IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 04.15

Moore's Law is dead.

Long live Moore's Law.

IEEE

ENGINEERS START HERE

Research. Design. Prototype.

1 800 463 9275

newark.com

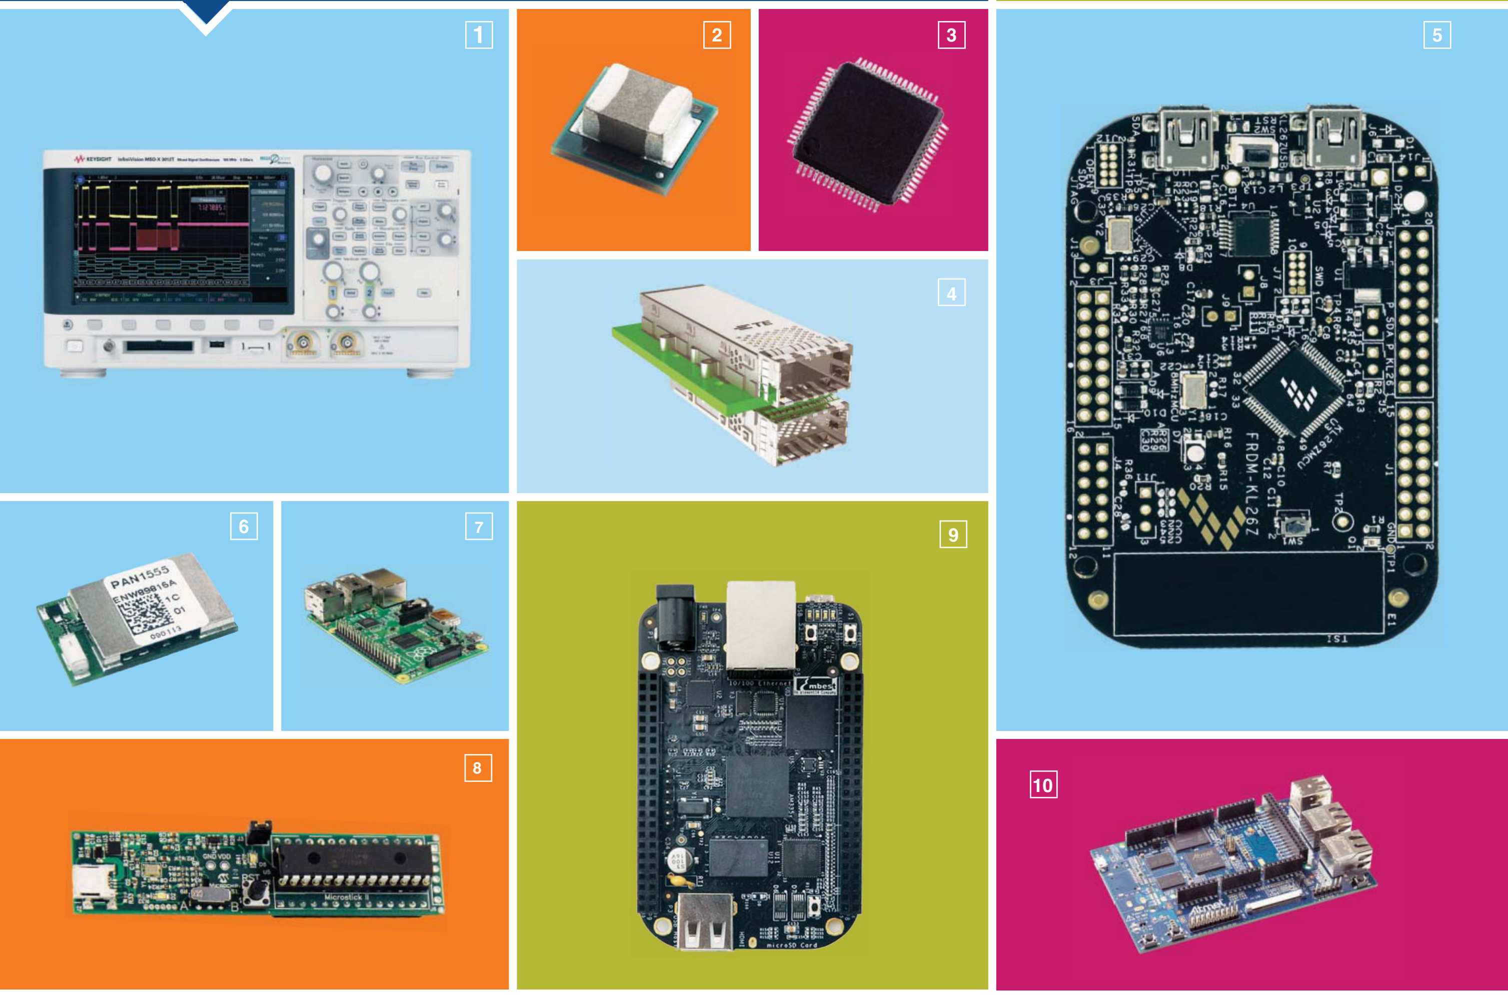

Atmel | beaglebone | element14 | freescale | KEYSIGHT TECHNOLOGIES | MICROCHIP | NXP | RaspberryPi | TE connectivity Authorized Distributor | TEXAS INSTRUMENTS

1 Keysight 3000T Oscilloscope

2 Texas Instruments Simple Switcher® Nano2 Power Modules

3 NXP LPC54100 Series Microcontrollers

4 TE Connectivity CXP Pluggable I/O Connector Assemblies

5 Freescale Freedom Platform

6 Panasonic Bluetooth Low Energy RF Modules

7 Raspberry Pi Single Board Computer

8 Microchip Microstick II

9 element14 BeagleBone Black (RevC)

10 Atmel Xplained Evaluation Boards

500,000 In-Stock Electronics Products | Guaranteed Same-Day Shipping | 500+ World-Class Manufacturers

Newark™ element14 ecia MEMBER

IEEE SPECTRUM

SPECIAL REPORT

50 Years of MOORE'S LAW

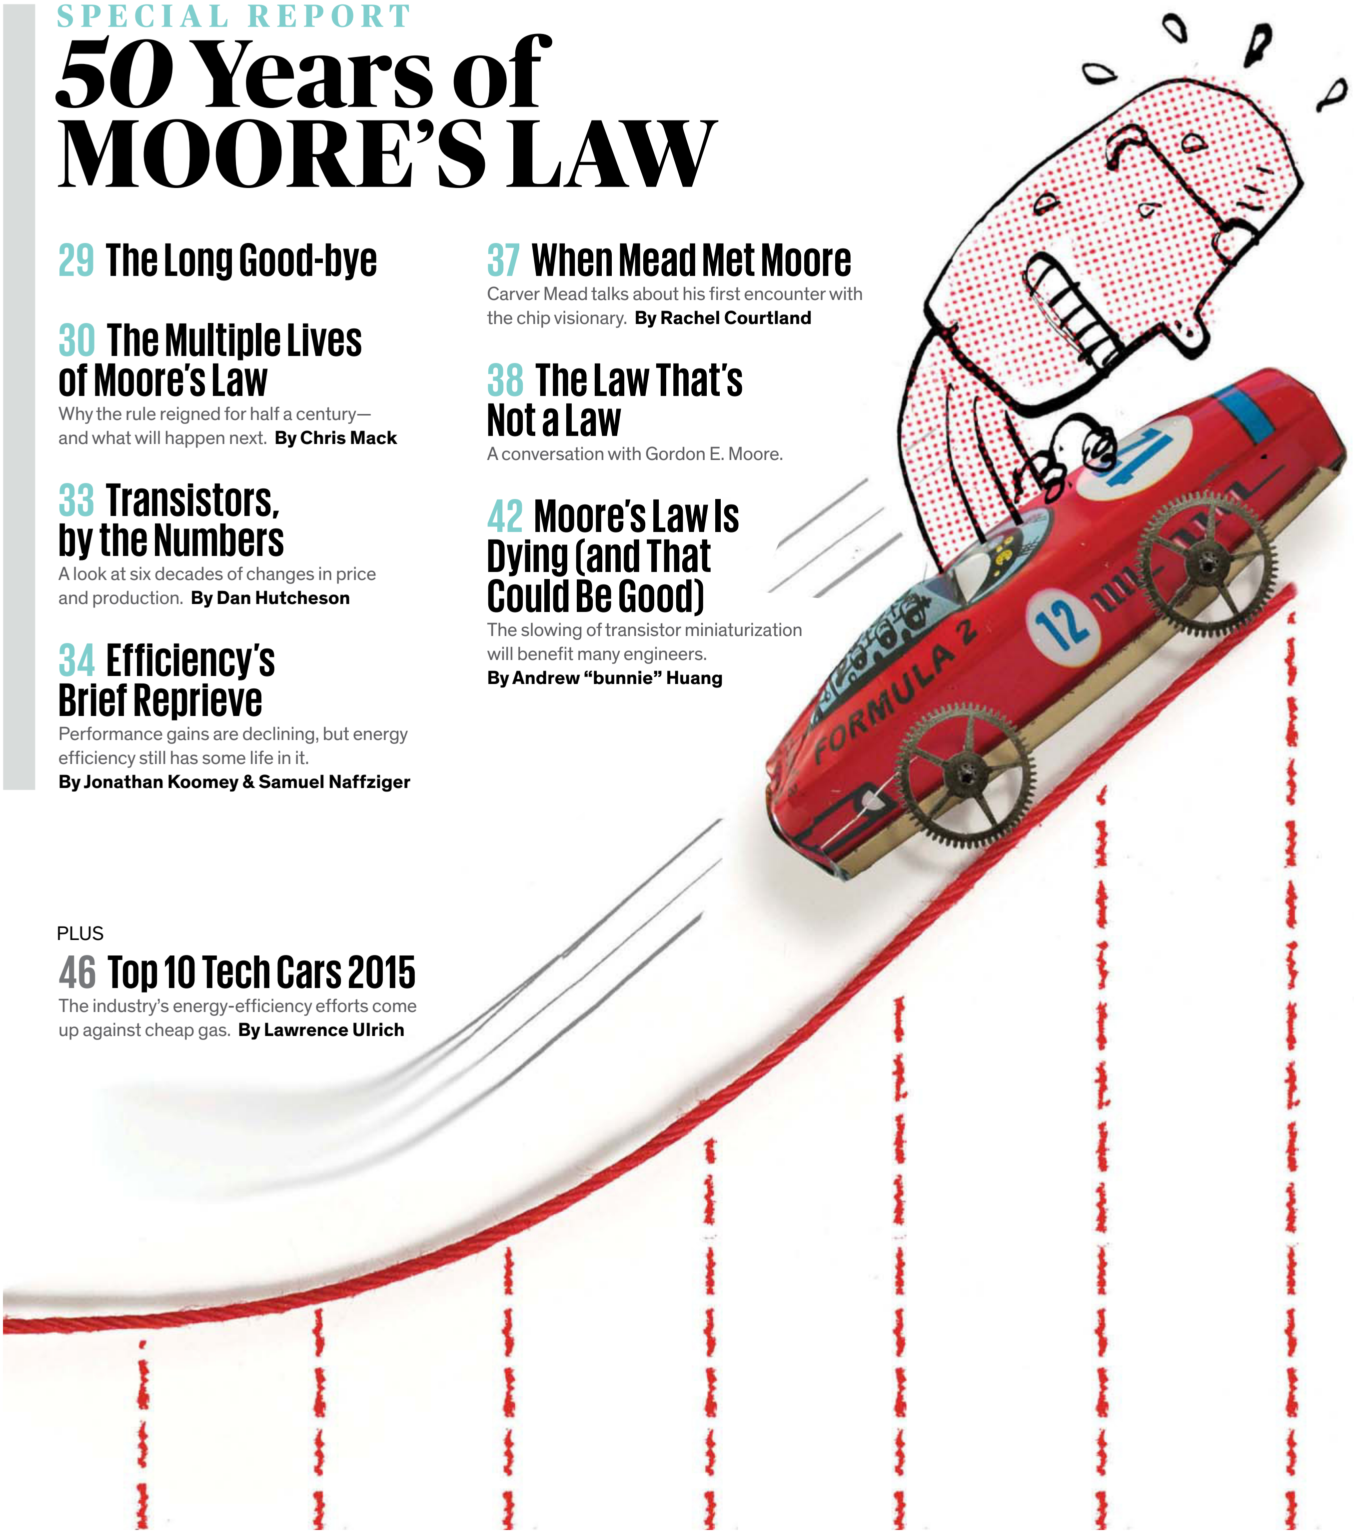

COVER PHOTO STYLIST: ROBIN FINLAY

On the Cover **Photograph for *IEEE Spectrum* by Adam Voorhes**

ILLUSTRATION BY Serge Bloch

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


Who will industrialize the Internet of Things?
You and NI will. From smart manufacturing to the smart grid, bridging the digital and physical worlds is one of the most significant engineering opportunities we face today. NI, together with other visionary companies, is ensuring the superior design and performance of the increasingly complex systems that will fuel the Internet of Things and drive industrial adoption. Learn more at ni.com.
NATIONAL INSTRUMENTS™
©2015 National Instruments. All rights reserved. LabVIEW, National Instruments, NI, and ni.com are trademarks of National Instruments. Other product and company names listed are trademarks or trade names of their respective companies. 19044

DEPARTMENTS_04.15

IEEE SPECTRUM

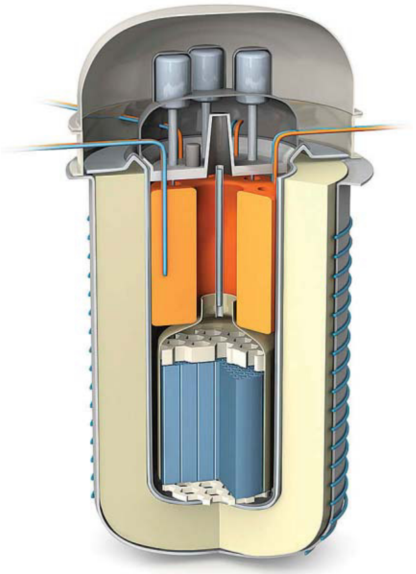

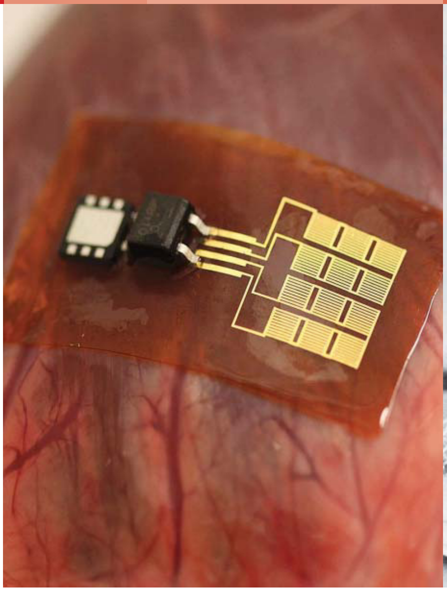

Online

Spectrum.ieee.org

More Moore
See more of our coverage of the Moore's Law semicentennial, including an extended interview with Caltech's Carver Mead and an interactive timeline: http://spectrum.ieee.org/50yearsofmooreslaw

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Modeling of Capacitive and Resistive Devices—9 April
- Get Your Head Out of the Cloud and Into Open-Source Software Security and Compliance—Now Available on Demand
- Get Ready for the Internet of Things: Speed Innovation With Continuous Engineering for Electronics—Now Available on Demand
- Open-Source Software Security—Mitigating Risks From Development, Integration, Distribution, and Deployment—Now Available on Demand
- **MASTER BOND RESOURCE LIBRARY** http://spectrum.ieee.org/static/master-bond-resource-library
- **NEW PRODUCT RELEASE LIBRARY** http://spectrum.ieee.org/static/new-product-release-library
- **WHITE PAPERS** http://spectrum.ieee.org/whitepapers

The Institute

Available on 6 April 2015

CONSUMER ELECTRONICS Inspired by the 2015 Consumer Electronics Show, this issue explores new applications designed to benefit humanity and save lives. One researcher is working on an augmented-reality system to help children with autism learn to socialize, while others are developing technology to prevent distracted driving.

ASSISTIVE TECHNOLOGY Many of today's consumer electronics can be used to help people with disabilities perform everyday tasks with much greater ease. Just ask IEEE Member Lama Nachman, who recently helped upgrade the device that Stephen Hawking uses to communicate with others.

A BURGEONING FIELD More and more companies are looking for engineers with experience in 3-D printing and additive manufacturing. Learn what it takes to get in on the ground floor.

FROM LEFT: EMILY COOPER; UNIVERSITY OF ILLINOIS COLLEGE OF ENGINEERING; INTEL

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2015 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 52, issue No. 4, North American 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.

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

BACK STORY_

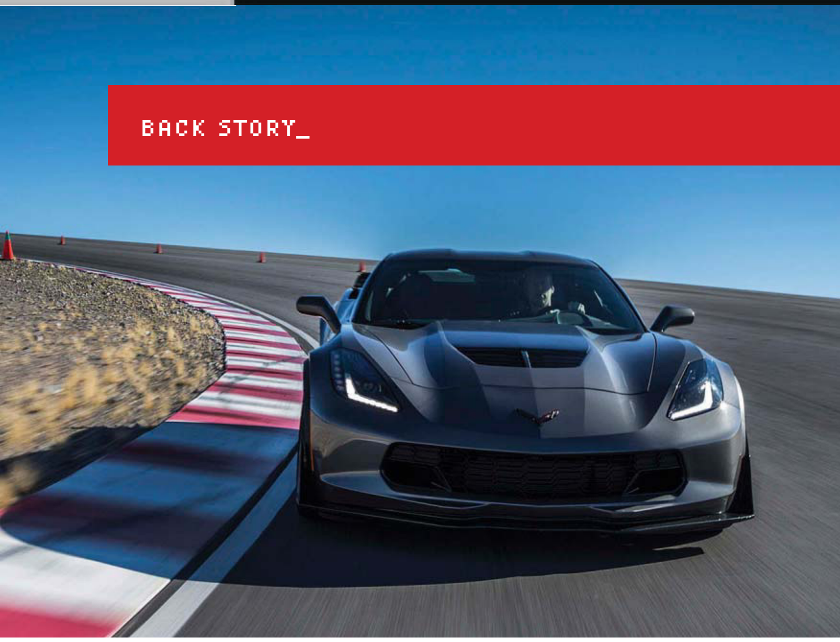

Sports Car Selfies

AUTO WRITERS, LIKE WINE CRITICS, must regularly try to put the ineffable into words. For years, our contributor Lawrence Ulrich has been doing just that—when he wasn't test-driving the world's hottest cars. Out of this lifestyle have come some alluring allusions, such as these:

- "a mellow, baritone blat" (on the BMW X6 M in 2010)
- "a silvery, shrieking blur" (on the McLaren MP4-12C in 2012)
- "The Tesla hunts down internal-combustion prey with eerily silent remorselessness" (2013)

Recently, on a test track in Nevada, he was able to meld his evocative prose with video and data, for the first time, after test-driving a car. It's a Corvette Z06 [above, with Ulrich at the wheel], one of our Top 10 Tech Cars for 2015.

He flipped on the Vette's performance data recorder—a new feature much like the training tools race car drivers use—and whipped around the track, giving second-by-second commentary. Meanwhile, the recorder used the car's own cameras and sensors to register everything that happened—Ulrich's voice, his view from the driver's seat, the workings of the engine, the scene in front of the car, and the car's position on a map of the track.

"It's a video game come to life," Ulrich told us.

His feature article is on display in this issue; his real-time test is up on our website. To see it, just search for "Watch This Corvette Z06 Test Us, as We Test It." Here's a still from that report [right] to tempt you. ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, And Now This... is in *IEEE Spectrum*, Vol. 52, no. 4 (INT), April 2015, p. 60, or in *IEEE Spectrum*, Vol. 52, no. 4 (NA), April 2015, p. 68.

