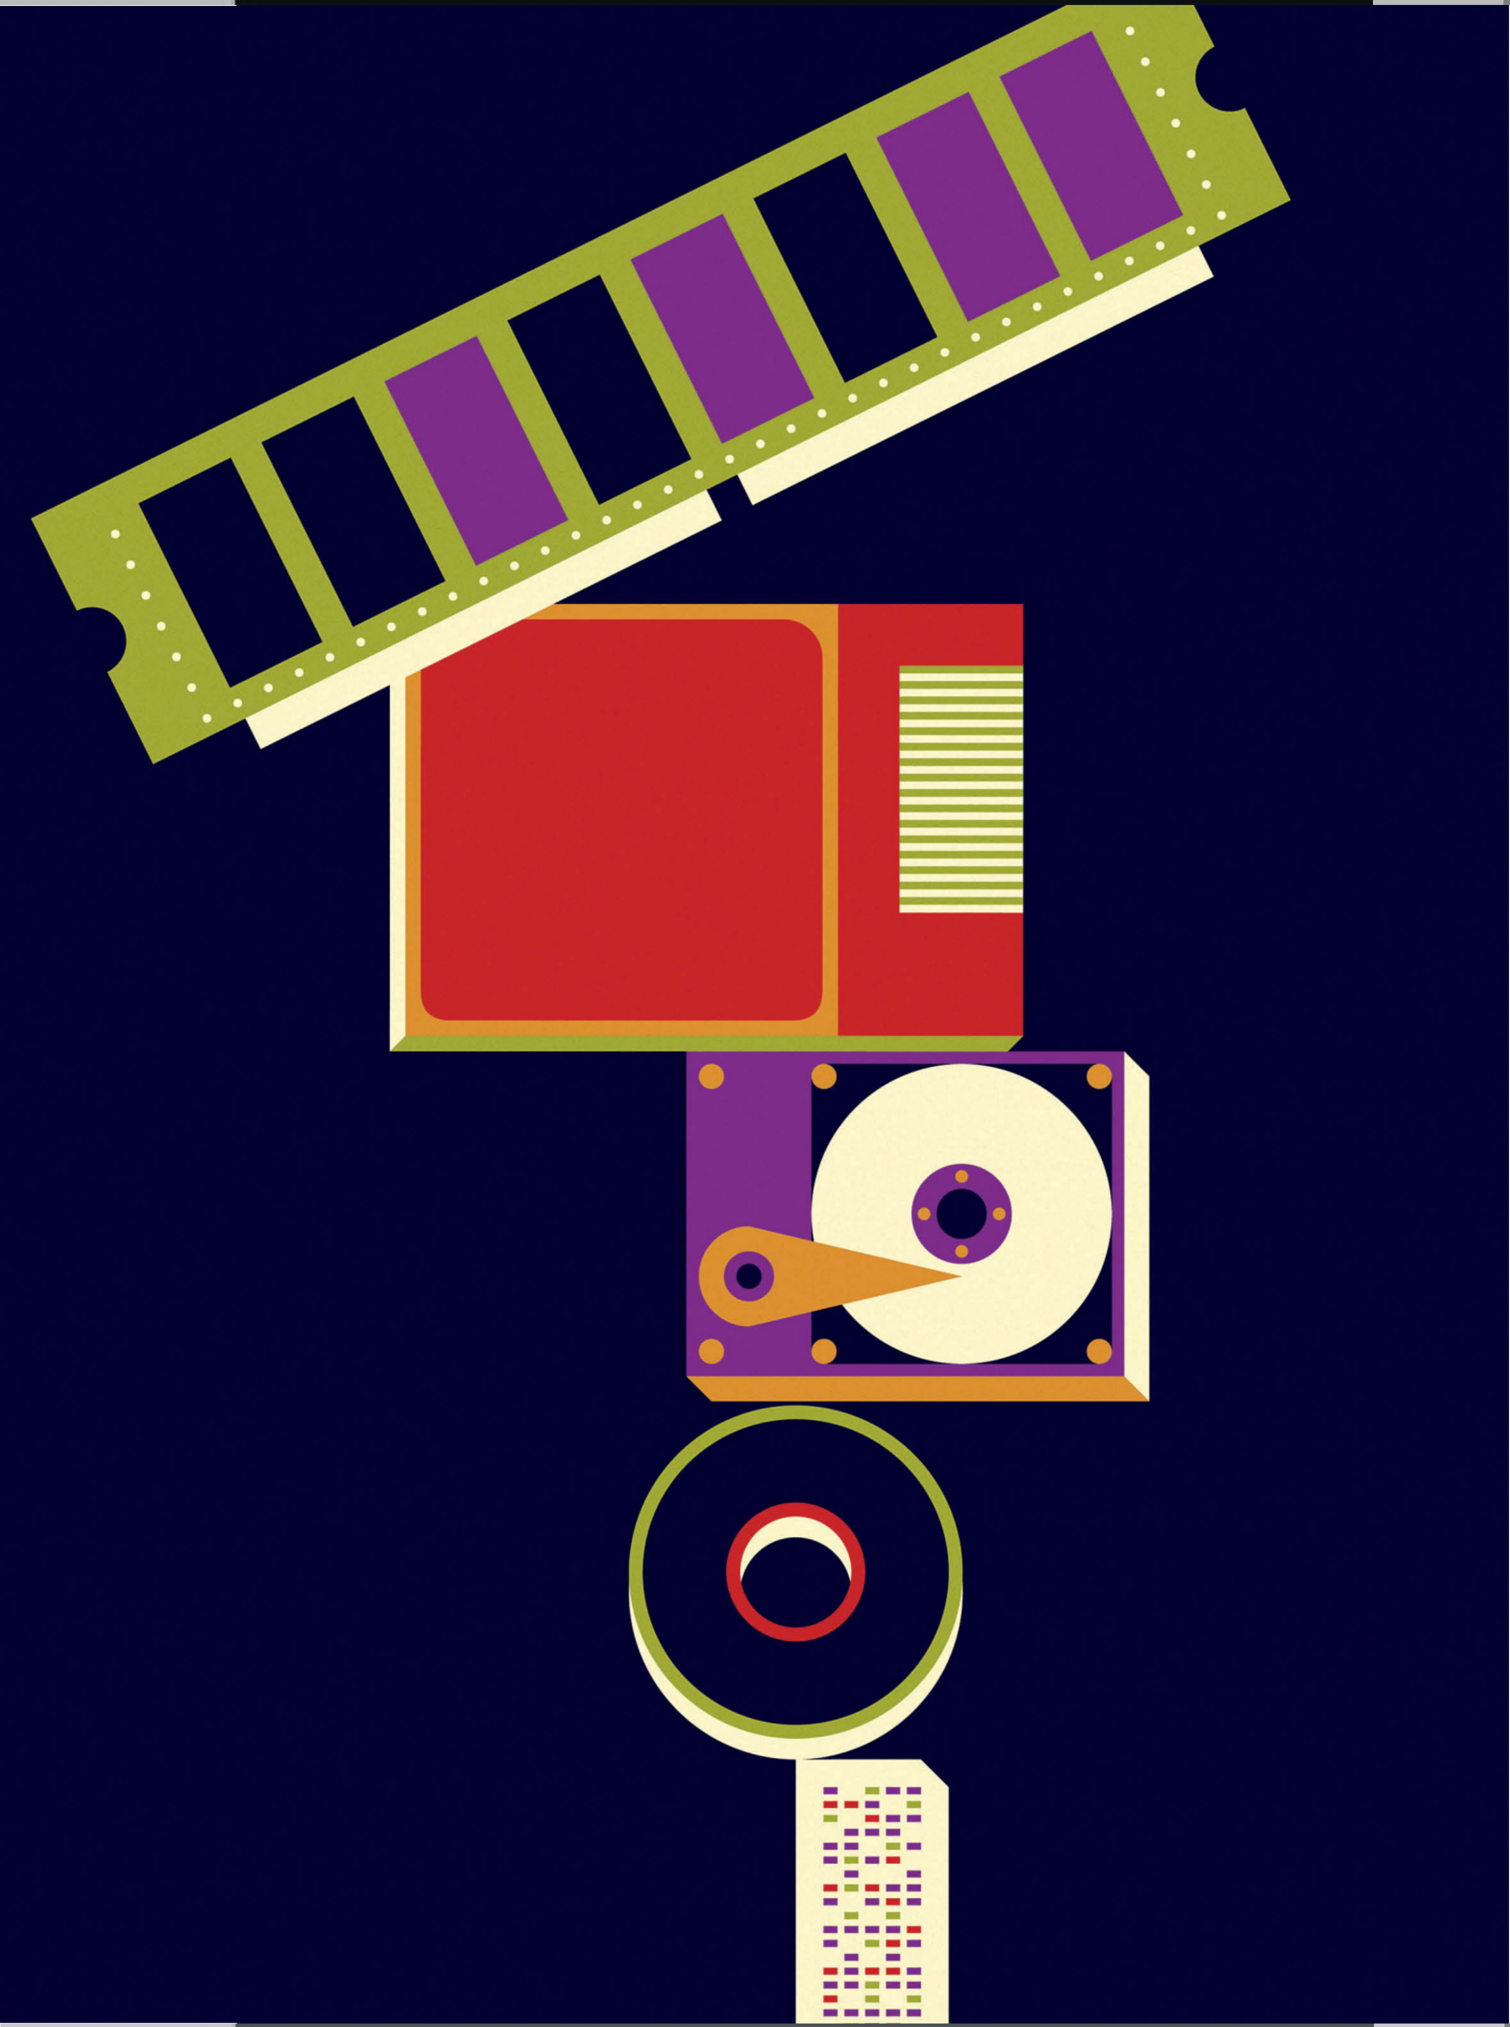

Now though, there is yet another alternative to disk: using dynamic random-access memory (DRAM) as the primary storage location for long-lived data. More and more applications, particularly large-scale Web applications, are keeping most or all of their data in DRAM. For example, all of the popular Web search engines, including Google, service people's queries entirely from DRAM. Also, Facebook keeps most of its social-network data in DRAM. And IBM's Watson artificial-intelligence system kept all of its data in DRAM when it won the "Jeopardy!" challenge a few years ago.

On the surface, this seems ridiculous. After all, DRAM was intended to hold information temporarily during active computations. Although it is about 1,000 times as fast as flash, it is also 100 times as expensive as disk, and it is volatile, which means that the data it holds will disappear if the computer loses power. Nevertheless, I believe that DRAM could soon become the primary storage medium for large-scale applications running in data centers. Here's why: If DRAM is backed up on disk or flash, users can enjoy the medium's huge speed advantage without worrying that data will be lost during the inevitable server crashes and power outages.

As a research project, my colleagues at Stanford University and I have constructed a general-purpose storage system we call RAMCloud, which keeps all of its data in DRAM at all times. RAMCloud aggregates the DRAM memories of a collection of servers—potentially hundreds or thousands in a typical data center. To work around DRAM's volatility, RAMCloud stores copies of data on disk or in flash memory, and it automatically recovers data from those backups after a server crashes. Our hope in the RAMCloud project is to make it as easy for developers to use DRAM-based storage as it is for them to use disk.

OUR MAIN MOTIVATION FOR MOVING FROM disk to DRAM comes from the evolution of disk technology. Although the storage capacity of disks has mushroomed over the years, their access speed has not improved as much. The reason for

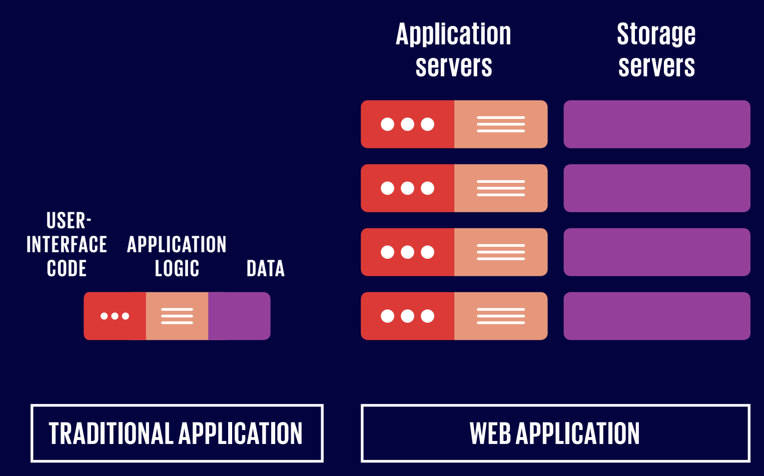

WEB-SCALE APPS

Traditional applications manipulate data held in the memory of a single machine [left]. But large-scale Web applications require a more complicated architecture [right], with many different computers in a data center running the application logic and the server-side user-interface code. Other computers act as storage servers. They keep the bulk of the data on disk with only the most recently accessed information held in memory.

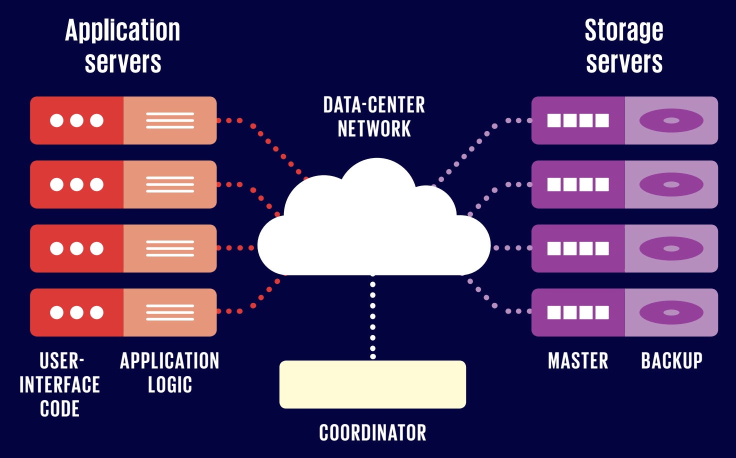

UP WITH DRAM

RAMCloud's storage servers have lots of DRAM (from 32 to 256 gigabytes on each machine)—enough in aggregate to hold all of the data. Those storage servers also contain hard disks and a small amount of battery-backed DRAM, together used for backup copies of data from other servers. A separate machine, called the coordinator, manages the communications between application servers and storage servers.

