IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 01.17

TOP TECH 2017

12 BREAKTHROUGHS COMING THIS YEAR, INCLUDING:

- Autonomous taxis
- SpaceX astronauts
- An internal artificial heart
- Flexible smartphones
- New AR gaming

IEEE

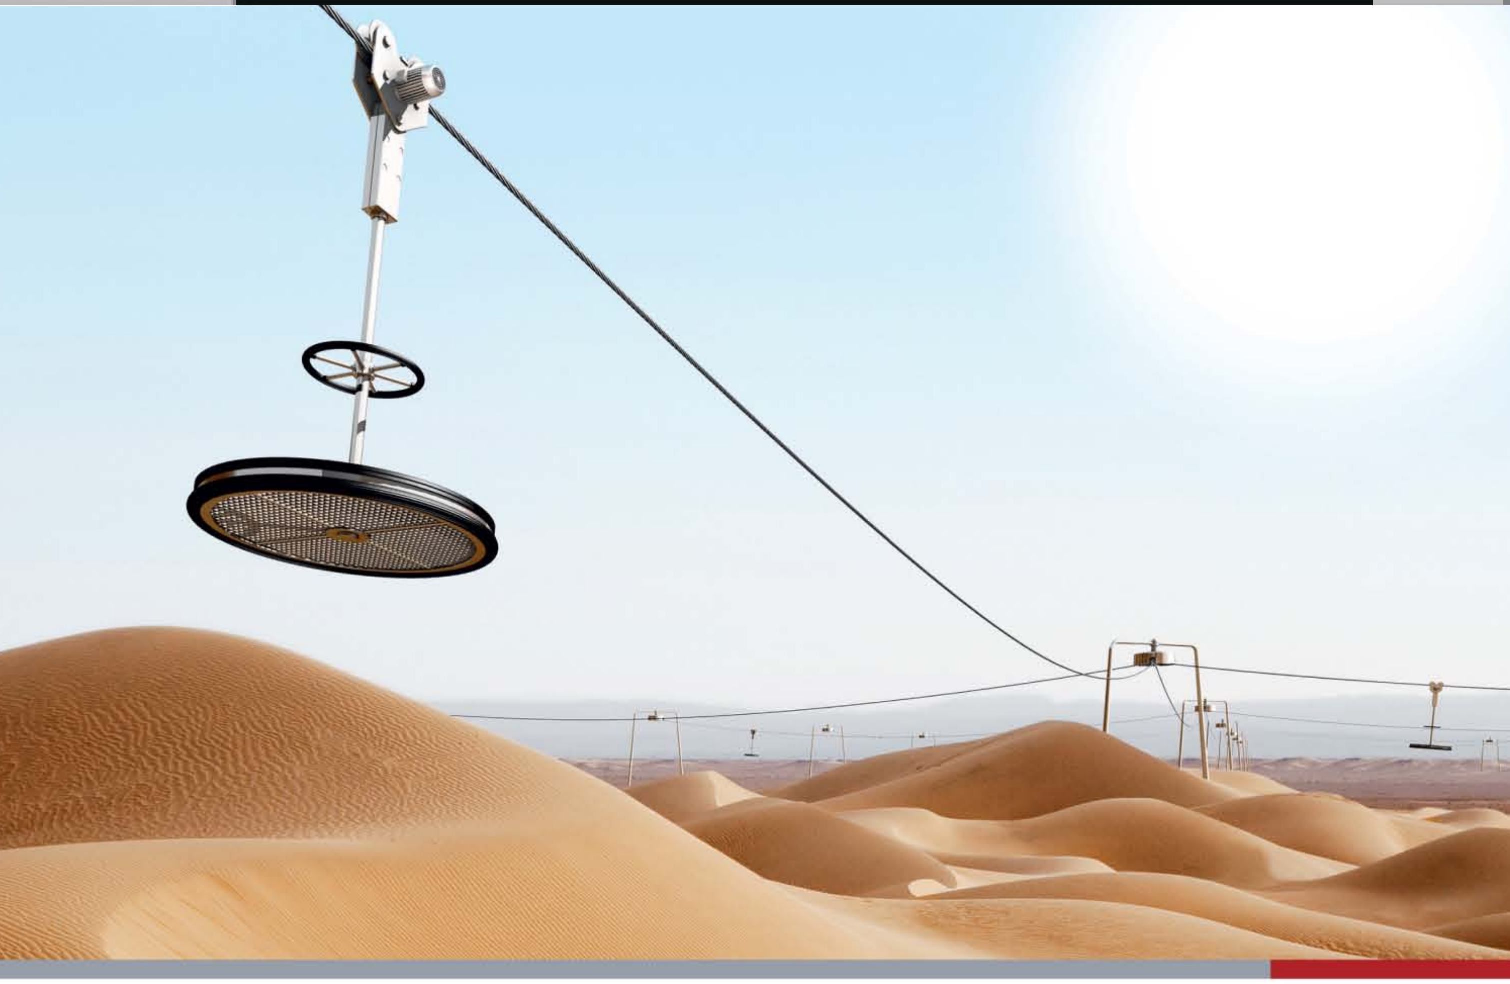

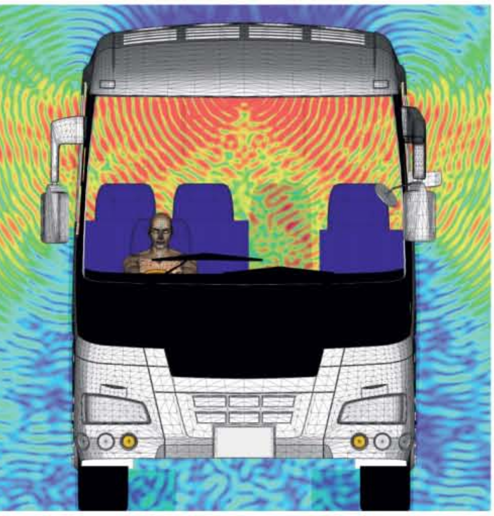

FEATURES_01.17

IEEE SPECTRUM

SPACEX

TOP TECH 2017

IEEE Spectrum's editors survey technological developments that you're sure to hear about in the coming year. P. 21

On the Cover **Illustration for *IEEE Spectrum* by Thirst**

DEPARTMENTS_01.17

IEEE SPECTRUM

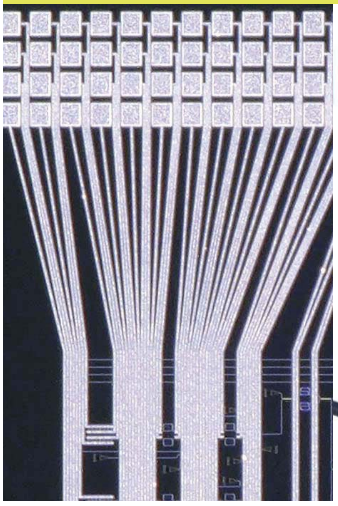

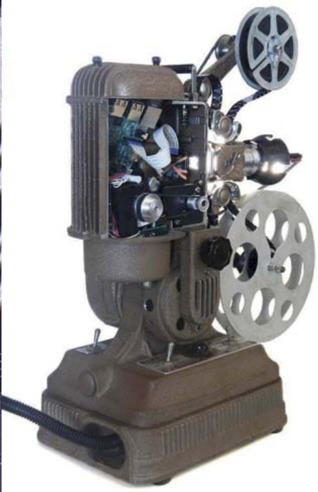

ADDITIONAL RESOURCES

White Papers

Available at spectrum.ieee.org/whitepapers

Engineering Lessons From the Dark Side
In this parody e-book, infamous *Star Wars* characters discuss how Arena's cloud-based PLM solution enabled them to design products that changed the galaxy. *Sponsored by Arena Solutions*

The Role of RTK GPS in Survey Result Accuracy
Explore how high-accuracy GPS equipment can improve the accuracy of survey results. *Sponsored by Swift Navigation*

High-Power Solid-State Pulsed Amplifiers
An attractive alternative to TWTA technology for pulsed-field generation: solid-state pulsed amplifiers that include various frequency ranges and output power levels. *Sponsored by AR*

The Institute

Available at theinstitute.ieee.org

GIVING BACK The IEEE Life Members Committee and several major tech companies support the IEEE Foundation's humanitarian programs, scholarships, and grants.

LIGHTING THE WAY An IEEE Foundation signature program, IEEE Smart Village is empowering entrepreneurs in impoverished areas by helping them start their own solar power companies to bring electricity to local communities.

EVOLUTION OF THE PHOTOVOLTAIC CELL In honor of *The Institute*'s 40th anniversary year and the work being done by IEEE Smart Village volunteers, we're taking a look at the developments of the PV cell over the last four decades.

FROM LEFT: HEWLETT PACKARD ENTERPRISE; JOE HERMAN; ISTOCKPHOTO

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2017 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 54, issue No. 1, 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_

EVAN ACKERMAN

Follow That...Robo-taxi

BEFORE CONTRIBUTING EDITOR EVAN ACKERMAN flew to Singapore to see a trial run of self-driving taxis by a startup company called NuTonomy, *IEEE Spectrum* photography director Randi Klett asked him to take a few pictures. "She told me not to stress out about it," Ackerman [above] says, "but if I could get a photo of the car in motion with an interesting background, that would be great. I thought a shot of the car at night would look cool, but the right effect was going to involve shooting a lot and hoping that some of the pictures would turn out the way I wanted."

On his last night, Ackerman set up his camera gear in the test area in the One-North district of Biopolis, Singapore's biomedical R&D hub. Then he sat and waited for the taxi to drive past. Engineers were making sure that the map in the car's memory corresponded to features of actual streets. This endeavor meant that the taxi would be driving around at random, and there was no way for Ackerman to know exactly where it would be, or when.

"I'd always see the car one or two streets over, or catch it turning away from me while I was running after it from half a block away," Ackerman says. "I started trying to ambush it at stop signs and intersections. Sometimes it would drive along the same road three times in a row, and I'd get in a good position, and then it would disappear for 20 minutes and show up somewhere else."

After chasing the robot car around One-North for several hours, Ackerman did manage to get several good night shots. "It was a lot of work," he says, "but I'm glad that I ended up with some pictures that I can be proud of. I just wish they'd decided to actually use them in the article, you know?" ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, Past Forward is in *IEEE Spectrum*, Vol. 54, no. 1 (INT), January 2017, p. 56, or in *IEEE Spectrum*, Vol. 54, no. 1 (NA), January 2017, p. 64.

01.17

IEEE
SPECTRUM

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

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

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

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

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

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

DEPUTY ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSOCIATE ART DIRECTOR Erik Vrielink, e.vrielink@ieee.org
SENIOR ASSOCIATE EDITORS
Rachel Courtland, r.courtland@ieee.org
Eliza Strickland, e.strickland@ieee.org
ASSOCIATE EDITORS
Celia Gorman (Multimedia), celia.gorman@ieee.org
Willie D. Jones (News), w.jones@ieee.org
Amy Nordrum, a.nordrum@ieee.org
SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

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

INTERN Andrew Silver

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

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

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

ADVERTISING PRODUCTION +1 732 562 6334
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org
SENIOR ADVERTISING PRODUCTION COORDINATOR
Nicole Evans Gyimah, n.gyimah@ieee.org

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

MANAGING DIRECTOR, PUBLICATIONS Michael B. Forster

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

ADVERTISING INQUIRIES
IEEE GLOBALSPEC
30 Tech Valley Dr., Suite 102, East Greenbush, NY 12061
+1 844 300 3098 Toll-free: +1 800 261 2052
www.globalspec.com

VP, DIGITAL MEDIA & ENGINEERING INSIGHT Don Lesem
+1 518 238 6514 don.lesem@ieeeglobalspec.com
VP, SALES & CUSTOMER CARE Peter Hauhuth
+1 303 594 8007 peter.hauhuth@ieeeglobalspec.com
SENIOR DIRECTOR, PRODUCT MANAGEMENT & MARKETING
Christian Noe
+1 518 238 6611 christian.noe@ieeeglobalspec.com
SENIOR PRODUCT MANAGER Linda Uslaner
+1 518 238 6527 linda.uslaner@ieeeglobalspec.com

REPRINT SALES +1 212 221 9595, EXT. 319
REPRINT PERMISSION / LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Managing Editor, IEEE Spectrum.

CONTRIBUTORS_

Edmon de Haro

De Haro, a graphic designer from Badalona, Spain, is known for his clever use of found photographs to create whimsical illustrations, as evidenced in his depictions of the top technologies for 2017, in this issue. De Haro studied at ELISAVA, the Barcelona School of Design and Engineering, where he rubbed shoulders with a lot of engineers. But he doesn't classify himself a technophile. "I'm very interested in technology and its possibilities, but somehow I try to live without being enslaved by it," he says.

Jeff Foust

Foust is a senior staff writer for *SpaceNews* and founder, editor, and publisher of *The Space Review*. To write "Boeing and SpaceX Vie to Fly Astronauts to the Space Station" [p. 42], he visited the two companies' facilities in California and Florida and watched launches of the Atlas V and Falcon 9. He's long dreamed of extraterrestrial adventures, but a career in journalism means it's unlikely he'll be able to afford a trip on even a suborbital vehicle.

Joe Herman

Herman is a Philadelphia-based software developer and actor. In this issue, he writes about a low-cost, high-quality way to digitally transfer a collection of old film reels [p. 13]. Herman's own collection of reels was filmed by his grandfather. "It was a thrilling experience to be able to integrate vintage film equipment, current affordable hardware, and open-source software. At the same time, it allowed me to connect with decades of my own family history," he says.

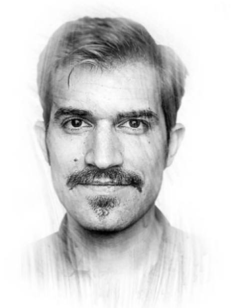

Lucas Laursen

A Madrid-based journalist, Laursen writes in this issue about new ways to fight an old problem: air pollution. Laursen saw his share of smog growing up in Southern California. "Air quality there has improved despite a growing population, thanks to public pressure and improved technology," he says. He hopes the technologies he describes in "Big Data vs. Bad Air" [p. 12] will help less developed regions "make the same changes, only faster."

Vaclav Smil

Smil is a distinguished professor emeritus at the University of Manitoba, in Canada. The latest of his 37 books is *Still the Iron Age: Iron and Steel in the Modern World* (Butterworth-Heinemann, 2016). Smil calls himself a "lifelong interdisciplinarian," whose interests include energy, manufacturing, and demographics. That background gives him plenty of fodder for writing *IEEE Spectrum's* Numbers Don't Lie column, which offers a data-centric look at technological trends. This month Smil takes on storage for the power grid [p. 19].

01.17

PHOTO-ILLUSTRATIONS BY **Gluekit**

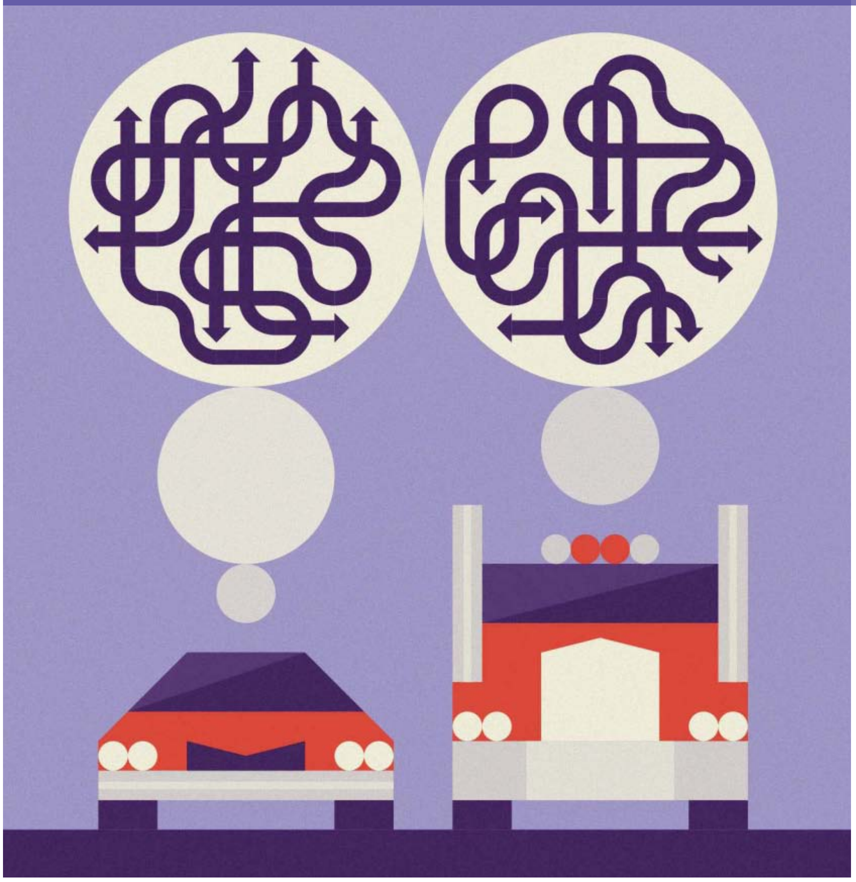

Self-Driving Cars and Trucks Are on the Move

How and when will they affect jobs and drivers?

Former GM R&D chief Larry Burns has likened it to an arms race. But whether you think the advent of self-driving vehicles is going to destroy our economic systems or save our cities, the total automation of driving is certainly going to transform the way we live.

Ford, Google, Mercedes-Benz, Tesla, and Uber, among others, have all boldly declared that they will get fully autonomous cars and trucks on the road in the United States by 2021. At the end of last year the Uber-owned company Otto sent a Budweiser beer delivery from Fort Collins, Colo., to Colorado Springs by autonomous truck. Chinese Internet company Baidu, partnering with Foton Motor Group, introduced its sleek semi-autonomous Super Truck. Daimler tested a driverless truck platoon in Germany. The only place driverless cars don't seem to be turning up anytime soon is India, where, according to Maruti Suzuki chairman R.C. Bhargava, autonomous cars will never be able to keep up with their make-it-up-as-you-go human counterparts.

In this year's Top Tech special report, you'll find a story by Evan Ackerman about nuTonomy, one of the up-and-coming companies jockeying for position in the race to build autonomous cars and fleets. The MIT spin-off has been experimenting with its software- and sensor-loaded Renault Zoes in Singapore since last year, and is also test-driving its nascent self-driving taxi fleet in the Boston area.

Singapore has to be proactive about self-driving cars, as Ackerman points out: The city-state's 5.6 million people are packed into just over 700 square kilometers, making it the third most densely populated country in the world. Roads in Singapore take up nearly as much land as housing does, and as the population keeps growing, building more roads is simply not an option. Boston has similarly run out of road room.

As with any disruptive technology, there's good news and bad news. Rice University's Moshe Vardi, interviewed in advance of the Humans, Machines, and the Future of Work Conference he organized in Houston last month, explained the many positive impacts self-driving vehicles will have on us and our environment, "until you start to think about what it will do to the job market." And that's because, in the United States at any rate, the most common job in more than 50 percent of U.S. states is that of driver. Some 3.5 million people drive for a living. Not to mention the tens of millions of people who work in the infrastructure that supports cars, trucks, and drivers–motels, restaurants, gas stations, you name it–some 50 million by Vardi's estimate.

Technological progress may not be inevitable, but technological change certainly seems to be. Questions about the impact of self-driving vehicles on the global workforce are already the subject of fierce and sometimes hand-wringing public debate, but mostly among sociologists and economists and Silicon Valley gurus. How interesting and useful it would be if the technologists and engineers who are assembling this new self-driving world helped frame these discussions about our driverless future.
—SUSAN HASSLER

ILLUSTRATION BY Greg Mably

NEWS

136 YEARS: TIME IT WOULD TAKE A 2005-ERA CPU TO SOLVE THE "TRAVELING SALESMAN PROBLEM" WITH 85,900 CITIES

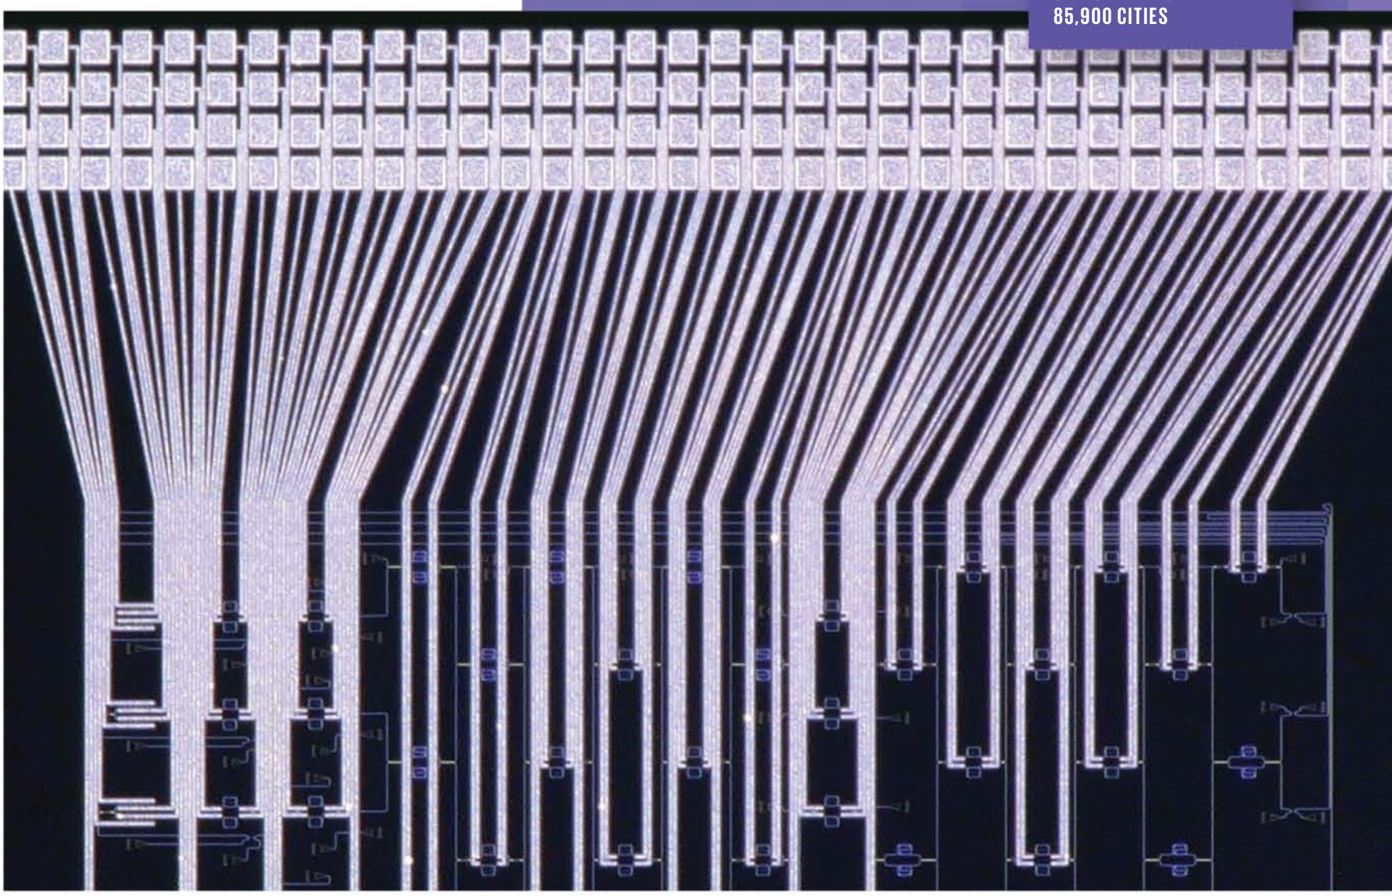

HEWLETT PACKARD ENTERPRISE

LIGHT AND HEAT: A lot of the real estate on this all-optical processor is taken up by wiring for on-chip heaters.