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

IMAGE: GENERAL MOTORS

PHOTOGRAPH BY **A.J. Mueller**

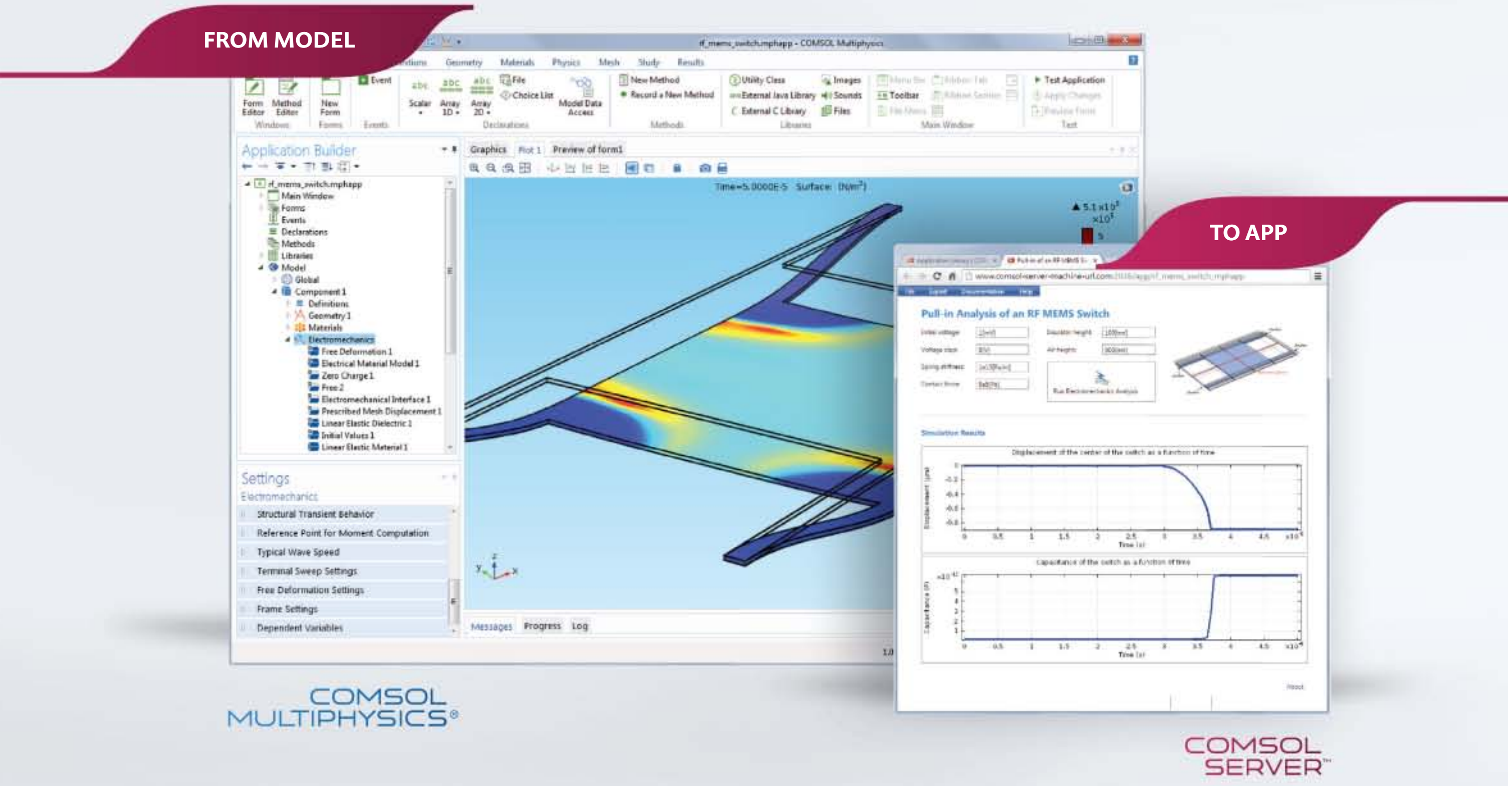

Verify and Optimize your Designs with COMSOL Multiphysics®

NOW FEATURING APPLICATION BUILDER & COMSOL SERVER™

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

Visit comsol.com/release/5.0

Product Suite

› COMSOL Multiphysics®
› COMSOL Server™

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

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

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

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

MULTIPURPOSE
› Optimization Module
› Material Library
› Particle Tracing Module

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

COMSOL

CONTRIBUTORS_

Serge Bloch

Bloch's artwork graces this month's special section on the 50th anniversary of Moore's Law. He was inspired to pursue his profession by the great French humorist-illustrators of an earlier generation, some of whom were slain at the offices of *Charlie Hebdo* in Paris. In a speech at the Society of Illustrators shortly after the attack, Bloch paid them tribute: "[T]hey often made us laugh and now we cry for them.... But drawing is our life, and life is stronger than violence."

Andrew "bunnie" Huang

Shortly after obtaining a Ph.D. in electrical engineering from MIT in 2002, Huang wrote *Hacking the Xbox: An Introduction to Reverse Engineering*. However, the need for such reverse engineering might soon decline, he argues in "Moore's Law Is Dying (and That Could Be Good)" [p. 42]. That's because as the pace of Moore's Law slows, manufacturers will increasingly adopt open designs. Huang's not worried that this change will hurt book sales, though: He's been offering free electronic copies of the book since 2013.

Chris Mack

For the 50th anniversary of Moore's Law this month, "gentleman scientist" Chris Mack examines its astounding and shifting history [p. 30]. He has worked in lithography—the technology used to pattern chips—for 32 years. For much of that, Mack has had the end of Moore's Law in his sights; he once wagered his beloved Lotus sports car on the future of long-delayed extreme ultraviolet lithography (the bet was declared a draw). Moore's Law has been a "truly remarkable" phenomenon, he says, definitely worth celebrating.

Martin LaMonica

LaMonica, a contributor to *IEEE Spectrum*'s Energywise blog, has long contemplated the potential of nuclear energy. "When you're looking for solutions to climate change, you have to think seriously about nuclear," he says. In this issue, he reports on startups that are aiming to get new nuclear technologies out into the real world faster than established companies can [p. 12]. He's now involved in a startup himself, as environment and energy editor at *The Conversation*.

Alexandra Ossola

A New York City–based journalist who writes frequently for *Popular Science* and Motherboard, Ossola first heard about the highly competitive fight to devise the best DNA search engine [p. 16] in passing while at a New Year's Eve party. It wasn't the usual way she gets story ideas, but she knew immediately that she wanted to pursue it. "When something like that falls into your lap, you pull at the thread and find something really cool at the end," Ossola says.

04.15

PHOTO-ILLUSTRATIONS BY **Gluekit**

You in Your Internet of Things

Should privacy and security measures be built into devices before they reach the market?

MY FITBIT SENDS ME ENCOURAGING NOTES throughout the day. My iPhone does a remarkable job of telling me where I am, where I should be, and how to get from here to there. And then there's my Honda, reminding me what music I last listened to. Would I like to hear something else?

For the most part, these devices work well, and I am pleased with the convenience that all this robust connectivity offers. On a more global scale, when my personal data is added to that of millions of others, I see the enormous value in all that aggregated information, culled and parsed by ever larger and more interconnected networks.

And yet, I'm also alarmed when I consider the security vulnerabilities present in nascent Internet of Things (IoT) systems. Digital privacy expert Christopher Soghoian, profiled in this issue, makes a point of never connecting his own devices directly to the IoT, for just that reason.

So why aren't these problems addressed during product development? It may be that today's pell-mell design cycles don't allow engineers to consider any issues that aren't absolutely essential to the day-to-day performance of a device. On the face of it, what does privacy have to do with the range of your Android's antenna? But that may well change in the next few years. Andrew "bunnie" Huang, in this issue, contemplates the demise of Moore's Law and concludes that the resulting slowdown in transistor-doubling times will give engineers more breathing room in which to be creative and responsive.

Or perhaps technology developers consider such matters beyond their purview. A while back, I attended a conference on nuclear weapons proliferation where a similar question arose: Do bomb designers have a moral responsibility for how their weapons are used? At that meeting the answer was no–the designers felt their job was to build the best possible weapons and someone else's job to determine what should happen with them. So are app developers and IoT builders absolved of responsibility for security and privacy? Who should decide these issues?

Now, my Fitbit tracker is not a nuclear weapon, but I can certainly imagine how the data it is collecting, along with the other information that is out there on the Web about me, could be used against me, in consequential or just absurd ways.

I notice, in passing, the programs that already track me online. In one recent creepy episode, I suddenly began receiving e-mails about burial-service insurance after looking at funeral home websites to make arrangements for a relative.

Which makes me wonder: Am I simply a collection of data? Luciano Floridi, Oxford professor and Google's information philosopher, seems to think so. Floridi has been exploring the newish field of the philosophy and ethics of information. Now he is helping Google come up with new ideas about what it means to be a person in an age when everyone and everything has an IP address. "For Floridi, you *are* your information, which comprises everything from data about the relations between particles in your body, to your life story, to your memories, beliefs, and genetic code," writes Robert Herritt in his recent profile of Floridi in *Pacific Standard* magazine. "Floridi's view can also help us think precisely about the consequential questions that today preoccupy us at a very practical level."

Just how we think about and answer these questions will define our tech-obsessed culture going forward. —SUSAN HASSLER

YAGI STUDIO/GETTY IMAGES

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

NEWS

26 GIGAWATTS: JAPAN'S PUMPED HYDROELECTRIC STORAGE CAPACITY, THE BIGGEST IN THE WORLD

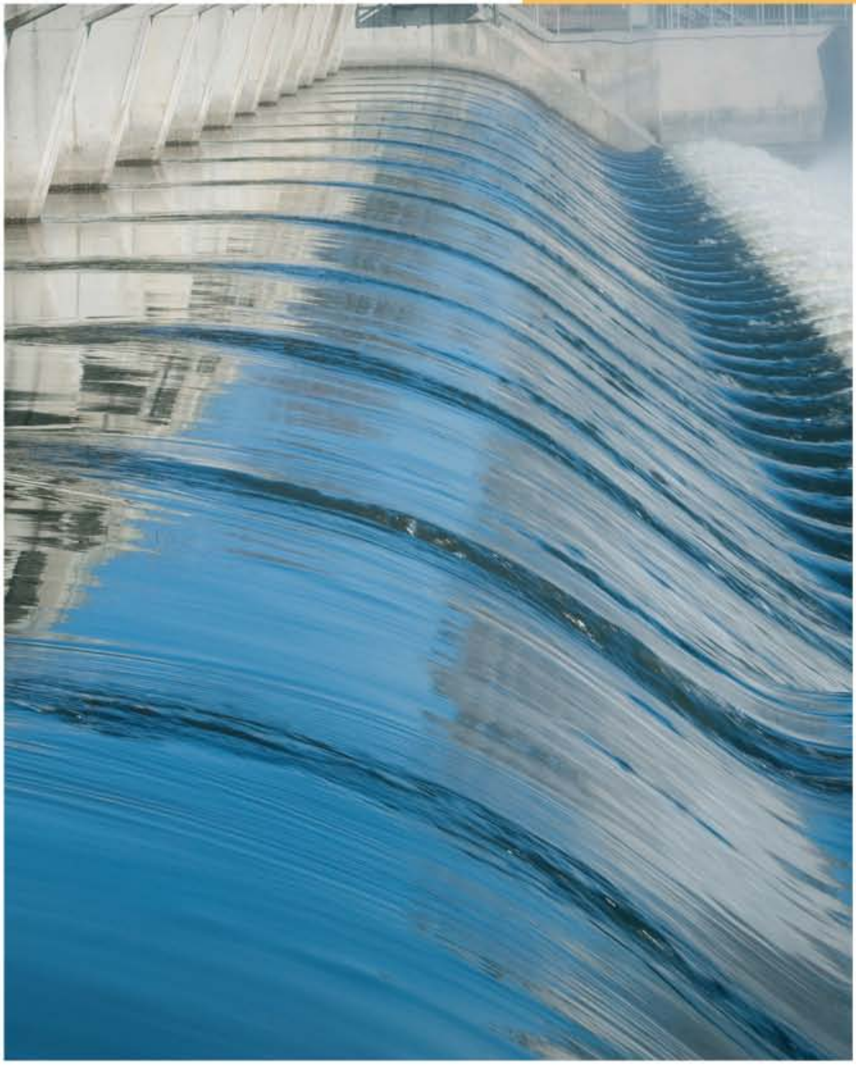

A RENAISSANCE FOR PUMPED HYDRO STORAGE

Energy-storage tech designed for nuclear power is being revamped for wind and solar

Pumping water uphill to store energy in hydropower reservoirs is an idea that, by power grid standards, is as old as the hills that such "pumped storage" plants are built on. But with the rise of intermittent solar energy and wind power, this technology could soon experience a revival, experts say.

Japan is positioned to lead this renaissance. It has the largest number of pumped storage plants, capable of absorbing and discharging 26 gigawatts of power. Those plants are also technology leaders whose variable-speed pumps, developed in the 1990s, are uniquely adapted to shifting energy flows.

But Japan is not alone. China's pumped storage capacity has more than doubled since 2008. And, in mature markets such as the United States and Europe, intermittent generation is accelerating pumped storage construction for the first time in several decades.

Ironically, what originally motivated pumped storage installations was the inflexibility of nuclear power. Nuclear plants' large steam turbines »

UPHILL BOTH WAYS: Renewable energy is driving construction of pumped hydro storage plants in Switzerland and elsewhere.

AXPO SERVICES

run best at full power. Pumped storage can defer surplus nuclear power generated overnight (when consumption is low) to help meet the next day's demand peak.

Japan's utilities chose to install variable-speed pumps so that the plants could also help stabilize their grids. Whereas single-speed pumps rotate synchronously with the grid, variable-speed pumps' asynchronous motor-generators can adjust a plant's charging and discharging to simultaneously balance power supply and demand, thus regulating the grid's AC frequency.

Variable-speed equipment is larger and costs more. But it reduced the need to ramp up and down Japan's oil-fired generators, a process that wastes pricey imported petroleum.

In post-Fukushima Japan, these nimble storage plants have a new opportunity to shine: managing the more than 10 GW of solar and other renewable energy capacity that has been installed in Japan since the March 2011 meltdowns and the subsequent idling of the country's nuclear reactors.

Concerns about the rising levels of intermittent generation prompted Japan's government to propose tougher rules for renewable installations in January. But international experts in integrating renewables say that Japan's utilities could manage their new grid challenges by redeploying the pumped hydro plants to absorb spikes in demand or surpluses of renewable power. "A substantial amount of Japan's renewable energy transients can be handled by the pumped storage," says Per Christer Lund, a principal consultant at the Singapore Clean Technology Center, run by Norway-based risk management and certification giant DNV GL.

Lund says that Japan's vertically integrated utilities appear to be reserving pumped hydro plants to serve the traditional and more parochial role of complementing nuclear power stations, which the utilities hope to begin restarting this year. But he predicts that Japan's shift to competitive power markets by 2020—which would separate generation from transmission and distribution—will drive more flexible use of the pumped storage plants.

Experience with competitive markets in China, Europe, and the United States, however, tells us they can also create new hurdles for expanded use of pumped hydro storage. New plants may be far cheaper than batteries—by as much as 95 percent per kilowatt-hour, according to a January 2015 report by Citibank. But most competitive markets do not currently pay for grid services, such as frequency regulation offered by advanced pumped storage plants. Ning Zhang, a power systems expert at Tsinghua University, in Beijing, says a lack of market payments for grid services explains why China has so far eschewed variable-speed equipment for its pumped storage.

Competitive markets can also bring unpleasant surprises. In Europe, a slide in anticipated profits poses a challenge for big, new pumped storage plants, such as those being built in Switzerland. Their backers plan to charge up the plants when power is cheap and sell when market prices peak. But surging solar power generation is blunting the demand peaks, thus narrowing daily fluctuations in wholesale power prices.

Greenhouse-gas mandates will ultimately force the use of storage, according to Scott Flake, director of power generation for the Sacramento (Calif.) Municipal Utility District (SMUD), whose planned US $700 million pumped storage plant is one of two approved last year by federal government regulators, after a two-decade hiatus.

Flake says that SMUD's variable-speed, 400-MW plant will take over grid-regulation tasks now performed by natural gas-fired plants, helping the utility comply with greenhouse-gas restrictions under California's carbon cap-and-trade program. "We see pumped hydro as a way to provide grid-management services in a carbon-free environment," says Flake.

Paris-based equipment supplier Alstom Power Systems, meanwhile, is developing a low-risk option for the pumped storage renaissance: retrofitting single-speed storage plants. That means finding a way to squeeze larger and heavier variable-speed equipment into existing subterranean vaults. In 2012 the European Commission awarded an Alstom-led consortium €13 million ($15 million) to prove that it's possible at a 36-year-old plant in France. The project is due to be completed by the end of 2017. —PETER FAIRLEY

> A substantial amount of Japan's renewable energy transients can be handled by its 26 GW of pumped storage

WHY NET NEUTRALITY ANALOGIES WILL ALWAYS FAIL

The Internet is a little bit like a lot of things

In late February, in a big victory for Net neutrality advocates, the U.S. Federal Communications Commission voted to reclassify broadband Internet as a telecommunications service subject to regulations like those used in the telephone industry. AT&T, Verizon, and other opponents of reclassification argued, unsuccessfully, that phone networks are a poor analogy for understanding the Internet. But what is the right analogy? Internet access is a complex tangle of technical, political, and commercial issues, and people have long tried to make it comprehensible through comparisons to other, more familiar systems. As a result, Net neutrality analogies have also become a powerful and effective rhetorical tool for rallying political action.

It goes without saying that the development of the Internet and the economy it supports has had an unprecedented global impact. But "unprecedented" means that nothing in history adequately prepared us to determine who should be paying whom and for what, exactly, in a world ruled by packets. Here's a quick reminder that some of the very things that make the Internet unique are precisely those that make it resistant to comparisons. —JOSHUA J. ROMERO

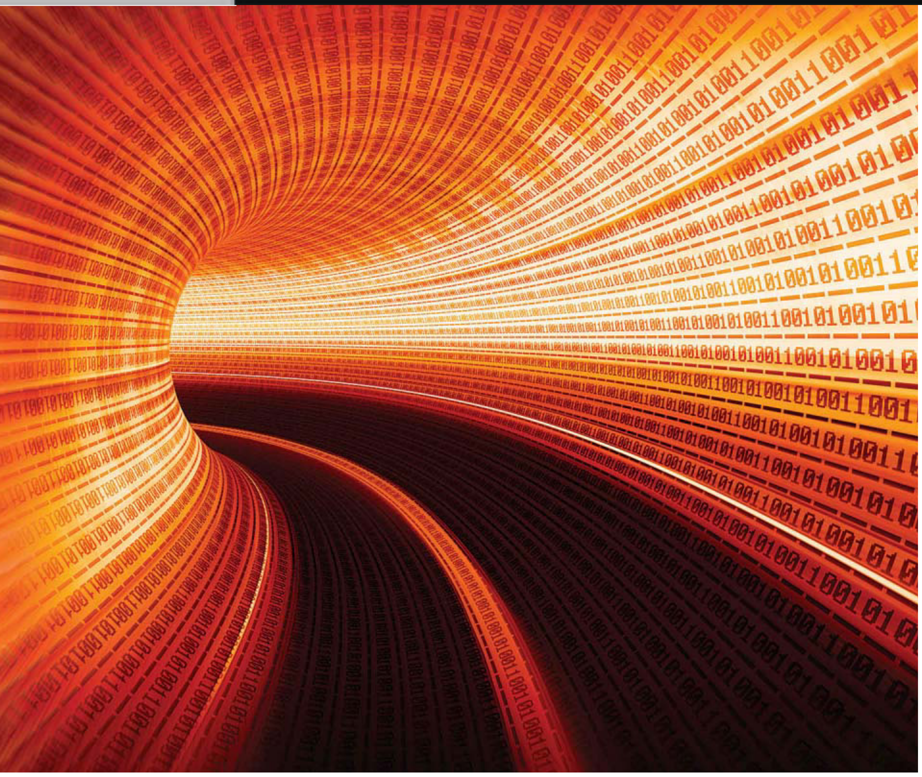

GETTY IMAGES

managed by private companies. Initial deployment is usually expensive, so the companies expect to profit by collecting recurring fees over time. When private companies do build private roads, tollbooths are common.

- **It's not a single-purpose network:** In a phone system, a call is a call, and the network is optimized for call quality. Similarly, cable TV infrastructure was purpose-built for delivering television channels. The beauty of the Internet Protocol, on the other hand, is in its ability to support a multitude of services. But one limitation is that the protocol itself doesn't guarantee timely delivery, and not all data streams are equally tolerant of delays. For e-mail or Web browsing, tiny delays in packet delivery have little impact, but when you're making a two-way voice call or playing an online game, even small amounts of latency can ruin the experience.

THE ANALOGIES:

- The FCC justified its latest decision by noting that when customers pay an Internet provider, they're primarily purchasing the ability to transmit (and receive) packets of data, just as customers pay **phone companies** to transmit and receive phone calls.

- That's a reversal from 2002, when the FCC said that wired Internet access is an information service, more like **websites** and **cable TV**. AT&T has argued that Internet service providers (ISPs) obviously do more than just transmit data; if they didn't provide a "computing functionality," paid prioritization, blocking, or throttling bandwidth wouldn't even be possible.

- It's also common for people to compare Internet data to **water flowing through pipes** or **electricity flowing through the power grid**. Internet law expert Lawrence Lessig, for one, has compared a nonneutral Internet to an electrical outlet that provides electricity with different prices, quality, and reliability depending on the brand of appliance you plug in.

- Ever since U.S. vice president Al Gore popularized the term **"information superhighway"** in the 1990s, the analogy of traffic on a multilane road has been ubiquitous. Over the past few years, the Net neutrality debate has largely focused on paid prioritization, the idea that ISPs could charge content providers for a higher level of service. Some opponents said this was like adding **tollbooths** to previously free roads. The ISPs argued that paid prioritization was the equivalent of adding new, **fast lanes** alongside the existing lanes; other opponents shot back that ISPs were more likely to add speed bumps to some of the old lanes and relabel the remaining lanes "fast."

WHY THE ANALOGIES FALL SHORT

- **It's privately owned infrastructure:** Unlike U.S. highways, which are built and maintained by the government, fiber connections are usually laid and

- **There's no such thing as generic data:** When you turn on your tap or plug in an appliance, all you really want is water or power. Where that water or power originated is irrelevant; once it's in the pipes or wires, it's all the same stuff. The Internet doesn't work that way. When consumers pay an ISP, they're purchasing the ability to access any specific content that they choose, not just the data that's convenient for the ISP to provide.

- **There are bottlenecks at interconnects:** When Columbia Law School professor Tim Wu first defined Net neutrality, he was arguing that network providers should not be allowed to discriminate against any of the bits flowing through their network. But what if the slowdown occurs outside of an ISP's network? When Netflix was battling with Comcast and Verizon over streaming-video speeds, the real issue revolved around who should foot the bill for upgrading the interconnects at the boundaries between those last-mile ISPs and the intermediate ISPs that Netflix paid to deliver its traffic. A better highway analogy might involve the financing of new on-ramps, rather than the arbitrary construction of tollbooths.

- **Traffic is two-way:** Many metaphors focus on content delivery. But on the Internet, of course, you're not just consuming packets, you're also sending them. Except for minimal interactivity (for example, on-demand movies) television is a one-way medium. Water and electricity are similarly services that are consumed (unless you're feeding solar power back onto the grid).