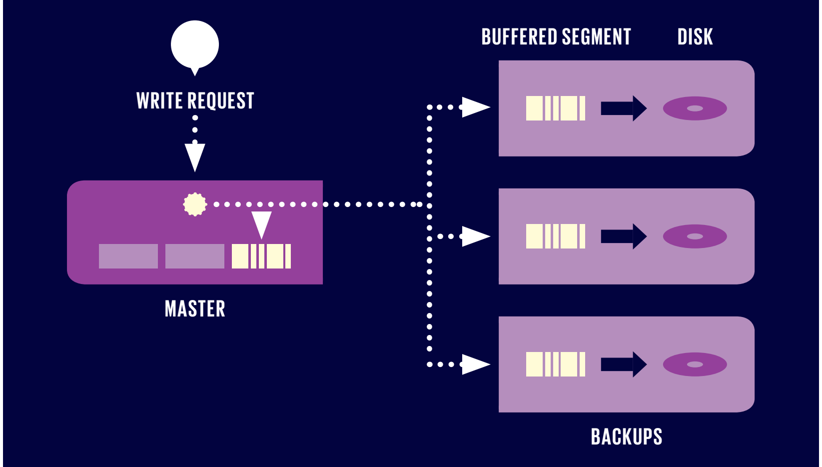

WRITE THREE DEEP

When a RAMCloud storage server receives a request to store a chunk of data, it adds the information to its master copy of the data in DRAM, which takes the form of an in-memory log. That storage server also backs up the newly recorded information on at least three other storage servers, which save redundant copies of the data, initially in battery-backed DRAM and then on disk.

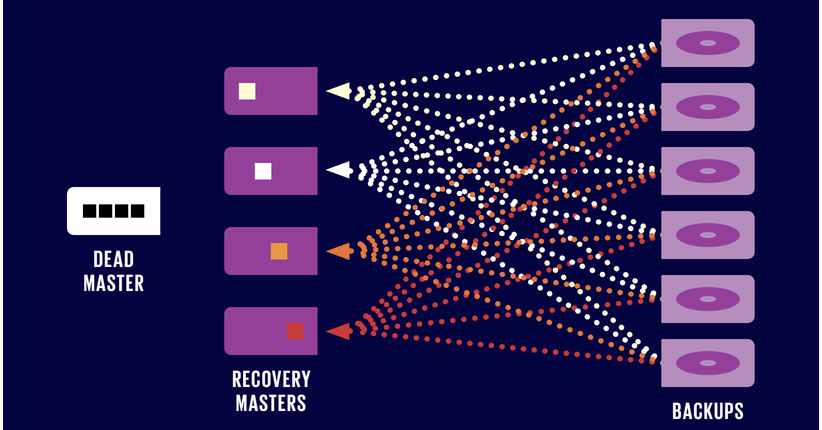

CRASH 'N' BURN

When a storage server crashes, the data it held in DRAM must be reconstituted—fast. For that, several existing servers are assigned to act as *recovery masters*, restoring different portions of the lost data and taking ownership of it. Each recovery master reads from multiple disk backups, speeding the recovery process. With enough servers working together, hundreds of gigabytes can be reconstituted in just a second or two.

this is simple enough. To read or write data on a hard disk, a physical mechanism must first position the read-write head over a particular track on a spinning platter. Then the system must wait until the desired information rotates under the head. These mechanical processes have proven difficult to speed up much. With the amount being stored on a disk increasing more rapidly than the access rate, the time needed to read or write all the data on a disk has increased. Indeed, if the information you need is stored in small blocks spread randomly on the disk (which is common in many applications), the time it takes to hunt down the data becomes significant. Even if you dedicated your computer to that task, to access each block in random order could take a few years!

This problem is making disks increasingly unsuitable for storing small pieces of data that must be accessed frequently. In a sense, the hard disk is putting itself out of business. This trend has been evident for many years and explains much of the interest in flash memory, with its higher access rates.

There is a second motivation for new storage technologies that's even more compelling: the rise of large-scale Web applications. These applications sometimes need to support tens or even hundreds of millions of users. Doing that requires a radically different operational structure than what has been used traditionally.

The old approach was to load the application's code and all of its data into the main memory of a single machine. The application could then access its data at DRAM speeds (typically 50 to 100 nanoseconds), which allowed it to manipulate its data intensively. The problem was that the total throughput of the application was limited by the capabilities of that one machine.

As the Web grew in popularity during the late 1990s, it quickly became clear that the traditional application architecture could not handle the loads generated by popular websites. Over the next 10 years, designers came up with a new architecture in which thousands of machines work together in huge data centers. One of the most important attributes of this arrangement is that application code and data are separated. One set of machines (called application servers) runs the applications, while a different set of machines (storage servers) holds the data. In this architecture, data is typically stored on disk, although frequently needed data may be held for quick access (cached) in DRAM.

This new architecture has enabled the creation of applications of a scale that was unimaginable 20 years ago–Facebook and Google Search come to mind–but it has also changed the relationship between an application and its data. In the Web architecture, each read or write access requires an application server to communicate over the data center's internal network to a storage server, and the storage server may additionally have to perform a disk access. The total time to fetch a chunk of data is now 0.5 to 10 milliseconds, four to five orders of magnitude longer than what it would be with the traditional one-machine-does-everything architecture.

As a result, Web applications cannot use their data very intensively. For example, when Facebook generates a Web page for a user, it can consult only a few hundred distinct pieces of data: Anything more would take intolerably long. This restriction limits the kinds of features that Facebook can provide, and it makes life difficult for the company's developers, who must labor to pack the most information possible into the smallest number of distinct chunks.

OUR AWARENESS OF THESE PROBLEMS spurred my Stanford colleagues and me to start the RAMCloud project in 2009. We figured that if we could speed up data access, it would make an enormous difference. Facebook is one example of an application that could benefit, but we believe there are many more things that are not even attempted today because no storage system can support them.

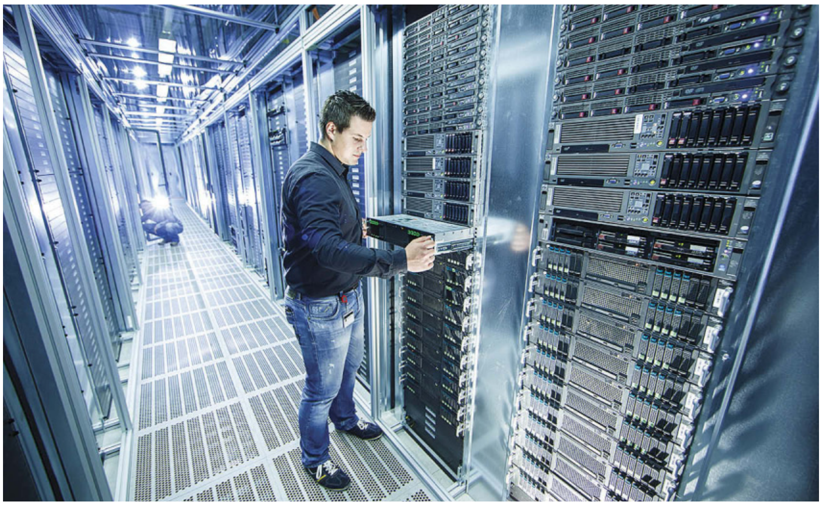

DESIGNED FOR SWAPPING: Servers can be replaced quickly. But even so, if DRAM is used for long-term storage, crashes remain problematic. RAMCloud automates data recovery, restoring within just a couple of seconds the information lost when a server fails, so users don't notice.

THOMAS TRUTSCHEL/GETTY IMAGES

To achieve that speed boost, RAMCloud keeps all data in DRAM at all times, but it also makes the data just as reliable as if it had been stored on disk. RAMCloud doesn't require any sort of exotic hardware to do so. Indeed, it's just a software package that runs on ordinary machines–RAMCloud servers, which take the place of today's storage servers.

Each RAMCloud server contains two software components: a *master* and a *backup*. The master code uses most of the server's DRAM to store RAMCloud data, and it makes that data available to the applications servers that request it. The backup code's job is to keep redundant copies of the data from other RAMCloud masters in the server's secondary storage (on disk or in flash memory).

A third software component of RAMCloud, the *coordinator*, runs on a separate machine. It manages the overall configuration of the RAMCloud cluster and coordinates recovery actions when storage servers fail. The various application servers running in the data center use a RAMCloud library package to access RAMCloud storage. The first time they do so, those application servers must contact the RAMCloud coordinator to find out which masters store which data. The application servers then cache this information locally, so subsequent RAMCloud requests can go directly to the relevant master.

While the RAMCloud servers are intended to take the place of normal storage servers, they can't entirely replicate what today's database servers do. RAMCloud's data model is a key-value store, which consists of a collection of tables, where each table can contain any number of objects. Each object consists of a variable-length key, which is unique within its table, and an associated value, which can be a blob of arbitrary data up to 1 megabyte in size. Applications read and write objects on a RAMCloud server simply by specifying a table and a key within the table. RAMCloud does not currently support all of the features of a database system, such as secondary indexes and transactions.

While some will see this feature deficit as a shortcoming, duplicating the functionality of today's database servers was not our goal. Our most important design imperative for RAMCloud was to achieve the fastest possible access times, defined as the total time, measured by a running application, to read or write small chunks of data that are stored in a RAMCloud server in the same data center. With the current version of RAMCloud, reads take about 5 microseconds (writes take about 14 μs) for 100-byte objects. That's roughly 1,000 times as fast as you could read data from a local hard disk and 10 to 20 times as fast as you could get that data from local flash memory.

When we began the RAMCloud project, access times like these were almost inconceivable. Even with data stored in DRAM, the networking infrastructure itself imposed significant delays. Each switch in the data-center network added 10 to 30 μs, and the request for data, and the supplied response, could each have to go through as many as five switches in a large data center. And data packets had to be processed by the operating system when entering or leaving a machine, which added roughly 60 μs. So typical round-trip times within a data center were hundreds of microseconds long.

Two improvements in networking have made RAMCloud's blazing speed possible. The first was a new kind of networking infrastructure that uses special-purpose switching chips with internal delays of less than 1 μs per switch. The second was a new generation of network-interface controllers (NICs), which have a capability called kernel bypass. This feature allows an application to communicate directly with the NIC to send and receive packets of data without involving the computer's core operating system.

The combination of fast-switching chips and kernel bypass makes ultralow-latency communication possible: less than 5 μs for requests that need to pass through only a few switches (such as in our test cluster), and around 10 μs for requests in a very large data center with 100,000 machines. With the additional improvement in networking technology we expect to see over the next 5 to 10 years, we believe that round-trip times in a large data center could be reduced to as little as 2.5 μs.

IN ADDITION TO PROVIDING FAST ACCESS, RAMCloud must also ensure that its data is stored as reliably as if it were held on disk. In particular, data must not be lost when a server crashes or the power goes down. Data centers typically lose power every few years, which can cause all the information in DRAM to be lost. As mentioned earlier, RAMCloud keeps backup copies of data on disk or in flash memory. (This scheme is analogous to the way that current disk-based storage systems keep backup copies on magnetic tape.)

But what if an individual RAMCloud server failed? To protect against such mishaps, RAMCloud keeps multiple backup copies (typically three) of every single piece of data, storing those on different servers. So when a write request arrives at a master, it updates its information in DRAM and then forwards the new data on to several backup RAMCloud servers. This need for replication explains why writing data takes longer than reading it.

In addition to these general measures, we had to solve two specific problems to make RAMCloud's backup strategy really bulletproof. The first was what to do if a server loses power before it has made backup copies. RAMCloud deals with that possibility using a small amount of nonvolatile memory on each machine. When new data is written to one machine, the backup machines associated with it collect that data temporarily in some form of fast but nonvolatile memory. That could take the form of battery-backed DRAM, for example. These backups then write the data to disk or flash in the background. Having the data in fast nonvolatile memory at the start ensures that the information can be recovered if a power failure occurs before the data is written to disk or flash.

The second problem is that a RAMCloud cluster with thousands of servers will inevitably experience frequent server crashes. RAMCloud keeps only a single copy of information in DRAM, so data that was stored on a crashed server will be unavailable until it can be reconstructed from information on the hard disks (or flash memories) of its backups. If all the crashed machine's data were held on one other disk, it would take several minutes to get it into working memory.

To avoid that long a delay, RAMCloud scatters the backup data for each master across all of the other servers in the cluster, which could amount to thousands of machines. After a crash, all of those machines work in parallel to reconstruct the lost data. With the work spread among so many computers, RAMCloud can recover from server failures in just 1 to 2 seconds—so fast that most Web users wouldn't even notice.

ALTHOUGH RAMCLOUD IS JUST A UNIVERSITY research project, our goal was to create a production-quality system that can be used for real applications. At the beginning of last year, the system reached version 1.0: The basic features were in place, and the software was mature enough for early adopters to begin experiments. Since then, several groups outside Stanford have begun experimenting with RAMCloud's open-source code, for such things as the Distributed Main Memory Database project at telecom giant Huawei.

One possible impediment is that RAMCloud does not support the relational data model that dominates information technology today. It does not provide the convenient facilities of a relational database management system, so many applications would have to be reprogrammed to use it. But that's perhaps not as big an issue as it would appear. The applications for which RAMCloud would be most advantageous haven't yet been written, simply because there is no storage system capable of supporting them.

So RAMCloud's early adopters are not likely to be those using existing enterprise applications. Instead, the first users will probably be developers who discover that today's storage systems are just not fast enough to meet their needs. These groups can design their applications around RAMCloud from the start, and they should be more willing than most to try it out. Desperate people are, after all, the friend of new technology!

If RAMCloud proves successful with its early adopters, it may not be long before it spreads into the mainstream. RAMCloud is a good match for cloud-computing data centers such as Amazon Web Services, Google Cloud, or Microsoft Windows Azure. In such an environment, the cost of a RAMCloud cluster could be amortized across many users, who could then take advantage of high-speed storage at very low cost, while retaining the ability to expand easily as their needs grow.