THE ISING ON THE COMPUTER CHIP

A 1,000-component optical processor explores a future after Moore's Law

We may use photons to carry our data, but we rely on the electron to put it to use. One day that division of labor might not be so stark. A team at Hewlett Packard Labs, in Palo Alto, Calif., has built a demonstration chip that could help push some particularly thorny computations into the realm of light, potentially boosting speed and saving energy in the process.

Silicon integrated circuits containing parts that can manipulate light are not new. But this chip, which integrates 1,052 optical components, is the biggest and most complex in which all the photonic components work together to perform a computation, says team member Dave Kielpinski, a senior research scientist at Hewlett Packard Labs (now a part of Hewlett Packard Enterprise, or HPE). "We believe that it is by a wide margin," he says.

The chip, which was developed through the U.S. Defense Advanced Research Projects Agency's Mesodynamic »

Architectures program and was still undergoing testing as *IEEE Spectrum* went to press, is an implementation of an Ising machine—an approach to computation that could potentially solve some problems, such as the infamous "traveling salesman problem," faster than conventional computers can.

The Ising approach is based on a century-old model for how the magnetic fields of atoms interact to give rise to magnetism. The model envisions every atom as having a property called "spin" that prefers to point either up or down. In a ferromagnetic material, above a certain temperature, these spins are oriented randomly and are flipped repeatedly by heat. But when the temperature falls below a certain threshold, the interactions between the atoms dominate, and most of the spins settle down to point in the same direction.

Computers inspired by this model—so-called Ising machines—use such settling actions to arrive at answers to optimization problems. The problem to be solved is entered into the machine by tuning the interactions between its computational elements. These elements, the spins, are designed to be in one of two states and interact with one another until they settle into an optimal configuration that corresponds to a low-energy state.

Stanford University's Yoshihisa Yamamoto pioneered an approach to building such a machine with light. The spins in his system are two phases of light that are 180 degrees out of phase of each other. In 2014, Yamamoto and his colleagues reported the construction of a four-spin machine based on this idea in the laboratory, built with mirrors, lasers, and other optical components.

But scaling up the machine was complicated by a macroscopic effect—ordinary acoustic noise, says Peter McMahon, a member of Yamamoto's group at Stanford. Even the vibrations created by someone emptying a nearby waste bin could cause a subtle expansion or contraction of a delay line, a part of the system that is used to impart delays so that spins can interact properly when they meet. This kind of expansion or contraction could alter the phase of the light enough to disrupt computation. While such perturbations can be corrected, McMahon says, scaling up the system was looking as though it was going to be an impractical and expensive prospect. So in the end, the team altered their approach, introducing electronic feedback into the mix. In October 2016,

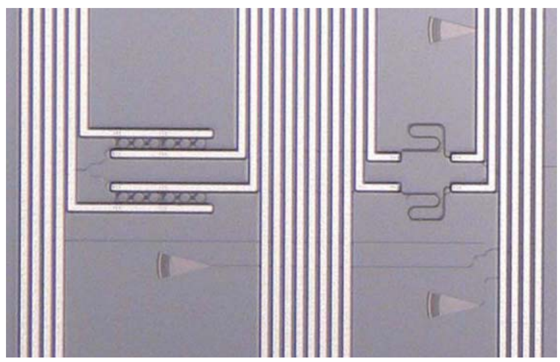

DOWN TO THE WIRE: A close-up of the Hewlett Packard Enterprise chip shows heater wires, microring resonators, wave guides, and optical input/output components.

HEWLETT PACKARD ENTERPRISE

McMahon and his colleagues reported in the journal *Science* that they had used this hybrid optical-electronic system to create an Ising machine with 100 spins. Research reported in the same issue by a number of the same scientists extended the technique to create a more specialized, 2,000-spin computer.

The HPE chip is designed to be a compact approach that doesn't need such electronic feedback. Four areas on the chip, called nodes, support four spins made of infrared light. After the light exits each node, it is split up and combined with light from each of the other nodes inside an interferometer. Electric heaters built into the interferometer are used to alter the index of refraction and physical size of nearby components. This adjusts the optical path length of each light beam—and thus its phase relative to the other beams. The heater temperatures encode the problem to be solved, as they determine how strongly the state of one spin is weighed against another when two beams are combined. The outputs of all these interactions are then condensed and fed back into the nodes, where structures called microring resonators clean up the light in each node so it once again has one of two phases. The light cycles over and over through the interferometer and the nodes, flipping spins between phases of 0 degrees and 180 degrees until the system equilibrates to a single answer.

McMahon says this system could eliminate the vibrational problem faced by the 2014 Ising machine. "If you integrate everything on a small chip," he says, the paths that light takes are etched in silicon. "Those are almost by definition very fixed things," he explains, such that any vibration or temperature swing will tend to affect all the paths that light can take equally. But McMahon notes it's still early days for this approach to computing; the basic physics by which these optical systems arrive at equilibrium still needs to be explored, he says, as does their potential competitiveness with conventional machines.

A key aim of HPE's chip project, Dave Kielpinski says, was to push the limits of photonic chip design. "One of the things we're proudest of is our computer-aided layout tools," he says. He presented the chip in October in San Diego at the inaugural IEEE International Conference on Rebooting Computing, which gathered researchers investigating a number of alternatives that could help keep improving computers as Moore's Law peters out.

In the future, Ising chips such as these might be able to act as accelerators, speedy specialists much like the graphics processing units used in many of today's machines. Kielpinski says the team is investigating designs that could be used to scale up to a larger number of spins. —RACHEL COURTLAND

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

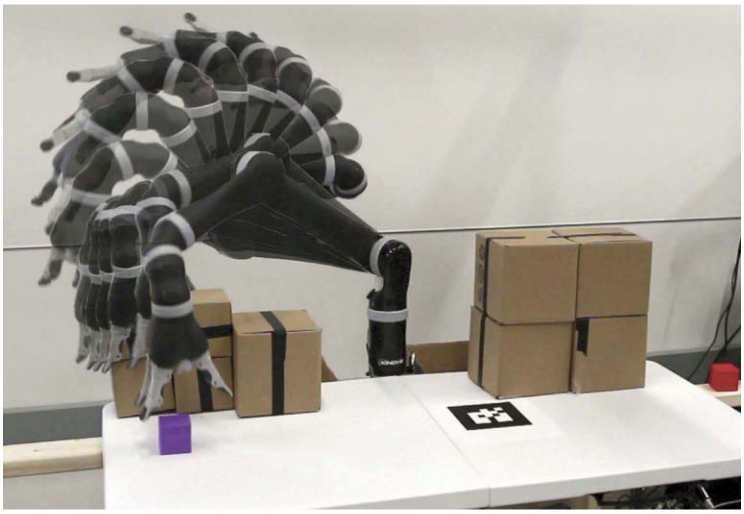

DUKE UNIVERSITY

PICK AND PLACE: A robot in an unfamiliar environment may need several seconds to plan a motion that won't send it crashing into something. A dedicated chip can reduce that calculation to microseconds.

MOTION-PLANNING CHIP SPEEDS ROBOTS

A programmable chip turns a robot's long pauses into quick action

If you've seen a robot manipulation demo, you've almost certainly noticed that the robot tends to spend a lot of time looking like it's not doing anything. It's tempting to say that the robot is "thinking" when this happens, and that might even be mostly correct: Odds are that you're waiting for some motion-planning algorithm to figure out how to get the robot's arm and gripper to do what it's supposed to do without running into anything. This motion-planning process is one of the most important skills a robot can have, and it's also one of the most time consuming.

Researchers at Duke University, in Durham, N.C., have found a way to speed up motion planning by three orders of magnitude while using one-twentieth the power. Their solution is a custom processor that can perform the most time-consuming part of the job—checking for all potential collisions across the robot's entire range of motion—with unprecedented efficiency.

Motion planning for something like a robotic arm usually involves first generating a probabilistic road map, or PRM. A PRM is a graph consisting of points in obstacle-free space, with lines called "edges" connecting points where direct movement between them doesn't result in a collision. Essentially, motion planning consists of picking a starting point and an end point in the PRM, and then figuring out the most efficient path of edges to follow to get from one to the other.

In practice, however, you also have to consider the fact that you're dealing with a physical robot arm, and that when the arm's gripper (the bit you care about) moves from one place to another, the rest of the arm (which you don't care about as much) may run into things. The area that a robot arm moves through is called the swept volume.

This is what motion-planning algorithms struggle with most: determining whether a swept volume will result in a collision with an obstacle, even if the path that the gripper takes doesn't. One study showed that collision detection consumed 99 percent of a motion planner's computing time.

To streamline this process, researchers at Duke used a combination of "aggressive precomputation and massive parallelism." Precomputation happens when you first set up your robot. You generate a single massive PRM that consists of something like 150,000 edges representing possible robot motions while avoiding self-collisions and collisions with things that don't change position, such as the floor.

Unfortunately, 150,000 edges were way too many for Duke's prototype system to work with. A more reasonable number is closer to 1,000, so the researchers had to find a way to prune the PRM down.

To do that, they first simulated some 10,000 scenarios with different numbers of randomly located obstacles of varying size, and then they checked to see which edges in the PRM are used for planning. Infrequently used edges are dropped from the PRM, and after a couple of iterations of pruning and retesting, the PRM in one example was reduced to fewer than 1,000 edges without affecting its ability to solve sample motion-planning problems.

Getting the number of edges in the PRM down to something manageable is important because of the limits of the processor that handles the planning. The processor, a field programmable gate array, is programmed with an array of collision-detection circuits, each one of which corresponds to one of the edges in the PRM. So the size of the PRM is limited to the number of such circuits that can fit on the FPGA—a few thousand at most.

NEWS

In operation, each circuit in the FPGA simultaneously accepts as an input the 3D location of a single pixel in a depth image. The circuit then outputs a single bit that reflects whether the edge that the circuit represents collides with the pixel's location. If there's a collision, the chip removes that edge from the PRM. Once the chip has run through all of the pixels in the image, what's left is a PRM consisting only of collision-free paths. The robot then just picks the shortest one.

The speedup is substantial. No matter how many edges you start with, the FPGA will take just 50 nanoseconds per pixel to determine all the potential collisions. In one particularly complex example, the FPGA took a little over 0.6 millisecond to come up with a plan, and a software-based planner running on a quad-core Intel Xeon processor clocked at 3.5 gigahertz took 2,738 ms–nearly 3 seconds.

Dinesh Manocha, a professor at the University of North Carolina at Chapel Hill who has been working on real-time motion planning with GPUs, agrees that FPGAs have the potential to be much more efficient at motion-planning tasks. "Currently, industrial robots do not use motion planners," Manocha explains. "As robots are increasingly used in new, uncertain environments, the role of motion planning will increase. I feel that this work is very exciting and provides a very practical solution."

The researchers at Duke, including professors George Konidaris and Daniel J. Sorin with grad students Sean Murray, William Floyd-Jones, and Ying Qi, are now exploring ways of applying similar techniques to the next bottleneck: finding the shortest path through the PRM.

A startup called Realtime Robotics will be commercializing the Duke technology. That will involve moving from FPGAs to application-specific integrated circuits (ASICs) that contain much larger PRMs (100,000 edges or more), allowing robots to handle a variety of different environments.

"We've been talking with many companies in the robotics space," Sorin says. "There's great interest in this. Motion-planning software has been a huge limiter to the adoption of robotics, and if you can do real-time motion planning, suddenly robots can now operate in dynamic, unstructured environments. That's what we're hoping to enable."

—EVAN ACKERMAN

Motion-planning software has been a huge limiter to the adoption of robotics. Real-time motion planning would let robots operate in dynamic, unstructured environments

JULIE MARKES

ELECTRIC FIELDS FIGHT DEADLY BRAIN TUMORS

Survival rates are boosted by an oscillating field that attacks dividing cancer cells

Jessica Morris was on a hiking trail in upstate New York last January when she suddenly uttered a line of gibberish and fell to the ground, her body shaking in a full seizure. A few hours later in a hospital she learned that she had glioblastoma, an aggressive brain tumor, and several days after that she was on the operating table having brain surgery. Since then, she's been fighting for her life.

She's grateful to have a radical new weapon in her arsenal, one that only became available to patients like her in 2015. She wears electrodes on her head all day and night to send an AC electric field through her brain, trying to prevent any leftover tumor cells from multiplying. She's been wearing this gear for about six months so far. "I think it's brilliant," Morris says. "I'm proud to wear it."

The Optune device from Novocure, an international company with R&D operations in Haifa, Israel, can't exactly be called convenient or unobtrusive. Morris goes about her business with a shaved head

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

plastered with electrodes, which are connected by wires to a bulky generator she carries in a shoulder bag. Every few days her husband helps her switch out the adhesive electrode patches, which requires reshaving her head, making sure the skin of her scalp is healthy, and applying the new patches.

Morris isn't complaining about the effort. "If you have a condition which has no cure, it's a great motivator," she says dryly.

Doctors typically combat glioblastoma with the triad of surgery, radiation, and chemotherapy. Optune's tumor-treating fields (TTFs) offer an entirely new type of treatment. Unlike chemo, this electrical treatment doesn't cause collateral damage in other parts of the body. Yet the technology has been slow to catch on. "The adoption rate has not been stellar to date," admits Eilon Kirson, Novocure's chief science officer.

He's hoping the most recent results of Optune's biggest clinical trial yet will make the difference: Two years after beginning treatment, 43 percent of 695 patients with glioblastoma who used Optune were still alive, compared to 30 percent of patients on the standard treatment regimen. Four years out, the survival rates are 17 percent for Optune patients and 10 percent for the others. "To patients, that's a big difference," Kirson says. "That's worth fighting for."

NOVOCURE

ELECTRIC LIFELINE: For nearly a year, Jessica Morris [left] has been wearing a system that keeps a brain tumor in check using an oscillating electric field. Patients like the one above must wear the system, which is carried in a bag, all day and night.

Many oncologists, however, still hesitate to prescribe Optune. Wolfgang Wick, a professor of neuro-oncology at the University of Heidelberg, in Germany, has written skeptically about TTFs and says the long-term results don't change his outlook. He draws a contrast with the chemotherapy drug temozolomide, which provides a clear benefit to a subset of patients who have a particular biomarker. Doctors don't know which patients will respond best to the electric fields, he says, and that makes Optune a less appealing treatment option. "If I listen to my patients, this is one thing missing with the TTF, and this has not changed," Wick says.

Novocure executive chairman William Doyle argues that every weapon should be deployed against brain tumors, which are notoriously tough to fight. The last advance in treatment came about 15 years ago when doctors introduced temozolomide. "Since then, every attempt to make an improvement in these patients' survival rates has failed," Doyle says.

Novocure's system uses electrodes stuck to the scalp to create a low-intensity oscillating electric field in the brain, which interferes with cancer cells as they try to divide and multiply. At one moment in the cell division process, the cell distorts into an hourglass shape. That's when the tumor-treating field has its impact because the cell's geometry concentrates the TTF at the center of the hourglass.

The TTF works on molecules inside the cell that are polarized and respond to electric fields. By pulling those polarized molecules out of their proper positions, the field interferes with the precise procedures of cell division. Or, as Doyle puts it, "all hell breaks loose inside the cell." The cells don't divide and may even go into a state of programmed cell death.

So why do the TTFs damage tumor cells while leaving normal cells unharmed? Doyle says the secret lies in the frequency of the electric field. Each different cell type has a membrane with specific filtering properties, allowing only certain frequencies to penetrate it. (You can think of the cellular membrane as a capacitor, Doyle says; at the right frequency, the field can go through it with very little impedance.) Optune uses a frequency of 200 kilohertz to get inside glioblastoma cells, but the frequency doesn't penetrate neurons and other normal brain cells.

It doesn't hurt that there's very little cell division going on in the brain. That's not an advantage Novocure will have as it looks to push Optune into treating cancer in other parts of the body, such as the pancreas, ovaries, and lungs, where normal healthy cells also divide frequently.

Jessica Morris hopes many other brain cancer patients will try the Optune, despite the uncertainties that come with a new medical technology. Her doctor originally recommended a nine-month treatment plan but recently revised that estimate, saying she might want to keep the gear on for two years. That change reflects her promising medical outlook: Her MRI scans haven't shown any new brain tumor growth, and Morris wants to keep it that way. "If I'm still doing well, why would I take it off?" she says. —ELIZA STRICKLAND

BIG DATA VS. BAD AIR

Physics simulations and AI combine to give pollution forecasts to city dwellers in Beijing and beyond

In mid-October 2016, officials from China's Ministry of Environmental Protection counted five illegal trash-burning sites and hundreds of thousands of vehicles exceeding emission standards in Beijing alone. For the first time since last winter's pollution high season, city officials issued a yellow air-quality alert, which required shutting down power plants and reining in Beijing's frenetic factories and road traffic. If this winter is anything like past winters, the city will have to pull out the yellow card again–and may even have to reach for its red card.

This winter, officials will be equipped with forecasting tools from IBM and Microsoft that they tested last year. IBM's tool, used by the city government, is designed to incorporate data from traditional sources, such as the 35 official multipollutant air-quality monitoring stations in Beijing, and lower-cost but more widespread sources, such as environmental monitoring stations, traffic systems, weather satellites, topographic maps, economic data, and even social media. Microsoft's system incorporates data from over 3,000 stations around the country. Both IBM's and Microsoft's tools blend traditional physical models of atmospheric chemistry with data-hungry statistical tools such as machine learning to try to make better forecasts in less time.

"Our advantage or differentiation is to combine all those together," says environmental engineer Jin Huang, who is project manager for the Green Horizon Initiative at IBM Research-China, in Beijing. IBM reports an accuracy of over 80 percent for 3-day forecasts and around 75 percent for its 7- to 10-day forecasts. Microsoft now provides China's Ministry of Environmental Protection with a 48-hour forecast that as of 2015 reached 75 percent accuracy for 6 hours and 60 percent for 12 hours in Beijing.

How best to combine physics models and machine learning for air-quality forecasts is "an active research area," says atmosphere scientist Vincent-Henri Peuch, the head of the European Copernicus Atmosphere Monitoring Service in Reading, England. He adds that blending is the right choice: Both types of models have something to offer and do not need to preclude each other. The market seems to agree so far. IBM now offers its combined model in New Delhi and Johannesburg, and the Beijing startup AirVisual also offers machine-learning-enhanced forecasts for private commercial use.

COULD'VE SEEN IT COMING: A heavy smog day in Beijing in November 2016 might have been predicted by IBM and Microsoft.

IMAGINECHINA/AP PHOTO

Beijing officials have been able to claim some success beating down their fine-particle pollution levels: They reported that 2015 levels were 6 percent below 2014 levels. And while governments are under pressure to reduce air pollution, they are also under pressure not to let economic growth slip. IBM's forecasting tool includes a simulator for measures such as shutting down factories upwind of the city or reducing road traffic for a day or two. "The tool estimates both emissions outcomes and the economic consequences of each proposed intervention," Huang says.

AirVisual, IBM, and Microsoft are all generalizing their software to work in different locations, which requires integrating different local physical models on the one hand but also tuning for differing types of input data and their changing parameters. Johannesburg, for example, has just 8 monitoring stations to Beijing's 35. Still, "there's an opportunity to reuse some of the assets they developed here in South Africa," says computer engineer Tapiwa M. Chiwewe, at the newly opened IBM Research lab in Johannesburg.

Each setting may require its own type of machine learning, a University of British Columbia team reported in 2016. In their study, they found that the computational expense of several types of learning depended on how much data they included up front versus how much data they fed into the program during its operation. The best solution for a place such as Beijing, with just a couple of years of historic air-quality data, may differ from what's best for a city with many more years of historical data, and that poses a challenge for officials trying to choose the right system for their city. It is difficult to compare different models without using the exact same data set at the same location, Peuch warns.

And cities around the world have a long way to go before they bring air quality down to levels recommended by the World Health Organization. In 2015, ambient particulate matter–which does not include tobacco smoke–cost 103.1 million disability-adjusted life years (a measure of the quality and length of human life), according to the 2015 Global Burden of Disease Study in *The Lancet*, making it the sixth most harmful disease risk factor. That makes it an important target for governments and companies. By one estimate, the market for monitoring air quality will grow 8.5 percent per year for the next five years, reaching US $5.64 billion. It seems safe to forecast that the market for air-quality forecasting will grow, too. –LUCAS LAURSEN

RESOURCES

1950: THE YEAR NITRATE-BASED FILM WAS COMPLETELY DISCONTINUED, DUE TO ITS HABIT OF EXPLODING

JOE HERMAN

FRAMES TO FILES: A film projector can be modified to capture old movies by reversing the path that light travels through it.