- **There's a history of peering arrangements:** In fact, the general business structure of the Internet was based on the idea that when customers on one network need to send data to customers on another network of similar size, it makes sense for those networks to exchange packets without charging each other, a practice known as "peering." Smaller networks paid to get access to larger networks, and at the top of the food chain were backbone providers, which found that it was mutually beneficial to exchange and deliver each other's packets without cost. With the rise of online video in particular, traffic has become much more asymmetric, undermining the historic definition of peers. It is difficult to work the concept of paid or nonpaid peering into a highway or utility metaphor.

- **Routing packets isn't like routing money:** Water and power companies charge based on consumption (in addition to connection fees), so that the more you use, the more you pay. And with a cable subscription, you're paying a flat rate for both the content itself and the transmission infrastructure that delivers it to your TV. Neither billing structure is suitable for the Internet, where nothing is actually consumed. Customers pay content providers for content and ISPs for access to that content. Companies in each business want to capture as much total revenue as possible while keeping their costs to a minimum. As Internet pioneer David D. Clark noted a couple of decades ago, "you cannot derive money flow from packet flow," so we can expect this battle of competing economic interests to continue for years to come.

IS THIS THE AGE OF ALTERNATIVE NUCLEAR POWER?

Smaller reactors and alternative tech are getting a serious look

There's more than one design for a nuclear reactor, but you wouldn't know it judging by the array of aging light-water reactors in operation around the world. A number of potentially better reactor types have been stuck on drawing boards or at the experimental stage for decades. Now a handful of small companies say the best way to commercialize these advanced concepts is to hitch nuclear to the distributed-energy bandwagon.

One vision for the future of nuclear involves smaller versions of today's gigawatt-scale behemoths. Building a power plant in smaller chunks, about one third or less the size of today's plants, should lead to lower up-front costs, site flexibility, and the ability to add more units as needed.

But others believe it's time to move past the light-water reactor. The idea is to modernize experimental reactor designs and make the leap from the lab to the marketplace.

"We're seeing a new innovation pathway emerging," says Richard Lester, the head of MIT's department of nuclear science and engineering. "Early-stage companies are getting into the business and seeking to develop new nuclear technologies—that's a big shift."

Terrestrial Energy, in Mississauga, Ont., Canada, for example, is seeking to build a molten salt reactor, one of which was built and operated at the Oak Ridge National Laboratory in the 1960s.

In a conventional light-water reactor, uranium rods are submerged in water. The nuclear reaction of uranium generates heat, and the hot water is continually circulated through the core to produce steam and turn a turbine. In Terrestrial's reactor, low-enriched uranium fuel is dissolved in molten salt. Aided by pumps and heat exchangers, the heat from that bath of fuel and molten salt produces steam to generate electricity. A molten salt reactor can use more of the available energy in the fuel and is safer than a water-cooled reactor, the company says. In the case of a loss of power, the heat from the reactor would dissipate, moving by convection from the core to an outer shell layer of solid salt, so it

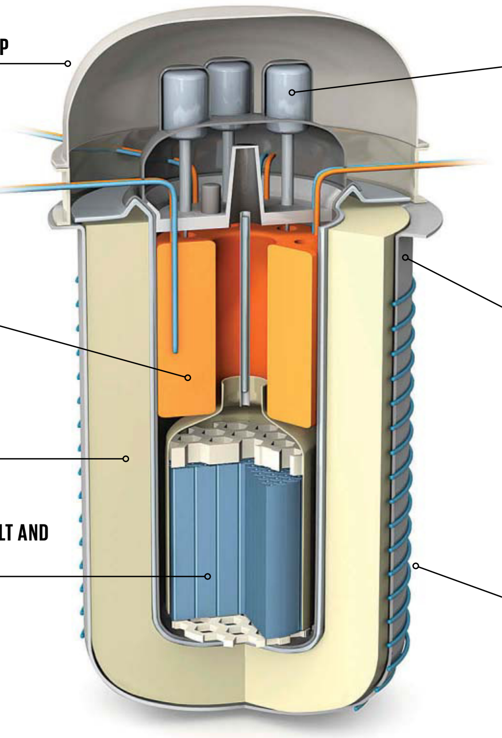

A FLUID FURNACE: In Terrestrial Energy's alternative nuclear reactor, uranium is dissolved in molten salt at the reactor's core. Heat exchangers move energy out of the hot core to turn a turbine generator. If power is lost to the pumps, heat passively dissipates into a solid buffer of salt.

wouldn't require pumps to cool it and keep it from melting down, says CEO Simon Irish.

Terrestrial's business plan is unconventional as well. Instead of only banking on building power plants for utilities, its management thinks it can get a commercial foothold in niche markets, including locations that now rely on diesel generators for electricity, such as remote locations in Canada or island nations. The reactor can produce both electricity and high-temperature heat, which is useful in many industrial settings, notes Irish.

The main component would be small enough to fit on a rail car or truck to simplify installation. And the reactor, able to generate 32.5 megawatts of electricity, would be sealed so that it could operate for about seven years before being decommissioned. Terrestrial plans to start the licensing process in Canada later this year and have designs ready for engineering blueprints by 2016.

"At the boundary between the laboratory–where we know it works–and the market, you have to have certain features, namely operational simplicity and a safety case you can demonstrate to a regulator to get an operating license," says Irish. "That's the essence of our innovation."

Some companies intend to make even smaller reactors for distributed energy. Boston-based UPower Technologies' nuclear generator would produce between 1 and 2.5 MW from a reactor the size of a shipping container. The core vessel would have a metal block with nuclear fuel at the base; tubes, partially filled with liquid, that fit into the block would transfer heat away through convection, rather than by pumping water through the core. As a first step toward seeking a license, UPower is designing tests of its heat removal system.

Quebec-based StarCore Nuclear intends to build a commercial high-temperature gas reactor, a variant of the pebble-bed, helium-cooled design developed in Germany in the 1960s. Fuel is contained in tennis-ball-size "pebbles" that generate heat, which is transferred via helium. Because it operates at high temperature, the design can be more efficient at generating electricity, and the reactor can cool itself by natural circulation in the case of an emergency or loss of power, say company executives.

StarCore's plans call for building a generator to produce 10 MW of electricity as well as heat. Like UPower, StarCore is initially targeting remote locations, such as mining operations, which pay a lot of money for heat and electricity. The company hopes to build and commission one installation in five years, "if we're lucky," says CEO David Dabney.

StarCore's founders see a future with many different nuclear technologies, including reactors that burn today's spent nuclear fuel. But to get there, new nuclear needs to start small, they say. "There are lots of other reactors that offer great capabilities," says StarCore chief technology officer David Ashley Poole. "We need to do that, but the first ones need to be as simple and safe as possible to get them out of the laboratories and into the ground, and get public acceptance."

It will take years for any new nuclear technology to be licensed and, most likely, years to get potential customers and partners to buy into the concept of mininuclear plants. A key time frame is the late 2020s and early 2030s, which is when many nuclear plants in the United States will likely go off-line, says MIT's Lester. "It is a long game, and it will hinge on whether these innovators and investors can find large companies to work with to commercialize the technology," he says. "Given the environment, it's not an unreasonable bet."
—Martin LaMonica

NEWS

ILLUSTRATION BY Emily Cooper

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

Keysight 89600 VSA software

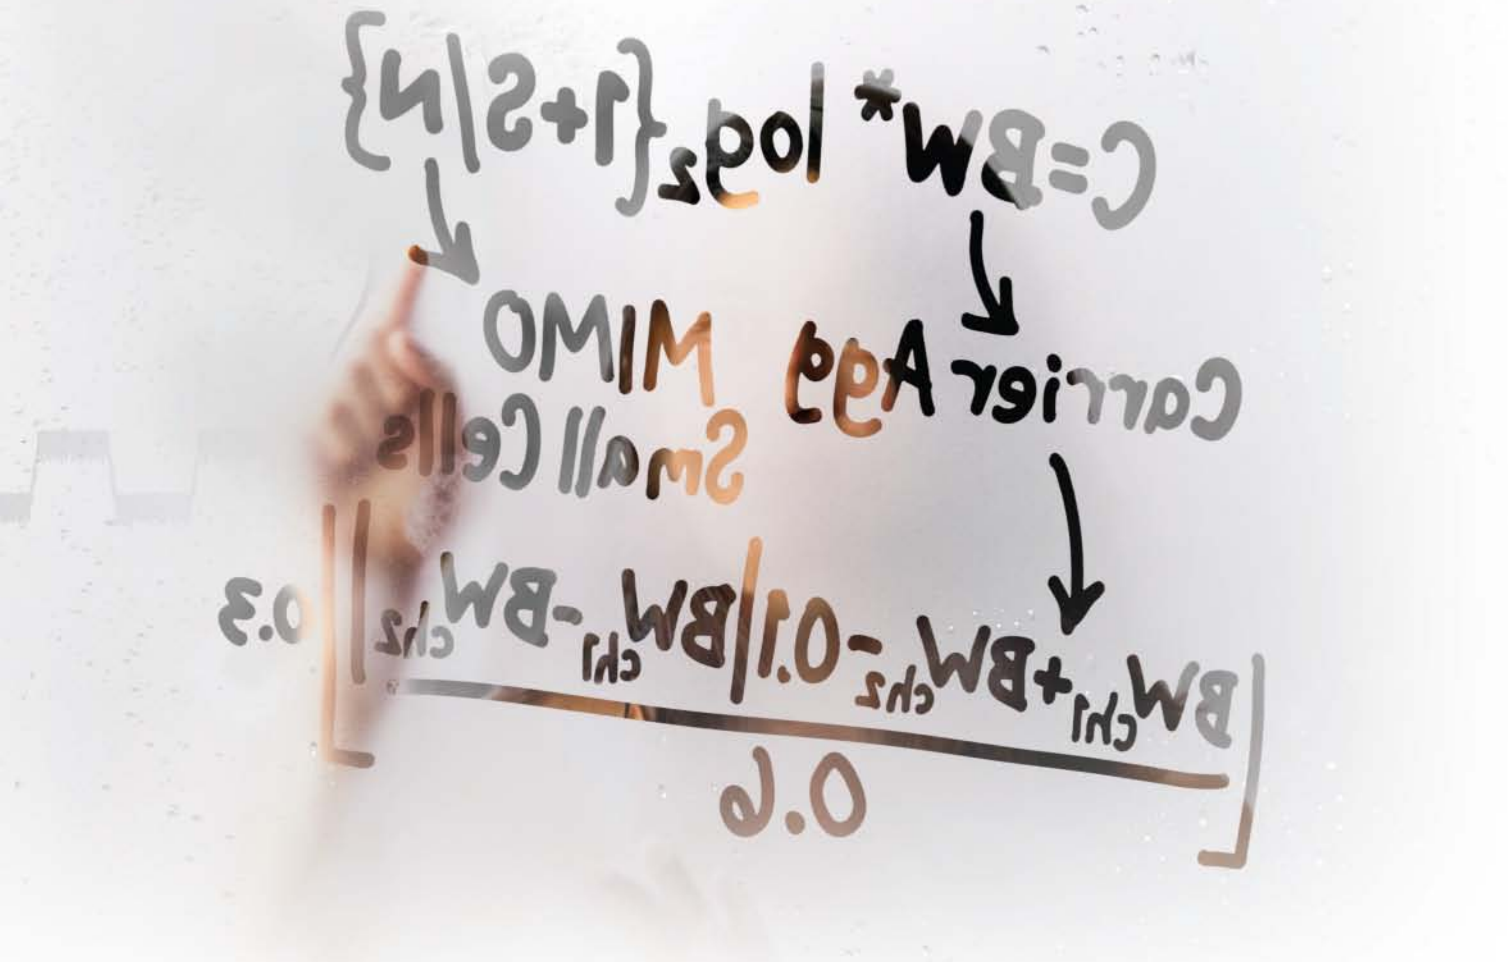

NEWS

A GOOGLE FOR DNA

Companies are vying to build the best genome search engine

In 2005, next-generation sequencing began to change the field of genetics research. Obtaining a person's entire genome became fast and relatively cheap. Databases of genetic information were growing by the terabyte, and doctors and researchers were in desperate need of a way to efficiently sift through the information for the cause of a particular disorder or for clues to how patients might respond to treatment.

Companies have sprung up over the past five years that are vying to produce the first DNA search engine. All of them have different tactics–some even have their own proprietary databases of genetic information–but most are working to link enough genetic databases so that users can quickly identify a huge variety of mutations. Most companies also craft search algorithms to supplement the genetic information with relevant biomedical literature. But as in the days of the early Web, before Google reigned supreme, no one company has yet emerged as the clear winner.

Making a functional search engine is a classic big-data problem, says Michael Gonzalez, the vice president of bioinformatics at one such company, ViaGenetics, which was expected to relaunch its platform in March. Before doctors or researchers can use the data, genomic data must be organized so that humans can read and search it. The first step toward that is to put it in a standard form called the variant call format, or VCF. As raw data, a person's complete sequenced genome would take up about 100 gigabytes, so a database that adds the genomes of even 10 patients per day would quickly get out of hand. But VCF files are more compact, requiring only a few hundred megabytes per genome, which helps researchers find the specific variants they want to search in a fraction of the time. Unlike a fully sequenced genome, VCF files point only to where a person's genetic data deviates from the standard, the genome originally compiled by the Human Genome Project in 2001.

With VCF, sifting the genomes themselves for pinpoint mutations isn't the challenge for search engine companies.

Most of these companies are allocating their resources toward efforts to seamlessly compile supplementary information about a specific mutation from other databases across the Web, such as the biomedical research archive PubMed or various troves of electronic medical records. Many of these tools have finely tuned algorithms that prioritize the results by credibility or relevance. "You want to be able to pull together the information known about a mutation in that position [of the genome] and quickly make an assessment," says David Mittelman, the chief scientific officer for Tute Genomics, based in Provo, Utah, another company designing a genetic-search engine.

In an effort to expand the information that can be attached to a genome under examination, ViaGenetics, based in Miami Beach, Fla., is making its newly updated platform useful for researchers who want to collaborate across institutions. With ViaGenetics' tools, researchers "can make their data available to other users, so other people can come across these projects, request access, and form a collaboration," Gonzalez says. "It helps people connect the dots between different researchers and institutions." This is especially helpful for smaller labs that may not have very extensive genome databases or for researchers from different universities working to decode the same mutation.

Although the genomic-search industry is now focused on serving scientists, that might not always be the case. Mittelman envisions that Tute Genomics could eventually serve consumers directly. People are already demanding information about their genomes just to understand themselves better, Mittelman says, but most companies don't yet consider the average person to be their primary customer. In order to make that shift, the tool will have to be even more intuitive and user-friendly. "Fire-hosing someone with data that's not easy to interpret, or using terminology that's not standardized, has the potential to confuse people," he says. Privacy is also a major concern for the average user; the information that Tute users upload isn't stored permanently, Mittelman says, but users will need extra reassurance if the platform becomes available to the lay public.

And a further evolution of the industry is in the offing. Both ViaGenetics and Tute are hoping to be able to run the entire process in-house–from the initial DNA sequencing to the presentation of final searchable results to users. "The market for analyzing and interpreting genomic data is very fragmented, like the computer industry in the 1990s, where you had to go to separate providers to buy a video card or a motherboard and then try to put it together," Mittelman says. "Soon this field will consolidate, as the computer industry did."

—ALEXANDRA OSSOLA

ANDREW BROOKES/CORBIS

PHOTOGRAPH BY Pim Hendriksen/ Daan Roosegaarde & Heijmans

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

LIGHT UNTO MY PATH

MUCH ATTENTION has been focused on making smarter vehicles. But what about the roads they ride on? Interactive and sustainable roads are the focus of a Dutch group called Smart Highway. For a bike path that runs through the city of Nuenen, Netherlands, where Vincent van Gogh painted his first masterpiece, the group used technology to create high art. The Van Gogh-Roosegaarde path is studded with thousands of stones that absorb solar energy during the day and glow at night. The stone arrangement–by Daan Roosegaarde, a Dutch designer–pays homage to *The Starry Night*, Van Gogh's most famous work of art.

RESOURCES

90,601: THE OFFICIAL NUMBER OF TARGETS SUBJECTED TO U.S. FEDERAL COMMUNICATIONS SURVEILLANCE IN 2013

RESOURCES_CAREERS

CHRISTOPHER SOGHOIAN

THE ACLU'S TECHNOLOGIST WANTS COMPANIES TO DO LESS SPYING FOR THE GOVERNMENT

Google, Facebook, and AT&T are just a few of the corporations that Christopher Soghoian has challenged in the name of protecting the privacy of ordinary citizens. Soghoian is the principal technologist of the American Civil Liberties Union. The job was created just for him, a computer scientist with a reputation for investigating digital invasions of privacy. • At the ACLU, Soghoian combines his investigative skills with public advocacy. His research often focuses on how telecommunications companies facilitate government surveillance, while his advocacy focuses on changing government policies and getting private companies to do a better job of guarding the interests of their customers. • "Half my day is spent in front of a computer, and 10 to 20 percent's talking to reporters," he says, referring to his ongoing media campaigns against alleged violators of privacy. Soghoian says his campaigns sometimes work, noting that in recent years major social networking companies have finally begun encrypting customers' data. "I had something to do with it—10 percent—by pushing them. I play the bad cop," he says. • Soghoian doesn't restrict himself to U.S. transgressions. Although the conventional wisdom is that Europeans are more attuned to privacy concerns, he rejects that notion: "Europeans are big into talking, but the »

PHOTOGRAPH BY M. Scott Brauer

debate over there on security and privacy is five years behind us."

Soghoian started out as a standard-issue computer nerd. After college he veered into computer security, getting a master's degree in the subject at Johns Hopkins University in 2005. "Then I found that the field was getting crowded, so I did a pivot into privacy," he says.

He went on to Indiana University Bloomington, where he earned a doctorate in informatics. His thesis focused on how corporations handled government surveillance requests. "I didn't do theoretical work, didn't do experimentation," Soghoian says. "My research methodology was to get lawyers drunk and telling me stories. My dissertation is filled with footnotes to anonymous sources—not your typical thing in computer science!" In 2009, while finishing his dissertation, he took a job watching over consumer privacy for the Federal Trade Commission. "They are the good side of the government," he says. "They don't kick down doors, they don't kill people—they protect people."

CHRISTOPHER SOGHOIAN

IEEE Member

Age: 33

Occupation: Principal technologist, American Civil Liberties Union

Location: Washington, D.C.

Hero moment: He tapped the wiretappers, then used their words to expose governmental invasions of privacy.

A few months after joining the FTC, he used his government affiliation to get into a closed conference run by Intelligence Support Systems for law enforcement and intelligence agents and telecom operators that Soghoian calls "the wiretappers' ball." There he tapped the tappers with a hidden recorder, causing a stir in the press when he caught one executive from Sprint Corp. saying that his company was so inundated with government requests to perform surveillance on customers that it had set up a self-service site for the police. The executive said the site got 8 million hits in the first year.

Exploits such as these led Soghoian, in 2012, to join the ACLU, an institution that pursues civil rights cases and is therefore primarily known for its lawyers. Meanwhile, he continues to submit papers, mostly to law journals, to leave open the possibility of an academic career: "I have one [paper] coming out on the technology used to intercept telephone calls."

> "My research methodology was to get lawyers drunk and telling me stories"

What does Soghoian do to protect his own privacy? For one thing, he never buys products that automatically hook to the Internet. Instead, he chooses those that connect via a detachable device, so that he can control when they connect (such as the smart meter that lets him adjust the air-conditioning in his Washington, D.C., home). "Otherwise, maybe years go by with no updates to [the product's] security, and someone hacks in and uses the built-in minicam to snoop on you," he says. For another, he has chosen his employer carefully, reasoning that it would be "pretty stupid of the FBI to target someone at the ACLU." **—PHILIP E. ROSS**

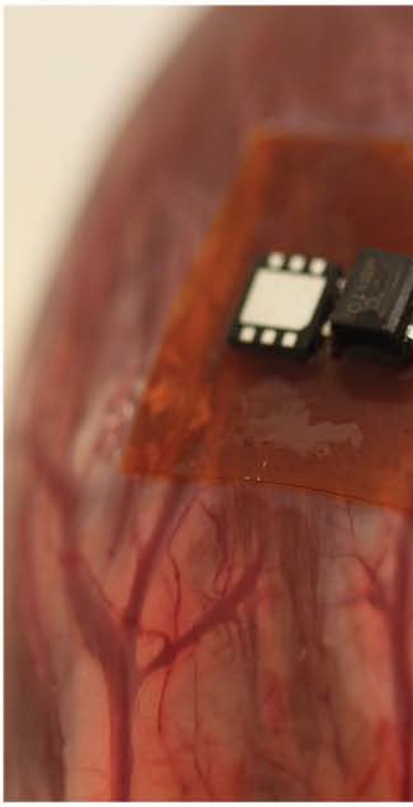

TOP: UNIVERSITY OF ILLINOIS COLLEGE OF ENGINEERING; BOTTOM: SOUTHEASTERN RAILWAYS

REAPING THE ENERGY HARVEST

CONSUMER DEVICES WILL RUN LONGER BY PULLING POWER FROM THE ENVIRONMENT

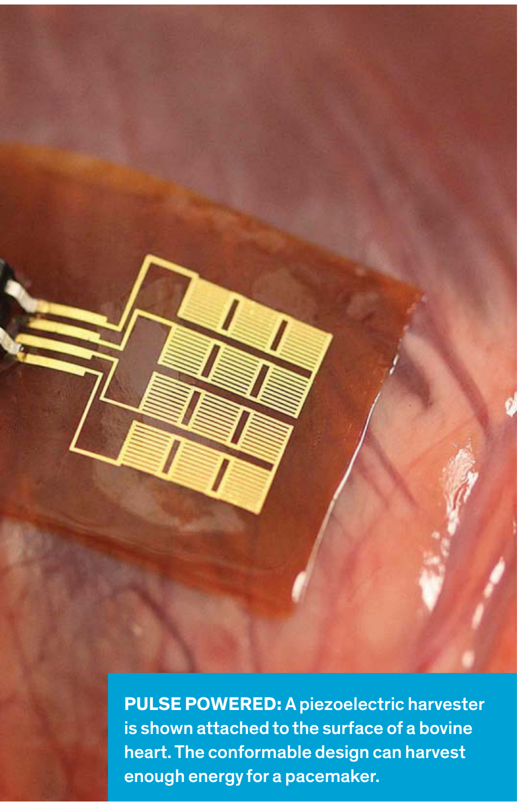

PULSE POWERED: A piezoelectric harvester is shown attached to the surface of a bovine heart. The conformable design can harvest enough energy for a pacemaker.