RAMCloud represents a new class of storage that simultaneously achieves vast capacity and low latency. It would let something like a million CPU cores within a data center work together on data sets of 1 petabyte or more, where any core can access any data item in 5 to 10 μs. That's one-thousandth of the time it typically takes now. So if you're already impressed by the speed and power of today's large-scale Web applications, fasten your seat belt and get ready for even bigger thrills in the future. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/ramcloud1115

CROWDSOURCING MUSTERED BOTH THE FUNDS AND THE DEVELOPER COMMUNITY NEEDED TO BUILD THE OPEN-HARDWARE NOVENA COMPUTER

BY **ANDREW "BUNNIE" HUANG** & **SEAN CROSS**

A Laptop With No Se

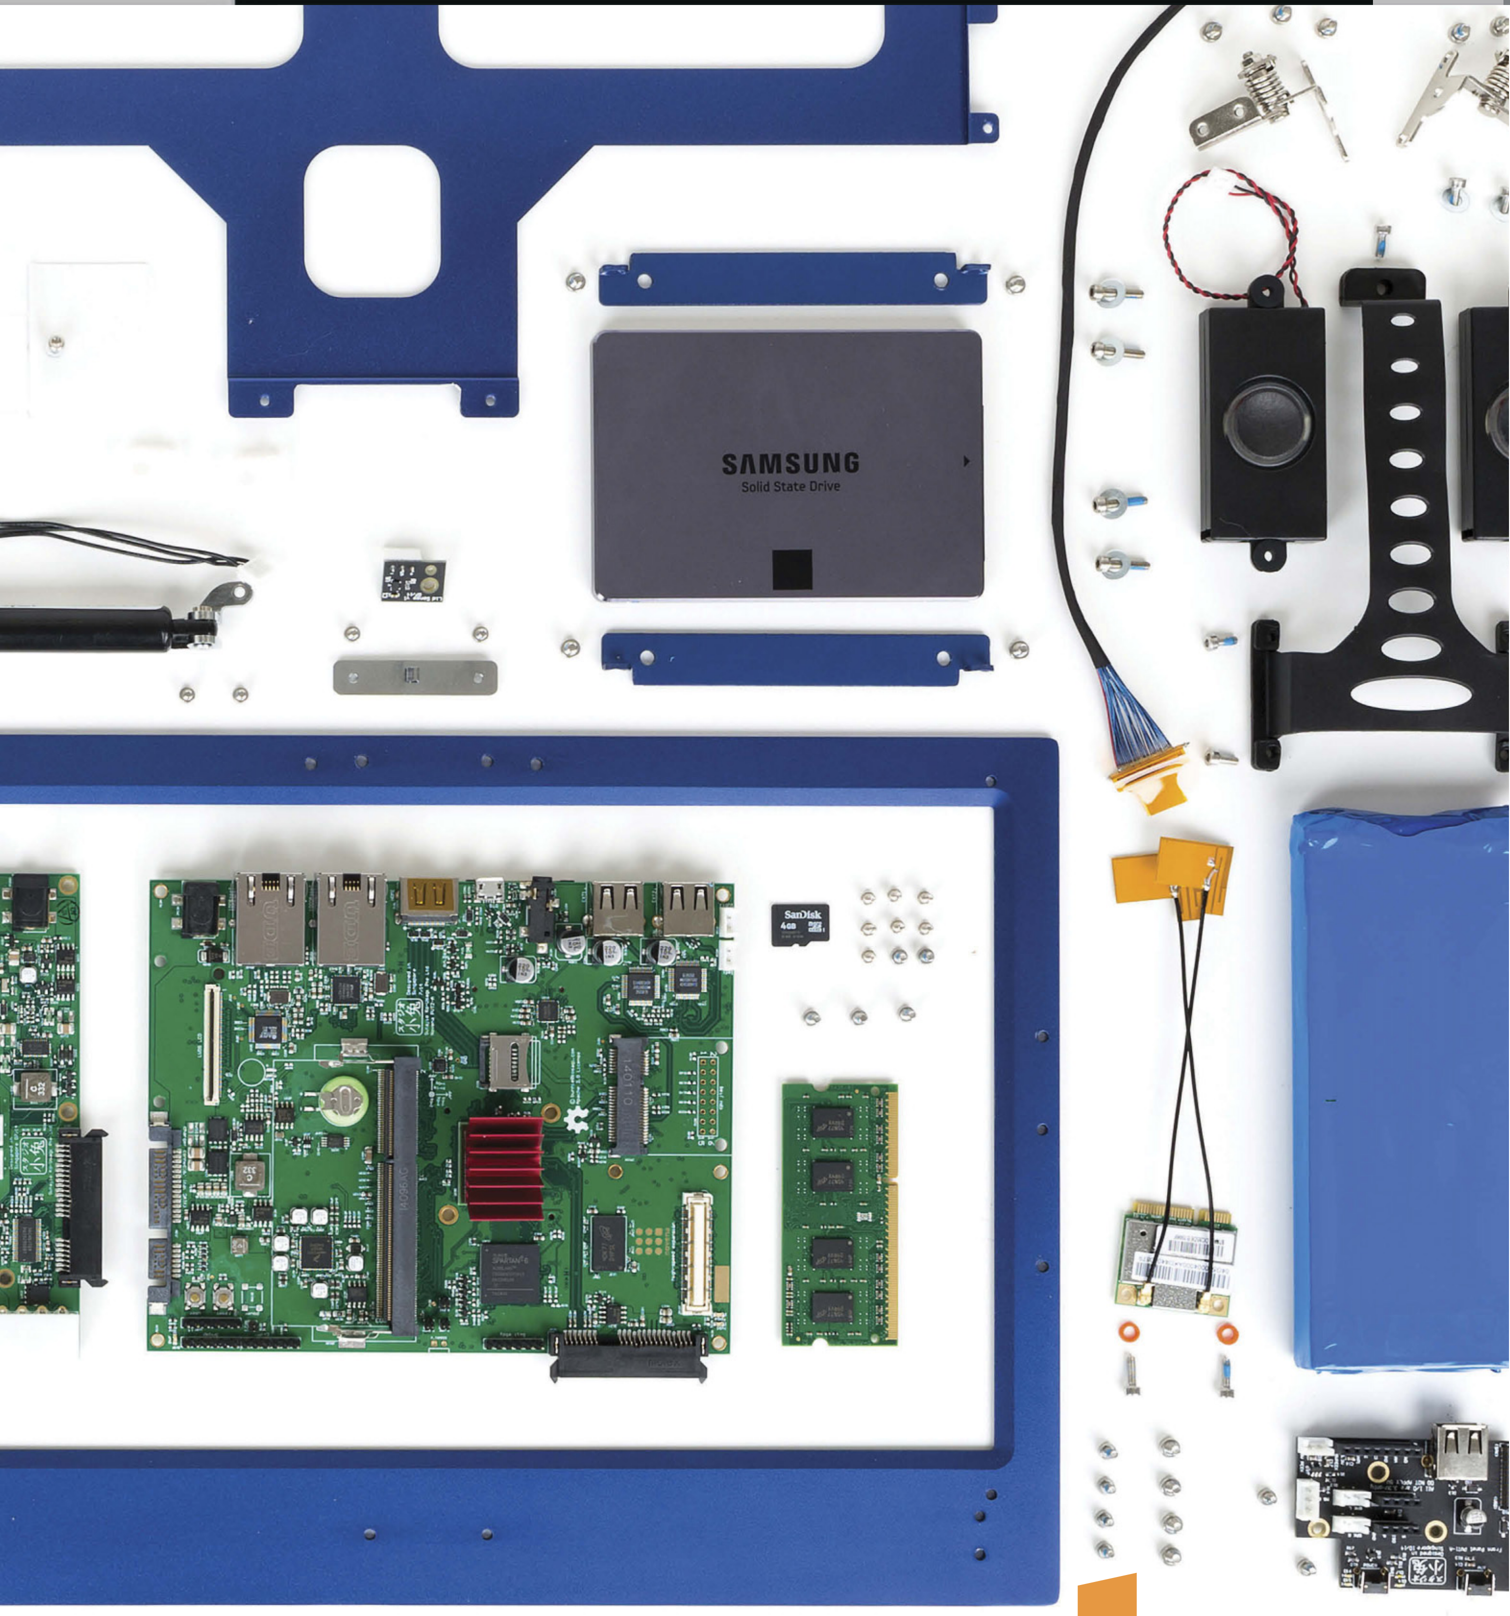

ecrets

Has the computer become a black box, even to experienced electrical engineers?

Will we be forever reliant upon large, opaque organizations to build them for us? Absolutely not, we say. And to prove our point, we built our very own laptop, from the circuit boards on up.

Admittedly, we did not delude ourselves that we could build a laptop that would be faster, smaller, or cheaper than those of Apple, Dell, or HP. However, we did set out to build a machine powerful and convenient enough to use every day. Fortunately, our dream inspired enough people to crowdfund the effort. Our laptop, which we call Novena, started shipping to backers in January 2015.

Events favored our quest. Because Moore's Law is slowing down, garage innovators can take a couple of years to develop a high-tech gadget without discovering that the completed version is obsolete. It has been three years since we started the Novena project and yet the 40-nanometer process on which our central microprocessor is based continues to occupy a sweet spot between cost and performance.

Also, the economic malaise of 2008 left a lasting mark on global supply chains. Even today, manufacturers are no longer too busy printing money to make the time for producing small, boutique projects like ours; instead, they see us as an opportunity to gain an edge over their competition. Owners of small factories in China turned out to be eager to learn our agile approach to engineering, so they, too, could take on the challenge of low-volume production and address the growing market of hardware startups.

We started by considering the constraints of the most complex and brittle pieces of any such system: the software. We wanted to build a complete computer, one with a long-term-support road map, and we had neither the resources nor the manpower to negotiate with vendors of proprietary hardware and software. We wanted to be able to inspect and understand as much of the system and its components as we could, so if we came across bugs or other anomalous behavior, we could rely on our wits to figure it out, rather than on the profit-motivated (and often empty) promises of a vendor's sales team. As a result, we decided to produce a laptop that was as free as possible of closed-source embedded firmware.

Firmware is basically software (such as drivers, kernels, and bootloaders), installed at the factory, that runs on "bare iron"—the computer itself, not its operating system. It's found not only on the main CPU in your laptop but also on about a dozen embedded controllers—small, special-purpose processors that take care of such things as managing the battery, keeping your hard disk free of errors, and maintaining your Wi-Fi connection. Each of these processors runs a bit of firmware; sometimes the firmware can be updated or modified, typically for the purpose of fixing bugs or adding features. But such updates can also introduce bugs or security flaws, and if you haven't got the ability to inspect the firmware, then you, the user, must depend utterly on the vendor to take care of security.

This open-source requirement of ours ended up influencing the selection of almost every piece of hardware, including the main CPU, the battery controller, and the Wi-Fi module. For example, we couldn't use Intel's *x*86 microprocessors because they can accept firmware updates that we cannot debug or inspect. Instead we chose an ARM-based Freescale i.MX6 system-on-a-chip, which has no such updatable code embedded. (A system-on-a-chip, or SoC, is similar to a microprocessor except it has more of the supporting hardware, such as memory and peripheral interfaces, needed to make a complete computer.) The i.MX6 does have some code burned into it to coordinate the computer's boot-up process, but this firmware can't be changed,

LAPTOP, OPEN WIDE: The Novena opens to reveal the circuit boards; when it's closed, the outward-facing display lets you use it like a tablet.

ALL PHOTOS BY SCOTT TORBORG/CROWD SUPPLY and its unencrypted binary code can be read out and analyzed for possible security problems.

Another advantage was Freescale's policy of distributing a very detailed reference manual covering most of the chip's real estate without requiring a nondisclosure agreement. That's important because such an agreement would have gotten in the way of the community involvement that our strategy for long-term support required.

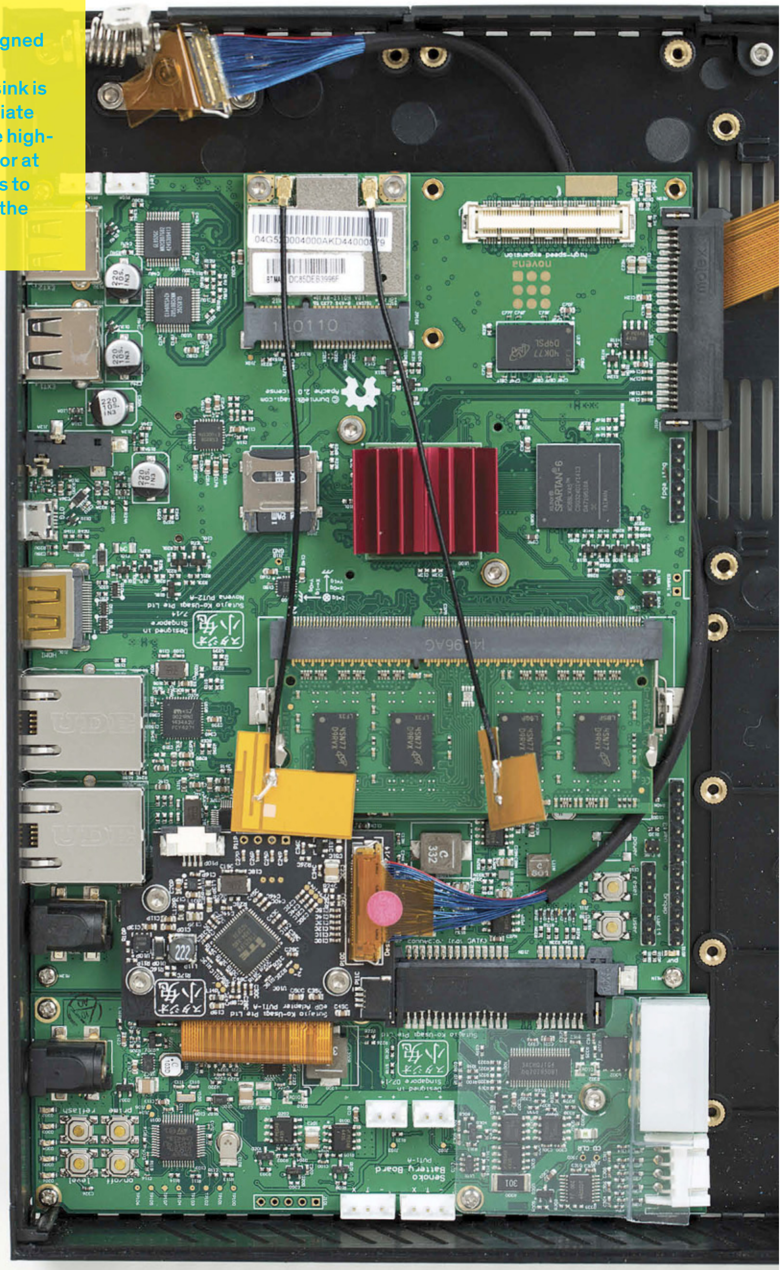

THE NOVENA'S HEART: A view of the custom-designed internal circuit boards. Underneath the red heat sink is the CPU, and to its immediate right, the FPGA. The white high-speed expansion connector at the top right enables users to easily expand and modify the hardware's capabilities.