DIY FILM DIGITIZER

A MODIFIED PROJECTOR CAN CONVERT OLD MOVIE REELS

My grandfather Leo was a self-taught electrical engineer and IEEE member who designed control systems for tire factories. He was also an avid photographer, and his eight children—and later, his grandchildren—were among his favorite subjects, right up to his death in 1974, when I was 5. Fast-forward to 2013: During a move, my uncle uncovered a trove of more than 130 reels of Leo's 8-mm and 16-mm home movies, some dating back to 1939. While commercial conversion services exist, converting so many reels would have been pretty expensive, so my cousin and I set out to preserve them digitally ourselves. ● First, a quick primer on film projection for digital natives. In a projector, a motor pulls film through a vertical "gate." Each film frame is held still and flat in the gate while a lamp illuminates it from within the projector's housing. Lenses on the other side of the gate focus the image so that it appears sharp on whatever surface the film is projected. Between the gate and the lamp, a rotating shutter wheel blocks the light while the next frame is sliding into position. (Without this shutter, the film would be one big blur.) ● For our initial movie conversion attempt, we tried to record digital video directly from a projector. ▸

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

MOVIE MODS

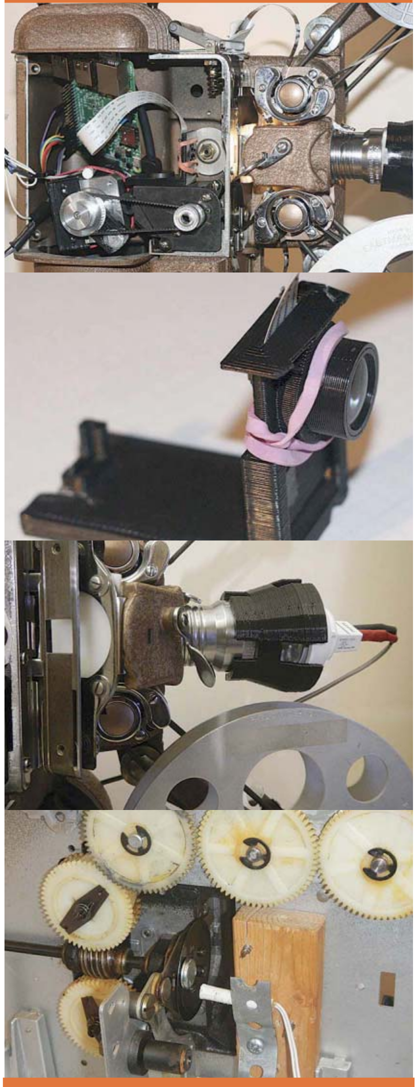

By replacing a projector's lamp and original motor with a Raspberry Pi and a stepper motor [top] frames can be recorded using a camera [second from top] located in the housing and illuminated from the outside via the projector's lenses [second from bottom]. A reed switch detects when the next frame of film is ready to be captured [bottom].

JOE HERMAN (4)

We mounted a DSLR camera to a series of lenses focused on the gate. We started the projector and camera, and everything seemed to be going great. The camera recorded video with excellent detail and color. However, it also highlighted the dirt, cracks, and scratches on many of the reels. Even more problematic, the speed of our projector—as with most—couldn't be precisely controlled. This meant that even under the best conditions, mismatches between the frame rates of the projector and the camera resulted in rolling dark areas, flickering, and other artifacts. These, along with the dirt and scratches, could not be easily remedied digitally.

Going back to the drawing board, we learned that a frame-by-frame transfer—capturing a still image of each film frame and then converting the collection of images into a video—gives far superior results. Still images are easily amenable to digital postprocessing. Such transfers are usually performed on specialized equipment for high fees. Fortunately, I remembered there was a spare Super 8 projector in my basement and an unused Raspberry Pi in my closet, and that a Pi camera was just US $30. Could we do a frame-by-frame transfer with these?

I removed the projector's motor and shutter wheel. I connected a small 60-rpm 12-volt motor to the projector's drive shaft using a section of thin fuel hose. Some basic interfacing electronics let me control the motor's speed and direction with the Pi. Then I attached a small magnet to the drive shaft so that a reed switch would close each time the film advanced one frame. This switch was connected to the Pi so that it would trigger the Pi camera.

The Pi camera's strength is that it's very easy to control using Python software libraries. This is a huge asset for a project requiring precise timing and color and exposure adjustments. But the Pi camera's fixed-focus lens isn't designed for close-up photography, and replacing it with a different lens produces severe color vignetting. Luckily, by carefully unscrewing the lens nearly all the way out of its housing, I was able to bring the Pi camera's focal distance down to about 20 mm. By adding an additional magnifier lens, I could get a nice sharp image of an 8-mm frame in the gate.

Positioning the camera close enough to the gate required mounting it in the projector's housing, so I placed my LED light source "backward," shining down the projector's original lens toward the film. Placing a diffuser (a bit of translucent white plastic from an old ceiling fixture) as close to the film as possible helped to hide scratches and age-related cracks in the film's coating.

Excited by this success, I started capturing films, but it soon became unbearably tedious. I had to run the system more slowly than I expected (under a frame per second) to ensure consistent behavior from the motor and accommodate delays caused by the Pi's slow file system. I also had no way of easily viewing images and adjusting settings midcapture, so I found myself having to recapture entire reels. Most worrying, after a few reels it was clear that the Pi camera's low dynamic range relative to that of the film was wiping out a lot of detail in high-contrast scenes.

A year (and lots of programming) later, though, I've managed to address each of these issues. To deal with the most critical—the high dynamic range of the film—I employ the same approach as smartphones do in similar circumstances. I shoot multiple images per frame with different exposures (two or three is usually enough) and combine them into a single image using functions available in the OpenCV library.

This requires significant processing, so instead of saving images on the Pi, I stream them via Ethernet to a PC running Linux desktop. This also eliminates the bottleneck of the Pi's file system. I further increased capture speed by swapping out my original motor for a more reliable NEMA 17 stepper motor. Using my fastest settings, I can now capture at more than two frames per second.

I built a GUI control panel for the Linux machine using PyQt, which lets me monitor captures on the fly and adjust settings.

Postcapture, I can quickly recombine a reel of images into a movie using one FFmpeg command. When advanced restoration is needed, I use AviSynth, a free video-processing tool for which others have written scripts tailored specifically for film restoration.

I have now built this setup three times: on 8-mm, Super 8, and 16-mm projectors. On the 16 mm I was able to embed the Pi, camera, and motor into the lamp housing to create a portable capture rig. I'm still developing the software, but so far I've captured about 7 hours of footage—nearly half a million frames—and the results are beautiful. It has been a long project, but being able to revisit—and share—such precious family history has made it more than worthwhile.

—JOE HERMAN

ENGINEERS WITH DISABILITIES: INVETERATE PROBLEM SOLVERS

DISABLED ENGINEERS MAKE GREAT CONTRIBUTORS—IF THEY CAN GET PAST THE INTERVIEW

After graduating, mechanical engineer Kurt Driscoll endured more than 100 interviews over 10 months before he was finally hired. A quadriplegic, he encountered some who told him that he couldn't do the job, while others simply claimed to be "going in a different direction."

He finally got hired by an engineering firm through a family connection. He worked there for three years until the company went bankrupt in 2001. His next job hunt was nine months long. "I tried the most direct, in-your-face approach I could think of," says Driscoll. To demonstrate his ability to do the job, he videotaped himself working at his desk. Finally, one interviewer at Faurecia's automotive seating group seemed more intrigued than put off by his disability. (He drives his wheelchair with his chin.)

"I told [the interviewer] I know how to solve problems. I know how to look for solutions. Despite what you're looking at, I'm a go-getter. If I don't know how to do it, I'll learn it and I'll get it done," he says. Faurecia made an offer, and Driscoll is still there.

People with disabilities are underrepresented in STEM (science, technology, engineering, and mathematics) jobs compared with their numbers in the overall population, according to the Bureau of Labor Statistics and the U.S. Census Bureau. But those who succeed share qualities of acceptance, tenacity, and resilience. By necessity, these engineers and coders have well-honed problem-solving skills.

Maggie Hauser, who works in data and information management for the London-based HSBC Bank, says that when she first graduated in 1992, information technology work was mostly solitary. "You went and coded it all in your cube," says the applied computer science major, who is hearing impaired and reads sign language. "It made communication easier."

Fortunately for Hauser, as workplaces became more global and conference calls and videoconferencing routine, technological advances that aid the hearing impaired kept pace. Today, she conducts her conference calls through a relay system, where an interpreter listens to the call and signs what people are saying to Hauser via a video connection on her laptop. In the United States, this service is free, paid for by a surcharge of a few cents on everyone's monthly phone bills.

Hauser had to convince her employer that this system wouldn't compromise company security and confidentiality. She used a separate laptop at work to access the Internet via public Wi-Fi to reach the interpreter, keeping the bank's data inaccessible to the service.

Colleagues were skeptical at first. "I have to come up with the ideas," she says. "Once you show them, they're like, 'That's great.' " Sometimes multiple people speak at once on conference calls. Hauser has to walk a fine line between being able to participate while not being perceived as weak. Often, others on the call will speak up when people are interrupting because they can't hear either.

To prevent miscommunication, she says, "I ask them to send me an email just so I know we're on the same page."

Electronics engineer Alison Kahn, of Boulder, Colo., doesn't miss a beat when new acquaintances underestimate her because of her limp. Some do a double take, some stare, and some are extremely cautious and solicitous. "I don't focus on what people think I can't do. If I did that, I wouldn't do anything," says Kahn, who has cerebral palsy and works on the nationwide public-safety broadband network for the National Institute of Standards and Technology. "I focus on knowing what I can do and showing people that. I'm not going to run a marathon, but the brain is there," says Kahn. One of her feet drags behind the other, so she says she's always tripping over her own feet. To put colleagues at ease, she jokes about it, saying, "If you can't laugh at yourself, you're kind of in trouble." **—THERESA SULLIVAN BARGER**

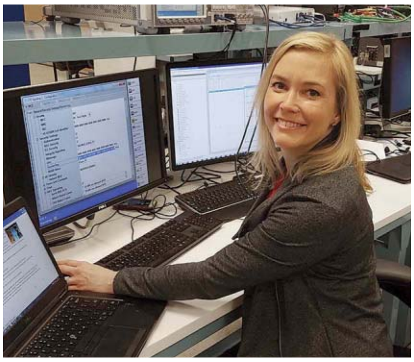

ON THE JOB: Alison Kahn [top], Maggie Hauser, and Kurt Driscoll have overcome skepticism about their abilities.

FROM TOP: NIST; SEAN HAUSER; AARON MAYBERRY/FAURECIA

RESOURCES_BEYOND THE DATASHEET

A CHIP TO PROTECT THE INTERNET OF THINGS

MICROCHIP'S AWS-ECC508 BAKES IN SECURE COMMUNICATIONS

ISTOCKPHOTO

The Internet of Things offers the promise of all sorts of nifty gadgets, but each connected device is also a tempting target for hackers. As recent cybersecurity incidents have shown, IoT devices can be harnessed to wreak havoc or compromise the privacy of their owners. So Microchip Technology and Amazon.com have collaborated to create an add-on chip that's designed to make it easier to combat certain types of attack—and, of course, encourage developers to use Amazon's cloud-based infrastructure for the Internet of Things.

The AWS-ECC508 is an add-on chip designed to make devices more secure—at least for developers using Amazon's IoT cloud. Cloud services are an integral part of the Internet of Things, which is built around the concept of connected objects becoming ubiquitous in our environment, and which must therefore rely on large-scale computing infrastructure.

For example, a smart lightbulb might upload its state to a cloud service operated by the lightbulb maker, and that information would update the lightbulb owner's smartphone app the next time the app is opened. The owner could then use the app to turn the lightbulb on or off as desired, sending the command via the cloud service.

The problem here is that this chain from device to owner and back again represents a potential opening for spoofing attacks: If an attacker successfully fools the cloud service into believing that a fake stream of information is coming from the lightbulb, then the owner could be fed incorrect information, and the attacker may be able to exploit the entry point into the cloud service for even deeper attacks. If the lightbulb is successfully fooled into believing a fake connection is coming from the cloud, then the attacker not only has control of whether or not the lightbulb is on or off but could also plant hostile malware inside the owner's network.

And unfortunately, IoT device manufacturers have been slow to address the problem, primarily "because they are always very sensitive about the cost" of adding better security, explains Microchip engineer Eustace Asanghanwa. "And this is especially true of manufacturers of products who do not see their products as critical. For example, if they are just making a lightbulb [they think] it's not a critical component. If it fails, nobody is going to be harmed. But the truth of it is that once the device is connected, it's not just the value of the device anymore that's at stake. It's the value of what that device is connected to," says Asanghanwa.

The AWS-ECC508 is designed to provide end-to-end security between the IoT device and the cloud infrastructure. It does this by leveraging Amazon's mutual authentication system, which verifies the identity of the cloud service and the device before any data or commands are accepted. The identities are based on cryptographic keys. Until now, creating such cryptographic identities relied on the original manufacturer—typically a contract manufacturer working for a device company—securely generating the keys and then passing the keys securely along the manufacturing chain. Instead, the AWS-ECC508 can generate its own keys that Amazon will accept as authentic.

As suggested by its name, the AWS-ECC508 relies on an "elliptic curve cryptography" algorithm rather than the better known RSA algorithm, which underpins much of the security on today's Internet. "In embedded systems, RSA is very expensive: The key size is very big, and the processes are expensive in terms of power consumption and the time it takes to complete a transaction.... ECC happens to be more efficient and uses less bits, which means less computing resources," says Asanghanwa. The chip is also designed to protect against hardware attacks, such as removing the casing to probe the circuitry or operating it outside normal operating voltages.

For IoT hardware creators looking to experiment with the AWS-ECC508, developer kits are available, and the chip costs around 68 U.S. cents in bulk. **—STEPHEN CASS**

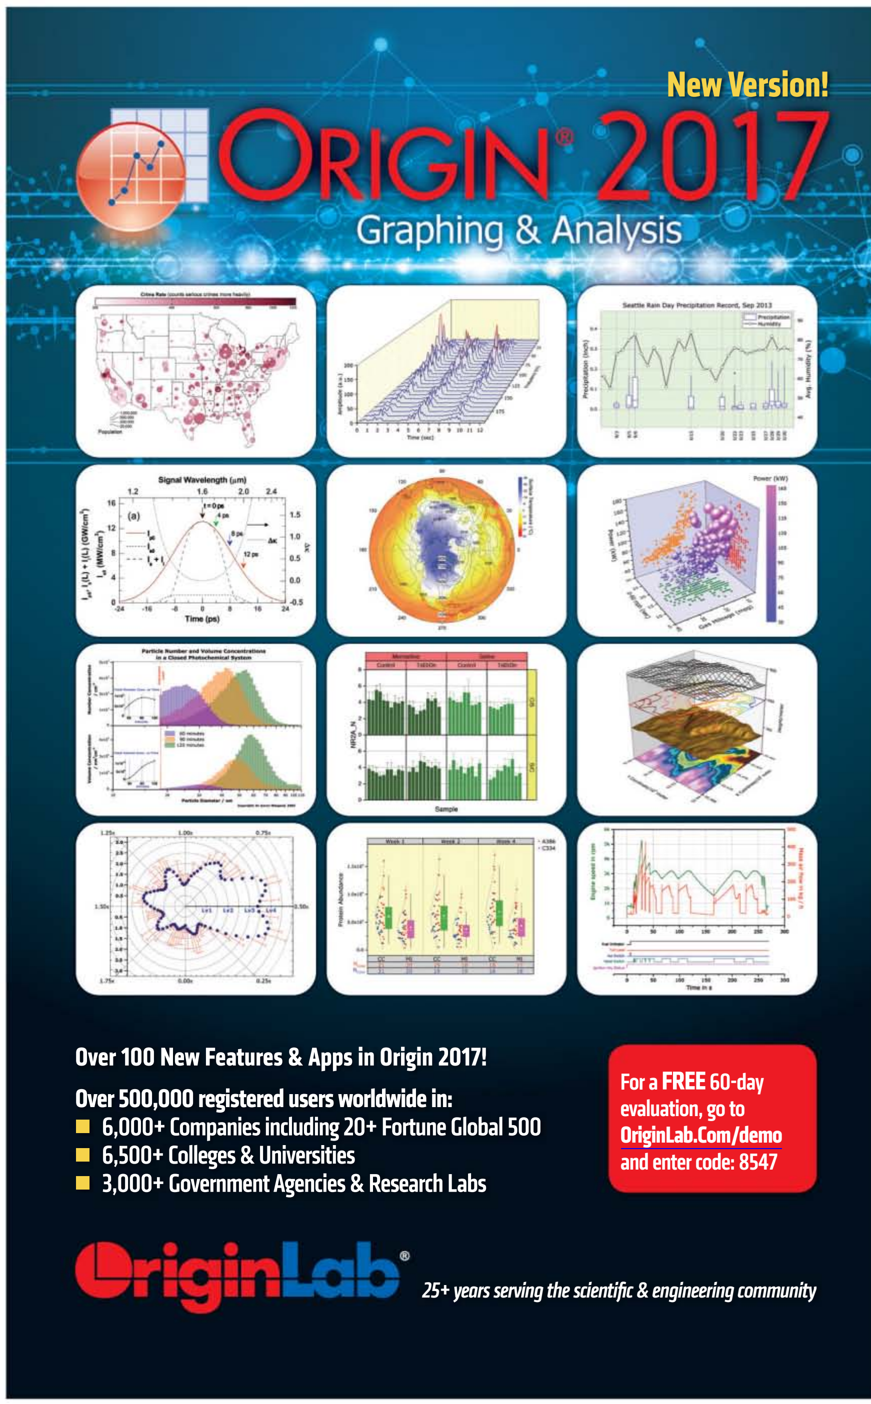

BOWERS & WILKINS P9 SIGNATURE HEADPHONES

OLD-SCHOOL DRIVER TECHNOLOGY STILL HAS WHAT IT TAKES

EMBRACE THE CELLPHONE: Paired with a good external digital-to-audio converter, the P9 headphones are terrific for music on the move.

These are odd times for busy audiophiles. You might have a US $5,000 system at home, but chances are it's gathering dust. Yet music remains a big part of your life—it's just that these days it's mostly coming to you from your phone.

Well, time to embrace the trend and get the best experience you can. The new Bowers & Wilkins P9 Signature headphones are a good place to start.

The $899 P9 is not a discreetly compact headphone of the sort typically designed for mobile use. But by the standards of high-end headphones, it's not large, and it also has a jointed headband that allows it to fold to a somewhat transportable size when not in use. The P9 is also very energy efficient, which means it can easily be driven by small and portable electronics.

To evaluate the sound, I was fortunate to have the help of Leonard Bellezza, who put the resources of his New York City store, Lyric HiFi & Video, at my disposal.

The P9 makes sound by means of dynamic-cone drivers, which means they are basically small versions of regular audio speakers. Though headphones incorporating dynamic-cone drivers are still the most common by far, the upper reaches of the market are becoming increasingly dominated by models based on a different technology, known as planar magnetic. Here, a leading manufacturer is Audeze, in Costa Mesa, Calif.

At the high end, "that's all people are buying these days," says Mike Deutsch, a salesman at Lyric. Planar magnetic headphones use a strong magnetic field to vibrate a very thin and relatively large diaphragm inside the ear cup. The advantage is speed: The diaphragm can respond very rapidly and precisely to changes in the magnetic field. In a well-designed headphone, this precision can translate into more natural "decay" of sounds and, in general, startlingly realistic audio. So I was particularly eager to hear how the P9 compared with comparably priced models from Audeze.

I paired the P9 with several different signal sources, but I found a real sweet spot with my DragonFly Red digital-to-analog converter (DAC)/headphone amplifier ($199). I used the DragonFly to play songs stored with lossless compression on my Samsung smartphone, going back and forth between the P9 and the Audeze EL-8TI, a fully enclosed planar magnetic model with a retail price of $799.

In general, the P9 had a slightly brighter sound. Its response emphasizes bass, but not overbearingly so, like many popular headphones at the moment (I'm looking at you, Beats). I was pleasantly surprised by how well the P9 conveyed the intimacy of spare, moody songs, such as Shelby Lynne's haunting version of "The Look of Love."

The P9 also did very well with large-scale orchestral pieces, such as Sir Simon Rattle's recording of Claude Debussy's *La Mer*. And the P9 really came into its own with rock music. Particularly on rollicking, thumping, hard-driving rock, such as Courtney Barnett's "Elevator Operator," the P9 came across as more gripping and compelling than the EL-8TI.

Greatly encouraged, I decided to see if I could push the P9 beyond its comfort zone. I found I could, but it wasn't easy (or cheap). I compared the P9 to the Audeze LCD-2 ($995), with the music coming from a Moon Neo 260D Transport/DAC ($3,000) playing through a McIntosh MHA100 headphone amplifier ($4,500). With this setup, the available power was for all intents and purposes unlimited, rendering the P9's efficiency advantage irrelevant. I also stuck to quieter, small-ensemble music, that played more to the LCD-2's strengths.

The P9 put up a valiant struggle, but I had to give this round to Audeze. The differences were not very stark: Listening to quiet and complex passages with the P9, the details are all audible and where they should be. But with the open-back LCD-2 those details have a lighter, more natural sound.