WE ARE ENTERING AN ERA WHEN SCAVENGING TINY AMOUNTS of power from the environment can power small devices to do extraordinary tasks. The needs of the Internet of Things and the advent of lower-cost components are moving energy-harvesting systems from niche applications to broad-scale practicality. ● This was the clear message of a conference sponsored by market research firm IDTechEx, held in Santa Clara, Calif., last November, where seven events took place at once. In addition to energy harvesting and storage, the topics covered were printed electronics, wearable technology, 3-D printing, supercapacitors, the Internet of Things, and graphene. These seemingly disparate fields are tightly linked; for example, as IDTechEx CEO Raghu Das pointed out, wearable devices were a strong driver for energy-harvesting technologies. Throughout the sessions and in the exhibit hall, other linkages were evident: Supercapacitors hold the potential to create highly efficient ways to store and release harvested energy, and 3-D printing plays a big role in creating prototypes and short-run custom production items.

LEVERAGING LOCOMOTION: Fitted to a train, this device uses the vibrations of the undercarriage to cause a magnet to oscillate in a coil, thus generating electricity for sensors.

According to IDTechEx's analysts, the energy-harvesting industry is projected to grow to US $2.6 billion by 2024. The harvesting show was built around three main sources of energy and the technology used to tap them: mechanical energy, heat, and electromagnetic emissions.

One means of harvesting mechanical energy discussed was a module created by the British company Perpetuum. Its module harvests energy from the vibration of railroad car wheels: The vibrations oscillate a magnet in a fixed coil, powering sensors that monitor the temperature of wheel bearings. Southeastern Railways, based in England, has installed Perpetuum's system on all 148 of its Electrostar trains to identify and address small problems before they become large ones.

Energy from motion can be harvested by other means. A variety of materials exhibit piezoelectric properties: Apply a force and the material generates an electrical current (and vice versa). And new piezoelectric materials are being created all the time. For example, a University of Illinois graduate student has created a thin, biocompatible piezoelectric film that can be implanted safely in living tissue. A multilayered device incorporating the film that is mounted directly on the surface of a human heart can power a pacemaker. This can eliminate the risk to patients that comes with surgically replacing a standard pacemaker every 5 to 10 years, when its batteries run low. Piezoelectrics can also work for larger-scale applications. During the 2012 Summer

Olympics, in London, part of an underground station's walkway lighting was powered by passengers' footsteps across energy-harvesting tiles. IDTechEx estimates that the piezoelectric market worldwide was about $35 million in 2014, nearly twice the size of the market for 2013.

IDTechEx's analysts were also bullish on the size of the market for thermoelectric systems that can convert unwanted heat energy to electricity. While thermoelectric systems have been around for decades, new solid-state thermoelectric devices can work with smaller temperature differences, increasing the number of places they can be used. At the show, the analysts forecast that the thermoelectric-generation market will more than double from 2014 to 2016, reaching nearly $95 million.

Automotive applications are an obvious target, but industrial applications are more likely to drive growth in this segment. Wireless sensors powered by excess heat in manufacturing settings can lower installation and maintenance costs by eliminating power and communications wires for each location.

Perhaps the most magical method of energy harvesting is drawing on ambient electromagnetic emissions. This isn't a new idea: Crystal radios were pretty popular once upon a time, and RFID tags and photovoltaic cells are pretty popular today. The concept can be taken further, however. It is possible to pull energy from the various radio frequencies that wash over us all the time. The Spanish textile research association Aitex, in collaboration with other researchers, has created cloth that contains tiny antennas woven directly into the fabric. This technology could be developed so that our clothing could capture enough energy to power biometric sensors and other wearable devices.

The IDTechEx conference was an excellent example of how the whole can be greater than the sum of its parts. There was an energy and enthusiasm among the participants that's been missing at many technology events in recent years. The excitement generated by the multiple markets, all experiencing rapid growth, was reminiscent of the heady early days of the personal computer, when anything was possible and amazing products were within reach. Energy harvesting is going to play an important role in changing the way we interact with everyday items, our data, one another, and even our own bodies. —**ALFRED POOR**

DIY WATER LEAK DETECTOR

TURN YOUR iPHONE INTO A MICROCONTROLLER

SARAH LEWIN

RECENTLY, I WAS AWAKENED BY A CRASH—MY BLINDS HAD fallen off my apartment's bedroom windows. Then, later, I saw that water was dripping from the holes left behind in the frame. I replaced the blinds with curtains, and the drip continued, on an annoyingly intermittent schedule, over the days that followed. • While I waited for the landlord to track down and repair the problem in the apartment upstairs, I needed something that would alert me when I was home and the dripping started. And for the times I wasn't home, I needed something that would automatically move the curtain as well in order to avoid water damage. Given that the situation was both urgent and temporary, I needed something I could throw together as a quick fix. Fortunately, I had just the thing to aid me in my time of need: the Nanospark. • Nanospark is a US $140 system for controlling electronics with your iPhone or iPad. A cable (owners of newer iOS devices will need an adapter) connects the phone to an interface board that has eight digital inputs and outputs, six analog inputs, and two analog outputs. A touch to an iOS app lets you read inputs or send commands to the board. Most interesting of all is that you can use an iPhone's sensors and messaging systems to control electronics. • My plan was to use the Nanospark with the Colour Detector sample app, which Nanospark offers as a free download. I wanted the device to spot when the ceiling started dripping, trigger a motor to shepherd the curtain out of the way, and grab my attention. In the daytime or when my room was lit, I wanted to be alerted via a loud buzzer and an LED, but when it was dark and my lights were off (when I—or my neighbors—were more likely to be asleep), I wanted to be alerted just by the LED.

I downloaded the detector app once I got an old iPhone set up. (My iPod Touch was too old to be updated to the required version of iOS.) I soon learned that while it could indeed turn the board's outputs on and off as color values changed in a designated part of the camera's view, the app was inflexible. For example, it could pay attention to only one area at a time in the camera's field of view, so I couldn't watch the dripping on one half of the screen and check for light on the other half. It was something I could surely program myself by writing my own iOS app, but this was supposed to be a quick fix. I decided to detect the level of light in the room with a photoresistor wired to the circuit controlling the buzzer, bypassing the Nanospark.

A quick word about the Nanospark ecosystem: There isn't one yet. There are a few short tutorials, and there are some videos that optimistically describe manufacturing tasks you could automate with Nanospark. The Controller app—more like a debugger, and the only existing app available for download besides the Colour Detector—lets you trigger the outputs individually or read out voltages from the inputs, but there's no "if my phone detects this, then do that" interface. If you're going to have to program all your software from scratch anyway, using Apple's XCode (which has a nontrivial learning curve), maybe Nanospark isn't the go-to manufacturing and tinkering solution it aspires to be. But the device is young; a database of user-created apps and tutorials could eventually fill that gap.

Back to apartmentland. To sense errant droplets, I made a crude color-changing water sensor by attaching a white piece

A VITAL SPARK

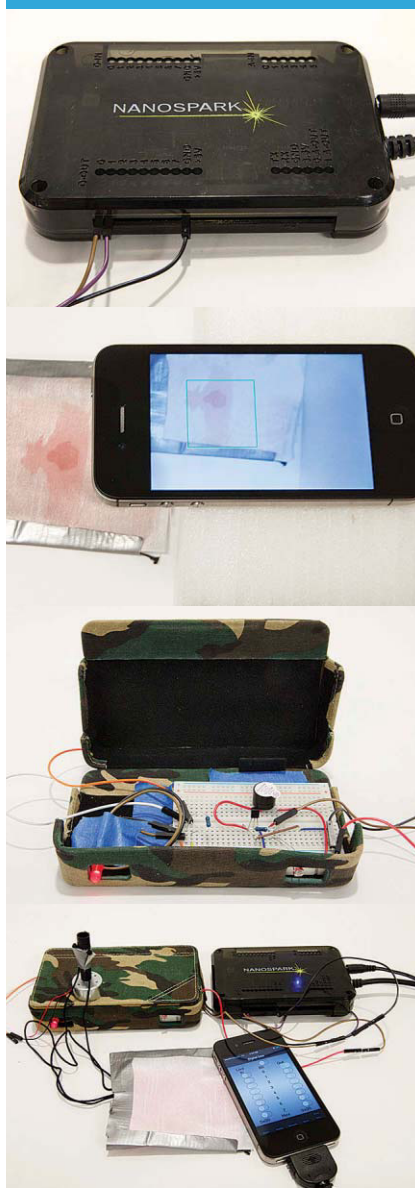

THE NANOSPARK controller comes in a plastic housing [top], which I connected to an iPhone that I'd set up to monitor a strip of fabric [second from top] for color changes caused by a water leak. When the iPhone detects a change, it sends a signal to some external circuitry [second from bottom] that controls an alarm and drives a motor [bottom], to pull a curtain out of the leak's way.

RANDI KLETT (4)

of fabric (a scrap from my newly installed curtains) atop a red piece of plastic placed on a jar. When wet, the fabric turns sheer and the camera can pick up the higher red values in its field of view. I tied this action to a digital output on the Nanospark, which delivers enough current to power a breadboard circuit with a small $5 five-volt motor and an LED. Then I connected another circuit with a photoresistor and the buzzer, all of which I'd borrowed from an Arduino starter kit. I housed all this in an old Nintendo DS case. It conveniently had holes in the front for the LED and photoresistor to peek out and a magnet in its top to hold the motor in place.

When the iPhone spots a drip, the motor reels in a string that draws the curtain aside and whips the fabric and plastic panel off the jar. Then, when the red is no longer visible, the motor stops and the buzzer and LED cut off. When it happens in the middle of the night, it's mildly alarming, even with the buzzer silenced, but when it happens later in the day, I trot off to investigate.

It'll all be over soon, I hope, when my landlord finally fixes the terrifying pipes-and-expanding-wood combo he found upstairs. But when that time comes, I have a long backlog of problems in my apartment that could be solved by a watchful eye and quick, simple motor response. With a relay and access to stronger power, the same setup could easily power something beefier, like a system that unlocks my door when I hold a colored card to the peephole or one that fine-tunes the light based on the height of someone sitting in a chair.

I could think of a million and one more complicated ways to use my iPhone's sensors—or a combination of those and remote messaging—to control and regulate my apartment. But most of those are beyond the amount of time I'm willing to spend without access to more customizable apps.

Until then, my trusty Nanospark, iPhone, and DS case full of parts will be the quickest of quick fixes. **—SARAH LEWIN**

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

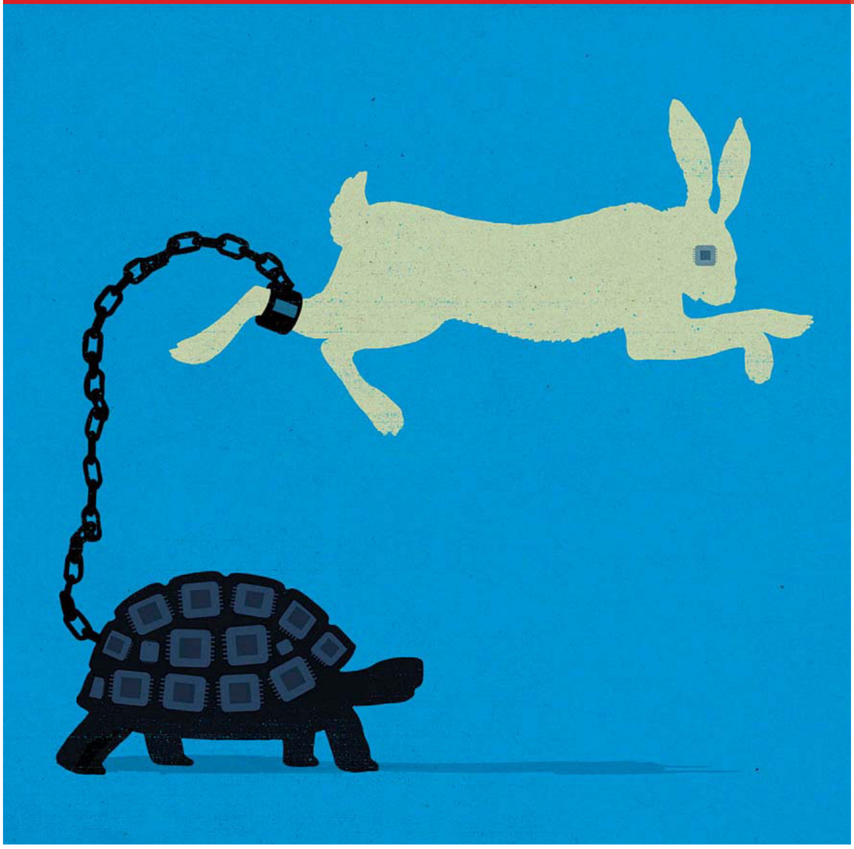

MOORE'S CURSE

IN 1965, THE YEAR IN WHICH THE NUMBER OF components on a microchip had doubled, Gordon Moore predicted that "certainly over the short term this rate can be expected to continue." In 1975 he revised the doubling rate to two years; later, it settled down at about 18 months, or an exponential growth rate of 46 percent a year. This is Moore's Law. ● As components have gotten smaller, denser, faster, and cheaper, they have increased the power and cut the costs of many products and services, notably computers and digital cameras but also light-emitting diodes and photovoltaic cells. The result has been a revolution in electronics, lighting, and photovoltaics. ● But the revolution has been both a blessing and a curse, for it has had the unintended effect of raising expectations for technical progress. We are assured that rapid progress will soon bring self-driving electric cars, hypersonic airplanes, individually tailored cancer cures, and instant three-dimensional printing of hearts and kidneys. We are even told it will pave the world's transition from fossil fuels to renewable energies. ● But the doubling time for transistor density is no guide to technical progress generally. Modern life depends on many processes that improve rather slowly, not least the production of food and energy and the transportation of people and goods. There is no shortage of historical data to illustrate this reality, and I have calculated representative rates for the decades coinciding with the development of transistors (the first commercial application was in hearing aids in 1952) and microprocessors, as well as the rates for the entire 20th century, or even longer.

Corn, America's leading crop, has seen its average yields rising by 2 percent a year since 1950. The efficiency with which steam turbogenerators convert thermal power to electricity generation rose annually by about 1.5 percent during the 20th century; if you instead compare the steam turbogenerators of 1900 with the combined-cycle power plants of 2000 (which mate gas turbines to steam boilers), that annual rate increases to 1.8 percent. Advances in lighting have been more impressive than in any other sector of electricity conversion, but between 1881 and 2014 light efficacy (lumens per watt) rose by just 2.6 percent a year, for indoor lights, and by 3.1 percent for outdoor lighting (topped by the best low-pressure sodium lamps).

The speed of intercontinental travel rose from about 35 kilometers per hour for large ocean liners in 1900 to 885 km/h for the Boeing 707 in 1958, an average rise of 5.6 percent a year. But that speed has remained essentially constant ever since—the Boeing 787 cruises just a few percent faster than the 707. Between 1973 and 2014, the fuel-conversion efficiency of new U.S. passenger cars (even after excluding monstrous SUVs and pickups) rose at an annual rate of just 2.5 percent, from 13.5 to 37 miles per gallon (that's from 17.4 liters per 100 kilometers to 6.4 L/100 km). And finally, the energy cost of steel (coke, natural gas, electricity), our civilization's most essential metal, was reduced from about 50 gigajoules to less than 20 per metric ton between 1950 and 2010—that is, an annual rate of about -1.7 percent.

Energy, material, and transportation fundamentals that enable the functioning of modern civilization and that circumscribe its scope of action are improving steadily but slowly. Gains in performance range mostly from 1.5 to 3 percent a year, as do the declines in cost. Outside the microchip-dominated world, innovation simply does not obey Moore's Law, proceeding at rates that are lower by an order of magnitude. ■

ILLUSTRATION BY Dan Page

Recording Data Has Never Been This Easy!

The all new TrendCorder TRC-18 from Dewetron is easy on the eyes, and even easier on your wallet. Set-up and operation is like a streamlined app, and includes multitouch pinch and zoom.

Our powerful DejaView® capability allows you to scroll back and zoom in on any part of the data even whill still recording at full speed! Even if that data was hours or days in the past, you don't have to stop recording to examine it at any resolution. And the price?

Only $19,950 for an 18 channel system

PLUS - you can add more fast channels by plugging in another TRION module – or add dozens of slow channels using our external EPAD 8-channel bricks. The system is 100% modular and field-upgradeable.

Toll-free: +1 877-431-5166
www.trendcorder.info

- 18 analog input channels
- Wide-aspect multitouch display
- Up to 250 kS/s/ch sample rate
- 1500 V isolation on all channels!
- 8th order filters on all channels!
- Multi-touch display with pinch and zoom operation.
- DejaView® provides unlimited scroll back and viewing resolution during recording

Please contact us to set up a demonstration today!

TRENDCORDER
+1 401-284-3750 www.dewamerica.com sales@dewamerica.com

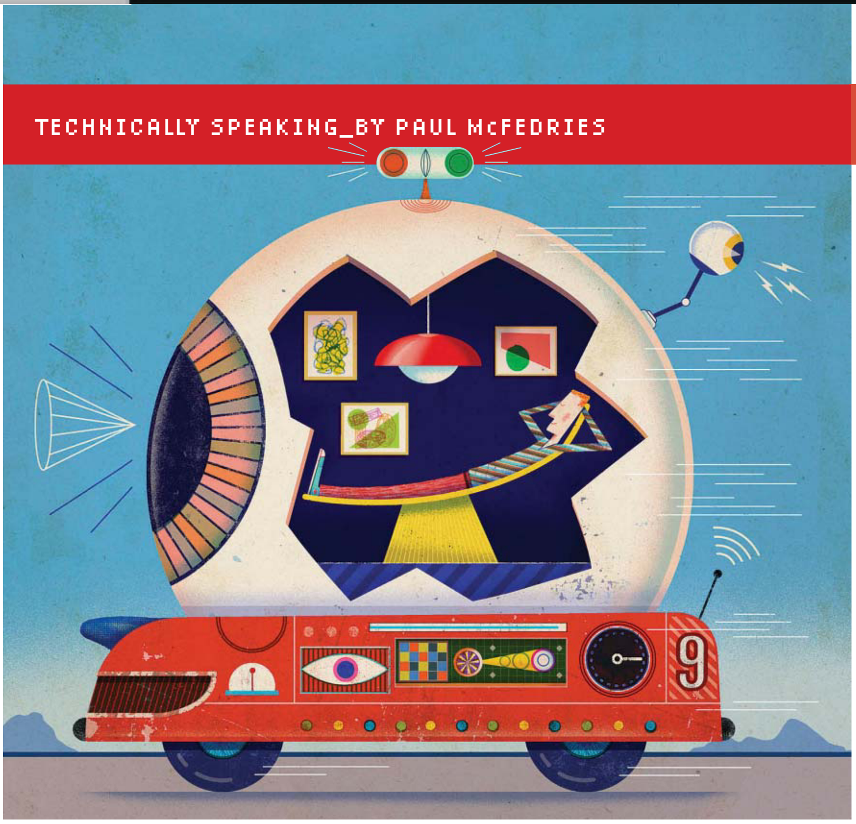

YOUR INEVITABLE ROBOCAR FUTURE

Hands-free driving, cars that park themselves, an unmanned car driven by a search-engine company. We've seen that movie. It ends with robots harvesting our bodies for energy. *—2011 Chrysler ad*

WATCH JUST ABOUT any movie or television show set in the not-too-distant future and you'll soon notice that although there are still plenty of cars on the road, nobody drives them. Or, I should say, no *human* drives them. From the **Johnny Cabs** in *Total Recall* to the sleek Lexus 2054 in *Minority Report*, in these imagined futures robots and computers have taken over the driving duties. ● Science fiction portrays **autonomous cars** as inevitable, but will we really see them anytime soon? We already have **self-parking cars** from companies such as BMW, Lexus, and Toyota, in which an onboard sensor detects a suitable parking spot and a built-in computer controls steering, accelerator, and brake to maneuver the car into the spot. But **automatic parking** (even the neat trick of **automatic parallel parking**) is a long way from **automatic driving**, although it might not be as long as you think. Google, famously, has been testing a fleet of cars that can operate in **auto-drive** mode. Google cofounder Sergey Brin was asked in 2013 when **driverless cars** might become mainstream, and he replied, "You can count on one hand the number of years until ordinary people can experience this." With nearly every major car company—including Audi, GM, Mercedes-Benz, Nissan, Tesla, and Volkswagen—testing **robotic car** prototypes, and with **robocars** already legal (albeit with restrictions) in California, Florida, Michigan, and Nevada, the fulfillment of Brin's prediction is probably not far off. ● Some are also predicting that we won't need to own cars when they're all **self-driving**. If you can summon a **robotaxi** in minutes with a few taps on your smartphone, who needs the hassle and expense of car ownership? (Since you in effect "whistle" for such a car, in a high-tech version of whistling for a taxi, **robocabs** and similar **self-delivering** vehicles are sometimes called **whistlecars**.) If this all sounds very Uber-like, know that Uber CEO Travis Kalanick has declared that replacing the company's current fleet of **human-driven vehicles** with **self-driven vehicles** is inevitable and will make the service incredibly cheap because there will be "no other dude in the car" to pay for. A similar idea is the **deliverbot** (also called the **robotruck**), an unmanned vehicle for delivering packages and other cargo.