Our next choice had social repercussions. When you adopt a CPU/operating-system combination, you also adopt its developers. We decided against Google Android because it's optimized for phones and tablets, its graphical display typically shows only one application at a time, and its touch-screen paradigm is too imprecise for computer-aided design work. Therefore, in order to create a system that our target market of developers and creators could use, we decided to run on our ARM chip a version of Linux called GNU/Linux. GNU, which authored both the OS libraries and the license that the Linux kernel uses, is a coder's organization, right down to the self-referential acronym itself (it stands for "Gnu's Not Unix").

Unfortunately, most Linux versions for ARM were not designed for personal computing but rather for routers and the set-top boxes that convert signals for viewing on televisions. Manufacturers that use Linux generally build their own, highly customized system around what's called the Linux kernel–the core framework of the OS. That way, they ensure that their customers, who integrate the software into a larger system, never see so much as a Linux command prompt. Examples include TiVo and airline in-flight-entertainment systems. On top of that, vendors use an old "snapshot" of Linux–that is, they copy whichever kernel is around at the time of chip release. Then they add a lot of patches to it to create a board support package (BSP) for their product, which is a kind of "quick start" kit for hardware developers who don't want to muck around too much with Linux kernel minutiae.

Frequently these BSPs contain proprietary bits that developers can't copy, that don't take advantage of features that came with the original kernel, and don't even follow the standard formatting procedures for Linux code. These undisciplined shortcuts may speed a product to market, but they make the kernel impossible to maintain, update, or improve. As a result, most systems that are based on Linux run kernels that are years out of date. This is not a plan for long-term customer support.

Therefore, we adopted the latest kernel available. We patched and extended Freescale's open-source BSP code to comply with Linux community standards, then submitted our changes to the managers of the Linux community at large. They reviewed our code and provided feedback over several months of iterations.

At that point, our code was finally ready to be "upstreamed" into mainline Linux. It's a time-consuming and laborious process, but when it's over the Linux community can adopt our version and maintain it. The main advantage of this upstreaming process is that features that come with future Linux kernels can be unlocked by simply downloading the latest version of the operating system and compiling it, without having to repeat the code-review process.

NOVENA'S PEEK ARRAY: The brass nuts in the open area at the left constitute the Peek Array, a number of points on which users may attach expansion boards of their choice. Standard-issue laptops lack this trick for making the hardware more hackable.

Hardware-accelerated graphics proved to be another challenge. Most modern desktop computers assign each application or graphical widget an off-screen memory buffer to render their output on the screen. Only then does the graphics processing unit (GPU) "composite" the various outputs to make the seamless image seen on the screen.

This two-step process decouples the application's drawing rate from the screen's refresh rate, preventing a "race condition," in which two input signals compete to be the first to affect the output. Such racing can produce artifacts such as "tearing," in which the attempt to drag a window ends up moving only the top part by the time the screen refreshes. The newly redrawn half of the window closest to the cursor thus looks torn off from the lower half. Also, the two-step decoupling process can be used to create subtle cues, such as transparency and blur, where foreground windows and title bars can take on the character of frosted glass, providing a hint of what content might lie underneath them.

Unfortunately, the integrated GPU on the Freescale i.MX6 chip is closed source. Remember how we said that the vendor-supplied BSPs are typically snapshots of older kernels? The code that Freescale provided for drivers—that's the software that controls various peripheral devices, like printers and GPUs—is compatible only with an out-of-date version of Linux. And this driver has an unfortunate quirk: It requires any application that talks to it to do all floating-point calculations in software! That's far slower than doing the calculations in hardware.

This example nicely illustrates why closed-source products can be so frustrating. If we'd used this closed-source driver, we would have locked ourselves into an outdated version of Linux and forced other applications to do floating-point calculations in software. But if we'd gone open source to exploit all the features and benefits of the latest Linux version and to allow floating-point calculations done in hardware, users would have had to reverse-engineer the GPU to create compatible open-source drivers.

We decided to stick to our guns and reject the closed-source GPU altogether, which meant we would have to render graphics in software.

Unlike Windows and Mac OS, Linux gives you a choice of how your computer appears to the user. This appearance is governed by a software utility called a desktop window manager. You can select this appearance by picking a particular "distribution" of Linux and coupling it with a particular window manager. We couldn't use the popular Ubuntu and RedHat distributions because they basically require GPU hardware acceleration, which we had rejected. Instead, we used a slightly older-looking but more open-source-friendly distribution called Debian. We coupled it with a desktop window manager option called Xfce4, which is explicitly designed to run well on systems without GPUs and thus is particularly good at software rendering.

But that arrangement, we hope, is merely a stopgap. The user community behind Novena is trying to create, through reverse engineering, open-source drivers that would allow the built-in GPU on the i.MX6 chip to render graphics directly. At the moment, we have drivers that can accelerate the drawing of 2-D figures on the screen, which makes scrolling and window dragging much smoother. We hope eventually to figure out enough of the GPU to let us do 3-D graphics with acceleration sufficient to produce a user experience much like that of any mainstream laptop.

As the software evolved, the hardware evolved with it. The most extraordinary step we took was to include a field-programmable gate array (FPGA), a type of processor chip that can be reconfigured by its user to change the chip's specs and capabilities. Basically, this reconfigurability allows the chip to do things in hardware that would otherwise have to be done in software.

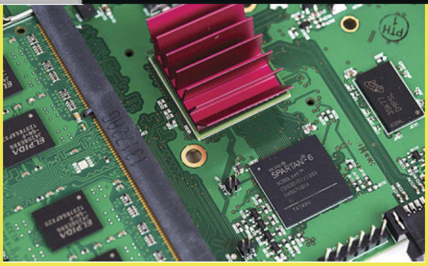

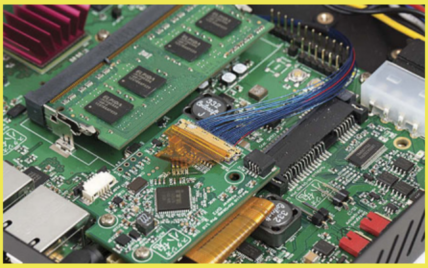

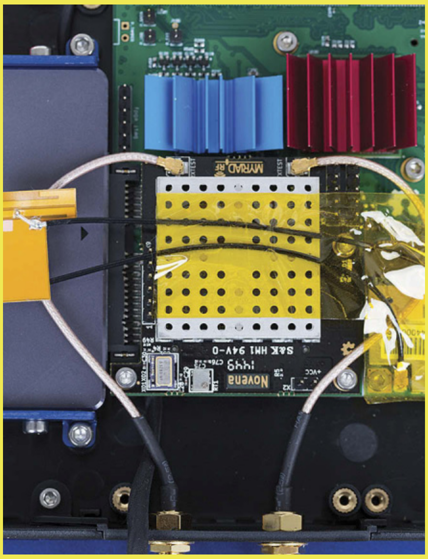

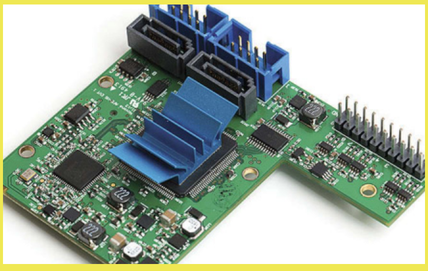

SOUP TO NUTS: At top are the CPU's red heat sink and the FPGA; next is the LCD signal adapter [lower left] and the battery board [lower right]; next ,the optional software-defined-radio module; then the port farm; at bottom is the digital sampler expansion card, which can turn the laptop into an oscilloscope.

For instance, if you wanted to accurately control a dozen motors, you'd have a lot of trouble doing that in software. Motion control requires exquisite real-time control over waveforms, and if you try to get your operating system to do that it would constantly be shifting from one task to another and back again, adding too much timing jitter. Imagine that stutter you see in your Web browser but in a self-driving car or a drone: instant wreck. However, it's a relatively simple matter to create hardware that runs the fussy timing and key control loops, and it's also trivial to replicate variations of that hardware again and again in an FPGA.

And at one point, we did in fact want to control a lot of motors. We had the wacky notion of mounting Novena in a quadcopter chassis so that our laptop could hover and follow us around the office; we figured we'd use the FPGA to interface with the requisite motors and sensors. Consequently, the FPGA's input-output connectors were originally targeted toward servos and motors.

Cooler heads eventually prevailed. We instead optimized the FPGA for data acquisition by adding a bank of local high-speed DDR3 RAM chips, a form of high-speed dynamic access memory. This upgrade allowed developers to implement a variety of expansion cards, including one for software-defined radio and another that worked as a digital sampler, which could allow engineers to use our laptop as a portable oscilloscope.

Because we developed the project on a shoestring budget and with no solid requirement other than to build the laptop we ourselves would want to use, the laptop evolved organically. Instead of designing circuits into the motherboard specific to a particular battery or LCD–typical for most consumer laptops–we modularized the system. We linked subsystems using generic connectors with pin-outs that were merely educated guesses as to what users would end up hooking onto those pins.

As a result, the final incarnation of the laptop includes specialized adapter boards for the battery pack, LCD, and front-bezel arrangement (the frame for the display). The good thing about this modularity is that users can adapt the system to meet their own needs. We're delighted to have received reports of users changing out LCDs and building custom battery packs. And users can add the keyboard of their choice. We don't include one.

Aside from the inclusion of an FPGA, our hardware decisions were fairly lackluster, even when compared with what comes in a low-end Intel laptop. Our i.MX6 SoC contains a quad-core ARM CortexA9 running at 1.2 gigahertz–a 32-bit processor, which means the system is limited to 4 gigabytes of RAM. We did, however, choose a fairly decent 1920-by-1080 display, which is the spec limit for the i.MX6 chip. Our battery life is acceptable, at a bit over 6 hours.

We included a few extra features just for the fun of it. We built in an accelerometer, although we really don't know what we're going to do with it. Accelerometers are cheap, and there are a lot of fun things you can do with your computer when it's aware of being tilted, hit, or dropped. For instance, you could make your laptop play a sound clip of "Goodbye, Cruel World!" when it detects that it is in free fall.

Most laptops have one Ethernet port at the most, which means that they can act only as endpoints in a network. But we gave Novena two ports to let it sit between nodes, filtering and monitoring traffic for security purposes. The ports also allow our laptop to function as a Wi-Fi router and as an Internet firewall for other wired devices–

ANDREW "BUNNIE" HUANG visits the Portland, Ore., workshop of collaborator Kurt Mottweiler, who is working on the Heirloom version of the Novena laptop. This special configuration features a hand-crafted wood-and-aluminum case.

a handy feature when on the road. And although 20 hertz is considered the low end of human hearing, one of us (Huang) has a broader hearing range, so we took the unusual measure of extending the analog frequency response of the headphone jack to below 20 Hz.

These little features are never included on a consumer laptop. But it's these customizations that make the project fun and unique. They add cost–but it's *our* cost overrun! The fully loaded laptop sold for US $2,000 during the crowd campaign. So far, we've sold about 500.

Initially, we didn't plan on going as far as building a mass-producible case. In fact, we made our first case by hand out of leather, bound like a book. The warm reception it got on social media led us to attempt a crowdfunding campaign via Crowd Supply, a site that takes the idea of Kickstarter and turbocharges it with a Web store and logistics services.

One thing we learned in that brush with the mainstream consumer was that most people were not prepared for the experience we were delivering. Our laptop won't run Windows or Mac OS, nor is it particularly fast or thin. So in order to differentiate our laptop and make it more attractive to true techies, we decided to add some of the things that developers prize. These included mounting bosses–an internal array of metal nuts–which we call the Peek Array (after Nadya Peek, of FabLab fame).

Typical laptops are unhackable because there's no place to put other...stuff. If you can find the empty space in a laptop, you'd have to glue your expansion board in place or drill a hole through the case to mount your custom hardware. Not so with the Peek Array! Its dozens of threaded nuts allow users to screw all manner of small project boards into the laptop case. Want to add a pulse oximeter to Novena so you can measure the level of oxygen in the blood running through the capillaries of your finger? Or maybe a barometer so you can monitor your airliner's cabin pressure? With just a few screws you can mount your customization inside Novena's laptop case.

We also designed the case to open easily under the impetus of an air spring, like the ones used to raise the seat on ergonomic office chairs. With the flick of a switch, the lid (with the screen) opens automatically. We preferred this to a conventional clamshell design on the theory that any laptop that exposes its guts to the user will particularly appeal to users prepared to hack the guts. (People who fear naked circuit boards shouldn't buy our laptop anyway!)

The case is designed to serve through generations of electronics hardware. For instance, to hold the liquid-crystal display in place we used computer numerically controlled (CNC) machining to make an aluminum bezel. Anyone with access to an entry-level machine shop can fabricate a custom bezel to accommodate a different LCD, as well as mount additional sensors (such as a camera or a microphone) or additional buttons and knobs.