Nevertheless, unless you are contemplating spending $7,500 on home-headphone electronics, I unhesitatingly recommend the B&W P9 headphones. The P9, together with the DragonFly Red DAC/headphone amp, are a true dynamic duo and an unbeatable combination. **—GLENN ZORPETTE**

BOWERS & WILKINS

GRID ELECTRICITY STORAGE: SIZE MATTERS

IT WOULD BE A LOT EASIER to expand our use of solar and wind energy if we had better ways to store the large quantities of electricity we'd need to cover gaps in the flow of that energy. • Even in sunny Los Angeles, a typical house roofed with enough photovoltaic panels to meet its average needs would still face daily shortfalls of up to about 80 percent of the demand in January and daily surpluses of up to 65 percent in May. You can take such a house off the grid only by installing a voluminous and expensive assembly of lithium-ion batteries. But even a small national grid—one handling 10 to 30 gigawatts—could rely entirely on intermittent sources only if it had gigawatt-scale storage capable of working for many hours. • Since 2007, more than half of humanity has lived in urban areas, and by 2050 more than 6.3 billion people will live in cities, accounting for two-thirds of the global population, with a rising share in megacities of more than 10 million people. Most of those people will live in high-rises, so there will be only a limited possibility of local generation, but they'll need an unceasing supply of electricity to power their homes, services, industries, and transportation. • Think about an Asian megacity hit by a typhoon for a day or two. Even if long-distance lines could supply more than half of the city's temporarily lowered demand, it would still need many gigawatt-hours from storage to tide it over until intermittent generation could be restored (or use fossil fuel backup—the very thing we're trying to get away from). Li-ion batteries, today's storage workhorses in both stationary and mobile applications, are quite inadequate to meet those needs. The largest announced storage system, comprising more than 18,000 Li-ion batteries, is being built in Long Beach for Southern California Edison by AES Corp. When it's completed, in 2021, it will be capable of running at 100 megawatts for 4 hours. But that energy total of 400 megawatt-hours is still two orders of magnitude lower than what a large Asian city would need if deprived of its intermittent supply. For example, just 2 GW for two days comes to 96 gigawatt-hours.

We have to scale up storage, but how? Sodium-sulfur batteries have higher energy density than Li-ion ones, but hot liquid metal is a most inconvenient electrolyte. Flow batteries, which store energy directly in the electrolyte, are still in an early stage of deployment. Supercapacitors can't provide electricity over a long enough time. And compressed air and flywheels, the perennial favorites of popular journalism, have made it into only a dozen or so small and midsize installations. We could use solar electricity to electrolyze water and store the hydrogen, but still, a hydrogen-based economy is not imminent.

And so when going big we must still rely on a technology introduced in the 1890s: pumped storage. You build one reservoir high up, link it with pipes to another one lower down and use cheaper, nighttime electricity to pump water uphill so that it can turn turbines during times of peak demand. Pumped storage accounts for more than 99 percent of the world's storage capacity, but inevitably, it entails energy loss on the order of 25 percent. Many installations have short-term capacities in excess of 1 GW—the largest one is about 3 GW—and more than one would be needed for a megacity completely dependent on solar and wind generation.

But most megacities are nowhere near the steep escarpments or deep-cut mountain valleys you'd need for pumped storage. Many, including Shanghai, Kolkata, and Karachi, are on coastal plains. They could rely on pumped storage only if it were provided through long-distance transmission. The need for more compact, more flexible, larger-scale, less costly electricity storage is self-evident. But the miracle has been slow in coming. ■

ILLUSTRATION BY Stuart Bradford

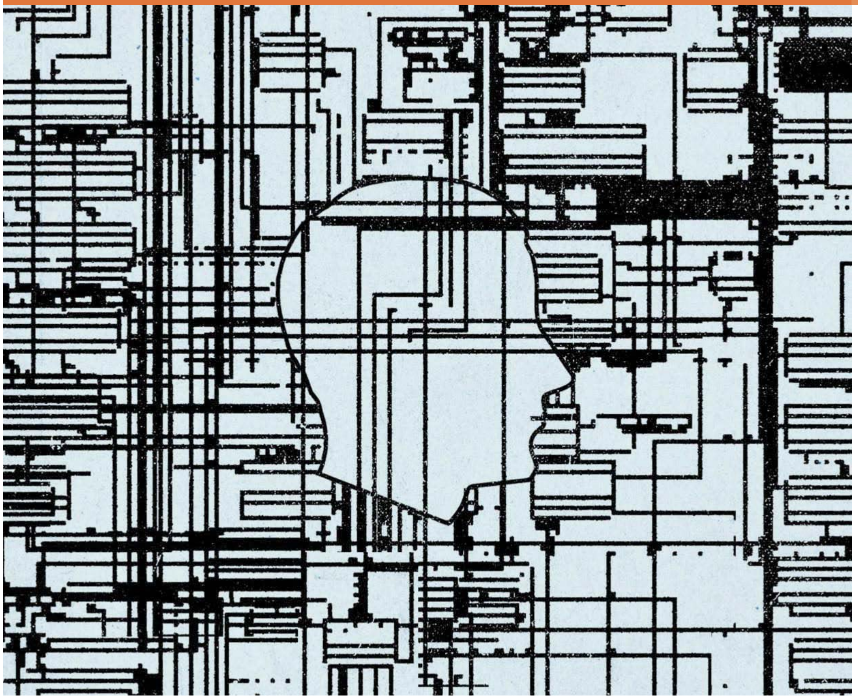

COZYING UP TO COMPLEXITY

I ENJOYED SAMUEL ARBESMAN'S FIRST BOOK, *The Half-Life of Facts*, which was a discussion of the exponential pace of change, as exemplified by Moore's Law, among other things. When I saw the title of his new book, *Overcomplicated*, I assumed that it would be a warning that we technologists had gone too far in creating complex systems. It would advocate moving to simpler systems, just as a doctor might advise an overweight person to go on a diet. I was prepared to argue against such a conclusion, but as I discovered upon reading the book, Arbesman does *not* say that complexity is necessarily bad or that we should seek simplicity. Instead, he maintains that systems are now unknowably complex, that they will become even more so, and we should...just get over it. • Much of the book is spent in discussing the reasons why complexity is inevitably increasing. Arbesman writes that "almost everything we do in the technological realm seems to lead us away from elegance and understandability, and toward impenetrable complexity and unexpectedness." He cites three main factors driving this increase—"accretion," "interconnection," and "edge cases." Accretion is the result of large systems being built on top of smaller and older systems, often via the incorporation of legacy code, producing what we call kludges. As these subsystems become interconnected, the resulting entanglement can change what was simply intricate to truly complex. Finally, complexity is exacerbated by the inevitable existence of edge cases—that myriad of individually negligible exceptions and rarities that yet constitute the long tail of cases that must all be accounted for in system design. The complexity resulting from these factors has passed a tipping point where no single person can fully understand a complete system. • At this point in reading the book I'm thinking that this is all familiar territory to us engineers. The question is: What do we do about it? Arbesman passes quickly by our usual strategies to manage complexity—the abstraction of subsystems to hide the complexity of lower levels, and "good hygiene" in coding. Good things to do, but they won't solve the ultimate problem, argues Arbesman. So what else is there?

Arbesman suggests two approaches. The first is that we need to create more generalists—ideally people with "T-shaped" knowledge bases, those with a deep specialty combined with more superficial knowledge of a broad area. This has become particularly important now that most systems involve divergent specialized fields. However, Arbesman acknowledges that generalists are relatively useless unless accompanied by narrow specialists. Moreover, the market does not at present support generalists. It was easier to be a renaissance man in the Renaissance!

The other approach, which is dealt with at some length, is that we need to think more like biologists than physicists. A physicist would be inclined to attempt a mathematical analysis of the system, hoping for an elegant solution that explains and predicts behavior. On the other hand, a field biologist, accustomed to the overwhelming complexity wrought by evolution, would work from the bottom up, cataloging the appearance and behavior of the systems found, hoping perhaps to identify a new species or ecosystem.

I've been thinking about this biological approach, but I'm not sure how much insight it would provide. Some of our bugs are like distinct new species, like the "Goto" trap in programming, but the overwhelming majority might be one-offs, with little general relevance.

I agree, however, with Arbesman's final conclusion that we should celebrate the functionality and sophistication of our creations. Our achievements will necessarily move beyond the understanding of a single human. After all, evolution has created us, creatures with astounding functionality and resilience alongside esoteric fragilities. Perhaps it will be the same in technology. ■

ILLUSTRATION BY Dan Page

HERE ARE SOME OF THE TECHNOLOGIES YOU'LL BE READING ABOUT THIS YEAR

Photo-illustrations by **EDMON DE HARO**

CARTOONIST BILL WATTERSON of Calvin and Hobbes fame once remarked, "The problem with the future is that it keeps turning into the present." And that is the essence of this annual challenge for *IEEE Spectrum*'s editors: Each January we offer you a preview of some key technological developments that we anticipate for the coming year. We do this knowing full well that, within just a few months, we can be proved completely wrong–as we were last year when NASA's Mars InSight mission was scrubbed. But overall, our track record has been quite good. And now we are pleased to unveil our dozen picks for this year, including artificial hearts, delivery drones, autonomous taxis, and more. We include two more online (at http://spectrum.ieee.org/toptech0117), where you can also tell us what you think of our choices. And, of course, feel free to tweak us if any fall short. ➜

ILLUSTRATION BY Thirst

Hail, Robo-taxi!

AFTER MASTERING SINGAPORE'S STREETS, NUTONOMY'S SELF-DRIVING CARS WILL TAKE ON NEW CITIES

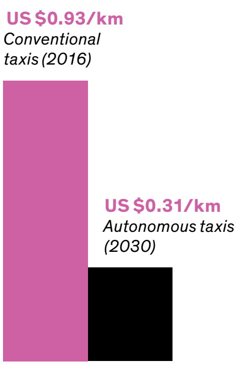

ROBO-TAXIS COME CHEAP

CHART SOURCE: MORGAN STANLEY

TAKE A SHORT WALK THROUGH Singapore's city center and you'll cross a helical bridge modeled on the structure of DNA, pass a science museum shaped like a lotus flower, and end up in a towering grove of artificial Supertrees that pulse with light and sound. It's no surprise, then, that this is the first city to host a fleet of autonomous taxis. ● Since last April, robo-taxis have been exploring the 6 kilometers of roads that make up Singapore's One-North technology business district, and people here have become used to hailing them through a ride-sharing app. Maybe that's why I'm the only person who seems curious when one of the vehicles–a slightly modified Renault Zoe electric car–pulls up outside of a Starbucks. Seated inside the car are an engineer, a safety driver, and Doug Parker, chief operating officer of nuTonomy, the MIT spinout that's behind the project.

The car comes equipped with the standard sensor suite for cars with pretensions to urban autonomy: lidars on the roof and around the front bumper, and radar and cameras just about everywhere else. Inside, the car looks normal, with the exception of three large buttons on the dashboard labeled Manual, Pause, and Autonomous, as well as a red emergency stop button. With an okay from the engineer, the safety driver pushes the Autonomous button, and the car sets off toward the R&D complex known as Fusionopolis.

BY THE END OF THIS year, nuTonomy expects to expand its fleet in Singapore from six cars to dozens, as well as adding a handful of test cars on public roads in the Boston area, near its Cambridge headquarters, and one or two other places.

"We think Singapore is the best place to test autonomous vehicles in the world," Parker tells me as the car deftly avoids hitting a double-parked taxi.

One-North offers a challenging but not impossible level of complexity, with lots of pedestrians, a steady but rarely crushing flow of vehicle traffic, and enough variability to give the autonomous cars what they need to learn and improve.

Riding in an autonomous car makes you acutely aware of just how many

PHOTO-ILLUSTRATION BY Edmon de Haro

potentially dangerous behaviors we ignore when we're behind the wheel. Human drivers know from experience what not to worry about, but nuTonomy's car doesn't yet, so it reacts to almost everything, with frequent (and occasionally aggressive) attempts at safety. If the car has even a vague suspicion that a pedestrian might suddenly decide to cross the road in front of it, it will slow to a crawl.

This mistrust of pedestrians as well as other drivers was designed into the software. "Humans are by far our biggest challenge," Parker says.

Over the course of 15 minutes, our car has to deal with people walking in the gutter, cars drifting across the centerline, workers repairing the road, taxis cutting across lanes, and buses releasing a swarm of small children. Even a human driver would have to concentrate, and it's unsurprising that the safety driver sometimes has to take over and reassure the car that it's safe to move.

TO HANDLE THESE complex situations, nuTonomy uses formal logic, which is based on a hierarchy of rules similar to Asimov's famous Three Laws of Robotics. Priority is given to rules like "don't hit pedestrians," followed by "don't hit other vehicles," and "don't hit objects." Less weight is assigned to rules like "maintain speed when safe" and "don't cross the centerline," and less still to rules like "give a comfortable ride."

The car tries to follow all of the rules all the time, but it breaks the less important ones first: If there's a car idling at the side of the road and partially blocking the lane, nuTonomy's car can break the centerline rule in order to maintain its speed, swerving around the stopped car just as any driver would. The car uses a planning algorithm called RRT*—pronounced "r-r-t-star"—to evaluate many potential paths based on data from the cameras and other sensors. (The algorithm is a variant of RRT, or rapidly exploring random tree.) A single piece of decision-making software evaluates each of those paths and selects the path that best conforms to the rule hierarchy.

By contrast, most other autonomous car companies rely on some flavor of machine learning. The idea is that if you show a machine-learning algorithm enough driving scenarios—using either real or simulated data—it will be able to figure out the underlying rules of good driving, then apply those rules to scenarios that it hasn't seen before. This approach has been generally successful for many self-driving cars, and in fact nuTonomy is using machine learning to help with the much different problem of interpreting sensor data—just not with decision making. That's because it's very hard to figure out why machine-learning systems make the choices they do.

"Machine learning is like a black box," Parker says. "You're never quite sure what's going on."

Formal logic, on the other hand, gives you provable guarantees that the car will obey the rules required to stay safe even in situations that it's otherwise completely unprepared for, using code that a human can read and understand. "It's a rigorous algorithmic process that's translating specifications on how the car should behave into verifiable software," explains nuTonomy CEO and cofounder Karl Iganemma. "That's something that's really been lacking in the industry."

Gill Pratt, CEO of the Toyota Research Institute, agrees that "the promise of formal methods is provable correctness," while cautioning that it's "more challenging to apply formal methods to a heterogeneous environment of human-driven and autonomous cars."

nuTonomy is quickly gaining experience in these environments, but it recognizes that these things take time. "We're strong believers that this is going to make roads much, much safer, but there are still going to be accidents," says Parker. Indeed, one of nuTonomy's test vehicles got into a minor accident in October. "What you want is to be able to go back and say, 'Did our car do the right thing in that situation, and if it didn't, why didn't it make the right decision?' With formal logic, it's very easy."

The ability to explain what's happened will help significantly with regulators. So will the ability to show them just what fix you've made so that the same problem doesn't happen again. Effective regulation is critical to the success of autonomous cars, and it's a challenging obstacle in many of the larger auto markets. In the United States, for example, federal, state, and local governments have created a hodgepodge of regulations related to traffic, vehicles, and driving. And in many areas, technology is moving too fast for government to keep up.

A handful of other companies are testing autonomous taxis and delivery vehicles on public roads, including Uber in Pittsburgh. The motive is obvious: When robotic systems render human drivers redundant, it will eliminate labor costs, which in most places far exceed what fleet operators will pay for their autonomous vehicles. The economic potential of

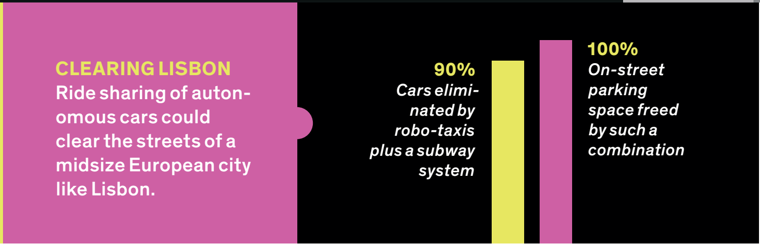

autonomous vehicles may be clear. But what's less clear is whether regulators will approve commercial operations anytime soon.

IN SINGAPORE, the city-state's government is both more unified and more aggressive in its pursuit of a self-driving future. "We're starting with a different philosophy," explains Lee Chuan Teck, deputy secretary of Singapore's Ministry of Transport. "We think that our regulations will have to be ready when the technology is ready." Historically, Singapore has looked to the United States and Europe for guidance on regulations like these, but now it's on its own. "When it came to autonomous vehicles, we found that no one was ready with the regulations, and no one really knows how to test and certify them," says Tan Kong Hwee, the director for transport engineering of the Singapore Economic Development Board.

Singapore's solution is to collaborate with local universities and research institutions, as well as the companies themselves, to move regulations forward in tandem with the technology. Parker says that these unusually close ties between government, academia, and industry are another reason nuTonomy is testing here.

Singapore has good reason to be proactive: Its 5.6 million people are packed into just over 700 square kilometers, resulting in the third most densely populated country in the world. Roads take up 12 percent of the land, nearly as much as is dedicated to housing, and as the population increases, building more roads is not an option. The government has decided to make better use of the infrastructure it has by shifting from private cars (now used for nearly 40 percent of trips) to public transit and car shares. Rather than spending 95 percent of their time parked, as the average car does today, autonomous cars could operate almost continuously, reducing the number of cars on Singapore's roads by two-thirds. And that's with each car just taking one person at a time: Shared trips could accommodate a lot more people.

Over the next three to five years, Singapore plans to run a range of trials of autonomous cars, autonomous buses, autonomous freight trucks, and even autonomous utility vehicles. The goal will be to understand how residents use autonomous vehicle technology in their daily lives. Beyond that, Lee says, Singapore is "about to embark on a real town that we're developing in about 10 to 15 years' time, and we're working with the developers from scratch on how we can incorporate autonomous vehicle technology into their plans." Building new communities from scratch, such as One-North, is a Singaporean specialty.

NO TIP, PLEASE: nuTonomy's test cars bristle with sensors, are always polite, and in principle should be inexpensive.

In this new town, most roads will be replaced with paths just big enough for small autonomous shuttles. For longer trips, on-demand autonomous cars and buses will travel mostly underground, waiting in depots outside the city center until they're summoned. It's a spacious, quiet vision, full of plazas, playgrounds, and parks–and practically no parking spaces.

TO BEGIN TO MEET this challenge, nuTonomy has partnered with Grab, an Asian ride-sharing company, making autonomous taxi services available to a small group of commuters (chosen from thousands of applicants) around One-North. Testing the taxis in a real application like this is important, but equally important is understanding how users interact with the cars once they stop being a novelty and start being a useful way to get around. "People very quickly start to trust the car," says Parker. "It's amazing how quickly it becomes normal."

If all goes well, Parker adds, the company should be ready to offer commercial service through Grab–to all customers, not just preapproved ones–around the One-North area in 2018. At first, each taxi will have a safety driver, but nuTonomy is working on a way to allow a human to remotely supervise the otherwise autonomous car when necessary. Eventually, nuTonomy will transition to full autonomy with the option for teleoperation.

"The whole structure of cities is going to change," Parker predicts. "I think it's going to be the biggest thing since the beginning of the automobile age."

–EVAN ACKERMAN

CHART SOURCE: OECD INTERNATIONAL TRANSPORT FORUM. PHOTO: EVAN ACKERMAN

PHOTO-ILLUSTRATION BY Edmon de Haro

A Make-or-Break Year for Artificial Hearts

MECHANICAL DEVICES THAT PERMANENTLY REPLACE HUMAN HEARTS MAY FINALLY BE READY

THE HUMAN HEART IS A MARVEL OF engineering. Inside the chest of the average adult, that hard-working muscle beats about 100,000 times per day, pumping blood through arteries that branch up toward the brain and twine down to the toes. ● So it's no wonder that biomedical engineers have had a tough time building a mechanical replica to keep patients with heart failure alive and well. Since the 1950s, ambitious researchers have tried to build artificial hearts but have always come up short. Now, four different companies think they've found the right technology, and they're out to prove it. In 2017, clinical trials and animal tests could finally demonstrate that permanent artificial hearts are ready for the clinic.

About 5.7 million people in the United States alone are currently living with a diagnosis of heart failure, meaning their hearts are gradually becoming less effective at pumping blood. Some of the worst-off patients join the waiting list for a heart transplant, but donor hearts are scarce and many people die while waiting.

The Total Artificial Heart from Arizona-based SynCardia Systems already has U.S. regulatory approval as a "bridge to transplant," and now the company is enrolling patients in a clinical trial that's testing the device as a permanent replacement. SynCardia CEO Michael Garippa says the trial is small—just 28 patients—because more than 1,600 temporary placements have already proven that the artificial heart is safe.

Garippa is confident that the device is durable, too, based on the simplicity of its design. "There's nothing electronic inside the body of the patient," he says. SynCardia's heart has two plastic chambers to mimic the heart's two pumping chambers, and each plastic chamber is

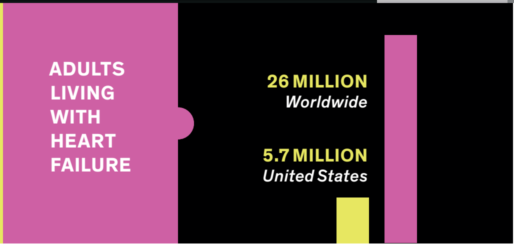

bifurcated by a membrane with air on one side and blood on the other. A patient with a SynCardia heart carries around a 6-kilogram air compressor attached to tubes that penetrate the abdomen to deliver air to the two chambers, pushing on their membranes to propel the blood on the other side. The compressor thumps loudly at a steady rate of 120 times per minute. "It's not a normal life," Garippa says, "but it's way better than these heart failure patients have ever had before."

In France, a company called Carmat is hoping to do better. "Our system is completely silent," says Piet Jansen, Carmat's chief medical officer. Like SynCardia's device, the Carmat heart also has two artificial chambers with membranes that press outward to pump blood. But instead of compressed air, it uses hydraulic fluid driven by an implanted pump. Carmat's heart is larger, heavier, and more complex than SynCardia's device, but its designers are proud of the sensors that determine the patient's exertion level and the microprocessor that calculates an appropriate and changeable heart rate. Wires emerge from the back of the patient's neck to connect to a 3-kg battery pack.