But a car that can drive itself is only the beginning. The long-term vision is to combine **smart cars** with **smart roads**. That is, an advanced **connected car** will drive autonomously, *and* it will tap into the sensors and beacons that will festoon future roads and highways, leading to the ideal of **crash-avoiding** or **crashless** vehicles. This will also enable **platooning**, in which cars drive at a steady speed and follow each other at a set distance. The resulting **car train** or **road train** will be completely controlled by **vehicle-to-vehicle** (**V2V**) communication and interaction with the so-called **automated highway system** (or **intelligent transport system**). The ultimate goal is an **autocar** that is so intelligent and so safe that you could fall asleep in it—the pie-in-the-sky, head-on-the-pillow **sleeper car**.

The cloud of words and phrases surrounding self-driving vehicles is interesting, and the technology itself offers a fascinating glimpse into the near future, but do we really need driverless cars and smart roads and vehicle platoons? Well, let's look at the facts. Drivers run stop signs, cut off pedestrians, tailgate the car in front, and race the car behind. And they do all this and worse in a fog of barely contained road rage. No wonder more than 32,000 people are killed in traffic accidents every year in the United States alone. Clearly, when it comes to robocars, the future can't come soon enough. ■

ILLUSTRATION BY Christian Northeast

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

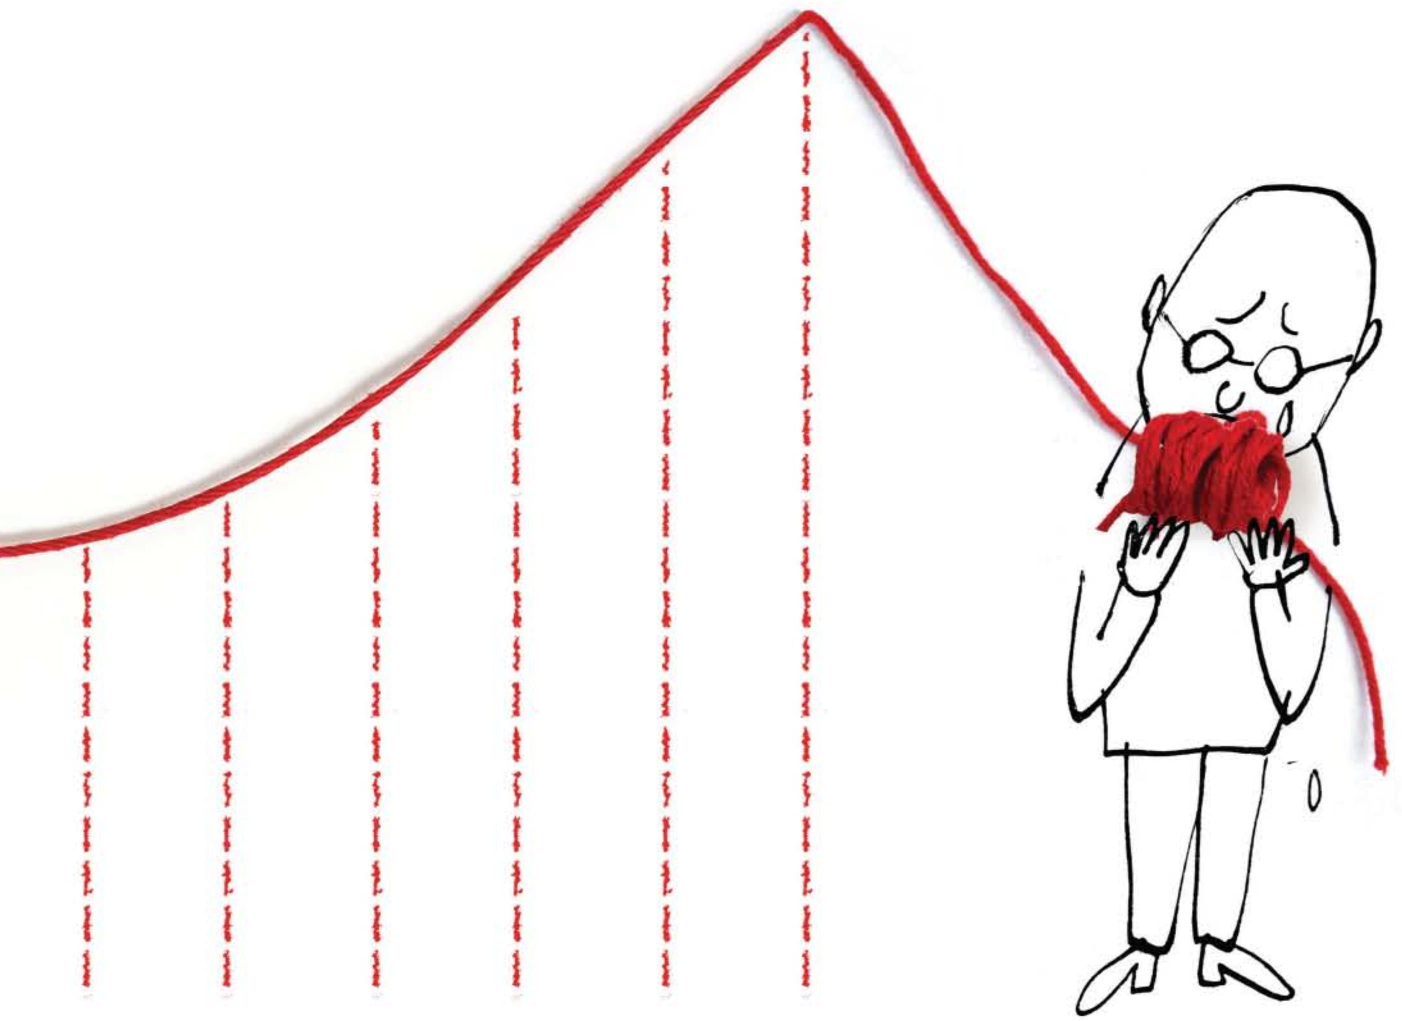

THE LONG GOOD-BYE

Moore's Law is going out not with a bang but a whimper

FIFTY YEARS AGO THIS MONTH, Gordon Moore forecast a bright future for electronics. His ideas were later distilled into a single organizing principle—Moore's Law—that has driven technology forward at a staggering clip. We have all benefited from this miraculous development, which has forcefully shaped our modern world. In this special report, we celebrate that glorious history and also consider the trend's inevitable decline. We find that the end won't be sudden and apocalyptic but rather gradual and complicated. Moore's Law truly is the gift that keeps on giving—and surprising, as well.

Illustrations by **SERGE BLOCH**

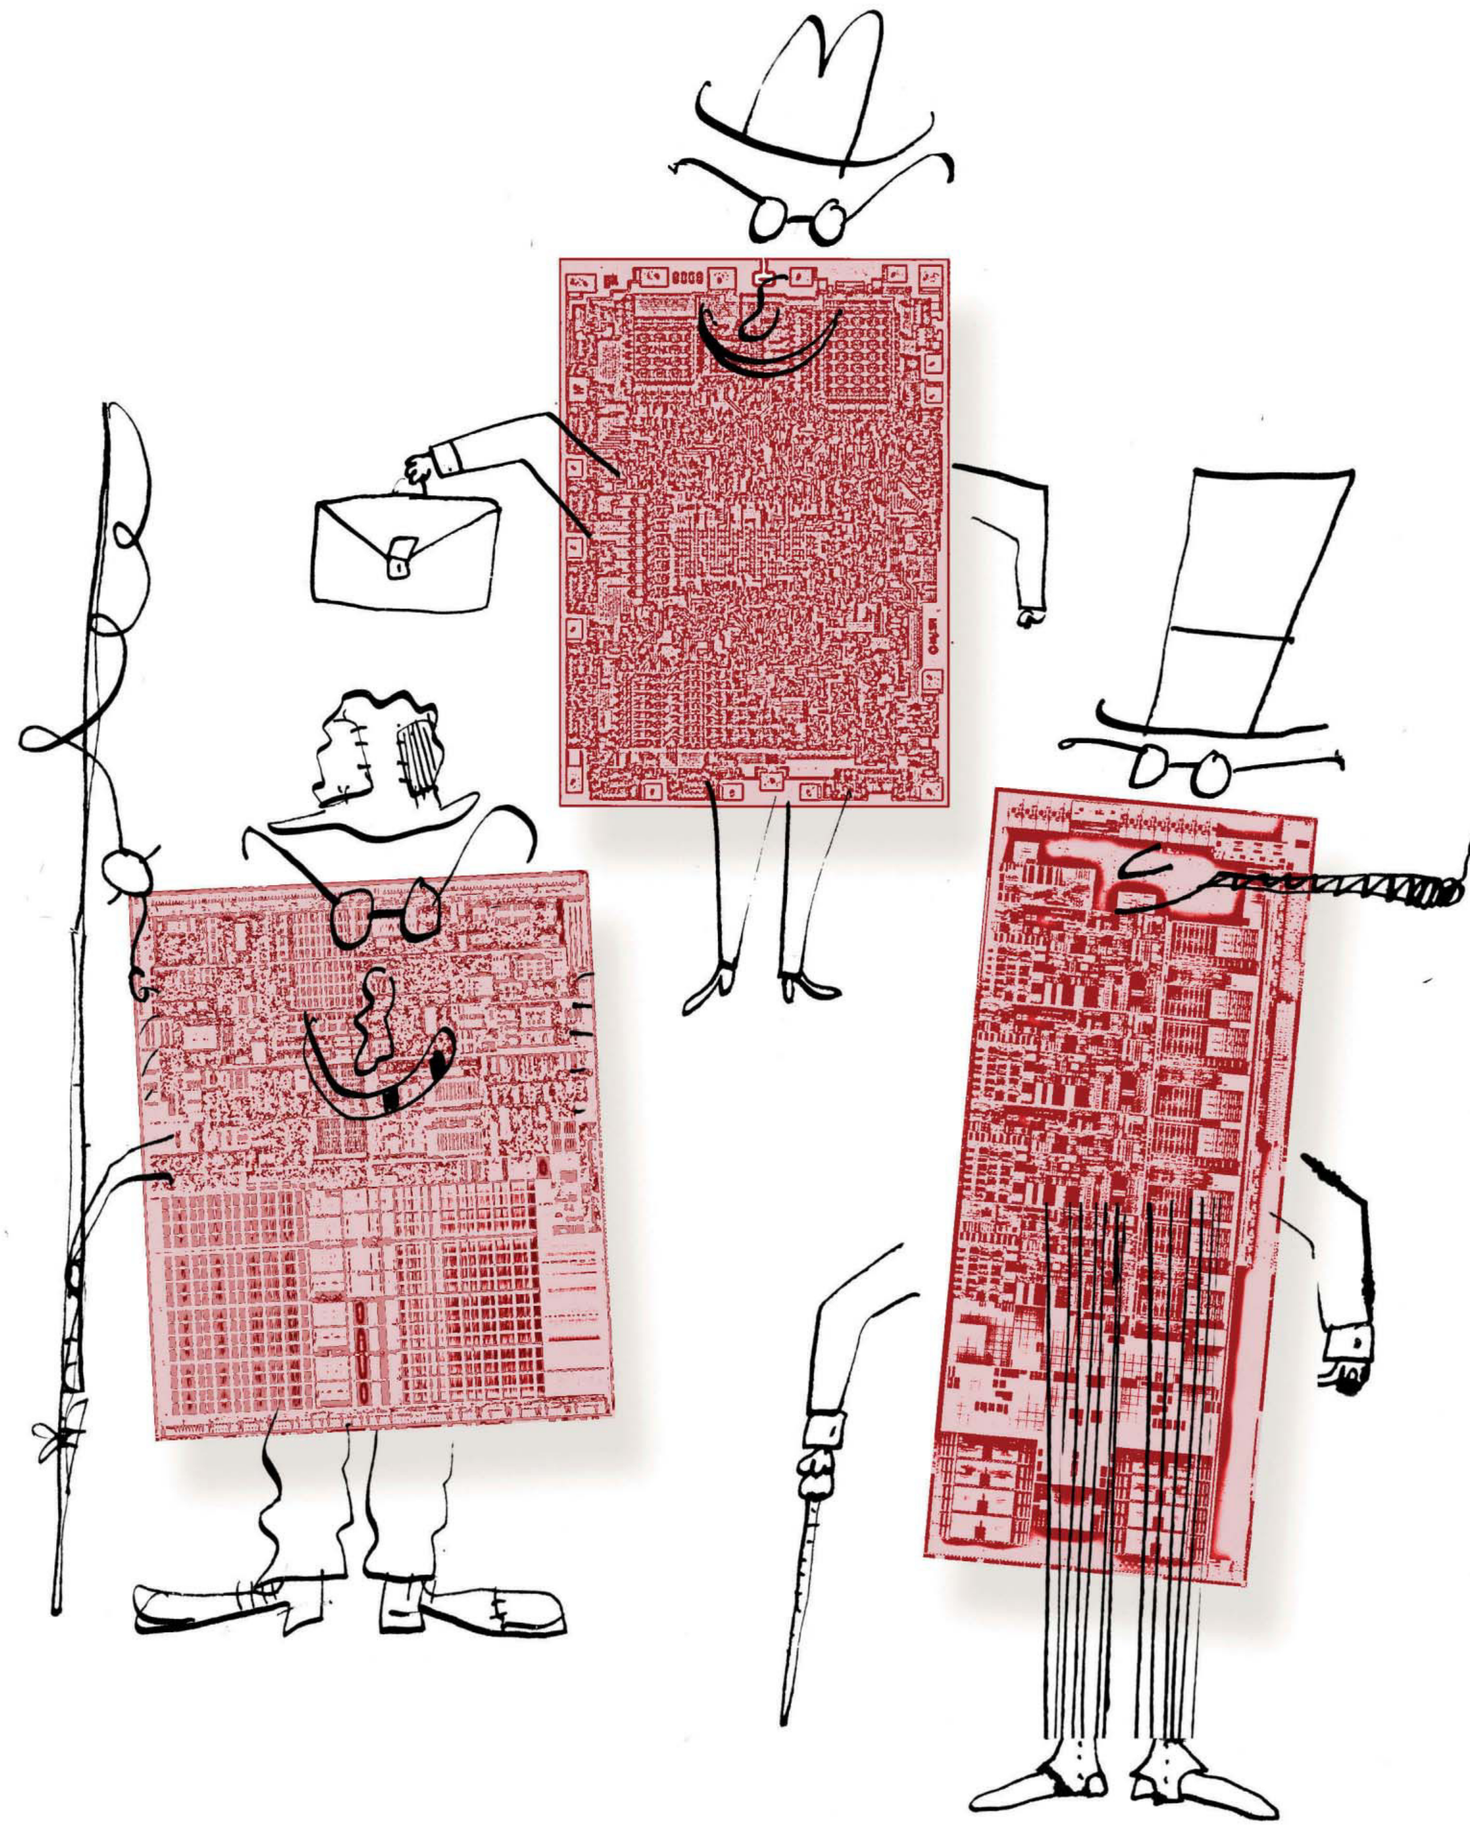

MOORE'S LAW 50 YEARS

The MULTIPLE LIVES of MOORE'S LAW

Why Gordon Moore's grand prediction has endured for 50 years

By **CHRIS MACK**

A HALF CENTURY AGO, a young engineer named Gordon E. Moore took a look at his fledgling industry and predicted big things to come in the decade ahead. In a four-page article in the trade magazine *Electronics*, he foresaw a future with home computers, mobile phones, and automatic control systems for cars. All these wonders, he wrote, would be driven by a steady doubling, year after year, in the number of circuit components that could be economically packed on an integrated chip. • A decade later, the exponential progress of the integrated circuit–later dubbed "Moore's Law"–showed no signs of stopping. And today it describes a remarkable, 50-year-long winning streak that has given us countless forms of computers, personal electronics, and sensors. The impact of Moore's Law on modern life can't be overstated. We can't take a plane ride, make a call, or even turn on our dishwashers without encountering its effects. Without it, we would not have found the Higgs boson or created the Internet. • But what exactly is Moore's Law, and why has it been so successful? Is it evidence of technology's inevitable and unstoppable march? Or does it simply reflect a unique time in engineering history, when the special properties of silicon and a steady series of engineering innovations conspired to give us a few decades of staggering computational progress?

I would argue that nothing about Moore's Law was inevitable. Instead, it's a testament to hard work, human ingenuity, and the incentives of a free market. Moore's prediction may have started out as a fairly simple observation of a young industry. But over time it became an expectation and self-fulfilling prophecy—an ongoing act of creation by engineers and companies that saw the benefits of Moore's Law and did their best to keep it going, or else risk falling behind the competition.

I would also argue that, despite endless paraphrasing, Moore's Law is not one simple concept. Its meaning has changed repeatedly over the years, and it's changing even now. If we're going to draw any lessons from Moore's Law about the nature of progress and what it can tell us about the future, we have to take a deeper look.

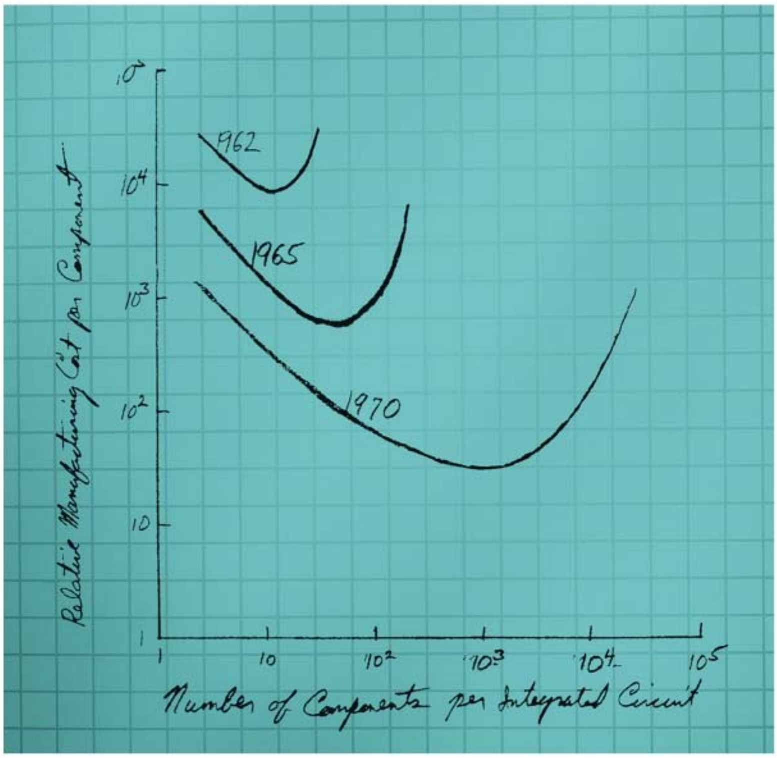

THE SWEET SPOT: Economics was at the core of Moore's 1965 paper. He argued that for any particular generation of manufacturing technology, there is a cost curve. The cost of making a component declines the more you pack onto an integrated circuit, but past a certain point, yields decline and costs rise. The sweet spot, where the cost per component is at a minimum, moves to more and more complex integrated circuits over time.

INTEL

In the early 1960s, before Silicon Valley became known as Silicon Valley, Gordon Moore was director of research and development at Fairchild Semiconductor. He and others had founded the company in 1957 after defecting from Shockley Semiconductor Laboratory, where they'd done some of the early work on silicon electronic devices.

Fairchild was one of a small group of companies working on transistors, the now ubiquitous switches that are built by the billions onto chips and are used to perform computations and store data. And the firm quickly started carving out a niche.

At the time, most circuits were constructed from individual transistors, resistors, capacitors, and diodes that were wired together by hand on a circuit board. But in 1959, Jean Hoerni of Fairchild invented the planar transistor—a form of transistor that was constructed in the plane of the silicon wafer instead of on a raised plateau, or mesa, of silicon.

With this configuration, engineers could build wires above the transistors to connect them and so make an "integrated circuit" in one fell swoop on the same chip. Jack Kilby of Texas Instruments had pioneered an early integration scheme that connected devices with "flying wires" that rose above the surface of the chip. But Moore's colleague Robert Noyce showed that planar transistors could be used to make an integrated circuit as a solid block, by coating the transistors with an insulating layer of oxide and then adding aluminum to connect the devices. Fairchild used this new architecture to build the first silicon integrated circuit, which was announced in 1961 and contained a whopping four transistors. By 1965, the company was getting ready to release a chip with 64 components.

Armed with this knowledge, Moore opened his 1965 paper with a bold statement: "The future of integrated electronics is the future of electronics itself." That claim seems self-evident today, but at the time it was controversial. Many people doubted that the integrated circuit would ever fill anything more than a niche role.

You can forgive the skepticism. Although the first integrated chips were more compact than their hand-wired brethren, they cost significantly more—about US $30 per component in today's dollars compared with less than $10 for stand-alone components. Only a handful of companies were making integrated circuits, and their only real customers were NASA and the U.S. military.