By putting most of the ports on a single edge, which we call the Port Farm, we made it easy to replace the panel covering it. That way the user can keep the case itself even as the motherboard and the Port Farm evolve.

We took into account the total costs of manufacturing tooling as well. For instance, though injection-molded plastic is very cheap on a per-unit basis, the steel tools that produce it are not. You need an oil-cooled block weighing about a ton, capable of handling pressures found at the bottom of the Mariana Trench, machined internally to a tolerance less than the width of a human hair, with a moving clockwork of dozens of ejector pins, sliders, lifters, and parting surfaces separating and coming back together again smoothly over thousands of cycles. It can easily cost $250,000, which is why the injection molding process can pay for itself only in very large manufacturing runs. We instead used a combination of CNC aluminum, optimized plastic design, and family molds (whose various parts each form a different aspect of a given product). We needed just 10 machine tools to produce the case, for an up-front investment of about $50,000.

As we write this, we've shipped all of our standard laptop and desktop units to our crowdfunding campaign backers. We've marshaled nearly a thousand components from dozens of vendors, from Shenzhen, China, to Fremont, Calif.

Thanks to the magic of crowdfunding and a slowdown of Moore's Law, a two-person team was able to have the funds and the time to implement an everyday-use laptop. Furthermore, this laptop's schematics, circuit board layouts, mechanical design files, kernel, drivers, and application programs are all open source and available for anyone to download.

Such transparency is unprecedented. We hope it will encourage other engineers to follow in our footsteps and help users reclaim their technological independence. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/novena1115

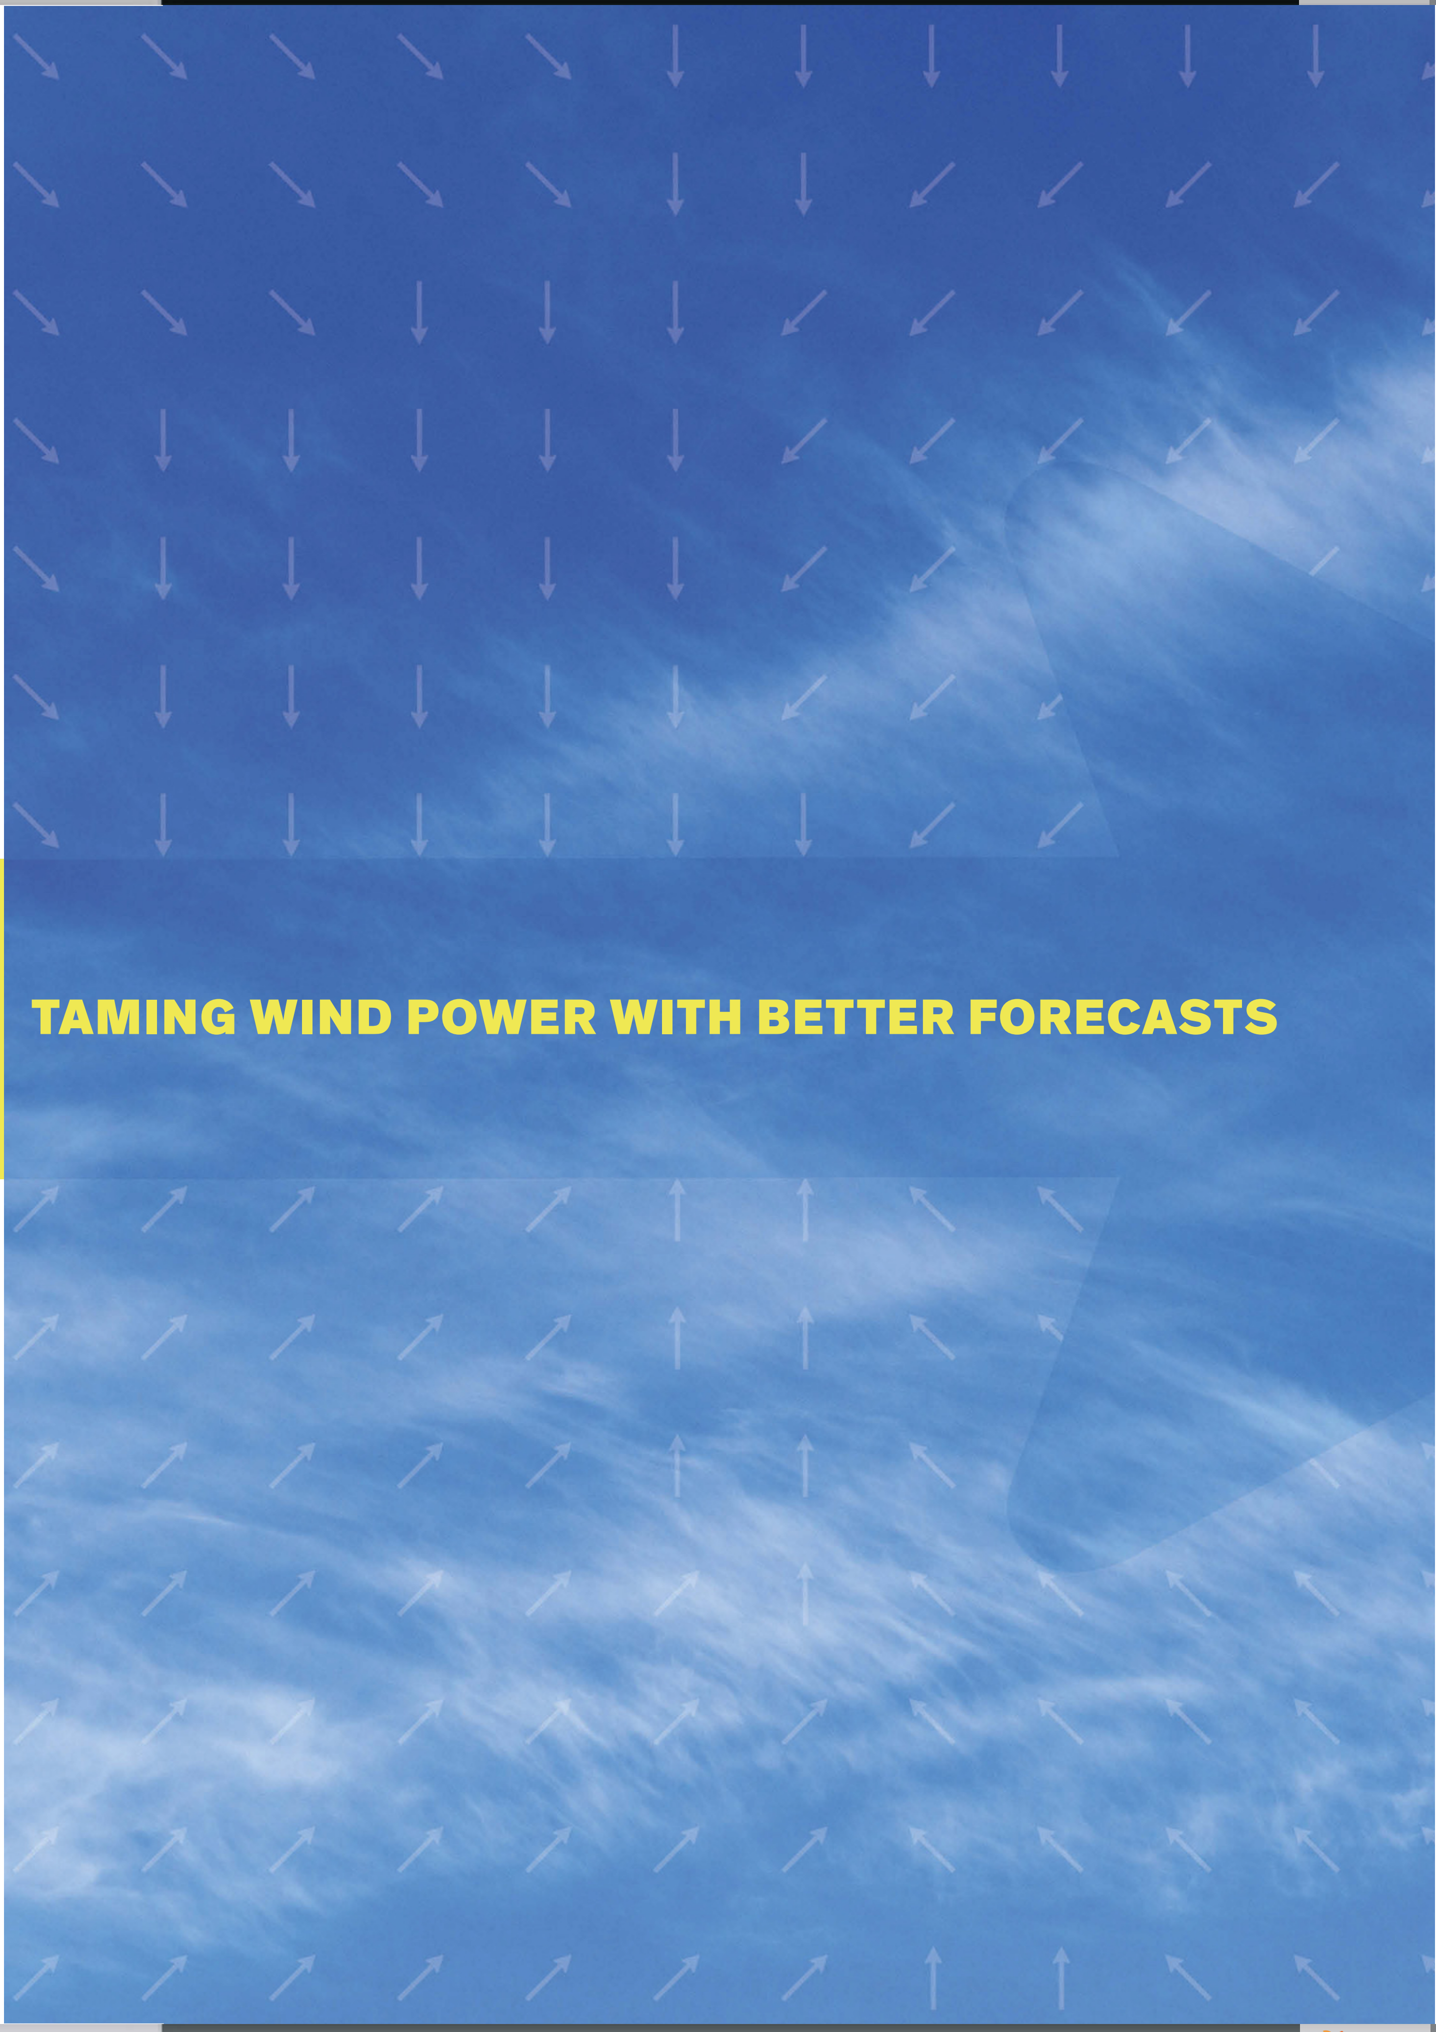

TAMING WIND POWER WITH BETTER FORECASTS

Sophisticated weather simulations are making wind power more grid friendly

BY **SUE ELLEN HAUPT** & **WILLIAM P. MAHONEY**

Wind energy. It's clean. It's renewable. Its potential is enormous. But to draw energy from the wind and send it to people's homes reliably and efficiently, you have to know when the wind will blow and when it won't. When it stops or changes rapidly, you have to be ready to substitute power from another source. And because such sources aren't always available at a moment's notice, you need this information many hours and even several days ahead.

None of this matters much if the wind supplies only a small percentage of the electricity going into the power grid. But several countries have already gone far beyond that, and more will soon, as wind is now the fastest growing source of energy in the world. Denmark already gets 28 percent of its electrical power from the wind and has at times drawn all of its electrical energy from wind turbines. The wind is already providing 20 percent or more of the electricity for several U.S. states, including Iowa, Kansas, and South Dakota. Colorado, for example, has at times obtained 60 percent of its electric power from the wind. And the United States as a whole will likely produce 20 percent of its electricity from wind power within 15 years.

Even though they can't control the wind, utilities in these regions are able to reliably manage wind power today and will be able to do so even more efficiently in the future. That's not because they have

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

installed massive warehouses of batteries or have a large number of natural-gas generators ready to power up instantly. It's because since the 2000s, scientists have been dramatically improving their ability to forecast the amount of power that utilities will be able to draw from the wind. And these forecasts are getting better all the time, allowing utilities to depend more and more on their wind resources and less on fossil fuels or nuclear power.

M**anaging the supply** of power to the electrical grid is tricky. Utilities must gauge power demand and constantly adjust generation to meet it. If you underestimate the need and produce too little, you get a brownout. Sudden changes in wind production may cause a deviation from the normal frequency and voltage of the grid, causing local problems like flickering lights and possibly damaging unprotected electronics, or at higher levels tripping circuit breakers in the grid and causing blackouts.

Grid operators are pretty good at formulating such estimates and can usually make adjustments quickly if the situation changes—at least when they have direct control of the generation systems, as is the case with coal, natural gas, nuclear energy, and much of hydropower. They have some ability to dial these sources up or down in a couple of hours—or minutes for many natural-gas plants—as long as the generators are already in operation.

Wind is different. Maintaining a steady flow of wind energy into the power grid is not so much a matter of controlling wind energy as it is of compensating for its fluctuations. But it takes 12 hours or more to fully start up a nuclear reactor and 6 hours for a coal plant. So the only quick fix is to use an expensive "spinning" reserve, typically a power plant running well below its capacity, like an idling car. Unless they want to keep a large number of such spinning reserves up and running—which is a wasteful proposition—utilities must be able to forecast the wind accurately.

Predicting winds at a specific site isn't easy. Terrain, nearby bodies of water, buildings, and vegetation dramatically influence local wind speeds. Traditional weather-prediction models are typically too coarse to pick up these details. It takes a specialized computer model that considers these local influences along with the effects of larger-scale weather systems.