Carmat's first feasibility study seemed rocky: Two out of four patients died within three months. But industry analyst Andrew Thompson, who recently authored a report on artificial hearts, says these patients were extremely sick–as might be expected of people who volunteer for an experimental treatment. "It was not so much a failure of the device as a failure of the body," Thompson says.

European regulators must have agreed, because they approved the major clinical trial that Carmat launched this past August. The company expects surgeons to implant its devices in about 20 patients by the end of 2017 and hopes that its artificial heart will be certified as a permanent replacement device for Europeans in 2018.

KEEPING THE BEAT: To do the job of the human heart, the artificial hearts from Carmat [top] and SynCardia will have to reliably beat about 37 million times each year.

Two other companies not yet at the clinical trial stage have embraced a technical approach that some experts find more promising. Both companies rejected pulsating membranes and instead use centrifugal pumps with whirling, fanlike blades that propel the blood forward, sending a constant flow through the arteries. A device from Cleveland Heart (based on technology developed at the Cleveland Clinic) kept two calves alive and healthy through a 90-day study in 2015. And in Texas, a company called Bivacor is currently conducting 90-day studies with calves in cooperation with the Texas Heart Institute. Both companies are still tweaking their designs and working toward human trials.

Gianluca Torregrossa, a cardiac surgeon who has implanted SynCardia devices and written about the progress of artificial-heart research, is eagerly watching these two companies. Torregrossa says their "continuous flow" designs have fewer points of failure. "If the device has fewer moving parts, you have better chances," he says.

When it comes to clinical trials, all of the technologies have to prove themselves under very tough circumstances. "Doctors don't want to refer a patient to a science project unless the patient has no options," says SynCardia's Garippa. If the technology works for these worst-off patients, the long wait for a reliable artificial heart may be over. The tryouts of 2017 could finally reveal an engineering marvel made by humans, not by biology.
—ELIZA STRICKLAND

SOURCES FOR CHART: CDC, EUROPEAN SOCIETY OF CARDIOLOGY. ILLUSTRATION, THIS PAGE: CARMAT; PHOTO: SYNCARDIA SYSTEMS

TOP TECH 2017

Air Traffic Control for Delivery Drones

ENGINEERS ARE FIGURING OUT HOW TO LET DRONES FLY BEYOND VISUAL RANGE

IN 2013, SHORTLY BEFORE CHRISTMAS, Amazon.com released a video depicting its plans to speed packages to their destinations using small drones. Some commentators said it was just a publicity stunt. But the notion began to seem less far-fetched when Google revealed its own drone-based delivery effort in 2014, something it calls Project Wing. And in the early months of 2016, DHL actually integrated drones into its logistics network, albeit in an extremely limited way–delivering packages to a single mountaintop in Germany that is difficult to access by car in winter.

"It started to get momentum after serious players came in," says Parimal Kopardekar, NASA's senior engineer for air transportation systems, who has been researching ways to work these buzzing little contraptions into an air traffic control system created for full-size aircraft. "We need to accommodate drones."

This past August, the U.S. Federal Aviation Administration (FAA) introduced Part 107, also known as the Small UAS Rule, which allows companies to use small drones in the daytime (or during twilight) and within visual line of sight of the pilot, so long as they are not flown over people who aren't participating in these operations.

This year promises to see the FAA's drone rules loosen even more. At the InterDrone conference in Las Vegas this past September, FAA head Michael Huerta explained that his agency was drafting rules to allow drones to be flown over random bystanders (the FAA calls them "nonparticipants") and that it plans to release proposed regulations to that effect by the end of 2016. "We're also working on a proposal that would allow people to fly drones beyond visual line of sight," he said. Such a move would open the door to the use of small

PHOTO-ILLUSTRATION BY Edmon de Haro

NUMBER OF WAIVERS
For some applicants, the FAA has waived certain parts of its Small UAS Rule (as of 1 December 2016).

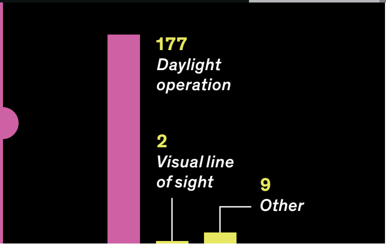

drones to deliver packages, among other things.

Of course, when you start flying drones where you can't see them, you need to put technology in place to be sure that they don't hit anything or injure anybody. While the details of how exactly to do that remain to be hammered out, there is no shortage of ideas.

One of the companies working on this challenge is PrecisionHawk, based in Raleigh, N.C. It's one of just two companies to have obtained a waiver from the FAA allowing it to fly small drones beyond the operator's visual line of sight. For such flights, the FAA does, however, require that an observer be posted to look out for full-scale aircraft.

Still, the waiver increases the range of the company's drone operations from how far away you can see a small aerial vehicle–typically a kilometer or less–to how far away you can see a full-size plane–6 to 7 kilometers. The waiver does not allow "a 200-mile straight-line flight from A to B," notes Thomas Haun of PrecisionHawk. Nevertheless, he's heartened by the "much broader area" the exception permits. The engineers at PrecisionHawk obtained that waiver in part because it had created a system to help drone pilots safely operate a vehicle that they can't directly see.

Avoiding a collision with a full-size aircraft is job No. 1, of course. But the more typical danger is much more mundane–running into a tree or a wall. To avoid that, PrecisionHawk uses satellite imagery to create a detailed terrain model, one of sufficient resolution to capture how high each tree and building is. Its system continually updates that model as new satellite imagery becomes available. Flight-planning software or even the autopilot on the drone itself can then use this information to avoid obstacles.

PrecisionHawk has also worked out a mechanism for drone operators to get updates through Verizon's cellular network on the location of full-size aircraft–the same sort of information that air traffic controllers have. And PrecisionHawk's drones report their positions over that same wireless network, so air traffic controllers and pilots can, in principle, know where these machines are. "What we're providing as a product is primarily the software and data," says Tyler Collins, the creator of this system, which goes by the acronym LATAS (Low Altitude Traffic and Airspace Safety). "We want LATAS on every drone."

PrecisionHawk's system mimics the strategy that is increasingly being used to manage full-size aircraft, whereby those aircraft determine their positions using GPS or some other form of satellite navigation and broadcast that information by radio to everyone else. The equipment for this form of air traffic management, called ADS-B (for Automatic Dependent Surveillance-Broadcast), will be mandatory on most U.S. aircraft by 2020.

WANT IT NOW: Amazon is testing a hybrid drone that takes off and lands vertically but flies like a plane to its destination. Such drones could one day speed packages to consumers.

While it might seem sensible to include small drones in the upcoming ADS-B regime, doing so could easily overwhelm that system, given the huge and growing number of drones–they're selling at a rate of about 2 million a year in the United States alone, according to the FAA. With those numbers growing so fast, an independent scheme for drone-traffic management seems inevitable.

NASA, Google, and Amazon have all been contemplating what such a system should entail. While the concepts that have been outlined vary in many ways, they are all similar in that they would restrict drones to the first few hundred feet above the ground and to locations that are well separated from any airports–that is, to parts of the sky full-size aircraft rarely visit.

At an airport in Reno, Nev., this past October, NASA and various industry partners carried out trials meant to help establish detailed technical requirements for a drone traffic-management system, one that would allow deliveries like the one depicted in that 2013 Amazon video. So whether or not it was a publicity stunt, perhaps this indeed is what the future holds. Haun of PrecisionHawk says, "We actually don't think the future is very far off."
–DAVID SCHNEIDER

CHART SOURCE: FEDERAL AVIATION ADMINISTRATION. PHOTO: AMAZON.COM

Employee of the Month. Every Month

A SAFE, SELF-REPLICATING INDUSTRIAL ROBOT CALLED FRANKA WILL SOON SHOW UP ON FACTORY FLOORS

SAMI HADDADIN ONCE ATTACHED a knife to a robot manipulator and programmed it to impale his arm. No, it wasn't a daredevil stunt. He was demonstrating how a new force-sensing control scheme he designed was able to detect the contact and instantly stop the robot, as it did. ● Now Haddadin wants to make that same kind of safety feature, which has long been limited to highly sophisticated and expensive systems, affordable to anyone using robots around people. Sometime in 2017, his Munich-based startup, Franka Emika, will start shipping a rather remarkable robotic arm. It's designed to be easy to set up and program, which is nice. But what makes it special is that, unlike typical factory robots, which are so dangerous they are often put inside cages, this arm can operate right next to people, assisting them with tasks without posing a risk.

And did I mention that it can build copies of itself?

The robot, also called Franka Emika—"It's like first and last name," Haddadin explains—is not the only one ever designed to operate alongside human workers. Indeed, this type of system, known as a collaborative robot, or cobot, is one of the fastest growing segments in the robotics market, with global sales expected to jump from US $100 million in 2016 to over $3.3 billion in just five years, according to one estimate.

All the big industrial robot makers are trying to develop their own cobots, but the most innovative designs have come from startups. Rethink Robotics introduced its Baxter dual-arm robot in 2012, and more recently it unveiled a single-arm robot called Sawyer. The cobot sector, however, is currently dominated by Danish company Universal Robots, which ships thousands of robots each year. Even so, such robots remain pretty rare.

PHOTO-ILLUSTRATION BY Edmon de Haro

WHO'S BUYING ROBOTS?
Unit sales estimates for industrial robots in 2016 show that Asia—in particular China, South Korea, and Japan—dominate this market.

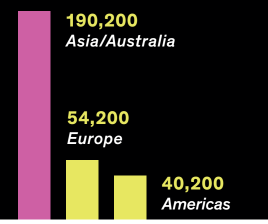

Expect that to change rapidly over the next few years as Haddadin's company—which is financially backed by a group of investors that include German robot maker Kuka—and other firms enter the market.

Haddadin, who's worked at one of Germany's top robotics labs and had a brief stint at the celebrated robotics company Willow Garage in Silicon Valley, says one thing that will set Franka apart from the competition is its manipulation skills. While some of its specs—seven axes of motion, 80-centimeter reach, 3-kilogram payload, and 0.1-millimeter accuracy—are comparable with those of other robots, Franka is designed to perform tasks that require direct physical contact in a carefully controlled manner. These include drilling, screwing, and buffing, as well as a variety of inspection and assembly tasks that electronics manufacturers in particular have long wanted to automate.

Franka has more dexterity than is typical for a robotic arm because it is what is known as a torque-controlled robot. It uses strain gauges to measure forces on all of its seven joints, allowing it to detect even the slightest collisions. In contrast, most industrial robots have no force-sensing capabilities at all—and that's why they are dangerous: They'll take you out and won't even notice it.

One prerequisite for torque control is an extremely detailed model of your robot's dynamics. You need to factor in even the smallest effects, such as elasticity, vibration, and friction in the components. That's because torque control works by comparing actual force measurements on the robot to reference values computed from a model in real time. So if your model is off, your control will be off too.

Haddadin saw that not as a hurdle but as an advantage. "The truth is, I model the hell out of everything I build," he says. Gerd Hirzinger, a pioneer of torque-controlled robots and one of Haddadin's mentors at the German Aerospace Center's Institute of Robotics and Mechatronics, called Franka a "long-yearned-for breakthrough."

Another factor that will make Franka stand out is cost. At the time of this writing, the robot was available for preorder at a yet-to-be-confirmed price of €9,900, or about $10,500. That's a startlingly low figure for such a capable robotic arm. For comparison, Rethink's Sawyer sells for $29,000, and Universal Robots' best-selling UR5 costs even more, at $35,000.

Henrik Christensen, director of the Contextual Robotics Institute at the University of California, San Diego, says Franka is "an impressive piece of hardware." But he adds that with cobots the main challenge is "not just the hardware but also the software to make it easily accessible to nonexperts." Universal Robots, he says, is "beating the competition by having by far the best user interface." So that's an area where Franka will need to prove itself.

Haddadin says his company devoted just as much attention to software as it did to the design of the robot itself. Users can program Franka by moving it with their hands and tapping on a touch screen, with a variety of preprogrammed motions readily available. And once you've created a program for one Franka, you can just copy it over the cloud to one or more other Frankas.

But perhaps the most ambitious part of Haddadin's plan is getting Franka to essentially clone itself. During initial production runs, the robot was performing about 80 percent of the work, but the goal is 100 percent, he insists. Looking further into the future, Haddadin envisions sending containers all around the world as mobile robot factories. "Inside there will be Frankas building Frankas," he says.

Hordes of self-replicating robots popping up everywhere? For whatever it's worth, it's probably a good thing Haddadin is making them very human friendly—even when holding a knife.

—ERICO GUIZZO

SAFE TO TOUCH: Sami Haddadin [left] wanted to create a factory robot that was safe to operate around people, and even be touched by them. The result is Franka Emika, which begins shipping this year.

CHART SOURCE: INTERNATIONAL FEDERATION OF ROBOTICS. PHOTOS: FRANKA EMIKA (2)

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

Augmented Reality: Forget the Glasses

POKÉMON GO IS SHOWING A FASTER ROUTE TO AR SUCCESS

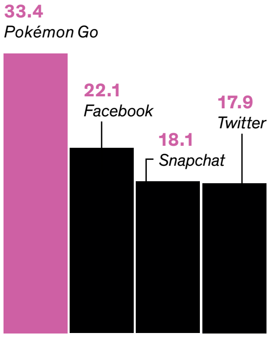

iPHONE USE, IN MINUTES
An analysis for 11 July 2016 shows that the average iPhone user spent more time on *Pokémon Go* that day than on any of the most popular social-networking apps.

CHART SOURCE: SENSORTOWER

IN MID-2014, MAGIC LEAP BEGAN teasing us with visions of realistic baby elephants playing in the palms of our hands, promising to soon unveil a mind-blowing augmented reality technology that would dramatically change the worlds of both entertainment and computing. Investors have ponied up an astounding US $1.39 billion so far to own a piece of this AR future, according to Crunchbase. ● We're still waiting. For a while, it seemed that 2017 was going to be the year of Magic Leap, but the company's technology does not appear ready for prime time, though AR fans are hoping for at least one public demo. ● Meanwhile, a funny thing happened on the way to beta testing. While we were waiting for Magic Leap to show us what's behind its curtain, another startup, Niantic, working with the Pokémon Company and Nintendo, launched a free mobile app in July 2016 that inserted the little critters of the decades-old Pokémon franchise into live scenes on the screens of mobile devices. *Pokémon Go* challenged fans old and young to go out and "catch 'em all!" And it worked: CEO John Hanke, speaking at an Apple event in September, reported 500 million downloads worldwide in just two months, with players collectively walking 4.6 billion kilometers while playing the game. In the United States, *Pokémon Go* quickly beat *Candy Crush Saga*, becoming the most popular game ever.

Since then, many companies have been scrambling in stealth to develop competing AR game apps, and dozens of these apps will likely roll out throughout this year. This is the start of something big, a new form of mass entertainment. It's a watershed moment as significant as the early days of video games.

Consider the episode widely referred to as the birth of today's video-game industry. In August 1972, a little startup called Atari put a prototype coin-op video game, *Pong*, in a bar in Sunnyvale, Calif. A few days later, the engineer who built the game went back to the bar to check out reports that *Pong* was broken. The problem turned out to be an overstuffed coin box.

Now consider last summer. For the first few days after Niantic put *Pokémon Go* out on app stores, its

PHOTO-ILLUSTRATION BY Edmon de Haro

servers around the world crashed regularly, and new users struggled to open accounts on the overloaded systems. Speaking at San Diego Comic-Con the month of the launch, Hanke admitted that the company hadn't been at all prepared for the app to become such a hit.

In the year after the introduction of *Pong*, companies like Sega and Taito announced that they, too, were getting into the video-game business, and with Atari, they built a new industry out of a revolution in entertainment technology.

That, in a PokéBall, is likely to be the story of AR in 2017.

LET'S START WITH a brief explanation of *Pokémon Go*, just in case you've been stranded on a desert island or something for much of the past six months. The game puts you in the role of a trainer, charged with catching little virtual critters—Pokémon—and then bringing them to virtual gyms where they can battle other critters. The virtual creatures are linked to real-world locations, and not randomly: You'll need to be near a body of water to find aquatic Pokémon; an amusement park crammed with flashing lights is a good place to hunt the electric type. To find them, you walk around using a sort of Pokémon "radar" that appears on your smartphone screen. Get close enough and you can see the Pokémon itself; the app combines the virtual image with whatever your phone's camera is seeing, so it looks like the creature is right in front of you, perhaps sitting on your car—or your dog—leading to hilarious photos shared on social media. You catch one by using your finger to toss a virtual ball at it, and it's not as easy as it looks.

SEEING PINK ELEPHANTS: Magic Leap's mysterious artificial-reality technology is the elephant in the room—or in the palm of your hand, as this early demo suggests—but it's unlikely to reach consumers this year.

MAGIC LEAP

That's the gist. There are also more complexities, but the main takeaways are that you "see" virtual creatures in the real world by looking at your smartphone's screen and that you can't do much sitting at home. You have to get out and walk and walk and walk.

Pokémon Go and its coming cohort represent a type of AR technology that some are calling "augmented reality lite." On the opposite end of the spectrum is "augmented reality heavy." That's where you'll find Magic Leap. Here, the ultimate goal is an augmented reality that passes the Turing test for AR: If you see two identical objects sitting on a table, you won't be able to tell which is real and which is virtual. This approach requires some kind of head-mounted display and lots of computing power.

Magic Leap has stated that its technology uses a "dynamic digitized light-field signal" to "generate images indistinguishable from real objects and... place those images seamlessly into the real world." Generally, people interpret this as meaning that the company is projecting images directly onto the retina in the form of a light field—that is, the technology models the direction from which each ray of light from the virtual objects would travel to the observer if the virtual objects were actually at the real-world locations they are supposed to occupy.

A light-field approach for AR would allow a virtual image to be more realistically mixed with real images than in a conventional display.

I have not tried out a Magic Leap prototype (and if I had, I couldn't admit it in print because the company's nondisclosure agreements are that tight). But I have gotten to one degree of separation: I talked to someone who has. He couldn't be specific but admitted that he was indeed blown away. "Before I saw it," he said, "I thought it was much further out. As soon as you see it you say, 'Oh, yeah, how could this not be...'" He trailed off, leaving what it could be to my imagination. A huge success? The future of computer games? The end of computing as we know it?

Also in the AR-heavy category, and already shipping to developers and some business customers, is Microsoft's $3,000 HoloLens, an AR headset with all its necessary computing performed on board. The trade-off Microsoft made to make this technology portable is in its field of view: At less than 45 degrees, it's like looking through a small window. One of Microsoft's competitors, Meta, is taking preorders from developers for a $950 headset with a 90-degree field of view—but the Meta headset requires tethering to an external computer to operate. Both project images

REVENUE IN 2020
By that year, the market for augmented or virtual reality is expected to be half the size of the film market.

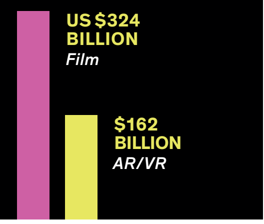

outward, not directly on your retina, as Magic Leap is expected to do. But Magic Leap may not end up being the only AR retina display out there. Kartik Hosanagar, professor of technology and digital business at the Wharton School of the University of Pennsylvania, says he believes that MicroVision, a pioneer in retina displays, may move into commercial augmented reality.

And then, inevitably, there is Apple. CEO Tim Cook told *The Washington Post* in August that the company was doing a lot of things "behind the curtain" with augmented reality. Apple last year purchased the AR startup Metaio, and now the company reportedly has hundreds of engineers working on the technology, including some hotshot researchers recently hired away from Oculus and Magic Leap.

But in spite of all of these efforts, for the average person without a big budget and the willingness to tinker with a technology still under development, AR heavy remains somewhere beyond the horizon. "Magic Leap is not going to be an order of magnitude lower in price than the $3,000 HoloLens," says Hosanagar, "which will make it prohibitive for the mass market."

Niantic chief technology officer Phil Keslin agrees. "I'm not going to wait for Magic Leap," he says in an interview. "I have products to deliver to the world, experiences to get out there. I'm going to use whatever tech is available." Right now, he adds, that's the phone, although he's also keenly anticipating improvements in smart-watch technology. Moving to that platform would let AR users move more seamlessly between the virtual and the real worlds, he says. He's also doing a lot of thinking about audio technology, looking at earbuds and other wearables (like the Bragi earbuds and the Oakley Radar Pace glasses) that have accelerometers and other sensors built in. "We haven't explored an audio-augmented experience as much as we could," he says.

THAT BRINGS US BACK to *Pokémon Go* and the soon-to-come wave of AR-lite competitors vying to be the next big hit. What exactly will that hit be? Wouldn't *you* like to know?

Janet Murray, professor of digital media and associate dean of research at Georgia Tech, whose students seem to be ahead of the curve—they demonstrated their *Harry PottAR* mobile phone game a semester before the launch of *Pokémon Go*—thinks the next big thing to exploit mobile-phone AR could be a time-travel application: As you walk through the real world, you will use your phone screen to peek back at the past. What did that building look like 10, 20, or 100 years ago? Who lived there? Can you follow a fictional character from another era through this parallel world, catching glimpses of intriguing past events, or maybe find the forgotten site of an old bank and collect a few virtual gold coins? I'd download that.

Niantic's Keslin is proud of what his company has achieved with *Pokémon Go.* "We exposed AR technology to a broader base of people. People had seen it in movies, but this made it real. And we demonstrated that there is a market for this, that people want this experience. That will help move the technology forward.

"The vision, when we started, was to get people off their couches and into the real world and meeting other people. We did that. But we were overwhelmed by its popularity this summer. It has been a humbling experience in many ways."

And a personally rewarding one for Keslin. "My wife has a good friend whose 14-year-old son is autistic, and this has changed his life. He used to have to wear noise-canceling headphones out in the world; now he doesn't have to do that—he can focus on the game. He is motivated to go out and play with others. And I've heard that this has had the same effect on other kids."