Compounding the problem was the fact that transistors were still unreliable. Of the individual transistors that were made, only a small fraction—just 10 to 20 percent, Moore later recalled—actually worked. Pack a half dozen of those devices together in an integrated circuit and you'd expect those small fractions to multiply, yielding a dismally small number of operational chips.

But this logic was flawed. It turned out that making a chip with eight transistors yields a fraction of operational chips similar to what you'd get by making eight stand-alone transistors. That's because the probabilities aren't independent. Defects take up space, and many types are distributed randomly, like paint splatter. If two transistors are placed close together, a single transistor-size flaw can take out both devices. As a result, putting two transistors side by side carries about the same risk of death by defect as one transistor by itself.

Moore was convinced that integration would ultimately prove economical. In his 1965 paper, as evidence of the integrated circuit's bright future, he plotted five points over time, beginning with Fairchild's first planar transistor and followed by a series of the company's integrated circuit offerings. He used a semilogarithmic plot, in which one axis is logarithmic and the other linear and an exponential function will appear as a straight line. The line he drew through the points was indeed more or less straight, with a slope that corresponded to a doubling of the number of components on an integrated circuit every year.

From this small trend line, he made a daring extrapolation: This doubling would continue for 10 years. By 1975, he predicted, we'd see the number of components on an integrated circuit go from about 64 to 65,000. He got it very nearly right. By 1975, Intel, the company Moore cofounded after leaving Fairchild in 1968, was preparing charged-coupled-device (CCD) memory chips with some 32,000 components–only a factor of two off from his thousandfold prediction.

Looking back on this remarkable paper, I'll note a few details that are often overlooked. First, Moore's prediction was about the number of electronic components–not just transistors but also devices such as resistors, capacitors, and diodes. Many early integrated circuits actually had more resistors than transistors. Later, metal-oxide-semiconductor (MOS) circuitry, which relied less on nontransistor components, emerged, and the digital age began. Transistors dominated, and their number became the more useful measure of integrated circuit complexity.

The paper also reveals Moore's focus on the economics of integration. He defined the number of components per chip not as the maximum or the average number of components but as the number for which the cost per component was at a minimum. He understood that the number of components that you can pack on a chip and the number that makes economic sense are not necessarily the same. Instead, there's a sweet spot for every generation of chip-fabrication technology. As you add more components, you drive the cost per component down. But past a certain point, attempting to pack even more transistors into a given space will raise the possibility of killer defects and lower the yield of useful chips. At that point, the cost per component will start to rise. The goal of integrated circuit design and manufacturing was–and still is–to hit this sweet spot.

DATA SOURCE: VLSI RESEARCH

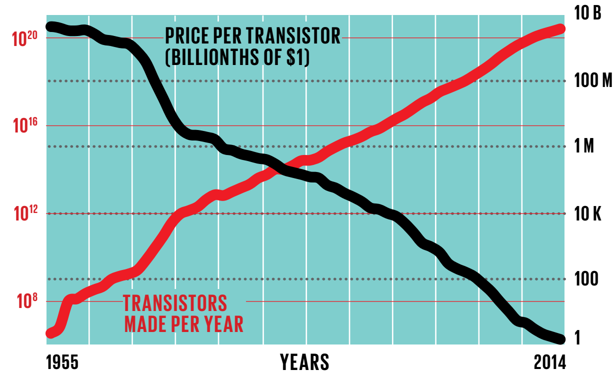

Transistors, by the Numbers

IN 2014, semiconductor production facilities made some 250 billion billion (250×10^{18}) transistors. This was, literally, production on an astronomical scale. Every second of that year, on average, 8 trillion transistors were produced. That figure is about 25 times the number of stars in the Milky Way and some 75 times the number of galaxies in the known universe.

The rate of growth has also been extraordinary. More transistors were made in 2014 than in all the years prior to 2011. Even the recent great recession had little effect. Transistor production in 2009—a year of deep recession for the semiconductor industry—was more than the cumulative total prior to 2007.

The collective pursuit of Moore's Law has driven this growth. For decades, manufacturing innovation and simple miniaturization have enabled engineers to pack more capability into the same area of silicon. The result has been a steady decrease in manufacturing cost per transistor (transistor price, which is easier to track, is plotted above).

This steady, predictable decline in prices was a self-reinforcing gift. Because electronics manufacturers could depend on Moore's Law, they could plan further ahead and invest more in the development of new and better-performing products. In ways profound and surprising, this situation promoted economic growth. It has been the ever-rising tide that has not only lifted all boats but also enabled us to make fantastic and entirely new kinds of boats. **—DAN HUTCHESON**

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

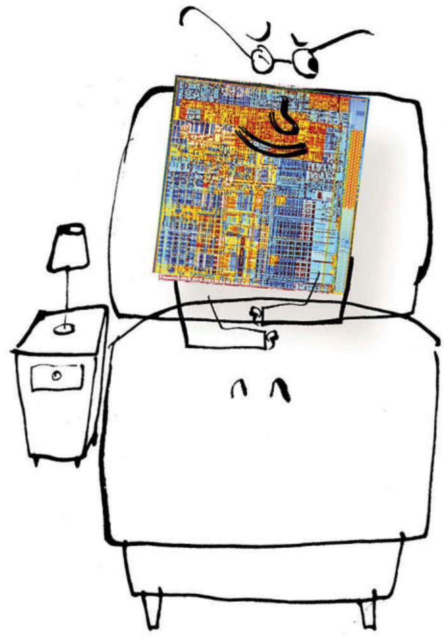

Efficiency's Brief Reprieve

Moore's Law slowdown hits performance more than energy efficiency

No one can say exactly when the era of Moore's Law will come to a close. Nevertheless, semiconductor experts like us can't resist speculating about that day because it will mark the end of an extraordinary period of history, with uncertain implications for one of the world's most important industries.

Here's what we do know. The last 15 years have seen a big falloff in how much performance improves with each new generation of cutting-edge chips. So is the end nigh? Not exactly, because even though the fundamental physics is working against us, it appears we'll have a reprieve when it comes to energy efficiency.

As chip-fabrication technology has improved, the sweet spot has shifted to larger numbers of components and lower costs per component. Over the last 50 years, the cost of a transistor has been reduced from $30 in today's dollars to a nanodollar or so. Moore could hardly have predicted such a dramatic reduction. But even in 1965, he understood that integrated circuits were about to change from an expensive, high-performance replacement for discrete components to a cheap, high-performance replacement, and that both performance and economics would favor integration.

Ten years later, Moore revisited his prediction and revised it. In an analysis he'd done for the 1975 IEEE International Electron Devices Meeting, he started by tackling the question of how the doubling of components actually happened. He argued that three factors contributed to the trend: decreasing component size, increasing chip area, and "device cleverness," which referred to how much engineers could reduce the unused area between transistors.

Moore attributed about half of the doubling trend to the first two factors and the rest to "cleverness." But when he considered the CCD memories that Intel was preparing to release, he decided that cleverness would soon go out the window. In CCD arrays, devices are packed together in tight grids with no wasted space to eliminate. So he predicted the doubling trend would soon be driven only by tinier transistors and bigger chips. As a result it would slow by half, doubling components once every two years instead of every year.

Ironically, CCD memory proved to be too error prone, so Intel never shipped any. But the prediction was nonetheless borne out in logic chips, such as microprocessors, which have grown at about a two-year doubling rate since the early 1970s. Memory chips, with their massive arrays of identical transistors, scaled faster, doubling every 18 months or so, mainly because they are simpler to design.

Of the three technology drivers Moore identified, one turned out to be special: decreasing the dimensions of the transistor. For a while at least, shrinking MOS transistors—the silicon-gated transistors we still use today—offered something that rarely happens in the world of engineering: no trade-offs. Thanks to a scaling rule named for IBM engineer Robert Dennard, every successive transistor generation was better than the last. A shrinking transistor not only allowed more components to be crammed onto an integrated circuit but also made those transistors faster and less power hungry.

This single factor has been responsible for much of the staying power of Moore's Law, and it's lasted through two very different incarnations. In the early days, a phase I call Moore's Law 1.0, progress came by "scaling up"—adding more components to a chip. At first, the goal was simply to gobble up the discrete components of existing applications and put them in one reliable and inexpensive package. As a result, chips got bigger and more complex. The microprocessor, which emerged in the early 1970s, exemplifies this phase.

But over the last few decades, progress in the semiconductor industry became dominated by Moore's Law 2.0. This era is all about "scaling down," driving down the size and cost of transistors even if the number of transistors per chip does not go up.

Although the Moore's Law 1.0 and 2.0 eras have overlapped, the dominance of scaling down versus scaling up can be seen

There are many ways to gauge a computer's efficiency, but one of the most easily calculated metrics is peak-output efficiency, which measures the efficiency of a processor when it's running at its fastest.

Peak-output efficiency is typically quoted as the number of computations that can be performed per kilowatt-hour of electricity consumed. And according to a peer-reviewed paper published in 2011 in the *IEEE Annals of the History of Computing*, it doubled like clockwork every year and a half or so for more than five decades.

This trend started well before the first microprocessor, way back in the mid-1940s. But it began to come to an end around 2000. Growth in both peak-output efficiency and performance started to slow, weighed down by the physical limitations of shrinking transistors. Chipmakers turned to architectural changes—such as putting multiple computing cores in a single microprocessor—but they weren't able to maintain historical growth rates.

These days, we've found, it takes about 2.7 years for peak-output efficiency to double. That's a substantial slowdown. Historically, a decade of doubling boosted efficiency by a factor of a hundred; at current rates, it would take 18 years to see a hundredfold gain.

Fortunately, the news isn't all bad. Our computing needs have changed. For years after Moore's landmark 1965 paper, computers were expensive, relatively rare, and regularly pushed to their computing peak. Now that they're ubiquitous and cheap, the emphasis in chip design has shifted from fast CPUs in stationary machines to ultralow-power processing in mobile appliances, such as laptops, cellphones, and tablets.

Today, most computers run at peak output only a small fraction of the time (a couple of exceptions being high-performance supercomputers and Bitcoin miners). Mobile devices such as smartphones and notebook computers generally operate at their computational peak less than 1 percent of the time based on common industry measurements. Enterprise data servers spend less than 10 percent of the year operating at their peak. Even computers used to provide cloud-based Internet services operate at full blast less than half the time.

In this new regime, a good power-management design is one that minimizes how much energy a device consumes when it's idle or off. And the better indicator of energy efficiency is how much electricity a computer consumes on average—not when it's operating at full blast.

We've recently defined a measure of efficiency that's more in sync with how chips are used nowadays, which we call "typical-use efficiency." Like peak-output efficiency, it's measured in computations per kilowatt-hour. This time, however, it's calculated by dividing the number of computations performed over the course of a year by the total electricity consumed—a weighted sum of the energy a processor and its supporting circuitry use in different modes over that same period. For example, a laptop might operate at peak power when its user is playing a game, but this only happens a tiny fraction of the time. Other common activities, such as word processing or video playback, might consume a tenth as much electricity, since only a fraction of the chip is needed for these functions, and smart power management can actively shut off circuitry between keystrokes and video frames.

Encouragingly, typical-use efficiency seems to be going strong, based on tests performed since 2008 on Advanced Micro Devices' chip line. Through 2020, by our calculations for an AMD initiative, typical-use efficiency will double every 1.5 years or so, putting it back to the same rate seen during the heyday of Moore's Law.

These gains come from aggressive improvements to circuit design, component integration, and software, as well as power-management schemes that put unused circuits into low-power states whenever possible. The integration of specialized accelerators, such as graphics processing units and signal processors that can perform certain computations more efficiently, has also helped keep average power consumption down.

Of course, as with any exponential trend, this one will eventually end, and circuit designers will have become victims of their own success. As idle power approaches zero, it will constitute a smaller and smaller fraction of the energy consumed by a computer. In a decade or so, energy use will once again be dominated by the power consumed when a computer is active. And that active power will still be hostage to the physics behind the slowdown in Moore's Law.

Over the next few decades, we'll have to rethink the fundamental design of computers if we want to keep computing moving forward at historical rates. In the meantime, steady improvements in everyday energy efficiency will give us a bit more time to find our way.

—JONATHAN KOOMEY & SAMUEL NAFFZIGER

ENERGY EFFICIENCY RELATIVE TO 1985 (1985 = 1.0)

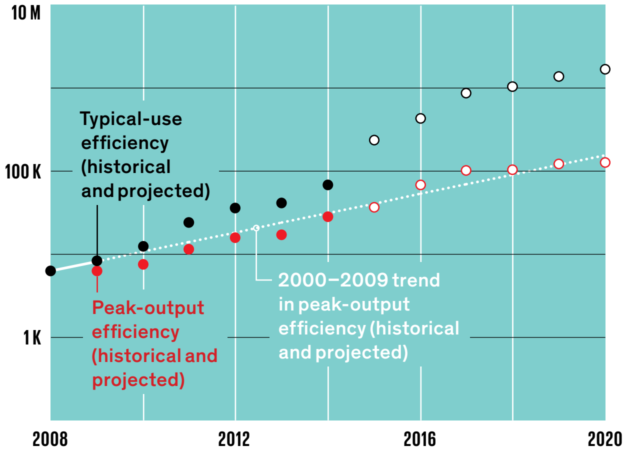

DATA SOURCES: AMD, KOOMEY ET AL. (2011)

in the way the semiconductor industry describes itself. In the 1980s and early 1990s, the technology generations, or "nodes," that define progress in the industry were named after dynamic RAM generations: In 1989, for example, we had the 4-megabyte DRAM node; in 1992, the 16-MB node. Each generation meant greater capability within a single chip as more and more transistors were added without raising the cost.

By the early 1990s, we'd begun to name our nodes after the shrinking features used to make the transistors. This was only natural. Most chips didn't need to carry as many transistors as possible. Integrated circuits were proliferating, finding their way into cars and appliances and toys, and as they did so, the size of the transistor–with the implications for performance and cost savings–became the more meaningful measure.

Eventually even microprocessors stopped scaling up as fast as manufacturing technology would permit. Manufacturing now allows us to economically place more than 10 billion transistors on a logic chip. But only a few of today's chips come anywhere close to that total, in large part because our chip designs generally haven't been able to keep up.

Moore's Law 1.0 is still alive today in the highest-end graphics processing units, field-programmable gate arrays, and perhaps a handful of the microprocessors aimed at supercomputers. But for everything else, Moore's Law 2.0 dominates. And now it's in the process of changing again.

This change is happening because the benefits of miniaturization are progressively falling away. It began in the early 2000s, when an unpleasant reality started to emerge. At that time, transistor sizes began to creep down below 100 nanometers, and Dennard's simple scaling rule hit its limit. Transistors became so small that it was quite easy for electrons to sneak through them even when the devices were supposed to be off, leaking energy and lowering device reliability. Although new materials and manufacturing techniques helped combat this problem, engineers had to stop the practice of dramatically lowering the voltage supplied to each transistor in order to maintain a strong electrical clamp.

Because of the breakdown of Dennard scaling, miniaturization is now full of trade-offs. Making a transistor smaller no longer makes it both faster and more efficient. In fact, it's very difficult to shrink today's transistors and maintain even the same speed and power consumption of the previous generation.

As a result, for the last decade or so, Moore's Law has been more about cost than performance; we make transistors smaller in order to make them cheaper. That isn't to say that today's microprocessors are no better than those of 5 or 10 years ago. There have been design improvements. But much of the performance gains have come from the integration of multiple cores enabled by cheaper transistors.

The economics has remained compelling because of an important and unheralded feature of Moore's Law: As transistors have gotten smaller, we've been able to keep the cost of manufacturing each square centimeter of finished silicon about the same, year after year after year (at least until recently). Moore has put it at about a billion dollars an acre–although chipmakers seldom think in terms of acreage.

Keeping the cost of finished silicon constant for decades hasn't been easy. There was steady work to improve yield, which started in the 1970s at around 20 percent and now sits at 80 to 90 percent. Silicon wafers–the round platters of silicon that are eventually cut into chips–also got bigger and bigger. The progressive boost in size lowered the cost of a number of manufacturing steps, such as deposition and etching, that are performed on a whole wafer at once. And crucially, equipment productivity has soared. The tools employed in lithography–the printing technology that's used to pattern transistors and the interconnections between them–cost 100 times as much today as they did 35 years ago. But these tools pattern wafers 100 times as fast, making up the cost increase while delivering far better resolution.

These three factors–improved yields, larger wafers, and rising equipment productivity–have allowed chipmakers to make chips denser and denser for decades while keeping the cost per area nearly the same and reducing the cost per transistor. But now, this trend may be ending. And it's largely because lithography has gotten more expensive.

Over the last decade, the difficulties of printing tiny features have raised the manufacturing cost per unit area of finished silicon about 10 percent per year. Since the area per transistor shrank by about 25 percent each year over the same period, the cost of each transistor kept going down. But at some point, manufacturing costs will rise faster than transistor area will fall, and the next generation of transistors will be more expensive than the last.

If lithography costs rise fast, Moore's Law as we know it will come to a quick halt. And there are signs that the end could come quite soon. Today's advanced chips are made with immersion lithography, which makes patterns by exposing water-immersed wafers to 193-nm, deep ultraviolet light. The planned successor is lithography using shorter-wavelength, extreme ultraviolet light. That technology was supposed to come on line as early as 2004. But it's suffered delay after delay, so chipmakers have had to turn to stopgaps such as double patterning, which doubles up some steps to fashion the finest features. Double patterning takes twice as long as single patterning. Nonetheless, chipmakers are contemplating triple and even quadruple patterning, which will further drive up costs. A few years from now, we may look back on 2015 as the year the tide turned and the cost of transistors stopped falling and started to rise.

I've been known for making grand pronouncements at lithography conferences about the coming end of Moore's Law. But the truth is, I don't think Moore's Law is over. Instead, I'd argue it's on the verge of morphing again.

Going forward, innovations in semiconductors will continue, but they won't systematically lower transistor costs. Instead, progress will be defined by new forms of integration: gathering together disparate capabilities on a single chip to lower the system cost. This might sound a lot like the Moore's Law 1.0 era, but in this case, we're not looking at combining different pieces of logic into one, bigger chip. Rather, we're talking about uniting the non-logic functions that have historically stayed separate from our silicon chips.

An early example of this is the modern cellphone camera, which incorporates an image sensor directly onto a digital signal processor using large vertical lines of copper wiring called through-silicon vias. But other examples will follow. Chip designers have just begun exploring how to integrate microelectromechanical systems, which can be used to make tiny accelerometers, gyroscopes, and even relay logic. The same goes for microfluidic sensors, which can be used to perform biological assays and environmental tests.

All of these technologies allow you to directly connect a digital CMOS chip with the outside, analog world. This could have a powerful economic effect if the new sensors and actuators can take advantage of the low-cost, mass-production approaches common to silicon manufacturing.

But this new phase of Moore's Law–what I call Moore's Law 3.0 and what others in the semiconductor industry call "more than Moore"–may not make economic sense. Integrating nonstandard components onto a chip offers many exciting opportunities for new products and capabilities. What it doesn't offer is the regular, predictable road map for continued success.

The path forward will be much murkier. Adding a new capability to a chip may make a company money today, but there's no guarantee that adding another will earn it more money tomorrow. No doubt this transition will be painful for some established semiconductor companies, with the winners and losers yet to be determined.

Still, I think Moore's Law 3.0 could be the most exciting rendition of the law yet. Once we get past our expectations for easily quantifiable progress, we could see an explosion of creative applications: bionic appendages that operate seamlessly with the body, smartphones that can sniff the air or test the water, tiny sensors that can power themselves from ambient energy sources, and a host of other applications we have yet to imagine. Moore's Law as we know it might be coming to an end. But its legacy will keep us moving forward for a long time to come. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/mooreslawfifty0415

When Mead Met Moore

DECADES AGO, Carver Mead performed some of the earliest work aimed at determining just how small transistors could ultimately get. A longtime colleague of Gordon Moore, he's widely credited with popularizing the phrase "Moore's Law." Their relationship dates all the way back to 1960—five years before Moore's famous article in *Electronics*, which kicked off popular attention to the trend. Recently, Mead told *IEEE Spectrum* about his poignant first encounter with the electronics legend.

"I was a brand new assistant professor at Caltech, just in my first year, and I was in my office working away on the results of some experiment I had done. This guy waltzed into my office and said, 'Hi, I'm Gordon Moore from Fairchild,'" Mead says, laughing. "Well, I had never heard of Gordon Moore, but I knew about Fairchild."

"We shook hands, and he said he was on campus to recruit some engineers, and would I like some transistors to teach my lab course? And I said, 'Oh, that would be absolutely great.' So he reached into the top of his briefcase, [and] the first thing he did was pull out a sock or a dirty shirt or something. . . . I was looking at him a little surprised. He turned around with this little grin on his face and said, 'I travel light.' "

Moore went on to pull out two big 8½-by-11 manila envelopes, each "bulging" with transistors, Mead says. One was full of devices from Fairchild's 2N697 line, one of the company's first, and the other zippy 2N706 switches.

"I had never seen so many transistors," Mead recalls. "I was completely blown away. In those days none of us had much budget for things like that for teaching. We were working with really cheap transistors that were about a dollar apiece in the stock room. For a student to shell out that for [a device] that might burn out on the first experiment was not easy. Having some transistors that the students could work with without having to break their budget was a great thing." It was the beginning of years of close collaboration.

—RACHEL COURTLAND

MOORE'S LAW 50 YEARS

The LAW That's NOT A LAW

Gordon Moore reflects on 50 years of technological progress

GORDON MOORE PIONEERED the integrated circuit and cofounded the chip giant Intel; in retirement he has focused on science- and technology-oriented philanthropy. But thanks to an article he published in April 1965 in *Electronics* magazine, he's known most widely for Moore's Law, the prediction that has reflected–and helped drive–steady and staggeringly fast progress in computing technology. In preparation for the 50th anniversary of Moore's prediction, *IEEE Spectrum* Associate Editor Rachel Courtland visited the man himself at his home on Hawaii's Big Island.