Utilities need accurate wind forecasts for two distinct time frames: one to five days ahead and 15 minutes to 6 hours ahead. They use their longer-range forecasts to plan the energy mix—both what the utility will produce from its different generation facilities and what it will buy if needed from other producers. Utility operators generally need this information a day ahead, say, by 6 a.m. for how much wind will be blowing from midnight that night through the following midnight. That gives them time to get coal, nuclear, or gas plants up and running and to place orders to buy electricity from other producers at an economical rate—purchases made any later would command premium prices. But one day isn't always enough lead time: Markets shut down on weekends, so utility planners need to have Monday's wind forecast on Friday morning. And over holidays they will need to plan their purchases even sooner.

Knowing *how* the wind will change just minutes to 6 hours ahead is important, too, because the weather can shift dramatically in that time, and with even a little warning of a big drop or increase in available wind energy, operators can adjust spinning reserves.

No single forecasting system performs optimally across both of these time scales. So we and our colleagues at the National Center for Atmospheric Research (NCAR) have been working to develop tools for both types of forecasts. We're doing this in partnership with Xcel Energy—a utility supplying energy to residents and businesses in Colorado, Michigan, Minnesota, New Mexico, North Dakota, South Dakota, Texas, and Wisconsin. Xcel has the largest wind-energy capacity of any utility in the United States, some 5,700 megawatts.

We call what we've built the NCAR Wind Power Forecasting System. Its development started in 2008. Xcel Energy began

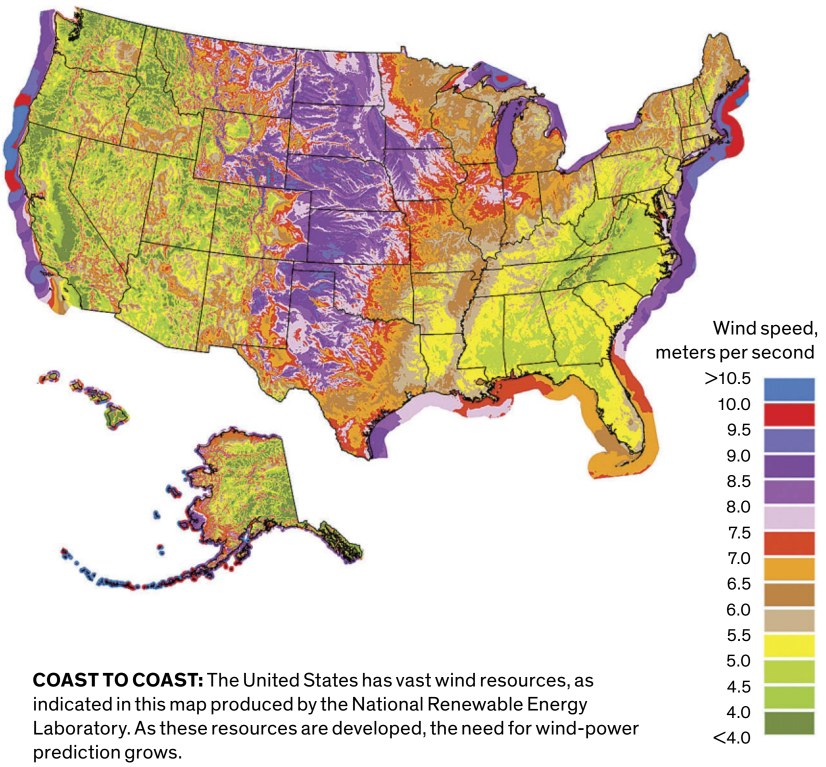

COAST TO COAST: The United States has vast wind resources, as indicated in this map produced by the National Renewable Energy Laboratory. As these resources are developed, the need for wind-power prediction grows.

PREVIOUS PAGES: ISTOCKPHOTO; THIS PAGE: AWS TRUEPOWER/NREL using an early version operationally in 2009, and NCAR has been improving it ever since. Denmark, Germany, Australia, China, and numerous private companies are developing similar prediction systems to enable them to reliably increase their use of wind power.

At the core of our forecasts—as with all weather forecasts today—are numerical models. Such systems start with measurements of temperature, barometric pressure, wind speed and direction, moisture, solar radiation, cloud cover, and more made by weather stations, weather balloons, satellites, and ground-based radars. These models assimilate the data into a grid of points covering the region of interest. They then compute changes in those measurements at each point using equations that govern the evolution of the atmosphere, taking many factors into account, such as changes in solar radiation, the effects of atmospheric turbulence, the physics of cloud formation, and land-surface usage. These numerical-prediction schemes are able to model everything from the large-scale weather systems circling the globe to the turbulent eddies that form when the wind flows over bumps in the local terrain. To do so, today's models typically include tens of millions of grid cells and require thousands of computer processors running in parallel to produce multiday forecasts as often as hourly.

Wind forecasts are publicly available from the U.S. National Weather Service (NWS), which runs its own numerical weather model. Ten years ago, when this was the only tool for wind forecasting, it provided weather forecasts at a resolution of 36 kilometers. At this level, smaller features of the local terrain are ignored. And although the model could generally indicate that a weather system would move through a region in the morning or the afternoon, it couldn't tell exactly when it would arrive at a particular spot.

Today, some NWS models generate predictions at a 3-km resolution over the United States—a huge improvement—and they can time the movement of weather fronts to within just a few tens of minutes.

But that's not quite good enough to tell us exactly what the winds will be doing at utility wind farms. So we're also running a specialized version of the Weather Research and Forecasting (WRF) model. We've specially configured it to provide information specific to Xcel Energy's needs. We call this the Real Time Four Dimensional Data Assimilation version of WRF. Like the new, high-resolution NWS model, it runs at a 3-km resolution. But for the vast majority of Xcel's 3,000 wind turbines, we've added a key piece of data—the wind speeds measured by the sensors sited on top of the turbines. We started using these additional data in 2008, and they have made a dramatic improvement in the accuracy of our wind forecasts.

A few wind farms have on-site weather stations mounted on special meteorological towers continuously measuring wind and temperature at different levels. This helps determine the stability of the atmosphere and can improve wind predictions, especially at nighttime, when the wind at the ground often weakens faster than the wind at turbine height, causing wind shear and making the turbines less efficient. We're using these data as well in the model we've adapted for Xcel.

To optimize our predictions, we combine results from this model with those from other forecasting models—including the NWS models—blending them all together with something called the Dynamic Integrated foreCast (DICast) system. NCAR started developing DICast almost 20 years ago for the Weather Channel. Since then NCAR, as well as some commercial weather

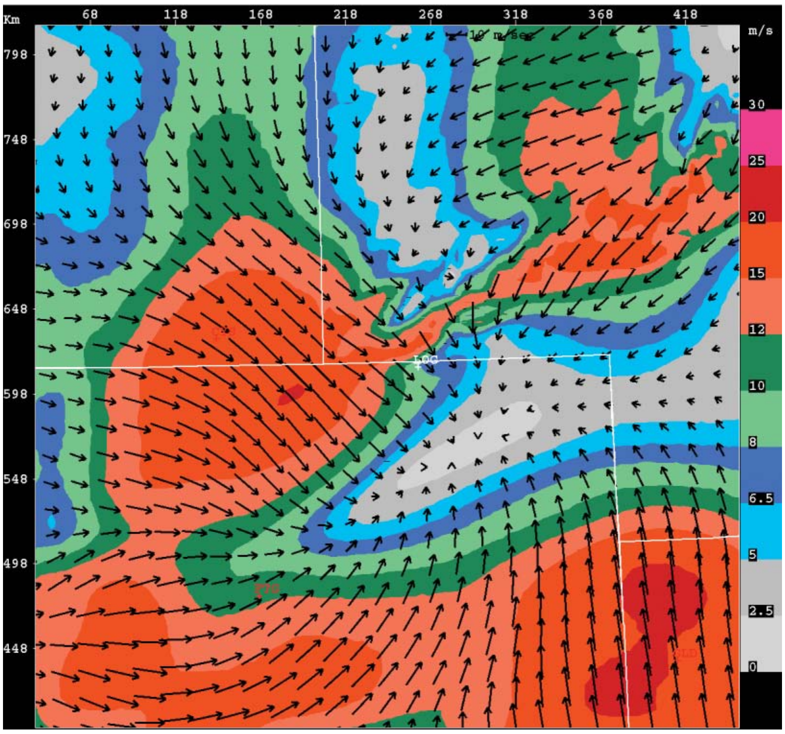

PICTURE THIS: Doppler radar returns blended into the Variational Doppler Radar Analysis System (VDRAS) can give utilities information about the current wind situation and likely near-term changes to come. In this example, covering northeastern Colorado and neighboring states, the gray band shows an area in which winds converge, causing lower wind speeds over wind farms locally and a likelihood of decreasing wind speeds over the next couple of hours. The red areas represent faster winds moving into the region that will increase the amount of power produced by wind turbines.

NCAR

IN THE WIND: Doppler weather radars, like the one at the National Center for Atmospheric Research [left], and weather sensors placed on or at the height of wind turbines [center] collect some of the data that are incorporated into numerical weather models to produce local wind forecasts. Methods to precisely predict extreme weather conditions like icing [right] are being developed.

companies, have refined and implemented it in various operational forecast systems.

We started using DICast for wind prediction in 2009. We train it by giving it information on how well each of the models we're blending performs under different weather conditions. We also give it past observations that allow it to determine the different biases and errors in certain models. One model, say, may always deviate from the correct forecast at a certain elevation by a certain amount; another may include too many clouds in a certain region, preventing adequate surface heating that may drive winds.

So now we run DICast in two stages. The first corrects the models for biases and errors; the second blends the weather models, giving them different weights depending on which models perform best under conditions that resemble the current weather for the given location and lead time. Although DICast works well enough with just 30 days of training data, it improves with time because it constantly checks its predictions against what actually happened and adjusts its bias corrections and weightings.

After our systems generate a wind forecast using these various components, we need to translate the predicted wind speeds into information that utilities can use, which is the amount of power the winds will generate at the local wind farm.

To do this, we first use the DICast system to predict the wind speed at the height of the hubs of the turbines—about 80 to 100 meters above the ground. Then our system translates those predicted wind speeds into power forecasts. Each type of wind turbine has a particular relationship between the wind speed and the power output, known as its power curve. But power production often deviates substantially from the turbine manufacturer's power curve, and a single wind speed can correspond to a variety of power outputs. Bad correlations between wind and power output can be caused by a variety of factors, such as times when the turbine is down for maintenance, snow or ice is on the blades, the utility is intentionally curtailing power, or the wind speed exceeds the design limits for safe operation and operators shut down the turbine. To avoid the worst anomalies in our wind-to-power conversions, we typically just consider the middle 50 per-

WIND TO POWER: In this plot of power produced by a wind turbine at different wind speeds [left], the manufacturer's data [green] differ slightly from the median of the observed data [yellow] because of site-specific features like terrain and air density. Taking the middle 50 percent of observations creates a usable site-specific curve [right].

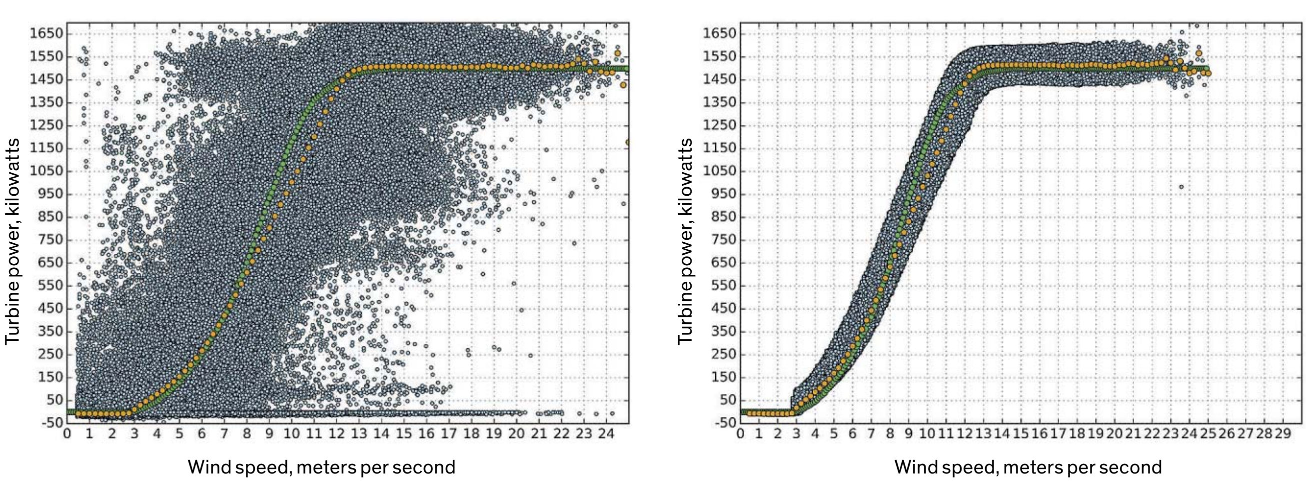

TOP, FROM LEFT: NCAR; ISTOCKPHOTO (2); BOTTOM: NCAR cent of possible power values for each wind speed at each turbine.

That's how we get from a basic weather forecast to detailed wind-power predictions at specific sites. There's one more thing we're doing to make those predictions even more accurate. It involves a unique approach to using ensembles of forecasts to predict weather.

In ensemble forecasting, about 10 to 30 different weather models are run concurrently depending on the available computing resources and the current weather situation. These differ in such factors as initial conditions, boundary conditions, and the use of physics packages—formulas that represent processes like land-surface usage or changes in solar radiation. From that set, the system generates an ensemble of possible outcomes. Sometimes the differences are slight, with all the models predicting, say, that a hurricane will follow a particular path. Other times the results vary widely, and it may be hard to predict with any confidence whether that hurricane is going to turn toward land or head out to sea.

At NCAR, we are refining the use of ensemble forecasting by identifying analogous forecasts—previous forecasts that predicted conditions similar to those of the current forecast. We then determine the errors in those analogous predictions by comparing them with how the weather actually evolved. If previous predictions tended to be off in some consistent way, we can correct the current forecast by compensating for the past errors.