The number of people playing *Pokémon Go* daily dropped off dramatically in late summer. The app lost some 15 million—about a third—of its daily users, according to a report by Axiom Capital Management. Keslin attributes part of that drop to the start of school. "School took a huge portion of our player base and sucked it into buildings," he says. Some of it is simply boredom. Keslin recognizes that and hopes to keep delivering "new and shiny" experiences to regularly draw players back in.

"People have said that *Pokémon Go* has peaked and is flaming out. I don't think so," says Matthew Szymczyk, CEO of AR content maker Zugara. "It's on track to make $1 billion, and that's just direct revenue, not including what retailers make by having *Pokémon Go* sites at their locations. It's a multibillion-dollar ecosystem."

"It's like a TV show," points out Tawny Schlieski, director of desktop research at Intel. "It's absurd to ask if it's lasting; that's not how content works. Content evolves. *Pokémon Go* will move forward in the same way as a TV show that comes back for another season. It is the start of a new form of gaming that is contextual for your world, that you can't play on the couch," she concludes.

"AR and geolocation is a whole frontier that is waiting to be mapped out and settled," adds Szymczyk. "This is just the beginning."
—TEKLA S. PERRY

CHART SOURCES: STATISTICA (FILM); INTERNATIONAL DATA CORP. (AR/VR)

Face Recognition Tech Goes On Trial

CLASS-ACTION LAWSUITS TARGET THE BIOMETRIC PRIVACY POLICIES OF SEVERAL INTERNET GIANTS

NIMESH PATEL, AGGRIEVED USER OF Facebook and Illinois resident, isn't naive: He well understands that the social networking company collects information about him. But Facebook went too far for his liking when it collected certain intimate details about his physiognomy, such as how many millimeters of skin lie between his eyebrows, how far the corners of his mouth extend across his cheeks, and dozens of other aspects of his facial geometry that enable the company's face recognition software to identify him.

Patel is a named plaintiff in a class-action lawsuit against Facebook alleging that the company's use of face recognition technology violates an Illinois law passed in 2008. The Biometric Information Privacy Act (BIPA) sets limits on how companies can store and use people's biometric identifiers, which the law defines as fingerprints, voiceprints, retina or iris scans, and scans of hand or face geometry. The case is scheduled for trial this October, and similar Illinois-based lawsuits are proceeding against Google and Snapchat. In the upcoming year, the courts will host a debate over who can keep our faces on file.

Civil liberties groups say that debate is long overdue. The Illinois law is a weird outlier in the United States, where face recognition is increasingly being integrated into surveillance systems and law enforcement databases. The technology has rapidly improved in recent years, says Jennifer Lynch, an attorney with the Electronic Frontier Foundation, and regulations haven't kept pace. "We could soon have security cameras in stores that identify people as they shop," she says.

The case against Facebook hinges on a handy photo-tagging fea-

PHOTO-ILLUSTRATION BY Edmon de Haro

IMAGES OBTAINED, BY SOURCE
The FBI's FACES face recognition database mostly contains images of law-abiding citizens taken from driver's license and passport photos.

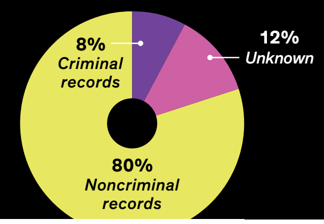

ture introduced in 2010: When a user uploads a photo, Facebook's system automatically picks out any faces in the shot, tries to match those faces to people it's seen in photos before, and offers up the names of any friends it has identified. According to the lawsuit, this "tag suggestion" system proves that Facebook collects and stores "face templates" for its American users. (The company turned off this feature in Europe in 2012 over privacy concerns.)

The Illinois law predates Facebook's introduction of the tag-suggestion feature and doesn't mention social networks. Instead, BIPA cites the potential use of biometric IDs in financial transactions, and notes that these identifiers differ significantly from PIN codes and passwords–if customers' biometric IDs are stolen by hackers, they can't be issued new fingerprints or faces. But the class-action lawyers who have recently seized on the law aren't going after banks; they're targeting tech companies. Yet another lawsuit, settled in April 2016 for an undisclosed sum, took aim at the photo storage site Shutterfly.

Under BIPA, private companies must develop written policies stating how long they will retain people's biometric information and when they will permanently destroy that data. "In a way, this is a modest law," says Claire Gartland, an attorney who works on consumer privacy issues at EPIC, the Electronic Privacy Information Center. "It just requires a disclaimer to the consumer."

By maintaining a database of Illinois users' face templates without a written policy in place, the suit says, Facebook has violated the law. A Facebook spokesperson declined to answer questions about the lawsuit, but notes that users can easily turn off the tag-suggestion feature for their accounts.

"WE COULD SOON HAVE SECURITY CAMERAS IN STORES THAT IDENTIFY PEOPLE AS THEY SHOP"

The legal wrangling has already begun. In late 2015 the company filed a motion to dismiss based on its interpretation of BIPA's list of biometric identifiers, which includes face scans and face geometries yet explicitly excludes photographs and physical descriptions. Facebook argued that the law refers only to physical face scanners that create biometric records based on flesh-and-blood faces. But the court called Facebook's argument "unpersuasive," saying that the law was intended to address all emerging biometric technologies, and allowed the suit to move forward. If Facebook loses the case, the company could be forced to pay damages to millions of Illinois users and change its policies in that state–or, more practically, throughout the United States.

In the courtroom, it's quite possible that the technical aspects of Facebook's face recognition technology will come into play. The courts may need to know whether the company uses the conventional approach to face-matching software, says biometrics expert Anil Jain, a professor of computer science and engineering at Michigan State University. Such systems build and store face templates based on thousands of measurements: "They extract landmark points by sampling across the contours of the face, the eyebrows, the nose, the points along the lips, the two ends of the mouth, and so forth," he says.

But Jain notes that Facebook researchers pioneered a new approach to face recognition that relies on machine learning, introducing their DeepFace system in a 2014 paper. In the report, the researchers describe training their system using a data set of 4.4 million labeled faces drawn from Facebook photographs. The system's deep neural network examined the faces based on millions of parameters, and derived its face-matching rules based on whatever mysterious lessons it learned. "It's more like a black box," Jain says.

Facebook won't say whether it now uses DeepFace, or something like it, for its standard tag-suggestion feature. If the company does employ this advanced method, however, its current technology might not violate the letter of the law. "The question is what they store in the database," explains Jain. As the DeepFace program analyzes raw photographs, the system might simply hold on to the analytic rules it has learned, and might not bother to store face templates that count as biometric identifiers. Therein lies the irony: If Facebook doesn't save faces in its database, it may save face in court.
–ELIZA STRICKLAND

CHART SOURCE: CENTER ON PRIVACY AND TECHNOLOGY, GEORGETOWN LAW

Deeper and Cheaper Machine Learning

SUPERCHARGED HARDWARE WILL SPEED UP DEEP LEARNING

LAST MARCH, GOOGLE'S COMPUTERS roundly beat the world-class Go champion Lee Sedol, marking a milestone in artificial intelligence. The winning computer program, created by researchers at Google DeepMind in London, used an artificial neural network that took advantage of what's known as deep learning, a strategy by which neural networks involving many layers of processing are configured in an automated fashion to solve the problem at hand.

Unknown to the public at the time was that Google had an ace up its sleeve. You see, the computers Google used to defeat Sedol contained special-purpose hardware—a computer card Google calls its Tensor Processing Unit.

Norm Jouppi, a hardware engineer at Google, announced the existence of the Tensor Processing Unit two months after the Go match, explaining in a blog post that Google had been outfitting its data centers with these new accelerator cards for more than a year. Google has not shared exactly what is on these boards, but it's clear that it represents an increasingly popular strategy to speed up deep-learning calculations: using an application-specific integrated circuit, or ASIC.

Another tactic being pursued (primarily by Microsoft) is to use field-programmable gate arrays (FPGAs), which provide the benefit of being reconfigurable if the computing requirements change. The more common approach, though, has been to use graphics processing units, or GPUs, which can perform many mathematical operations in parallel. The foremost proponent of this approach is GPU maker Nvidia.

Indeed, advances in GPUs kick-started artificial neural networks back in 2009, when researchers at

PHOTO-ILLUSTRATION BY Edmon de Haro

DEEP-LEARNING SOFTWARE REVENUES (US $, BILLIONS)
Revenues from deep-learning software should soon exceed $1 billion.

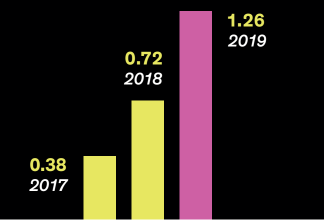

Stanford showed that such hardware made it possible to train deep neural networks in reasonable amounts of time.

"Everybody is doing deep learning today," says William Dally, who leads the Concurrent VLSI Architecture group at Stanford and is also chief scientist for Nvidia. And for that, he says, perhaps not surprisingly given his position, "GPUs are close to being as good as you can get."

Dally explains that there are three separate realms to consider. The first is what he calls "training in the data center." He's referring to the first step for any deep-learning system: adjusting perhaps many millions of connections between neurons so that the network can carry out its assigned task.

In building hardware for that, a company called Nervana Systems, which was recently acquired by Intel, has been leading the charge. According to Scott Leishman, a computer scientist at Nervana, the Nervana Engine, an ASIC deep-learning accelerator, will go into production in early to mid-2017. Leishman notes that another computationally intensive task–bitcoin mining–went from being run on CPUs to GPUs to FPGAs and, finally, on ASICs because of the gains in power efficiency from such customization. "I see the same thing happening for deep learning," he says.

A second and quite distinct job for deep-learning hardware, explains Dally, is "inference at the data center." The word *inference* here refers to the ongoing operation of cloud-based artificial neural networks that have previously been trained to carry out some job. Every day, Google's neural networks are making an astronomical number of such inference calculations to categorize images, translate between languages, and recognize spoken words, for example. Although it's hard to say for sure, Google's Tensor Processing Unit is presumably tailored for performing such computations.

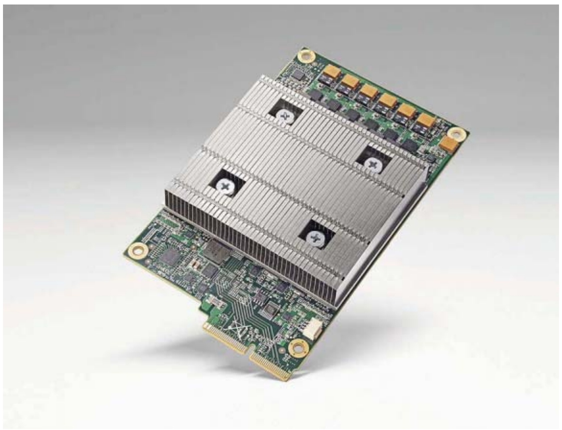

PEDAL TO METAL: Google's Tensor Processing Unit accelerates deep-learning calculations on the company's servers.

Training and inference often take very different skill sets. Typically for training, the computer must be able to calculate with relatively high precision, often using 32-bit floating-point operations. For inference, precision can be sacrificed in favor of greater speed or less power consumption. "This is an active area of research," says Leishman. "How low can you go?"

Although Dally declines to divulge Nvidia's specific plans, he points out that the company's GPUs have been evolving. Nvidia's earlier Maxwell architecture could perform double- (64-bit) and single- (32-bit) precision operations, whereas its current Pascal architecture adds the capability to do 16-bit operations at twice the throughput and efficiency of its single-precision calculations. So it's easy to imagine that Nvidia will eventually be releasing GPUs able to perform 8-bit operations, which could be ideal for inference calculations done in the cloud, where power efficiency is critical to keeping costs down.

Dally adds that "the final leg of the tripod for deep learning is inference in embedded devices," such as smartphones, cameras, and tablets. For those applications, the key will be low-power ASICs. Over the coming year, deep-learning software will increasingly find its way into applications for smartphones, where it is already used, for example, to detect malware or translate text in images.

And the drone manufacturer DJI is already using something akin to a deep-learning ASIC in its Phantom 4 drone, which uses a special visual-processing chip made by California-based Movidius to recognize obstructions. (Movidius is yet another neural-network company recently acquired by Intel.) Qualcomm, meanwhile, built special circuitry into its Snapdragon 820 processors to help carry out deep-learning calculations.

Although there is plenty of incentive these days to design hardware to accelerate the operation of deep neural networks, there's also a huge risk: If the state of the art shifts far enough, chips designed to run yesterday's neural nets will be outdated by the time they are manufactured. "The algorithms are changing at an enormous rate," says Dally. "Everybody who is building these things is trying to cover their bets."
—DAVID SCHNEIDER

CHART SOURCE: TRACTICA. PHOTO: GOOGLE

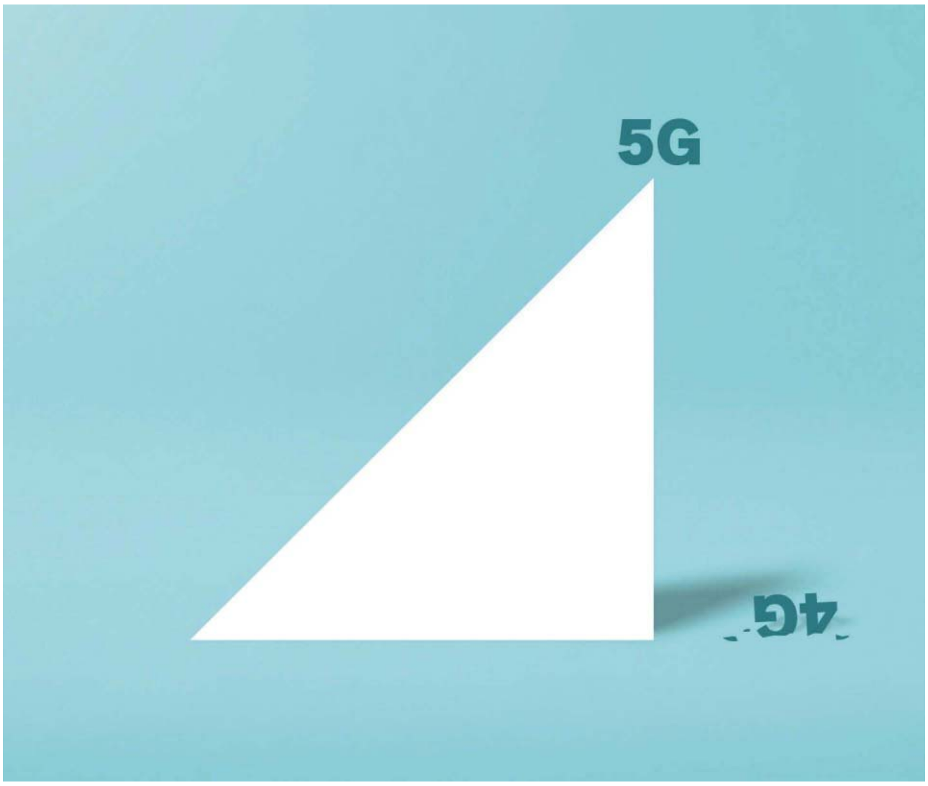

Here Comes 5G—Whatever That Is

CONTROVERSY FLARES AS CARRIERS RUSH TO BRING 5G SERVICES TO MARKET

UNLIKE MOST EPOCHS, THOSE OF the wireless age have come and gone with convenient numerical designations. And with each came marvelous new capabilities: 2G let us all text for the first time, for example, and 3G empowered us to surf the Web.

So now, as phone makers and Internet service providers begin hashing out the details of what 5G will be and how it will work, speculation is high about what new marvels will be possible when carriers can deliver data to smartphones at rates as high as 100 megabits per second. Soon, we'll start to find out: This year, Verizon and AT&T plan to deliver broadband Internet to select homes or businesses using fixed wireless networks built with early 5G technologies. These 5G pilot programs will give the public its first glimpse into a wireless future that isn't due to fully arrive until the early 2020s.

With 5G, carriers hope to deliver data to smartphone users at speeds 10 times as fast as on today's 4G networks, and with only 1 millisecond of delay. Right now, standards groups including the International Telecommunication Union (ITU) are in the middle of an eight-year process to set technical specifications and performance parameters for 5G that should be ready around 2020. "At this time, everything is open," says Colin Langtry, who oversees 5G for the ITU.

For many invested in the 5G rollout, that extended timetable could be a problem. By the time standards are finalized in 2020, carriers and governments will have already poured US $5 billion into 5G development. With so much cash on the line, and facing pressure from data-hungry customers, carriers are moving fast. Over the past year, companies have

PHOTO-ILLUSTRATION BY Edmon de Haro

GLOBAL CELLULAR ACCESS
Ninety-five percent of the world's 7.4 billion residents have access to a cellular network. Among them, quality of service varies considerably.

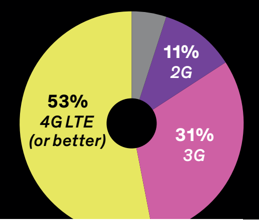

completed a flurry of lab tests and trials to figure out what types of radios, antennas, and signal processing techniques will work best to deliver 5G in hopes of bringing those technologies and their capabilities to market as soon as possible.

"5G is really an amorphous term," says Thyaga Nandagopal, a program director for the National Science Foundation who oversees wireless initiatives. "But it's increasingly going to become very narrowly focused on a few technologies, very fast."

Both Verizon and AT&T say their 2017 fixed wireless networks will rely on millimeter waves, which are arguably the hottest new 5G technology. Millimeter waves are officially defined as waves transmitted at frequencies between 30 and 300 gigahertz, and they are between 1 and 10 millimeters in length. That's much shorter than traditional cellular signals, which have frequencies below 6 GHz, typically far below, with wavelengths in the tens of centimeters.

Millimeter waves could solve a big problem for today's carriers: Customers want to transmit more data at once, but the traditional, sub-6-GHz bands are more crowded than they've ever been. In contrast, millimeter waves are part of a wide-open section of spectrum, which offers intriguing possibilities.

In early tests, Verizon homed in on one band in particular—28 GHz, which has become a clear favorite for 5G after the U.S. Federal Communications Commission opened it and a handful of other high-frequency bands for commercial use in July 2016. AT&T, too, has tested millimeter waves at 28 GHz and on another popular band, 15 GHz, at its trial site in Austin, Texas. In the lab, AT&T has achieved peak data speeds of 14 gigabits per second at 15 GHz.

Other companies have also had encouraging results. Last year, Sprint demonstrated peak data speeds of 4 Gb/s—good enough to stream a virtual-reality demo for spectators—on that same 15-GHz band at the 2016 Copa América Centenario soccer tournament, in Philadelphia. In Japan, NTT Docomo has incorporated another 5G technology called MIMO (multiple-input, multiple output), in which dozens of programmable antennas are made to send and receive signals at once from a single base station, to reach a blazing 20 Gb/s at 15 GHz. At that speed, a complete 2-hour, 1080p, high-definition movie can be transmitted in a second and a half.

Now, Verizon and AT&T are eager to share with customers the speedy data rates they have seen in early tests, despite the fact that comprehensive 5G standards are still years away. But though these two companies may be the first to bring millimeter waves to customers, their fixed wireless networks will fall short of what many experts consider "true" 5G.

Unlike the networks that connect smartphones, fixed wireless systems send a focused beam to connect one stationary point to another, such as a base station to a rooftop antenna. From there, carriers run Ethernet cables from the antenna to deliver broadband Internet to offices or apartments within a building.

To fully achieve 5G, carriers and smartphone manufacturers must also figure out how to deliver high-speed data to mobile users who are riding in cars or trains or walking on sidewalks. And the grandest vision for 5G extends far beyond mobile devices—to autonomous cars, connected appliances, and industrial robots.

There are other problems, too. Millimeter waves don't easily penetrate obstacles such as buildings, and they are more readily absorbed than traditional microwave cell signals by water and oxygen molecules in the air. So they require more power to travel the same distances as the signals from today's smartphones. Given these issues, using millimeter waves to serve mobile users will require more sophisticated signal processing and a greater density of base stations than are available today.

Some critics think fixed wireless deployments are an unfortunate distraction at a time when companies should be focused on developing these other capabilities. "Effectively, that's delaying mobile 5G," says Paul Struhsaker, chief technical officer for the investment group Carnegie Technologies.

Competitors also warn that Verizon and AT&T may have to forfeit their early fixed wireless investments if other companies turn up even better technologies in time, and industry experts fret about fragmentation across carriers. "The problem with prestandard implementation has always been you risk that your investment after three years is outdated," says Günther Ottendorfer, Sprint's chief operating officer for technology.

Still, Sanyogita Shamsunder, Verizon's director for network planning, says the company will press ahead with its bold 2017 plan for 5G through fixed wireless. "We see a reasonable use case that we think we can address with this technology, so we're going to do it," she says.
—AMY NORDRUM

CHART SOURCE: INTERNATIONAL TELECOMMUNICATION UNION

Boeing and SpaceX Vie to Fly Astronauts

THE WINNING CONTRACTOR WILL BUILD NASA A SPACECRAFT—AND FLY IT

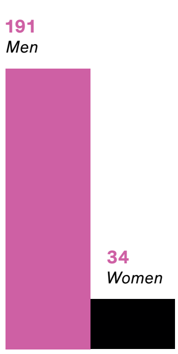

VISITORS TO THE INTERNATIONAL SPACE STATION More than 200 astronauts have visited the International Space Station (as of 29 October 2016).

CHART SOURCE: WIKIPEDIA

WHEN CHRIS FERGUSON commanded the final space shuttle mission in July 2011, he and his crew members left behind something on the International Space Station: a small American flag that had flown on the first shuttle mission three decades earlier. "It will hopefully maintain a position of honor until the next vehicle launched from U.S. soil brings U.S. astronauts up to dock with the space station," Ferguson explained in a call with President Obama during the mission. ● "I understand it's going to be sort of a 'capture the flag' moment here for commercial spaceflight," Obama responded. ● As the director for crew and mission operations at Boeing, Ferguson is one of the leaders of that company's effort to develop and operate a spacecraft to carry NASA astronauts to the space station—and capture that flag. Boeing is competing with SpaceX for that privilege, and both are racing the clock to get those vehicles flying before a NASA contract with Russia runs out.