PHOTOGRAPH BY Olivier Koning

IT'S GOOD TO BE GORDON: Since retiring from Intel, Gordon Moore has focused on philanthropy through the Gordon and Betty Moore Foundation. He stands here in the backyard of his home in Hawaii.

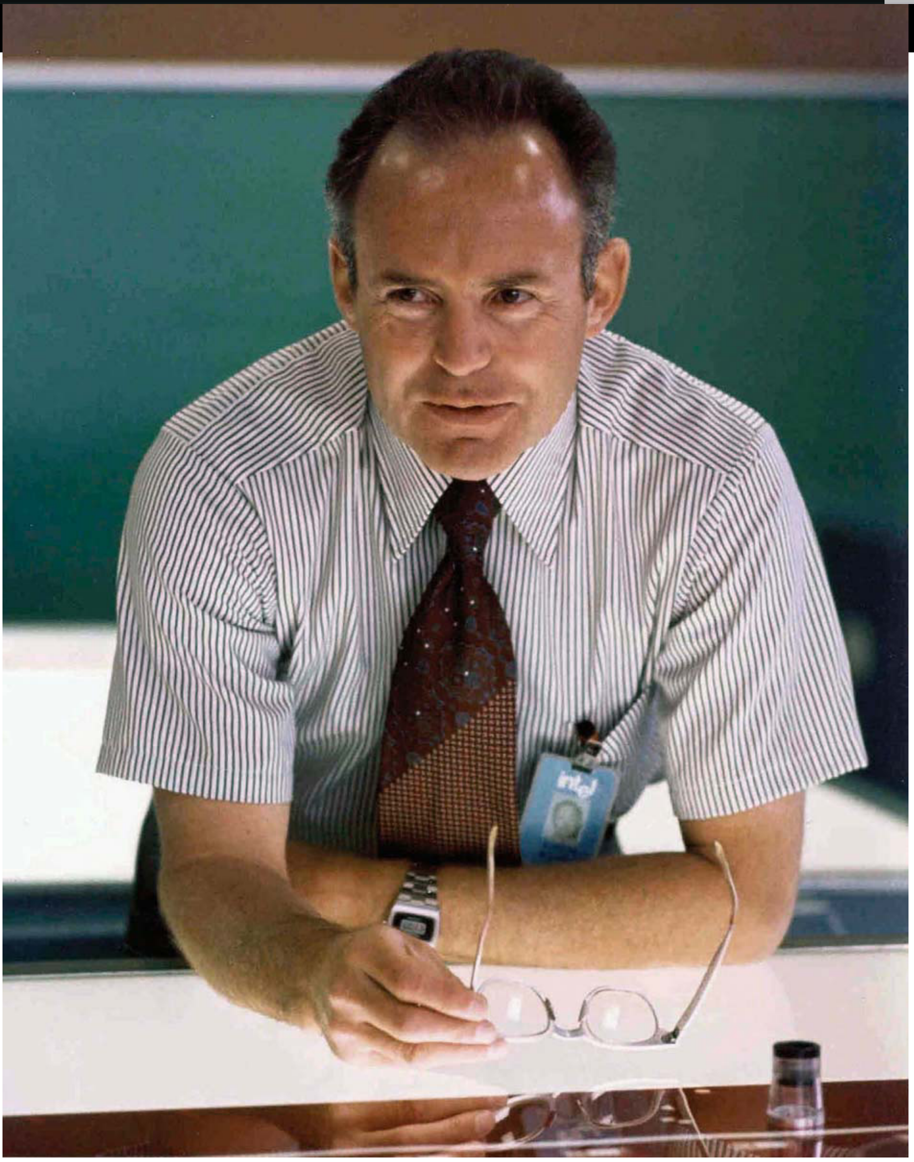

Rachel Courtland: It's been 50 years since the article came out.

Gordon Moore: It's hard to believe. I never would have anticipated anyone remembering it this far down the road.

R.C.: Why is that?

G.M.: At the time I wrote the article, I thought I was just showing a local trend. The integrated circuit was changing the economy of the whole [electronics] industry, and this was not yet generally recognized. So I wrote the article to try to get the point across—this is the way the industry is going to get things really cheap.

R.C.: At that point, the integrated circuit was still fairly new.

G.M.: The integrated circuit had been around a few years. The first few had hit the market with as many as about 30 components on the chip—the transistors, resistors, and so forth. I looked back to the beginning of the technology I considered fundamental—the planar transistor—and noticed that the [number of components] had about doubled every year. And I just did a wild extrapolation saying it's going to continue to double every year for the next 10 years.

And it proved to be amazingly correct. I had a colleague who saw that and dubbed this Moore's Law. It's been applied to far more than just semiconductors. Sort of anything that changes exponentially these days is called Moore's Law. I'm happy to take credit for all of it.

R.C.: You spoke to a colleague of mine after winning the 2008 IEEE Medal of Honor, and I believe you told her that you didn't want Moore's Law to be your legacy. You'd moved on to other things.

G.M.: Well, I couldn't even utter the term "Moore's Law" for a long time. It just didn't seem appropriate. But as it became something that almost drove the semiconductor industry rather than just recording its progress, I became more relaxed about the term.

R.C.: How long did it take to come to terms with having a law named after you?

G.M.: Oh, 20 years or so. It really took a long time. But [now] it is well established. A while back I googled "Moore's Law" and "Murphy's Law" and discovered that Moore's Law had more references than Murphy.

R.C.: Did that feel like a coup of some sort?

G.M.: I think so. It's about as profound a law [as Murphy's Law] too.

R.C.: Coming from a science background, when I think of laws, I think of ironclad, mathematically grounded laws of nature. And Moore's Law is…

G.M.: It's not a law in any real respect. It was an observation and a projection.

R.C.: Technological improvements are nothing new, but the rapid progress that's been made under Moore's Law has been pretty special. Is there something fundamentally different about the nature of silicon?

CONTINUED ON PAGE 56

MOORE'S BIG MOVE: Moore [above] wrote his seminal 1965 paper while working at Fairchild Semiconductor. Just three years later, he and colleague Robert Noyce left to cofound Intel.

INTEL

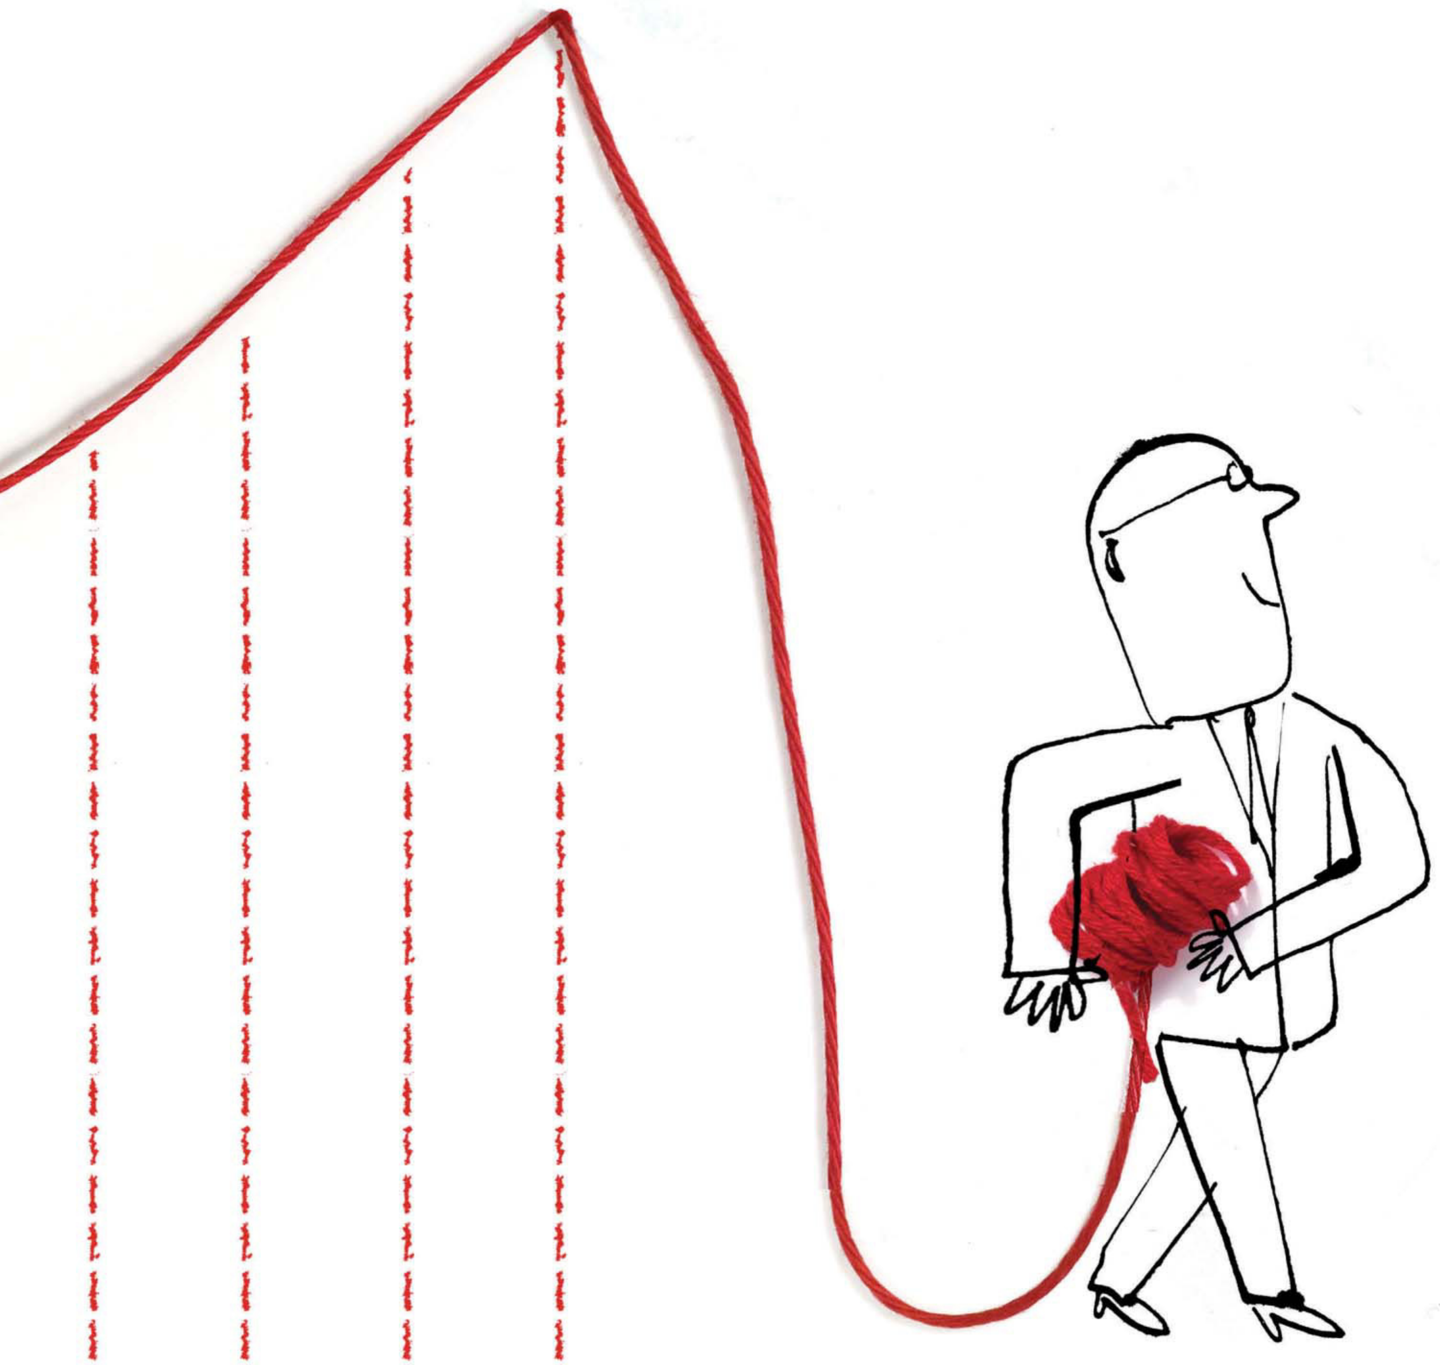

MOORE'S LAW 50 YEARS

MOORE'S LAW is DYING (and That Could Be Good)

Open-source hardware will benefit as transistor shrinkage slows

By **ANDREW "BUNNIE" HUANG**

COMPANIES THAT PRODUCE open-source hardware are few and far between. At least, they are if you define them in the usual way: an enterprise that provides documentation and permission sufficient for others to re-create, modify, improve, and even make their own versions of the devices it sells. And although open hardware has made strides in recent years–including an increasing number of companies adhering to these practices along with the establishment of the Open Source Hardware Association–it remains a niche industry. ● You might guess the reason to be simple–such companies must be set up and run by idealists who lack any hardheaded business sense. Not true! What's held back the open-source hardware movement is not a lack of business acumen; it's the rapid evolution of electronic technology. ● The reasons for this are subtle, but as I will explain below, swift advances in electronic technology inherently favor large "closed" businesses at the expense of small teams or individual innovators, who benefit most by working with open systems. At least that's the way things have been. But there are changes coming that I expect will tilt the balance the other way. ● The relevant shifts are fundamentally tied to the pace of microelectronic miniaturization, which for decades has been well described by

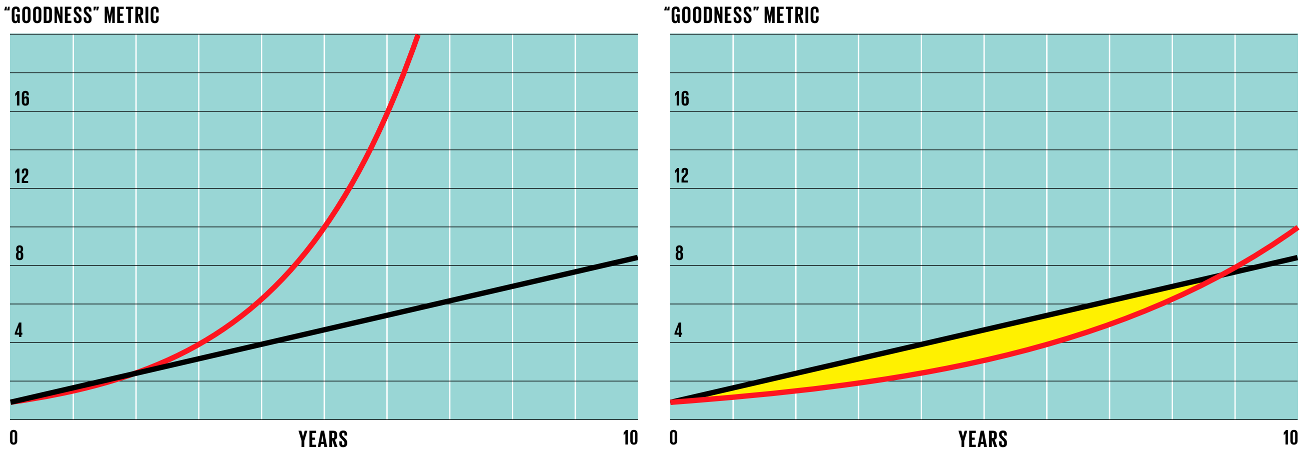

SLOW IT DOWN: The gains a small company could reasonably make [black line, a 75-percent-per-year, noncompounding rate] can't compete with Moore's Law [red line] if the doubling time is just18 months [left]. But a doubling time of 36 months [right] provides ample opportunity [yellow].

Gordon Moore's eponymous law. So it seems fitting that on the 50th anniversary of his seminal publication on the topic, we pause to consider how the impending end of Moore's Law–or at least its slowing–could, in fact, be tremendously beneficial to many manufacturers and consumers. But before I present this contrarian forecast, I should review for younger readers a little bit about how we got where we are.

In the beginning, there was the vacuum tube. And just about all the hardware based on it was what today we would call open. Early consumer electronic products, such as radios and television sets, often shipped with user manuals that contained full schematics, a list of replacement parts, and detailed instructions for servicing. Very little changed when the transistor was first introduced.

Indeed, into the 1980s, computers often came with schematic diagrams of their circuit boards. The Apple II, for example, shipped with a reference manual that included a full schematic of the computer's main board, an artifact that strongly influenced me to pursue a career in electronic design.

Modern user manuals lack such depth. The most complex diagram in a recent Mac Pro user manual instructs the purchaser on how to *sit* at the computer: "lower back supported," "thighs tilted slightly," "shoulders relaxed," and so forth.

What happened? Did electronics just get too hard to service or improve?

Not really. In fact, improving electronic products has become too easy–particularly for the system integrators. For decades, they have, in essence, been able to sit and wait for the ICs populating their circuit boards to get better rather than put in the hard work needed to hone their existing product designs. For example, throughout the 1990s and into the new millennium, programmers were encouraged to abandon hand-optimized assembly language in favor of cramming in more features using ever higher-level languages. Snappy performance, if it wasn't there on release day, would come soon enough with the next generation of CPUs.

You can see this effect clearly if you graph the "goodness" of electronic gadgets over the years. Pick virtually any metric–performance, feature set, whatever–and arrange your graph so that the plotted parameter doubles every 18 months following Moore's Law. But do that on a chart with a linear vertical axis. Most diagrams depicting Moore's Law use a logarithmic vertical scale, which flattens the curve's sharp upward trend into a much more innocuous-looking straight line.

Now compare the sharply climbing Moore's Law curve with what might come from a design team at a small-scale manufacturer working diligently to improve its latest product by polishing the firmware or tweaking the way memory and other components are connected. Assume that such a team can achieve a constant but respectable rate of progress. Let's say the first year of such work improves the product by 75 percent over its initial performance, and that improvements continue to accrue by this same amount year after year. How do the results compare to what the design team gains just by sitting and waiting for Moore's Law to do its magic?

Not well. Indeed, if you plot both lines, linear improvement versus exponential Moore's Law curve, you'll see only a tiny sliver of opportunity for diligent work on the part of a small-scale manufacturer to make for a better gizmo. And that opportunity expires within two years of product launch, because at that point higher-performance chips become available, requiring the company to create an entirely new product that uses them lest it lose out to the competition.

The two curves I've just described highlight the central challenge that small innovators have faced during the past few decades. Almost universally, it has been more rewarding for them to sit and wait rather than innovate. In particular, if it takes two years to conceive and implement changes to a design that doubles its performance, you and your customers are better served by doing nothing and upgrading to the latest available ICs two years down the road. It's been a Sisyphean exercise for the engineers at many small businesses to try to race against Moore's Law.

Indeed, the exponential growth of Moore's Law works against them while it favors large businesses, which have the resources to work on three or four generations of products simultaneously. And even for large companies, this is hard.

The race manufacturers are in to keep up with the pace of Moore's Law has several pernicious effects. The most relevant one here is that the small edge (of perhaps a few months' time) created by keeping a design proprietary and forcing the competition to reverse engineer your products constitutes a significant advantage, one that manufacturers are understandably hesitant to squander. Their reluctance to give up even the slightest edge on the competition helps explain why they rarely share schematics, code, or other technical details about their products.

Thankfully, Moore's Law is slowing, and the dynamics I've just described are changing.

Before I say how, though, let me defend my assertion that Moore's Law is slowing. Lots of analysts and commentators have warned recently that the era of exponential gains in microelectronics is coming to an end. But I do not really need to hang my argument on their forecasts. The reduction in size of electronic components, transistors in particular, has indisputably brought with it an increase in leakage currents and waste heat, which in turn has slowed the steady progression in digital clock speeds in recent years. Consider, for example, the clock speeds of the various Intel CPUs at the time of their introduction. After making dramatic gains, those speeds essentially stopped increasing about a decade ago.

Since then, CPU makers have been using multicore technology to boost performance, despite the difficulty of implementing such a strategy. [See "The Trouble With Multicore," by David Patterson, *IEEE Spectrum*, July 2010.] But engineers didn't have much of a choice: Key physical limits prevented clock speeds from getting any faster, so the only way to use the increasing number of transistors Moore's Law made available was to create more cores.

Transistor density continues to increase exponentially, as Moore predicted, but the rate is decelerating. In 1990, transistor counts were doubling every 18 months; today, that happens every 24 months or more. Soon, transistor density improvements will slow to a pace of 36 months per generation, and eventually they will reach an effective standstill.

When will that be? The short answer is, nobody really knows. But one study suggests that it will stop at an effective gate length of about 5 nanometers sometime around 2020 or 2030. Five nanometers is about the space between 10 silicon atoms, so even if this guess is wrong, it can't be wrong by much.

The implications are profound (at least for the kind of person who reads this magazine). Someday in the foreseeable future, you will not be able to buy a better computer next year. The flash drive you purchase next will cost the same and store the same number of bits as the one you're replacing. And you'll have to stop looking forward to your next phone being more powerful and doing more amazing things than your last one.

Rather than expecting that within a few years of its purchase you'll just throw out whatever electronic gizmo you buy, you'll be anticipating keeping it for a good long while.

Under such a regime, you'll probably want to purchase things that are more nicely made to begin with. The idea of an "heirloom laptop" may sound preposterous today, but someday we may perceive our computers as cherished and useful looms to hand down to our children, much as some people today regard wristwatches or antique furniture.