We typically identify hundreds of analogous forecasts, then use the 10 to 20 that are most similar to the current forecast to generate an ensemble. We then look carefully at how well those past predictions played out so that we can adjust the current forecast to compensate for any consistent bias. And we use the range of errors to help estimate the uncertainty in today's forecast.

Over the past year, this approach has improved the accuracy of our forecasts and allowed us to give the utilities probabilities to help them with their planning. A utility would be more likely to keep additional spinning reserves on line when the possibility of strong winds is 20 percent, rather than 80 percent, for example.

All of these efforts are aimed at getting one- to five-day forecasts. For our very short range forecasts, we need a different approach—for two reasons. Most important is that these computationally intensive weather models take longer to run than there is time available. Also, we want to take advantage of local information, which the larger-scale models miss.

At the shortest time scales, it is currently difficult to beat what is called a persistence forecast. For example, 5 minutes from now, the wind speed is likely to be much as it is at the moment. So it's appropriate to predict no change at these very short time scales. But in 30 minutes to several hours, wind speeds can shoot up and down dramatically.

To forecast these rapid changes, we need to include observations near the wind farms. One approach that we've been evaluating blends radar observations into a rapidly updating weather model. It uses data from Doppler weather radars called Next Generation Weather Radar, which the NWS has deployed around the United States. (If you live in that country, you likely see images from these radars when you tune in to a local TV weather forecast.) The radars send out pulses of microwave energy and measure the amount that is reflected off various sorts of particles entrained in the air, including raindrops and hailstones. Changes in the frequency of the signal—the Doppler shift—indicate the speed of the particles' horizontal movement toward or away from the radar antenna, and thus the speed of the wind. Combining results from multiple radars allows you to ascertain the direction of the wind.

Mapping radar echoes lets you identify general weather patterns—patterns such as fronts that tend to persist even as the overall weather system moves along horizontally. So we can use our knowledge of the overall pattern and how it is moving to predict how the wind will soon change in one spot. This allows us to produce much better short-term forecasts.

How are we doing so far? Each of Xcel Energy's service regions has different challenges, and as a consequence the reliability of our forecasts varies. But even in the most difficult areas, we're doing well. For example, Xcel's Public Service Company of Colorado subsidiary operates in a region with very complex terrain, where it is notoriously difficult to provide good local wind forecasts. In 2008, the year we started working on this problem, the forecast error at a typical wind farm in this region was about 18 percent for 18 to 42 hours ahead, calculated as a percentage of the capacity of each farm. But by 2014, the mean forecast error had fallen to 10.8 percent. Xcel estimates that improvements in forecasting have saved customers US $49 million so far.

That's pretty good, but we think we can do better. This past winter, we added the ability to forecast turbine icing. Sometimes icing happens gradually, and the turbines slow before stopping. In such situations, we can sometimes use the drop in power to diagnose the icing. Other times, such as with freezing rain, icing can happen quickly, and the turbines need to be shut down immediately to avoid damage from the weight of the ice. Some of the most advanced work in icing prediction is in Scandinavia, where this phenomenon is a constant threat, and some wind turbines there and in other similarly cold, wet regions are even beginning to include mechanisms to remove ice from the spinning blades.

While we are excited to incorporate methods to predict icing in our model, and while we've had a few successes with them, we still have a long way to go. The United States didn't see many ice-producing storms this past winter, so we have less data than we expected to test our modeling.

Having even more accurate wind forecasts—ones that give day-ahead wind-speed projections with the accuracy of today's 6-hour forecasts—is essential to moving the world toward a largely renewable-energy future. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/forecast1115

COLLEGE OF ENGINEERING
ELECTRICAL & COMPUTER ENGINEERING
UNIVERSITY OF MICHIGAN

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

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

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

Applications will be considered as they are received. However, for full consideration applications must be received by **December 7, 2015.**

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

THE OHIO STATE UNIVERSITY

The Ohio State University invites applications for multiple tenure track faculty positions in the Department of Electrical and Computer Engineering. All areas and ranks will be considered. We are especially interested in (i) control systems or robotics for intelligent vehicle and transportation systems, (ii) biomedical applications including cancer detection and imaging, tumor cell biology, or the genomics of cancer, (iii) electronic materials and devices, emphasizing low energy electronics or wide bandgap semiconductors, (iv) neural electronics, and (v) senior candidates in electromagnetics, microwave theory, RF systems, or remote sensing. All positions may involve joint appointments with other engineering departments. Applicants must have a Ph.D. and outstanding academic credentials. Successful candidates will be expected to develop a vigorous externally-funded research program, show excellence and leadership in academic and scholarly activities, and demonstrate outstanding teaching at the undergraduate and graduate levels. Applicants are requested to send a letter of application, curriculum vitae, statement of research plans, brief statement of teaching philosophy, and name, address, and email of four references to Professor John L. Volakis at search@ece.osu.edu. Several of these positions are partially funded by Ohio State's Discovery Themes Initiative, a significant faculty hiring investment in key thematic areas that build on the university's culture of academic collaboration to make a global impact. The Ohio State University is committed to establishing a culturally and intellectually diverse environment, encouraging all members of our learning community to reach their full potential. We are responsive to dual-career families and strongly promote work-life balance to support our community members through a suite of institutionalized policies. We are an NSF Advance Institution and a member of the Ohio/Western Pennsylvania/West Virginia Higher Education Recruitment Consortium.

上海科技大学
ShanghaiTech University

ShanghaiTech Faculty Search

ShanghaiTech University invites highly qualified candidates to fill multiple tenure-track/tenured faculty positions in the School of Information Science and Technology. Candidates should have exceptional academic records or demonstrate strong potential in cutting-edge research areas of information science and technology. English fluency is required and overseas academic connection or background is highly desired. ShanghaiTech is built as a world-class research university for training future generations of scientists, entrepreneurs, and technological leaders. Besides establishing and maintaining a world-class research profile, faculty candidates are also expected to contribute substantially to graduate and undergraduate education.

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/integrated signal processing systems, computational foundations, big data, data mining, visualization, computer vision, bio-computing, smart energy/power devices/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 with international practice for performance evaluation and promotion.

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

Applications:
Submit (in English, PDF) a cover letter, a 2-page research plan, a CV plus copies of 3 representative publications, and names of three referees to sist@shanghaitech.edu.cn by December 31, 2015. For more information, visit **http://www.shanghaitech.edu.cn**.

FACULTY POSITIONS
Electrical Engineering
NYU SHANGHAI

NYU Shanghai invites outstanding applications from all ranks to apply for tenure-track or tenured faculty appointments in Electrical and Computer Engineering (ECE) with special emphasis in Computer Engineering, Multimedia Analytics, Neural Engineering, and Robotics. Exceptionally strong candidates in other areas and academic ranks in ECE will also be considered. The applicants must have demonstrated abilities in both research and teaching. In describing teaching experience, candidates should identify courses they could teach both within and outside their specialty. Candidates must have completed a Ph.D. or equivalent by the time of appointment.

The terms of employment in NYU Shanghai are comparable to U.S. institutions in terms of research start-up funds and compensation, and include housing subsidies and educational subsidies for children. Faculty may also spend time at NYU New York and other sites of the NYU global network, engaging in both research and teaching.

Applications are due no later than February 1, 2016 and will be reviewed until the position is filled. To be considered, candidates should submit a curriculum vitae, separate statements of research and teaching interests (no more than three pages each), and electronic copies of up to five recent, relevant publications. To complete the online process, applicants will be prompted to enter the names and email addresses of at least three referees. Each referee will be contacted to upload their reference letter. Senior candidates are not required to submit references at this time. Please visit our website at **http://shanghai.nyu.edu/about/work/faculty-positions** for instructions and other information on how to apply. If you have any questions, please e-mail **shanghai.engineering.recruitment@nyu.edu**.

About NYU Shanghai:
NYU Shanghai is the newest degree-granting campus within the NYU Global Network University. It is the first Sino-US higher education joint venture to grant a degree that is accredited in the U.S. as well as in China. All teaching is conducted in English. A research university with liberal arts and science at its core, it resides in one of the world's great cities with a vibrant intellectual community. NYU Shanghai recruits scholars who are committed to NYU's global vision of transformative teaching and innovative research.

New York University has established itself as a Global Network University, a multi-site, organically connected network encompassing key global cities and idea capitals. There are three degree-granting campuses in New York, Shanghai, and Abu Dhabi, and complemented by eleven additional academic centers across five continents. Faculty and students circulate within the network in pursuit of common research interests and cross-cultural, interdisciplinary endeavors, both local and global.

上海纽约大学
NYU SHANGHAI

EOE/Minorities/Females/Vet/Disabled

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

JIE Joint Institute of Engineering

Faculty Positions available in Electrical and Computer Engineering

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

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

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

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

JRI SHUNDE INTERNATIONAL Joint Research Institute

Research positions available in Electrical and Computer Engineering

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

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

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

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

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

SUN YAT-SEN UNIVERSITY   Carnegie Mellon University

Technical University of Denmark DTU

PROFESSOR

DTU Electrical Engineering invites applications for a position as professor (with special responsibilities) in power components in distributed energy systems. The research focus of the professor shall be within use of power electronics in smart grids. The professor should already have a strong international standing in the field and will be responsible for further strengthening our research activity as well as developing the collaboration within the field with external academic and industrial partners.

Application deadline: **29 November 2015**

DTU is a technical university providing internationally leading research, education, innovation and public service. Our staff of 5,800 advance science and technology to create innovative solutions that meet the demands of society; and our 10,300 students are being educated to address the technological challenges of the future. DTU is an independent academic university collaborating globally with business, industry, government, and public agencies.

Further details: career.dtu.dk

SCIENCE AND TECHNOLOGY AT A GLOBAL SCALE
– SET THE STANDARDS FOR THE FUTURE
See our PhD-programmes at dtu.dk/phd

THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Faculty Positions (Autonomous Systems and Robotics)

Founded in 1991, the Hong Kong University of Science and Technology (HKUST) is a world renowned international research university in Asia's most vibrant city. The University comprises more than 500 faculty members whose research ranges from science to engineering, business, humanities and social sciences. It promotes interdisciplinary studies, dedicated to educating all-rounded students to acquire a strong entrepreneurial spirit and innovative thinking with a global outlook. HKUST has been ranked among the top universities in Asia by /QS Asian University Rankings/. It ranks 16th in the latest /Global Employability University Survey/ reported by the New York Times, and remains No. 1 in Greater China area. English is the medium of teaching, research and administration.

With the approaching of its 25th anniversary, the University will conduct a clustered hiring campaign aimed at advancing the University's academic strengths in selected cross-disciplinary areas that are strategically important to HKUST's future development. We are seeking renowned scholars, academics, leaders or potential leaders in the core cross-disciplinary area of Autonomous Systems and Robotics.

Applications/nominations are invited for a number of substantiation-track faculty positions at all levels of Professor, Associate Professor and Assistant Professor in the above area. The successful candidate will hold a joint appointment in two or more departments.

Successful candidates are expected to work closely with our newly founded HKUST Robotics Institute to (i) conduct fundamental and cutting-edge research in the area of autonomous systems, robotics, unmanned aerial vehicles (UAVs), manufacturing automation, and medical robotics; (ii) provide a world class educational platform for autonomous engineering; (iii) transfer knowledge related to autonomous systems to the community, industry and government; and (iv) create an entrepreneurial environment for students and faculty.

Applicants/nominees for the positions should have a PhD degree, demonstrated leadership abilities, extensive teaching and research experience as well as an ability to interact effectively with students, faculty, industry and the Government.

Salary is highly competitive and will be commensurate with qualifications and experience. Fringe benefits include annual leave, medical and dental benefits. Housing benefits will also be provided where applicable. Appointment for Professor will be on substantive basis. Initial appointment for Assistant Professor/Associate Professor will normally be on a three-year contract, renewable subject to mutual agreement. A gratuity will be payable upon successful completion of contract.

Application Procedure
Applications/nominations, together with full curriculum vitae and the names and addresses of three referees, recent publications, and research papers should be sent to the Search Committee to asrsearch@ust.hk. Review of applications and nominations will begin immediately and continue until the positions are filled.

For further information about the University, please refer to: http://www.ust.hk.

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

UTSA

Open Faculty Positions in Electrical and Computer Engineering

The Department of Electrical and Computer Engineering at **the University of Texas, San Antonio (UTSA)** invites applications for multiple faculty positions starting Fall 2016 at the Assistant, Associate, or Full Professor ranks in: (1) VLSI Design and FPGA Systems, (2) Microelectronic, Mixed Signals and Advanced Materials, (3) Smart Grid, Power and Energy Systems, and (4) Computational Biology, Neuroscience, and Medical-Health Informatics. Outstanding candidates from all research areas will be considered, but preference will be given to candidates with primary research interest in VLSI design, FPGA systems, mixed signal integrated circuits, solid state devices and materials, power and energy systems, neuroscience, or big data applications. Please read the entire advertisement on the Department website at **http://ece.utsa.edu/about/faculty-openings.html**. The search committee will begin review of applications immediately and will continue until the positions are filled.

The UTSA is an Affirmative Action/Equal Employment Opportunity Employer, committed to diversity in its faculty and its educational programs. Women, minorities, veterans, and individuals with disabilities are encouraged to apply.

McMaster University

Tenure-stream Faculty Position in Power Electronics

The Department of Electrical and Computer Engineering at McMaster University invites applications for a tenure-stream position in Power Electronics. The successful applicant will have an earned PhD in a relevant discipline, and the potential for excellence in both research and teaching.

For full position details, please visit **http://www.eng.mcmaster.ca/postings_fulltime.html**

All qualified candidates are encouraged to apply. However, Canadian citizens and permanent residents will be given priority for these positions. McMaster University is committed to employment equity and to recruiting a diverse faculty and staff.

IEEE JobSite IEEE
The Right Candidate – Right Now!

Take the next steps to finding your ideal job.

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

One of the most influential reference resources for engineers around the world.

For over 100 years, *Proceedings of the IEEE* has been the leading journal for engineers looking for in-depth tutorial, survey, and review coverage of the technical developments that shape our world.