Much more than scheduling is at stake. If spaceflight can escape the hothouse of government contracting to become a truly commercial enterprise, price competition may at last bring costs down to earth. And 2017 may see some crucial steps taken along that path.

NASA'S COMMERCIAL Crew Program dates back a decade, when the agency began helping aerospace companies to develop commercial cargo vehicles. Rather than doing this under a typical government contract, with NASA designing the spacecraft and a contractor building it, NASA sought to have the companies create spacecraft on their own. NASA would provide financial and technical support, and the agency would later buy cargo-delivery services rather than purchase the spacecraft themselves.

That effort, known as Commercial Orbital Transportation Services (COTS), included an option to provide ferry services for crew. Although SpaceX described how it could satisfy that option in its COTS proposal, NASA chose not to select it, deciding to focus on the cargo services it knew it needed once the shuttle retired. Under the COTS program, SpaceX and Orbital Sciences Corp. (now Orbital

PHOTO-ILLUSTRATION BY Edmon de Haro

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

ATK) both eventually developed vehicles, which are now transporting supplies, experiments, and other equipment to and from the station.

As part of a broader shake-up of the agency's plans for human spaceflight, NASA announced in 2010 that it was moving ahead with development of crew transportation systems using the same approach as it did with COTS. The agency made a series of awards to several companies, including Boeing and SpaceX, to support early development of key technologies for those vehicles.

In September 2014, NASA issued the final set of contracts to Boeing and SpaceX–contracts with a combined value of US $6.8 billion. Besides completion of the vehicles, those contracts also include two test flights for each company, one with and one without a crew, and as many as six missions each to ferry two people at a time to and from the station. And both companies have been hard at work.

Boeing's CST-100 spacecraft, called the Starliner, and SpaceX's Crew Dragon spacecraft, a.k.a. Dragon v2, look remarkably similar. Unlike the shuttle, both are capsules that will launch atop rockets. Both are designed to carry up to seven people, although NASA anticipates sending only four at a time; the companies may use the additional seats for other, commercial missions.

There are, however, considerable differences. The CST-100 will launch atop an Atlas V from United Launch Alliance (a joint venture of Lockheed Martin and Boeing). This same rocket has flown some 60 consecutive successful missions since its introduction in 2002.

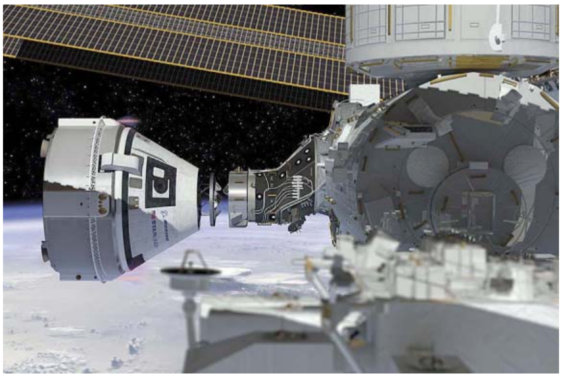

BOEING

YOUR PIZZA, CAPTAIN: Boeing's CST-100 will touch down where no NASA flight has ever gone: dry land.

"It's not unlike a stock Atlas V," Ferguson says of the CST-100's launch vehicle. The rocket will feature a few differences, such as the use of an upper stage with two engines instead of one, and the addition of a system to detect imminent failures, which will be linked to the capsule's abort system, allowing it to escape in the event of a problem.

On its return to Earth, the CST-100 will go where no NASA capsule has gone before: dry land. While the Mercury, Gemini, and Apollo capsules all splashed into the ocean, the CST-100 will touch down on land, using retrorockets, parachutes, and airbags.

That strategy is important for Boeing because it intends to reuse each CST-100 spacecraft up to 10 times. "We felt we really needed to land on land to get reusability," says John Mulholland, Boeing's CST-100 program manager. Splashing down in the ocean, he says, creates problems, such as salt water getting into the capsule and corroding components.

In addition, touching down on land makes it easier and faster to reach the crew, using trucks and helicopters, much as Russia does on its Soyuz landings.

Unlike Boeing's capsule, SpaceX's Crew Dragon will launch on a rocket of the company's own design, the Falcon 9, which has flown nearly 30 times since entering service in 2010. A Dragon cargo spacecraft failed during a launch, and a Falcon 9 exploded on the launchpad this past September.

The Dragon spacecraft itself, however, has performed well on its cargo missions to the International Space Station. Moreover, SpaceX designed the Dragon from the beginning to be capable, with modification, of carrying people as well as cargo.

"A big difference between Crew Dragon and the cargo Dragon is that we've added these SuperDraco thrusters," says Benjamin Reed, director of crew mission management at SpaceX. Those thrusters, mounted on the side of the capsule, will serve as the abort system for the spacecraft. The company tested this equipment in May 2015, when a Dragon capsule flew off its Florida launchpad using only its SuperDraco thrusters, splashing down a short distance offshore.

But SpaceX has additional plans for those thrusters. Initial Crew Dragon missions will splash down in the ocean, just as the cargo version of Dragon does today. According to Reed, splashdowns shouldn't prevent the capsule from being reused. "Our team has some of the world's experts in parachute landing and water recovery," he says. SpaceX plans to start reusing the cargo version of the Dragon spacecraft in the coming year.

Ultimately, though, SpaceX plans to use the Crew Dragon's SuperDraco thrusters to allow a powered landing on a pad after descending by parachute. Reed says SpaceX will first

CHEAPER THAN CAVIAR
The cost of placing 1 kilogram in orbit using SpaceX's Falcon Heavy rocket is less than the cost of 1 kilogram of caviar.

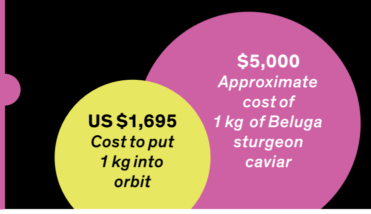

demonstrate that "propulsive landing" capability on later cargo missions: "We'll get very comfortable with flying and doing propulsive landing with cargo first, and then crew."

The two companies run their programs differently. Boeing makes extensive use of partners. It is leasing a hangar at the Kennedy Space Center previously used to refurbish space shuttles, turning it into an assembly plant for CST-100 capsules. Boeing also hired NASA's Mission Operations Directorate, which runs the iconic Mission Control Center in Houston, to take on a similar role for CST-100 flights.

SpaceX, on the other hand, is doing just about everything itself. It's building both the capsule and rocket, and it will also handle mission control for its crewed missions, using a facility at its headquarters in Hawthorne, Calif., which currently runs cargo Dragon missions.

"We're not just building a crewed Dragon," says Reed. "We're building a whole system that needs to be certified to fly humans safely."

SINCE NASA AWARDED contracts to Boeing and SpaceX in 2014, both companies have been busy developing their respective spacecraft. That work has included parachute and other landing-system demonstrations, SpaceX's pad abort trial in 2015, and various other tests.

But schedules have slipped. NASA originally hoped to have commercial crew vehicles ready to fly by the end of 2015. Spending cuts by Congress, though, prompted both by skepticism about the program as well as tightening budgets, pushed that target to the end of 2017.

And this past September NASA's Office of Inspector General warned that more delays were likely. The problem was no longer caused by money issues–Congress was now fully funding the program–but by technical ones.

"Notwithstanding the contractors' optimism, based on the information we gathered during our audit, we believe it unlikely that either Boeing or SpaceX will achieve certified, crewed flight to the ISS until late 2018," the report concluded.

When the report came out, Boeing's schedule had already slipped by a few months, with a crewed test flight planned for early 2018. "We'll be ready for flight services by the middle of 2018," Ferguson said in mid-September.

A month later, Boeing announced it was delaying its schedule by half a year. That crewed test flight, previously planned for February 2018, is now scheduled for August 2018, with the first operational mission expected in December 2018. Boeing blamed the delay on problems with aerodynamic loads on the Atlas V created by the CST-100 capsule, as well as other snafus.

"We've gone through a lot of hurdles on our technical development," Mulholland said in mid-October, after Boeing announced the schedule change, adding he was optimistic that the company could keep to the new schedule.

SpaceX has been more circumspect about its timetable. This past July, according to official NASA schedules, the company said it planned to carry out an uncrewed test flight of its Crew Dragon in May 2017, followed by a crewed test in August 2017. That schedule, though, predated the Falcon 9 pad explosion that is delaying many missions, including those test flights.

And in mid-December, NASA announced that SpaceX was delaying its uncrewed demonstration flight until November 2017 and its first crewed flight until May 2018. So no astronauts will fly in the coming year.

NASA is feeling the pressure. It needs to have at least one company ready to start flying astronauts to the space station by the end of 2018, when its contract to use Russia's rockets for those flights expires.

Extending the contract with Russia is not an option, says William Gerstenmaier, NASA associate administrator for human exploration and operations. "They require a three-year lead time," he says, meaning it's too late now to purchase seats on Soyuz flights for missions in early 2019.

Will Boeing and SpaceX be ready to fly by the end of 2018? "It's still feasible," Gerstenmaier says. The future of the station depends on one or the other capturing that flag. The commercialization of spaceflight requires that both companies, and others besides, join the fray.

–JEFF FOUST

FLYING OUR WAY: SpaceX built both the Crew Dragon capsule and its rocket alone, without partners.

CHART SOURCE: CHRISTIAN SCIENCE MONITOR. ILLUSTRATION: SPACEX

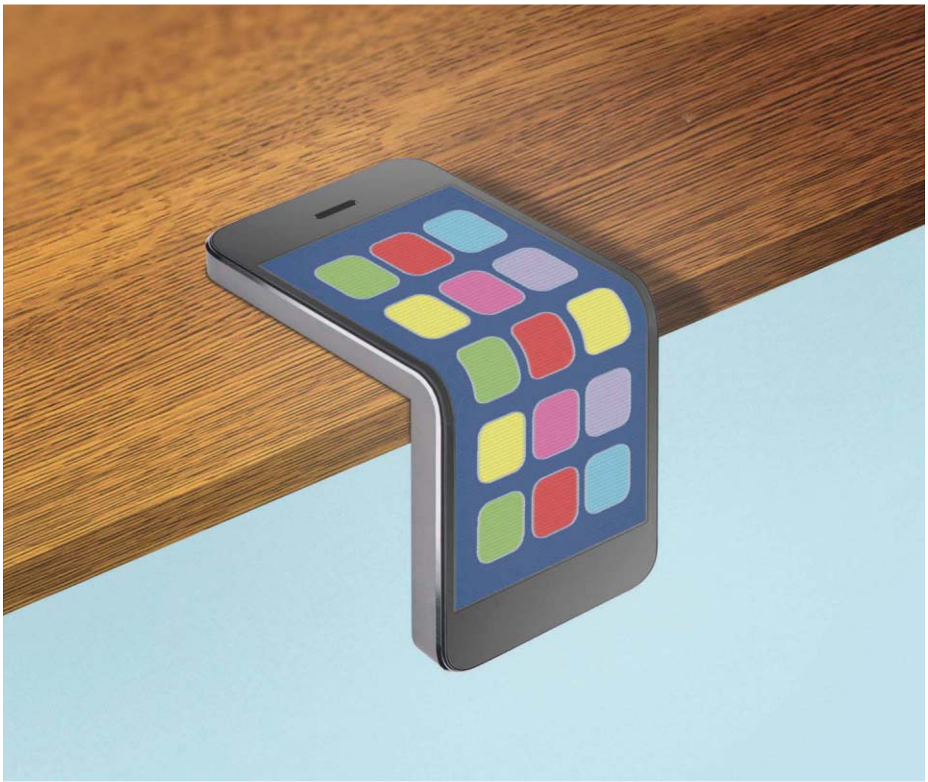

Fold-Up Screens Could Make Their Big Debut

IN 2017, SAMSUNG WILL LIKELY RELEASE A SMARTPHONE THAT TRANSFORMS INTO A TABLET

THE RUMORS HAVE BEEN swirling for months. Though they couldn't be confirmed, their persistence suggests that something significant may be coming from Samsung, possibly as early as this year: a foldable mobile. ● Today, the world of mobiles consists of two major realms–tablets and smartphones. Tablets are good for reading magazines and books, typing long messages on a linked keyboard, looking at pictures, and surfing the Web. Smartphones are good for texting and talking. Engineers have long dreamed of merging the two.

Such a device would morph from one to the other by folding: Open, it's a tablet, but by bending or folding it in half you'd transform it into a phone. "You can expect to open up your phone and double the screen real estate," says Roel Vertegaal, a computer scientist at Queen's University in Ontario. Besides the versatility, you'd have interesting new possibilities–imagine bending your phone to flip ahead in an e-book, just as you would flex a novel's covers to jump ahead a few pages.

Samsung has pursued flexible designs for at least four years, going so far as to develop "artificial muscles" that push and pull a smartphone's components into new positions to prevent damage as it bends. Now, according to media reports, the company may finally be ready to share those technologies with the world and save users the hassle of carrying both a phone and a tablet.

"Having that bimodality in a device would, I think, be really game changing," says mobile analyst Wayne Lam at IHS Markit. "You're not only creating a new form factor for the phone, but you're also cannibalizing other product categories."

Competitors are thinking along similar elastic lines. At a trade show last summer, Lenovo showed

PHOTO-ILLUSTRATION BY Edmon de Haro

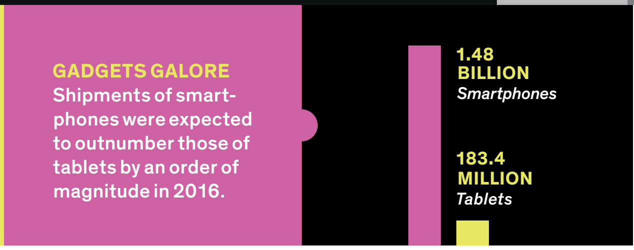

off a concept product for a smartphone that folded around a user's wrist into a wearable device. Throughout 2016, a Chinese manufacturer named Moxi Group promised a limited release of its own high-end flexible smartphone. But Samsung would be the first of any major company to debut a device with a truly flexible screen.

If Samsung does release such a phone, it would signal the first major departure from the flat, rectangular form that has defined smartphone designs since Apple released the first iPhone in 2007. Manufacturers have experimented with curved glass and adopted larger screens, but essentially all smartphones today are design descendants of that original iPhone.

The simple, rigid smartphone has endured partly because the challenges of building a foldable screen that is rugged and dependable are great. Stiff components such as the battery must be made to either bend along with the screen or be situated away from the fold.

Vertegaal himself built a flexible smartphone in his lab last year and tested hundreds of screens before settling on one that worked–a high-definition organic light-emitting-diode screen produced by LG Display. OLED screens contain a thin film of organic compounds that produce light from an electric current right at the surface of the device. They are a favorite of designers working on flexible TV and mobile units because they do not require the bulky backlight and filters found in LCD screens.

Samsung happens to be the largest global supplier of OLED panels. In 2013, the company showed off a concept product with a bendable OLED screen at the CES electronics show. It set off a frenzy in the tech blogosphere and led to speculation that the company would release a smartphone based on it.

Vertegaal says the biggest challenge in building his own flexible phone was powering all the pixels in his LG display with connectors that could withstand repeated bendings. To keep it simple, he used a relatively primitive screen that had only 720 pixels. He realized that the rigid materials found in conventional smartphones are, unfortunately, quite delicate. "Circuits are made out of metals, and these metals break under stress," he says. "While it's possible to make these bendable screens, it's difficult to make them in a way that they don't break."

One solution may be to use printed electronics to integrate razor-thin circuits and flattened antennas along the surfaces of a smartphone. In theory, this technique could make phones more flexible by reducing the number of large components and fragile attachments within the device. However, the easiest way to create such products is with injection molding, a process that is seldom used in smartphone manufacturing.

Right now, only two companies in the world have the expertise and production chops to manufacture a smartphone with a bendable display for the mass market: Samsung and LG, says William Stofega, a mobile analyst at International Data Corp. Just last year, at CES, LG exhibited an OLED screen, less than 1 millimeter thick, that could roll up like a newspaper. But Stofega says the time, complexity, and expense of manufacturing means that any flexible products that debut this year will likely be pricey, high-end devices.

FLEXIBLE DEVICES: This smartphone designed by Chinese manufacturer Moxi Group can wrap around a user's wrist.

Samsung needs a hit to regain momentum after last year's Galaxy Note 7 fiasco, in which it coped with reports of dozens of the smartphones catching fire. Ultimately, the problems prompted a recall that slashed profits by 17 percent, or US $4 billion, in that quarter. A flashy line of foldable phones could help the company rebuild its reputation. However, it would be a high-risk strategy, Stofega notes. "No one wants to risk coming out with a device that looks pretty cool and then, after about 2,000 bends, just cracks right in half," he says.

Samsung wouldn't comment on its plans for 2017. So we'll all have to wait and see if the company dazzles us this year with a couple of flexible smartphones–or leaves the many design headaches and teething pains for its rivals to endure. –AMY NORDRUM

CHART SOURCE: INTERNATIONAL DATA CORP. ILLUSTRATION: MOXI GROUP

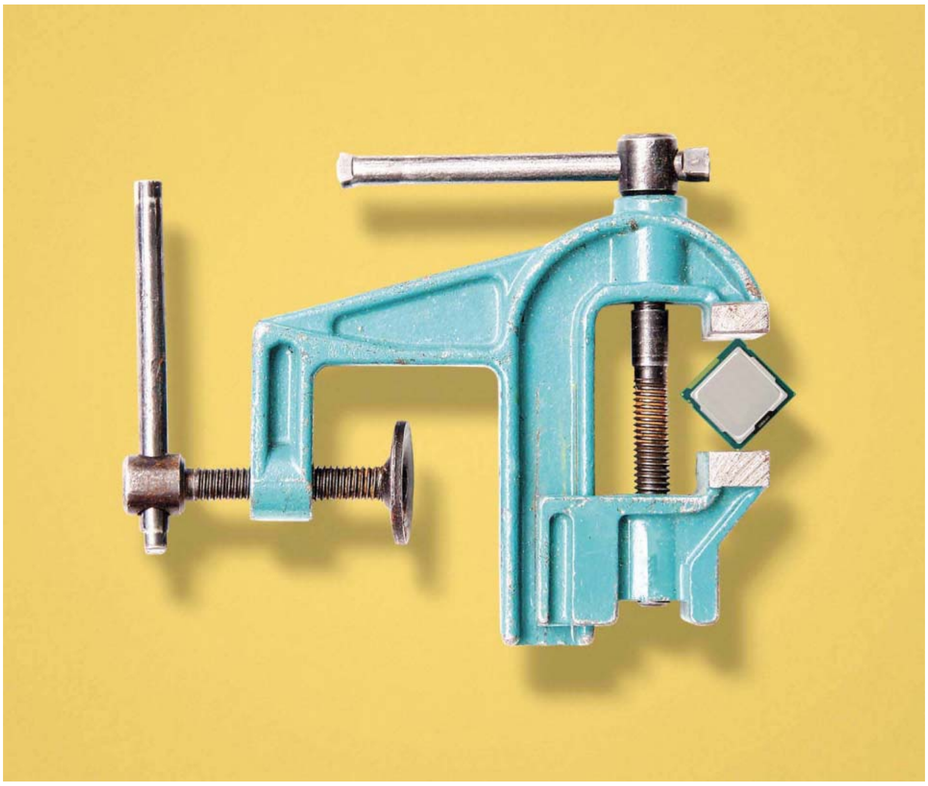

Moore's Law's Next Step: 10 Nanometers

IN 2017, INTEL HOPES ITS MANUFACTURING EDGE WILL KEEP IT AHEAD OF THE COMPETITION

THESE DAYS, FORECASTS ABOUT the future of Moore's Law tend to look quite gloomy. But Intel's outlook—at least for the next few years—is decidedly bright.

Sometime in 2017, Intel will ship the first processors built using the company's new, 10-nanometer chip-manufacturing technology. Intel says transistors produced in this way will be cheaper than those that came before, continuing the decades-long trend at the heart of Moore's Law—and contradicting widespread talk that transistor-production costs have already sunk as low as they will go. In the coming years, Intel plans to make further improvements to the design of these transistors. And, for the first time, the company will optimize its manufacturing technology to accommodate other companies that wish to use Intel's facilities to produce chips based on ARM architecture, which is nearly ubiquitous in modern mobile processors.

Although it wasn't always the case, today there is little about the name of a chip-manufacturing generation, or "node," that aligns with the dimension of any particular feature on a chip. But the transistors on Intel's 10-nm generation will still be denser than those on today's 14-nm chips—as well as other companies' 10-nm offerings, says Intel senior fellow Mark Bohr. Intel has released few figures to date on the dimensions of the new generation. But both the length of transistor gates, which turn transistors on and off, and the distance from one gate to the next—a figure known as gate pitch—will be smaller this time around. The minimum gate pitch will go from 70 nm to 54 nm. And logic cells—transistor combinations used to perform standard logic functions—can be less than 46 percent

PHOTO-ILLUSTRATION BY Edmon de Haro

ARM ON THE RISE
Annual sales of chips based on ARM technology exceed world population by a wide margin. No wonder Intel is gearing up to fabricate such chips.

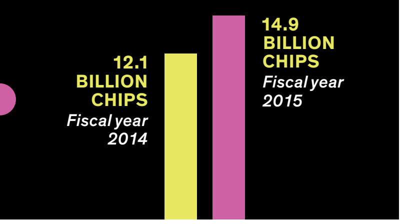

the size of those built on 14-nm technology.

This is a more aggressive level of miniaturization than in years past, says Bohr, and it helps counteract a recent trend: a slower cadence to the introduction of new chip-manufacturing generations. "One important message is that this node, and the products that we'll be making on it, will hopefully dispute some of the concerns of the industry that Moore's Law is slowing down," Bohr says.

Even though the cost of producing a wafer full of chips will be higher at 10 nm than at 14, Intel says the cost per transistor will be lower. The 10-nm chips are also expected to deliver improvements to switching speed and energy consumption. As has been the case for years already, clock speed isn't liable to increase, though. "It's really power reduction or energy efficiency that's the primary goal on these new generations, besides or in addition to transistor cost reduction," Bohr says. Improved compactness and efficiency will make it more attractive to add more cores to server chips and more execution units onto GPUs, he says.

In the past, Intel upgraded its transistors once every couple of years, with the introduction of a new manufacturing generation. But at 14 nm, the company offered what might be called a deminode: a set of improvements to its 14-nm transistors before the debut of 10 nm. At 10 nm, the company aims to introduce two of these deminodes (10 nm+ and 10 nm++) before it introduces its next manufacturing generation at 7 nm. This is in part a reaction to the increasing challenges in moving to the next chip generation, Bohr says: "If it's going to take you longer to [get to] the 10-nm or to the 7-nm generation, then the next best thing to do is to find ways to enhance what you have today and keep delivering better products every year."

Samsung and TSMC are also making 10-nm chips; in October, Samsung announced that it is the first to reach mass production, with devices bearing its 10-nm chips launching in early 2017. GlobalFoundries is opting to go straight to 7 nm beginning in 2018. All are aggressively pushing the limits of miniaturization, says Dan Hutcheson, CEO of VLSI Research, a firm that analyzes developments in the semiconductor industry. "Everybody's making these leaps," Hutcheson says. But they are not marching entirely in lockstep: If you compare 10-nm offerings, he says, "Intel is still about a couple of years ahead" in terms of feature size.

In addition, Hutcheson says, Intel is distinguished from other chipmakers by the control and uniformity of its manufacturing process. This could offer the company a significant advantage when it comes to its foundry services, which use Intel production lines to fabricate chips for other companies.

Manufacturing chips for others is a new approach for Intel. And adding the ability to manufacture ARM processors is an important step. "If you don't have ARM," says Hutcheson, "it's really hard to be a foundry, because it is the standard architecture."

Intel says the first 10-nm chip to ship will be one of its own processors, with chips for others to follow, but perhaps not until 2018.
—RACHEL COURTLAND

"IT'S REALLY POWER REDUCTION OR ENERGY EFFICIENCY THAT'S THE PRIMARY GOAL ON THESE NEW GENERATIONS"

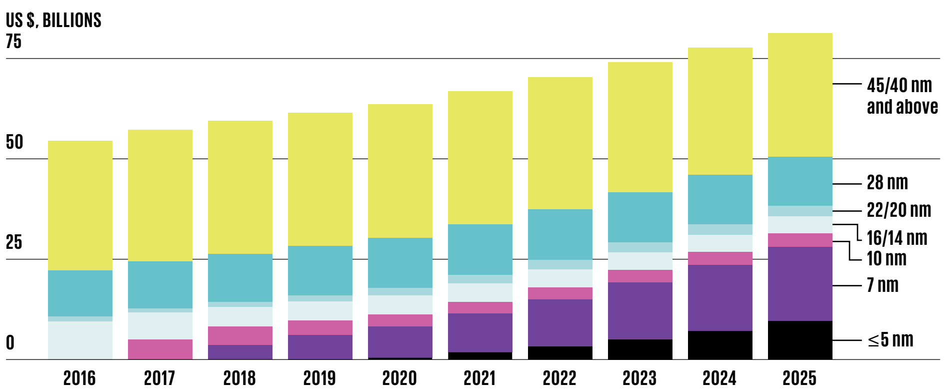

THE EVOLVING FOUNDRY MARKET: Chips built with 10-nanometer technology will come first. But International Business Strategies projects that Apple and others will be drawn to the next node in line: 7 nm.

TOP CHART SOURCE: ARM; BOTTOM: IBS

Fly Robotic?

AUTONOMOUS AIR TAXIS WILL TAKE OFF IN 2017, BUT WON'T GO FAR

IN THE FUTURE, THE JOKE GOES, airliners will each have a pilot and a dog. The dog will be there to bite the pilot if he touches the controls, and the pilot will be there to feed the dog. ● It's no joke, though, when NASA scientists begin entertaining the idea of replacing the copilot with a wideband connection to a ground controller. Who will take over the plane should the pilot become incapacitated? Nor is it a joke to carry the argument to its logical conclusion and do away with the pilot altogether.

It's a beguiling vision. An autonomous airplane reliable enough to be trusted by passengers and air-safety regulators could save not just on salaries but also on the cost of managing the glitch-prone minuet by which well-rested flight crews are united with the planes they're supposed to fly. That logistical problem will get harder as the pilot shortage worsens, and it will be hardest of all for short-hop air service, where the pilot-to-passenger ratio is high.

Now comes a slew of startups that propose to serve that very niche with tiny, autonomous aircraft. Most would be powered by electricity, use multiple propellers or ducted fans, take off vertically or nearly so, and range perhaps a few tens of kilometers.

Vahana, a tilt-wing, autonomous air taxi that's been developed by Airbus's Silicon Valley outfit, A3, is supposed to begin testing later this year. The German company e-Volo, already known for lofting a pilot in an elaborate multicopter, says it's gearing up to make an autonomous version. Zee Aero, reportedly personally funded by Google cofounder Larry Page, offers another example, and Uber yet another. And Terrafugia, a veteran in the flying-car space—at last, a proper use for that bit of

PHOTO-ILLUSTRATION BY Edmon de Haro

EHANG AIR TAXI'S QUICK TRIP
The maximum flight time for the Ehang 184 is just 23 minutes.

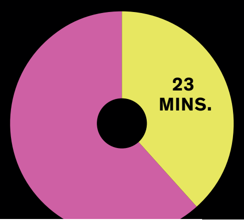

biz lingo!—is also talking about making a model that's autonomous.

When so many new startups are pursuing the same goal, it's tempting to think there must be something there. But hope springs eternal in tech land, and so does the propensity to promise big. All these companies have proven tight-lipped (not one returned inquiries from *IEEE Spectrum*), which suggests that there might be less here than meets the eye.

Spectrum reported on Terrafugia, the one company that has a real history, back in 2007, in our January special issue. We called the company a "loser" for describing a flying car it said it was about to bring to market. It didn't happen.

"It can be done—we could be flying around in pilotless planes, just as we could be living on cities on Mars—but is it worth the cost and the effort?" asks Patrick Smith, author of the Ask the Pilot column, which ran for years in *Salon* magazine. "I fly airplanes for a living, and my jaw drops when I hear people say that flying is already mostly automated. Even the most 'automated' flight is still subject to so much human input and subjective decisions."

So why then are all these startups starting up? "It'd be a novelty, not necessarily meant even for profit, but as a way to prove and build the technology," Smith suggests.

And should one of these outfits ever offer seats to the paying public, would you entrust your life to a robotic pilot? "People want a pilot in the cockpit, to know there's someone in charge who shares their fate," says Missy Cummings, a former U.S. Navy fighter pilot, now a professor of mechanical engineering and materials science at Duke University. "I don't think we'll ever have a passenger airliner be a drone—there will always be some version of a Captain James T. Kirk on board." But, she adds, things are different for hops of, say, 50 miles (80 kilometers), where for some people, at least, convenience might outweigh fear.

"It's technologically achievable in the near term; as for the regulatory problem, I think we'll see it in China first," Cummings says. "Ehang [in Guangzhou, China] is supposedly doing a test in March."

The company claims that its roboplane has already carried a passenger, and if it performs the feat in public, we'll let you know. And if it doesn't.
—PHILIP E. ROSS

LOOK, MA—NO PILOT: By dispensing with the pilot, short-hop electric air taxis would save money and restore the thrill to air travel. Here are artist's impressions of A3's Vahana [top] and Terrafugia's TF-X [middle and bottom].

CHART SOURCE: EHANG. ILLUSTRATIONS, FROM TOP: AIRBUS; TERRAFUGIA (2)

UNIVERSITY OF CAMBRIDGE

www.jobs.cam.ac.uk

Professorship of Information and Communications

Department of Engineering

The Board of Electors to the Professorship of Information and Communications invite applications for this Professorship from persons whose work falls within the general field of the Professorship to take up appointment on 1 October 2017 or as soon as possible thereafter.

Candidates will have an outstanding research record of international stature in the field of Information and Communications, and the vision, leadership, experience and enthusiasm to build on current strengths in maintaining and developing a leading research presence. They will hold a PhD or equivalent postgraduate qualification.

Standard professorial duties include teaching and research, contribution to clinical service delivery, examining, supervision and administration. The Professor will be based in Cambridge. A competitive salary will be offered.

To apply online for this vacancy and to view further information about the role, please visit: http://www.jobs.cam.ac.uk/job/12166.

Further information is available at: http://www.hr.admin.cam.ac.uk/professorships or contact Human Resources, University Offices, The Old Schools, Cambridge, CB2 1TT, (email: ibise@admin.cam.ac.uk).

Applications, consisting of a letter of application, a statement of current and future research plans, a curriculum vitae and a publications list, along with details of three referees should be made online no later than Tuesday 10 January 2017.

Informal enquiries may be directed to Professor Simon Godsill, Convenor of the Board of Electors; e-mail: sjg@eng.cam.ac.uk or tel: 0044 (0) 1223 332604.

Please quote reference NM10774 on your application and in any correspondence about this vacancy.

The University values diversity and is committed to equality of opportunity.

The University has a responsibility to ensure that all employees are eligible to live and work in the UK.

FACULTY POSITIONS

Electrical and Computer Engineering

NYU SHANGHAI

NYU Shanghai is currently inviting applications for tenured/tenure track positions at the rank of assistant, associate, or full professor. We will consider applicants in all areas of electrical and computer engineering, and have a particular interest in applicants in the fields of Multimedia Analytics, Neural Engineering, Systems Neuroscience, Robotics, Smart Cities and Smart Vehicles. The applicants must have demonstrated abilities in both research and teaching. 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 ***January 15, 2017*** *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), electronic copies of up to five recent, relevant publications and the names and addresses of three or more individuals willing to provide letters of reference. Please visit our website at **http://shanghai.nyu.edu/en/about/work-here/open-positions-faculty** for instructions and other information on how to apply. If you have any questions, please e-mail **shanghai.faculty.recruitment@nyu.edu**.

About NYU Shanghai:
NYU Shanghai is the newest degree-granting campus within New York University's global network. It is the first higher education joint venture in China authorized to grant degrees that are 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 of the highest caliber who are committed to NYU's global vision of transformative teaching and innovative research and who embody the global society in which we live.

NYU's global network includes degree-granting campuses in New York, Shanghai, and Abu Dhabi, 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.

NYU Shanghai is an equal opportunity employer committed to equity, diversity and social inclusion. We strongly encourage applications from individuals who are under-represented in the profession, across color, creed, race, ethnic and national origin, physical ability, and gender and sexual identity. NYU Shanghai affirms the value of differing perspectives on the world as we strive to build the strongest possible university with the widest reach.

EOE/AA/Minorities/Females/Vet/Disabled/Sexual Orientation/Gender Identity Employer

上海纽约大学 NYU SHANGHAI

INL

INTERNATIONAL IBERIAN NANOTECHNOLOGY LABORATORY

2 Tenure Track Staff Researcher openings at the International Iberian Nanotechnology Laboratory – INL, (Braga, Portugal):

The INL (**http://www.inl.int**) is an international organisation that carries out interdisciplinary research to deploy and articulate nanotechnology for the benefit of society. INL aims to become the world-wide hub for nanotechnology addressing society's grand challenges with specific emphasis on Aging & Wellbeing, Mobility & Urban Living and a Safe & Secure Society.

INL recruits from all over the world. We find the diversity and multiculturalism of our teams as well as family friendly working practices and benefits to be a great asset and an essential element in cultivating an attractive and inspiring workplace.

INL encourages researchers with a distinctive profile at the highest international level in their field of competency to participate and apply for these 2 positions:

Ref.11.16.82 Staff Researcher – Hybrid Sensing Devices

Ref.11.16.83 Staff Researcher – Energy Harvesting for Sensor Networks

The job description, the Candidates's profile, application procedure and employment conditions can be found at **http://inl.int/job_offers** - Staff Researchers.

Department of Electrical and Electronic Engineering
Academic Positions in Electrical Engineering
South Kensington Campus
Competitive Salary

Imperial College London

Consistently rated amongst the world's best universities, Imperial College London is a science-based institution with a reputation for excellence in teaching and research. We are seeking to make up to three faculty appointments in Electrical and Electronic Engineering at Lecturer, Senior Lecturer, Reader or Professorial level.

Our department has a record of fundamental research of international reputation that delivers lasting impact through industrial engagement and innovation. We expect appointees with an outstanding research record as demonstrated by publications, raising research funding and delivering impact, commensurate with the level of appointment. The standard of our education and student cohort is equally high. We look for a passion in teaching and teaching innovation, and a track record in supporting the next generation of innovators. Appointees with a vision for integrating their research into their teaching strategy are particularly favoured.

We will consider appointments in all research areas in which the department is active. More information on these can be found at: **http://www3.imperial.ac.uk/electricalengineering/research**.

Appointees are particularly sought in the area of Energy Systems Integration and the use of smart technologies to support Low-Carbon Electricity Systems. Topics of interest include whole-system and multi-vector energy modelling; stochastic modelling and control; decentralised control; data-driven model development; analysis of risk and planning under uncertainty. We also seek an appointee in the developing area of Big Data science and technology; aspects of interest include massive-scale sensing, network intelligence, machine learning, data processing techniques, and information theory for application domains such as smart cities, Internet of Things, digital economy, future communications and healthcare.

Our preferred method of application is online via **http://www3.imperial.ac.uk/employment** (select "Job Search" then enter the job title or vacancy reference number EN20160427SA). Alternatively, if you are unable to apply online, please contact Miss Anna McCormick: email a.mccormick@imperial.ac.uk.

Closing date: 28 February 2017

Committed to equality and valuing diversity. We are also an Athena SWAN Silver Award winner, a Stonewall Diversity Champion, a Disability Confident Employer and are working in partnership with GIRES to promote respect for trans people. Imperial College is a Family Friendly Employer.

IEEE **JobSite** IEEE

The Right Candidate — Right Now!

Take the next steps to finding your ideal job.

Take advantage of this unique member benefit.

Post your resume on the IEEE Job Site and start your job search now!

Visit www.ieee.org/jobs

BEIJING JIAOTONG UNIVERSITY

HANERGY SCHOOL OF RENEWABLE ENERGY

JOB OPENINGS for A Vice Dean and Faculty Members

BEIJING JIAOTONG UNIVERSITY. Established in 1896, Beijing Jiaotong University (BJTU), is a national key university, under the direct administration of the Ministry of Education of China and is also jointly supported by the Beijing Municipal Government and the China Railway Corporation.

HANERGY SCHOOL OF RENEWABLE ENERGY.
Hanergy School of Renewable Energy (HASRE) was jointly established by the Chinese State Administration for Foreign Expert Affairs, the Chinese Ministry of Education, and the Hanergy Holding Group, one of the world's leading companies in renewable energy technologies in November 2015. HASRE is one of only a handful of national "International Demonstration Schools" selected through a rigorous and competitive process in China over the past two years. Offering cutting-edge research opportunities and bachelor, masters and doctoral degrees. The inaugural class began in September 2016. All major courses are taught in English. As it pursues cutting edge cross-disciplinary research in renewable energy, HASRE will build upon core research areas including:

- Renewable Generation Systems
- Control Techniques
- Renewable Generation Reliability
- Inspection and Maintenance Strategy for Renewable Generation
- Asset Management on Renewable Generation
- Integration of Renewable Energy
- Modeling of Renewable Generation
- Reliability Assessment of Renewable Generation
- Economic Assessment of Renewable Generation
- Short-term Renewable Energy Forecast for System Operation
- Online Stability and Control of Renewable Generation
- System Restoration with consideration of Renewable Generation
- Economics and Environment Regulation and Policies
- Energy Market and Economics
- Environment and Regulation Policies

CURRENT OPENINGS:
HASRE plans to build a group of international faculty members with the following academic background:

- Renewable Generation (Wind, Solar)
- Power Systems (planning and operation)
- Power Electronics
- Energy Storage Technologies
- Control Theory
- Information Technologies (Big Data, Cloud, Mobility, Internet of Things)
- Energy Economics

By the end of 2017, HASRE expects to recruit 6-7 faculty members at all levels—Assistant, Associate, and Full Professor for full-time basis; 3-5 visiting faculties at all levels in semester basis;
Vice Dean for Planning, Finance & Internationalization Affairs. For candidates of the vice dean, an academic administration experience in university with fluent English will be highly respected.
Note: Applicants are expected to be fluent in English.

All of these faculty members will have a unique opportunity to build the foundation of what is envisioned to be the leading international school of renewable energy in China. In addition to maintaining a dynamic teaching and research program, successful candidates will contribute leadership in building HASRE by participating in curriculum development, developing internship and interdisciplinary opportunities, professional interactions with colleagues and industry, student club and professional and development activities, etc. Junior researchers who already have research partners at BJTU are especially welcome to apply. Competitive salaries commensurate with experience will be offered at each level, for example, with starting salary US$ 48,000 for (foreign) associate professors.

APPLICATION DETAILS

- Applicants should submit:
- Cover Letter
- Curriculum Vitae;
- Research & Teaching Statement;
- Publication List;

At least 3 Letters of Recommendation (sent directly by referees) to: Dean Jin Ho Kwak, **Email Address:** jinkwak@bjtu.edu.cn. **Postal Address:** Hanergy School of Renewable Energy, Beijing Jiaotong University, No.3 Shangyuancun, Haidian District, Beijing 100044
P. R. China
All applicants are recommended to submit all necessary materials by **April 30, 2017** (or by **September 31, 2017** for the 2nd round recruiting).

Applications are invited for faculty positions in the Department of Electrical and Computer Engineering for tenure, research, and teaching-tracks at multiple locations.

The Department of Electrical and Computer Engineering (ECE) invites applications for tenure, research, and teaching-track positions at its main campus in Pittsburgh (www.ece.cmu.edu), as well as Silicon Valley (www.cmu.edu/silicon-valley) and Rwanda (www.cmu.edu/rwanda) locations. Silicon Valley faculty will have a unique opportunity for research collaborations with industry and involvement in entrepreneurial activities emerging from research projects. Rwanda faculty will collaborate with regional and multi-national industry in technology innovation and entrepreneurship and deliver innovative, interdisciplinary graduate teaching and research programs within the African context.

We are strongly committed to all members of our community: students, faculty, and staff. Our vision is to be a creative driving force, within the university and worldwide, of the highest scholarly and entrepreneurial quality. Our mission is to inspire and educate engineers capable of pursuing fundamental scientific problems and important societal challenges. We strive to accomplish this with the highest commitment to quality, integrity, and respect for others. We are particularly interested in applicants who are committed and have passion for a culturally diverse environment in research and/or teaching, and demonstrate a willingness to nurture the uniquely inclusive Carnegie Mellon environment. We take pride and active steps in considering a diverse applicant pool in terms of gender, race, veteran status, and disability. Carnegie Mellon University seeks to meet the needs of dual-career couples and is a member of the Higher Education Recruitment Consortium (HERC) that assists with dual-career searches.

We are looking to hire in strategic thrust areas (www.ece.cmu.edu/research/index.html):

- Theoretical and Technological Foundations serve as a bedrock for our work;
- Systems and Technologies, connecting system level to physical and cyber levels (beyond CMOS, compute/storage systems, cyber-physical systems, data/network science systems, secure systems), with particular attention to nanoscale systems and integration of heterogeneous devices, sensors and materials; and
- Application Domains, where our fundamental and system level work makes significant impact (energy, healthcare & quality-of-life, mobile systems, smart infrastructure).

In Silicon Valley, our emphasis on experimental computing systems, platforms for mobile supercomputing, human, vehicle and environment sensing, and sensor signal processing and inference. In Rwanda, our focus is on software engineering, mobile computing, cloud computing, big data analytics, communications, information and cyber security, intelligent infrastructures, wireless networking, mHealth, energy systems, eLearning, and mobile applications.

- Tenure-track faculty carry a moderate teaching load that allows time for quality research and close involvement with students. We expect you to establish and grow a strong research program, contribute to our teaching mission, and show your passion for mentoring and advising students.
- Research-track faculty are not required to teach, but do so occasionally when of clear benefit to the faculty and the Department; you will be compensated for both teaching and advising Ph.D. students. You will typically focus on developing leadership within your area of research, developing research collaborations, and supervising Ph.D. students.
- Teaching-track faculty typically focus exclusively on teaching and service, but may do research as well. We will rely on you to help strengthen our teaching and mentoring mission.

For all tracks, we are seeking individuals who hold a Ph.D. in a relevant discipline and have demonstrated commitment to our core values: scientific truth, creativity, quality, innovation, and engineering solutions, all within a diverse and tight-knit community guided by respect and joy of doing. Faculty positions are primarily at the Assistant Professor level; however, appointments may be made at the rank of Associate Professor or Professor depending on the qualifications. Our Department and the College of Engineering are ranked among the top programs in the United States both at the undergraduate and graduate levels. We house and have ties to several multidisciplinary institutes and centers. We collaborate with colleagues around the world through a number of formal research and educational programs. We have extensive experimental and computing infrastructure, including state-of-the-art nanofabrication facilities.

Please submit an online application at www.ece.cmu.edu/faculty-staff/employment/index.html. We will begin evaluation of applications immediately and will continue throughout the academic year until positions are filled; we encourage you to submit early. Carnegie Mellon is an EEO/Affirmative Action Employer -- M/F/Disability/Veteran.

Electrical & Computer ENGINEERING

Carnegie Mellon University College of Engineering

LAWN CARE FOR THE LAZY

The most miraculous thing at the 1960 Miracle Garden Exhibition in Paris had to be the H.C. Webb & Co. radio-controlled electric lawn mower. The machine, which ran off an electric motor powered by two 12-volt car batteries, was by all accounts quite successful. Somehow, though, it failed to inspire a new era of sedentary lawn care, much to the disappointment of suburban teenagers everywhere. ■

GAMMA-KEYSTONE/GETTY IMAGES

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


MATLAB SPEAKS
MACHINE
LEARNING
With MATLAB you can use clustering, regression, classification, and deep learning to build predictive models and put them into production.
mathworks.com/machinelearning
@2016 The MathWorks, Inc
MathWorks®