To learn more and start your subscription today, visit
ieee.org/proceedings-subscribe

IEEE

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) at the University of Toronto invites applications for a continuing Teaching Stream appointment at the rank of Assistant Professor, Teaching Stream. The expected start date is **July 1, 2016**.

Candidates are expected to demonstrate potential for outstanding teaching, which will include developing and delivering undergraduate courses and laboratories and supervising undergraduate design projects. While our primary objective is teaching excellence across a range of fundamental and advanced ECE courses, our most pressing teaching needs are in the areas of computer software and energy/power systems, as well as engineering design. The successful candidate should have a Ph.D. in Electrical or Computer Engineering by the time of appointment or shortly thereafter and should be eligible to apply for Professional Engineering registration in Ontario.

The successful candidate will have demonstrated excellence in teaching and pedagogy, have a strong engineering background, and have excellent organizational and communication skills. Evidence of excellence in teaching will be demonstrated through teaching evaluations, letters of reference and the teaching dossier submitted as part of the application. Salary will be commensurate with qualifications and experience.

The ECE Department consistently ranks among the top 10 in North America and attracts outstanding undergraduate and graduate students, has excellent facilities, and is ideally located in the middle of a vibrant cosmopolitan city. Additional information is at **www.ece.utoronto.ca**.

All qualified candidates are invited to apply by clicking on the link below. For further information about the application process, please see the submission guidelines at **http://uoft.me/how-to-apply**.

Applicants should provide a curriculum vitae and a teaching dossier (including a statement of teaching philosophy and interests). They should arrange to have three confidential letters of recommendation sent on their behalf to:

Professor Farid N. Najm, Chair,
The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
University of Toronto
10 King's College Road, Toronto, Ontario
Canada M5S 3G4

All application documents must be received by **December 1, 2015**.

The University of Toronto is strongly committed to diversity within its community and especially welcomes applications from visible minority group members, women, Aboriginal persons, persons with disabilities, members of sexual minority groups, and others who may contribute to the further diversification of ideas.

All qualified candidates are encouraged to apply; however, priority will be given to Canadian Citizens and Permanent Residents.

Rochester Institute of Technology
Kate Gleason College of Engineering

Position: Instructional Faculty
Tenure Status: Tenure Track
Faculty Rank: Assistant Professor
Discipline: Electrical Engineering

R·I·T

POSITION ANNOUNCEMENT: KGCOE Electrical & Microelectronic Engineering

The Department of Electrical and Microelectronic Engineering at the Rochester Institute of Technology invites applications from candidates in the broad area of Power and/or Controls. Appointment to this tenure-track position is intended to be at the Assistant Professor level. However, candidates at other levels are welcome to apply. In order to be considered for this position, applicants must have a Ph.D. in Electrical Engineering or a closely related field, with specialization in Power and/or Controls, and must have published in refereed journals based on their Ph.D. dissertation and any subsequent research work. Candidates must have strong aptitude and interest in undergraduate and graduate teaching; strong potential for establishing and conducting sponsored research; satisfactory written and oral communication skills. We are seeking an individual who has the ability and interest in contributing to a community committed to Student Centeredness; Professional Development and Scholarship; Integrity and Ethics; Respect, Diversity and Pluralism; Innovation and Flexibility; and Teamwork and Collaboration. Select to view links to RIT's core values, honor code, and statement of diversity.
To apply to this position, go to: http://careers.rit.edu/faculty, search openings, then Keyword Search: 2012BR.

RIT does not discriminate. RIT is an equal opportunity employer that promotes and values diversity, pluralism, and inclusion. For more information or inquiries, please visit RIT/TitleIX or the U.S. Department of Education at ED.Gov

BAYLOR
UNIVERSITY

Faculty Position

The Electrical and Computer Engineering Department of Baylor University seeks faculty applicants for a tenured/tenure-track Faculty Position at any level. Any area of expertise will be considered but applicants in computer engineering will be given special consideration. Applicants for assistant professor must demonstrate potential for sustained, funded scholarship and excellent teaching; applicants for associate or full professor must present evidence of achievement in research and teaching commensurate with the desired rank. The ECE department offers B.S., M.S., M.E. and Ph.D. degrees and is rapidly expanding its faculty size. Facilities include the Baylor Research and Innovation Collaborative (BRIC), a newly-established research park minutes from the main campus.

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

Applications will be considered on a rolling basis until the **January 1, 2016** deadline. Applications must include:

1) a letter of interest that identifies the applicant's anticipated rank,

2) a complete CV,

3) a concise statement of teaching and research interests,

4) the names and contact information for at least four professional references.

Additional information is available at **www.ecs.baylor.edu**. Send materials via email to Dr. Keith Schubert at keith_schubert@baylor.edu. Please combine all submitted material into a single pdf file.

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

A VACUUM TUBE FOR THE AGES

The humble 807, still beloved of ham-radio enthusiasts and audiophiles, outlived the company that created it

In the 1930s, as the United States recovered from the Great Depression, most people couldn't afford a new radio. Even so, the Radio Corporation of America earned a tidy profit on vacuum tubes like the 807, introduced in 1937. The 807 proved especially popular with amateur radio operators, as evidenced by this ad from the March 1940 issue of *QST*, a ham-radio magazine. And in 1947, English engineer D.T.N. Williamson described the circuit for his eponymous high-fidelity audio amplifier, variations of which used the 807. Although RCA closed its electron tube operation in 1976, 807s are still made in Russia and China, and audiophiles continue to debate the technical and aesthetic qualities of 807s from different companies, eras, and countries.

—ALEXANDER B. MAGOUN, IEEE History Center

PHOTOGRAPH BY Randi Klett

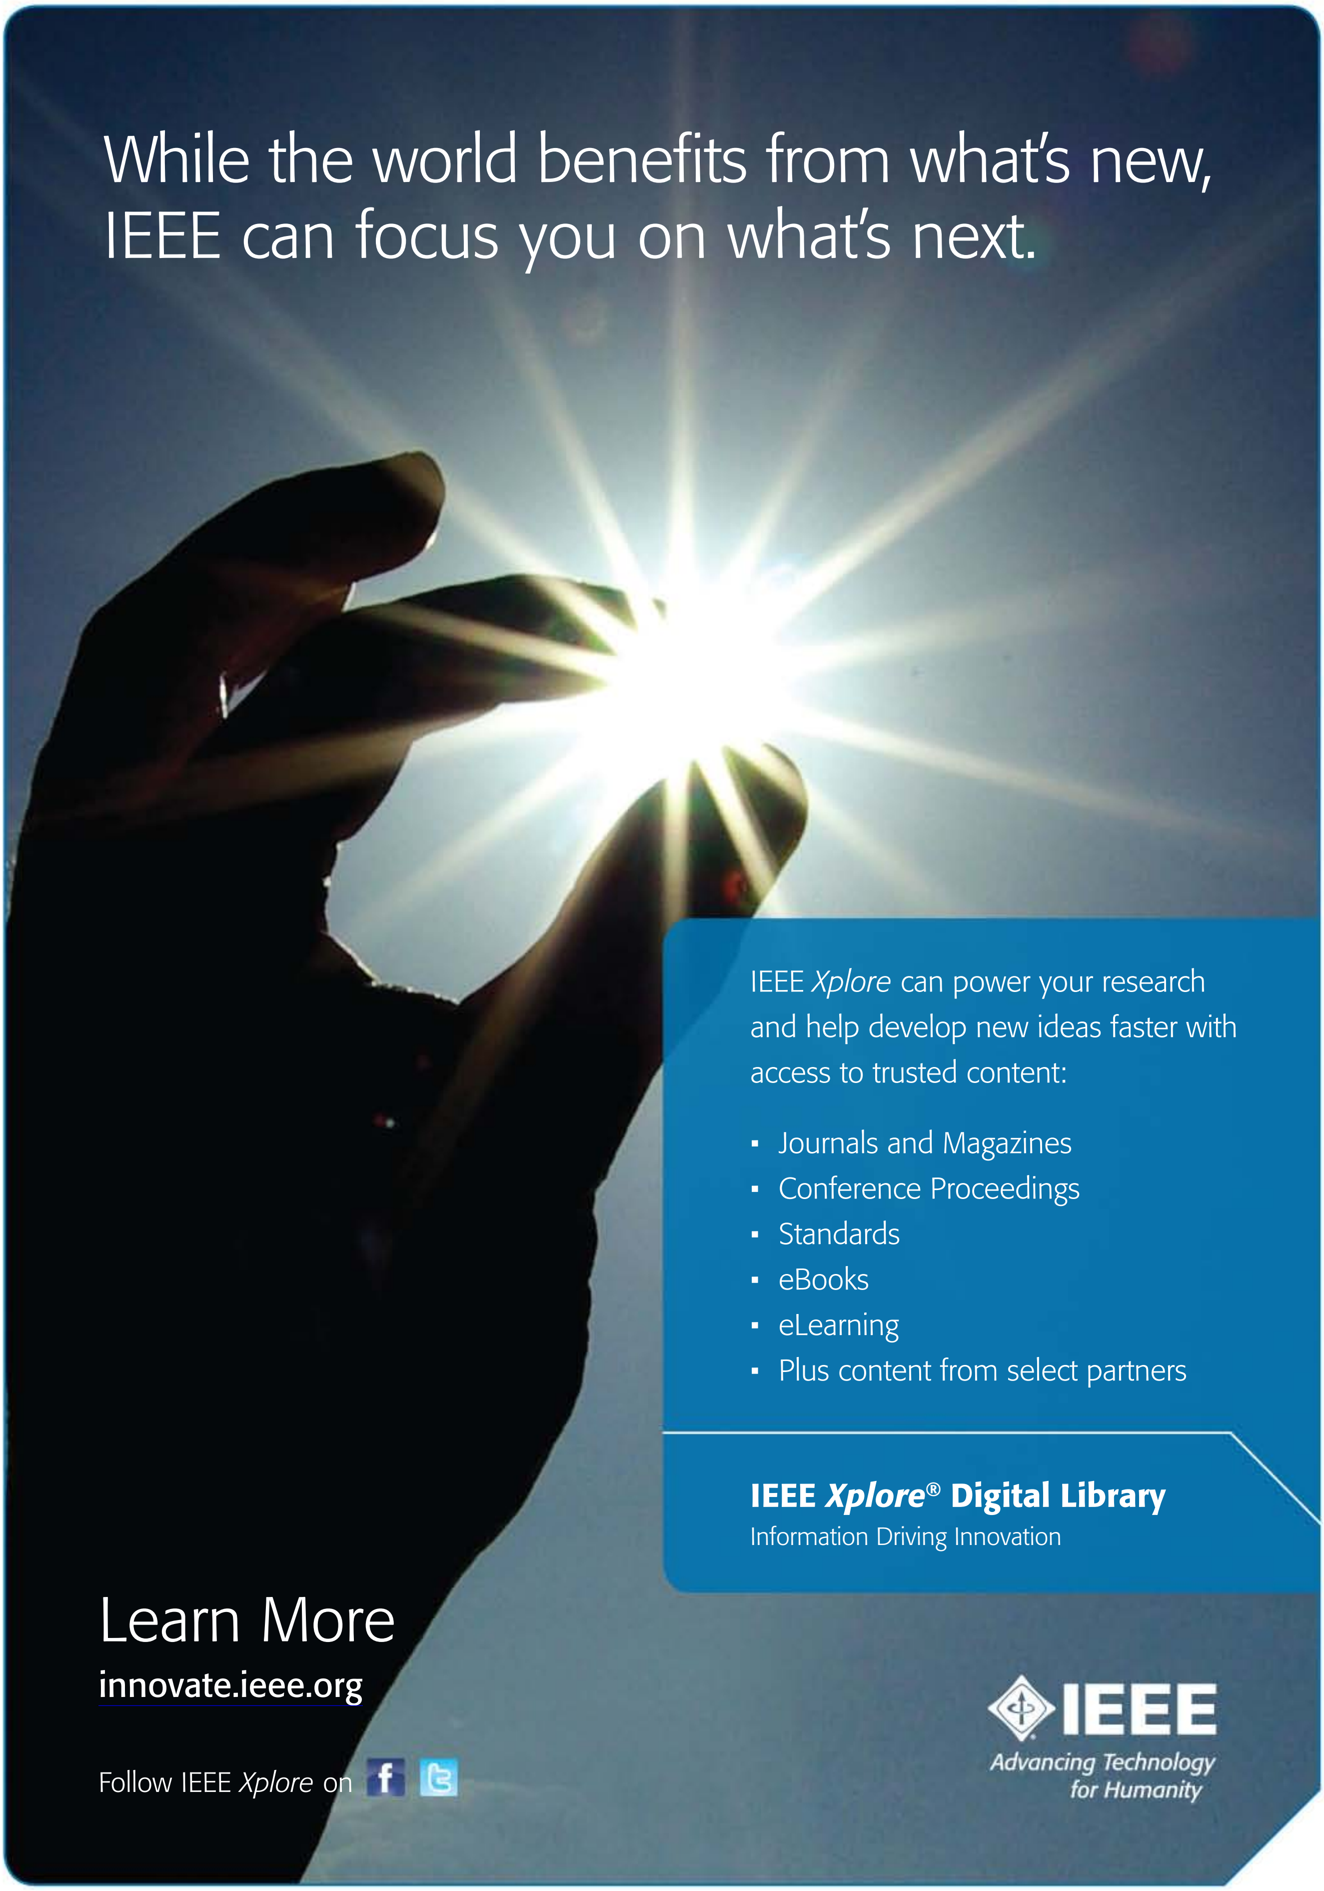
While the world benefits from what's new, IEEE can focus you on what's next.
IEEE *Xplore* can power your research and help develop new ideas faster with access to trusted content:
- Journals and Magazines
- Conference Proceedings
- Standards
- eBooks
- eLearning
- Plus content from select partners
IEEE *Xplore*® Digital Library
Information Driving Innovation
Learn More
innovate.ieee.org
Follow IEEE *Xplore* on
IEEE
Advancing Technology for Humanity