This slowing of Moore's Law portends a bright future for many small businesses–and likewise | CONTINUED ON PAGE 58

Top 10 Tech Cars 2015

CHEAP FUEL CATCHES CAR DESIGNERS OFF GUARD

BY LAWRENCE ULRICH

FOR THE AUTO INDUSTRY, predictions have been as reliable as a moth-eaten Yugo: Global oil prices are at a five-year low, sales of pickup trucks and SUVs are booming, and purchases of gas-electric hybrids have fallen. Yet automakers still face a monumental challenge to boost fleetwide fuel economy: In the United States, they must reach 4.3 liters per 100 kilometers (54.5 miles per gallon) by 2025, from approximately 7.6 L/100 km today. In the European Union, meanwhile, automakers face other headwinds, flowing from a requirement to cut carbon emissions and fuel consumption, even as sales remain mired in a vicious slump. • This year's Top 10 Tech Cars reflects on the effects of those competing demands. Consider the Tesla Model S: Only three years ago this electric sedan dazzled pundits, who predicted that Tesla would revolutionize automobiles. But any such revolution depended on a lower-price follow-up—the Model X crossover—which has been delayed *again.* And while

WORK IN PROGRESS: Workers at BMW's plant in Leipzig, Germany, assemble an all-electric i8.

long-range electric vehicles and plug-in hybrids remain very much in play, the world is still waiting for one of them to go beyond a mere plaything of the wealthy to become the Model T of its age.

General Motors has paid homage to that ideal with the Chevy Bolt, an all-electric hatchback with a 300-km range. GM aims to sell it for roughly US $30,000—though the Bolt remains a lightly fleshed-out concept ahead of its anticipated 2017 arrival. Another drawing-board conceit was the Mercedes-Benz F 015, a car so automated that its driver would be able to swivel to socialize with backseat passengers—another car we'll dissect if and when it reaches production.

One showroom car that did make our list, the Volkswagen Golf, shows that the technical contest among fossil-fueled vehicles, hybrids, and pure electrics is far from over. Among the world's best-selling cars, this hatchback for the common man is being offered in four propulsive flavors: turbocharged gasoline and diesel models, the high-performance gasoline GTI, and an electric e-Golf. Even hydrogen cars, which had fallen back to pipe-dream status a couple of years ago, are on the upswing again, with the advent of the Toyota Mirai.

Ford chief executive Mark Fields is one of a group of industry leaders convinced that a return to pricey gasoline is only a matter of time. With Americans buying more cars than during any year since 2006, the whiz-kid models of 2015 face a tough challenge: to provide enough thrills to keep the party going, even as they face, someday, a last call on fossil fuels. ■

BMW i8

Rich in carbon, not in carbon emissions

As a splashy billboard for a plug-in future, the BMW i8 is unmatched. Environmentalists who shake their fists at guzzling Lamborghinis fall in love. Speed freaks who scoffed at electrified sports cars change their tune in 4.4 seconds, the time it takes the i8 to howl to 100 kilometers per hour (62 miles per hour). Design mavens thrill to the BMW's thermoplastic-clad, wind-cheating teardrop shape.

The i8 isn't the first carbon-fiber-intensive car. But along with its cousin, the i3, the i8 is the clearest indication yet that the future belongs to cars made from material that's ultrastrong yet weighs half as much as steel.

Whereas the US $42,300 i3 is a toaster-shaped urban runabout, the i8 is a two-plus-two sports car that can sip fuel, cruise on electricity for short hops, or scorch the asphalt anytime you like. Channeling the spirit of Messrs. Jekyll and Hyde, I took test drives in New York and the Midwest, including thrilling laps at GingerMan Raceway, in western Michigan.

The carbon-fiber passenger cell of an i8, its technical heart, is molded in a patented process that takes minutes rather than the hundreds of hours needed to handcraft the material in recent supercars that cost $1 million. A 7.1-kilowatt-hour lithium-ion battery pack nestles below the floor. A 96-kW (129-horsepower) electric motor drives the front wheels through a full automatic two-speed transmission. The rear wheels get their urge from a 1.5-liter three-cylinder gasoline engine borrowed from BMW's Mini Cooper but turbocharged to 170 kW (228 hp). Drivers manage the combined 266 kW (357 hp) and 569 newton meters (420 foot-pounds) of torque through a superlative six-speed, paddle-shifted automatic transmission.

Toggle up the console's eDrive mode and the BMW will cover up to 35 km entirely on front-drive electricity, albeit at a leisurely 9.5-second pace to 100 km/h and a 120 km/h top speed. But cleverly, those 35 km don't have to tick off sequentially the instant you unplug from a wall charger. The all-wheel-drive Sport mode replenishes the battery on the fly, adding about 1 km of pure electric-vehicle range for every 17 km driven. That feature enables drivers to preserve or add electric miles for urban use. A few cities, notably London, charge hefty entry fees

PREVIOUS PAGES AND THIS PAGE: BMW; OPPOSITE PAGE: ROLLS-ROYCE

Price: US $136,650 **Power train:** 96-kW (129-hp) AC electric motor with 1.5-L 170-kW (228-hp) three-cylinder gasoline engine **Overall fuel economy:** 8.4 L/100 km (28 mpg) on gasoline; electric equivalent of 3.1 L/100 km (76 mpge)

for internal-combustion vehicles—and may eventually ban them outright.

Traveling up New York's Hudson River, I spent some hours in the BMW's Comfort mode, which maximizes electric efficiency but will seamlessly blend in gasoline power when you goose the throttle. Here, gentle highway driving produced a knockout 6.2 liters per 100 km (38 miles per gallon), roughly twice the efficiency of a similarly sized luxury GT, the Aston Martin DB9. Credit in part the BMW's svelte 1,567-kilogram (3,455-pound) curb weight, some 220 fewer kilograms than the aluminum-chassis Aston and 540 fewer than the battery-stuffed Tesla Model S sedan. Overall, the BMW showed me 7.6 L/100 km (31 mpg) over several hard-driving days. That's 3 mpg better than the car's U.S. federal estimate and about 50 percent better than that of the latest Corvette Z06 or Porsche 911.

Solid-citizen duties complete, it was time for some fun. Tromp the i8's throttle and you're whisked into motion like protons in a collider. The instant signature shove of electric torque helps. The miniature gas engine chimes in with a backup duet voice. It's canned but uncanny: a synthesized simulation of engine sound piped through door speakers to complement the actual exhaust note. You'd never know that a three-cylinder engine was at work; instead, the BMW's chesty rumble sounds like a Porsche flat-six mating with a flying saucer.

Admiring that performance is easy, but admiring the electric motor is hard: As with the iPhone, another lowercase techno marvel, the i8's hood is sealed, as it's meant to be opened only by certified technicians. Other drawbacks include a toddler-size backseat and stingy storage.

But blasting down the road, bypassing the gas stations, you really won't care. The i8 might have ended up as another Fisker Karma, a poorly integrated stew of hybrid technology. Instead, the BMW is real, and it's here. Just in time for the future. ■

Rolls-Royce Wraith

Satellite to Earth: Downshift a bit

For 111 years, the Rolls-Royce experience has been all about wafting, often in chauffeur-driven comfort. Now the stunning Wraith coupe adds a computerized wingman, using GPS data to automatically select the perfect gear for the road ahead. The Wraith's Satellite-Aided Transmission (SAT) is a production-car first for the BMW-owned Rolls-Royce.

Rolls-Royce engineer Phil Harnett, formerly with BMW's Formula One team, saw a colleague working on SAT as a predevelopment project. "I saw how perfect it was for Rolls-Royces—dynamic yet effortless—and brought it into the Wraith," Harnett says.

This yachtlike fastback coupe marries a 465-kilowatt (624-horsepower) V-12 to the silken, eight-speed automatic transmission produced by the ZF Group, based in Germany. But the Wraith takes the ZF unit to the next level, adapting to a driver's individual style and applying GPS route and location data to anticipate and change gears. The engine's electronic control unit adds an algorithm that responds to upcoming road patterns.

Rolls-Royce spokesman Gerry Spahn says the system works especially well on curvy terrain, off-ramps, or roundabouts. As the car heads into a series of curves, for example, the SAT will automatically hold a lower gear to prevent midcorner upshifts and to avoid upsetting the car's occupants with needless gear changes. In testing on public roads, the system has reduced the number of automated shifts by up to 30 percent, Spahn says. Besides reducing wear and tear on the transmission, the system also improves responsiveness and fuel economy. Because it requires no additional components beyond existing sensors, controllers, and the navigation system, no weight is added to the vehicle.

The SAT isn't mapped for altitude and topography, so it can't anticipate hills and change gears accordingly. But that kind of functionality could be baked into a future version.

"Customers have asked about altitude," Spahn says. "It's not technically impossible; it just hasn't been applied." ■

Price: US $298,225 **Power train:** 465-kW (624-hp) 6.6-L twin-turbo V-12 **Overall fuel economy:** 15.7 L/100 km (15 mpg)

Formula E Spark-Renault SRT_01E

A green machine on the racing scene

Power train: Rechargeable Energy Storage System (RESS) with maximum 200 kW (268 hp)

Better batteries, better electric vehicles, and a pox on fossil fuel: That's the ethos of Formula E, a revolutionary international race series and test bed for potential EV technology for the street. This year's 10-race series began in Beijing and will conclude in London in June. Backers are convinced that validating electric performance in harsh race conditions can spin off benefits—greater energy density, shorter charging times, new battery chemistries—to street EVs.

Governed by Formula One's sanctioning body, the single-spec Formula E car is designed and built by a who's who of racing heavyweights. McLaren Electronic Systems created the AC "motor-generator unit," with a maximum 200 kilowatts (268 horsepower) and a roughly 3.0-second sprint from 0 to 100 kilometers per hour (62 miles per hour). A driver-operated "push to pass" feature delivers temporary bursts of an additional 67 kW (90 hp).

Italy's Dallara built the carbon-fiber-and-aluminum honeycomb chassis. Williams Advanced Engineering, which helped introduce KERS (kinetic energy recovery systems) to F1, designed in just six months what it calls one of the most sophisticated traction batteries in history. Limited to 200 kilograms (441 pounds), the liquid-cooled, lithium-ion Rechargeable Energy Storage System passed the same ultrarigorous fire and impact standards of F1 components.

Formula E rules say the battery can deliver no more than 28 kilowatt-hours of juice per race, or less energy than in 1 gallon of gasoline. Speeds are limited to 225 km/h to conserve that energy, but there's no recharging during the event: Instead, each driver needs two cars to finish the 1-hour race, "refueling" by jumping from a depleted racer into a fully charged car in the pits.

Prior to January's Buenos Aires event, Marco Andretti, the IndyCar driver and the third generation of the famous racing clan, told *IEEE Spectrum* that he was eager for his first electron-aided competition: "I'm going in cold turkey, which I always like. I'll drive anything with wheels on it." Earlier races in Beijing and Malaysia were exciting and closely contested, with wrecks and even controversy: As with the much-reviled, four-cylinder F1 cars, fans grouse that the eerily silent Spark-Renaults take the ear-splitting joy out of racing. Backers considered adding artificial, synthesized engine sounds, but rejected that idea because it would add weight.

Andretti counts himself "a big critic" of wimpy-whiny F1 cars, but he has decided to cut Formula E some slack. "I'm all about the sound as well. But I think fans have to get over that hump and look forward. Being more green is the key here, and it's unbelievable what they've been able to do." ■

FIA

Volkswagen Golf

Choose your power: gas, diesel, or electric

The Golf recently marked its 40th anniversary as an affordable overachiever. But the seventh-generation VW has outdone itself, winning worldwide accolades for design, performance, and technology. It all starts with the Modularer Querbaukasten, or MQB, which translates to "Modular Transverse Toolkit." Consider it a set of building blocks in a single platform that will underpin more than 40 corporate models, from tiny Czech-made Skodas to Audi luxury sedans.

Mercedes-Benz C-Class

Luxury-car brains at mass-market prices

Within any automaker there is a slow but steady trickle down of technology from high-end to middle and entry class. But for Mercedes's newly redesigned C-class, that trickle was more like a torrent. Appropriately, the Mercedes demonstrated its social-climbing skills on the Côte d'Azur, in France. The test convoy included the all-wheel-drive C400 4Matic, which surges from 0 to 100 kilometers per hour (62 miles per hour) in 4.7 seconds, spurred by a wicked Biturbo V-6. The engine's 245 kilowatts translate to 329 horses—81 more than the departing 3.5-liter V-6, despite that model's extra half-liter of displacement.

Price: US $41,325 **Power train:** 245-kW (329-hp) 3.0-L Biturbo V-6 **Overall fuel economy:** 8.4 L/100 km (28 mpg)

But the more important part of what's trickled down—from the US $95,000 to $200,000 S-Class to this $41K cousin—involves not speed but smarts. The Intelligent Drive suite, priced at the same $2,800 as in S- and E-Class models, lets the car manage its own speed and brakes and even steer itself along gentle highway curves. Radar and a windshield-mounted stereoscopic camera provide a 360-degree enclave of sensory safety. The Mercedes can react to other vehicles, pedestrians, and bicyclists, even the traffic crossing at an intersection.

At a test facility in Salon-de-Provence, I was invited to try to smack a pedestrian—actually a mannequin dressed nattily in Mercedes attire—and a slow-rolling trailer that mimicked a car ahead. No dice: The Mercedes performed a partial braking, then a full-on stop to avoid a collision, even when I kept my foot on the gas.

Active Blind Spot Assist doesn't simply warn drivers about looming vehicles but will apply individual brakes to pop you safely back into a lane should you veer toward danger. Attention Assist monitors multiple parameters, including twitchy movements of the steering wheel, to recognize drowsy drivers and urge them to pull over for a break. Inside the surprisingly lush cabin, a new touch-pad console controller allows drivers or passengers to navigate contact lists, destinations, and the like by drawing letters or numbers or by using the swipes and gestures familiar from smartphones.

The C63 AMG version will debut just in time for spring in the northern hemisphere. This BMW M3 killer is armed with the same ripping 4.0-L Biturbo V-8 as the new AMG GT sports car, with up to 375 kW (503 hp). And there's more: A plug-in hybrid C-Class goes on sale in the United States this fall, followed by a C-Class diesel in 2016. ■

TOP: MERCEDES-BENZ; BOTTOM: VOLKSWAGEN

Price: US $18,815 **Power train:** 112-kW (150-hp) 2.0-L four-cylinder diesel **Overall fuel economy:** 6.5 L/100 km (36 mpg)

The MQB's standardized common core allows every imaginable power train—gasoline, diesel, electric, hybrid—to plug and play smoothly in the same front location. The electric version, for example, doesn't detract a whit from passenger or cargo space. There's even a compressed-natural-gas Golf.

Start with the turbocharged TDI model, whose clean-diesel technology is familiar in Europe but still rather novel in the States. The TDI clocks in at 5.2 liters per 100 kilometers (45 miles per gallon) on the highway in its federal fuel-economy estimate. But I saw a remarkable 3.9 L/100 km (60 mpg) over a 110-km (68-mile) highway run, eclipsing the Toyota Prius hybrid and setting my own personal record in a U.S.-market car. With its 2.0-liter, 112-kilowatt (150-horsepower) diesel, the Golf sipped fuel at well under 5.6 L/100 km during our test drive while cruising at 110 kilometers per hour.

Gasoline Golf models include the brilliant 157-kW (210-hp) GTI. With a 2.0-L direct-injected four that's a near twin of the Audi A4's engine, the GTI hustles to 97 km/h in 5.8 seconds and steers like a junior Porsche—but for just US $26,000 to start. GTI and TDI models get a six-speed, dual-clutch direct-shift gearbox, a seven-speed version of which does duty in the roughly $250,000 Lamborghini Huracán.

For zero-emission fans, the e-Golf is the handling standout among affordable plug-ins like the Nissan Leaf and Chevrolet Volt. Ten percent more aerodynamic than the standard Golf, the $35,445 electric VW gets its go from an 85-kW synchronous electric motor that maxes out at 12,000 rpm, driving the front wheels through a single-speed gearbox. The e-Golf does take about 10 seconds to reach 100 km/h, because it's weighed down by a 318-kilogram, 24.2-kilowatt-hour lithium-ion battery. But VW claims the e-Golf scoots from a standstill to 30 km/h in just 4.2 seconds—faster than the GTI, thanks to a heady 270 newton meters (199 foot-pounds) of torque. ■

Chevrolet Corvette Z06

The best bang for your buck

Finally, Chevy has managed to put it all together. The supercharged, 485-kilowatt (650-horsepower) version of the new Corvette Z06 combines a race-bred aluminum space frame, carbon-fiber bodywork, an adaptive magnetic suspension, superb aerodynamics, and a brilliant limited-slip differential. (It's called "limited slip" because it reduces power to the wheel that's not getting traction, thanks to an algorithm based on real-time parameters.)

Coupe or convertible, it's simply one of the fastest, best-performing production cars in history, beating track times of virtually any rival that sells for less than US $500,000, at a price that reads like a typo: Just $78,995 to start.

On a test drive near Death Valley, in California's Mojave Desert, the Z06 coupe peaks at 185 kilometers per hour (115 miles per hour) in third gear and 233 km/h (145 mph) in fourth. I stop there, before the 314-km/h (195-mph) top speed becomes too tempting. Besides, a test track beckons at Spring Mountain Motorsports Ranch, in Nevada.

The Z06 hurtles to 97 km/h (60 mph) in 2.95 seconds, then stops in 30 meters, thanks to mighty Brembo carbon-ceramic brakes. It pulls an insane 1.2 *g*'s of lateral cornering force, thanks in part to its Michelin-designed Super Sport Cup tires, which use a compound developed exclusively for the Z06. Those cornering and braking numbers may well set production-car records. Yes, the Ferrari 458 Speciale is one such production car.

Magnetic shocks and the Driver Mode Selector deliver a pleasingly wide envelope of driver settings, allowing the Z06 to cruise in all-day comfort or drive straight onto the track. With its 20,000-rpm supercharger nestled between its cylinder banks, the LT4 engine is only a couple of centimeters taller than the standard Stingray's 6.2-liter engine, but it generates nearly 37 percent

Price: US $78,995 **Power train:** 485-kW (650-hp) 6.2-L eight-cylinder **Overall fuel economy:** 13.0 L/100 km (18 mpg)

more power and a fearsome 881 newton meters (650 foot-pounds) of torque. Its V-8 even operates in four-cylinder mode to save fuel. Buyers choose a seven-speed manual transmission or a new eight-speed, paddle-shifted automatic that executes commands 160 times per second.

On Spring Mountain's sinuous track, I feel serenely in control. It's a testament to the work of the seven-time-LeMans-winning Corvette team that transferred technology to this street Corvette, which in turn informed the development of the new C7.R racer.

The coup de grâce is the Performance Data Recorder, which captures high-def video and in-car audio of track laps or Dairy Queen runs and overlays the video with telemetry data. Drivers can play back video in the car, share it online, or apply track-analysis software to improve their skills. If you've got $80,000 to $95,000 to spend on a Z06, it sure beats PlayStation. ■

GENERAL MOTORS

Ford F-Series

An aluminum truck for the masses

Aluminum on a six-figure Aston Martin may be no big deal, but when it goes on the Ford F-Series pickup—America's best-selling vehicle for 33 consecutive years—the world takes notice. The bodacious 2015 model must be considered one of the biggest product gambles in Ford's history, being the product of a multibillion-dollar investment in Ford's assembly lines, tooling, and engineering.

The F-150 has roughly 450 kilograms (1,000 pounds) of aluminum in its body, suspension, and other components, and it saves nearly 320 kg versus the outgoing model—a 15 percent loss. The switch from tried-and-true welded steel for an aluminum cab and cargo box alone saves 204 kg. Structural adhesives and about 4,000 rivets replace the roughly 7,000 spot welds of the old model, boosting rigidity.

Inside the engine compartment is another advance. It's a pair of optional, twin-turbocharged Eco-

Lamborghini Huracán LP 610-4

Fast, but it won't run away from you

Lamborghinis have always been fast. But they've often been infuriatingly balky and volatile. Enter the Huracán, which has finally managed to harness technology to become Lamborghini's best, most civilized supercar yet. The latest Lambo to be named after a famous Spanish fighting bull, the Huracán showed me its snorting stuff at Spain's Ascari racing circuit.

Price: US $240,245 **Power train:** 455-kW (610-hp) 5.2-L V-10 **Overall fuel economy:** 14.7 L/100 km (16 mpg)

The successor to the Gallardo shares its all-wheel-drive chassis with the upcoming Audi R8 supercar (both brands benefit from technology developed within the expanding VW empire). The chassis is basically an erector set of bonded-and-riveted aluminum interspersed with carbon fiber that's 10 percent lighter—yet 50 percent more rigid—than the Gallardo's. An alluring, modern wedge of a body is tasteful and functional, developing 50 percent more aerodynamic downforce than its predecessor.

Inside, the Lamborghini is the first car to feature the Audi-engineered "virtual cockpit," in which drivers can reconfigure a 12.3-inch LCD touch screen with all manner of data, such as a virtual tachometer or a full-screen navigation display. Meanwhile, on terra firma, power flows to all four wheels from a brand-new, 455-kilowatt (610-horsepower) V-10, good for a concussive 3.2-second run to 100 kilometers per hour (62 miles per hour). The new, fighter-jet-inspired Inertial Platform plops a trio of gyroscopes and accelerometers at the car's center of gravity to analyze physical forces and optimize systems for performance.

What really tames the Huracán are the whip-crack shifts of the Lamborghini Doppia Frizione. A dual-clutch, seven-speed automated gearbox is a first for a Lamborghini and a huge improvement over the clunky, lurching unit of old. Even as it scorched the Andalusian countryside, the Huracán pampered me and a passenger, adjusting its performance personality via the Anime switch on the steering wheel. Push the car beyond its lofty limits and you can feel the AWD system divvying power between front and rear wheels to keep the car on its swift and sure-footed course. More than any previous Lamborghini, the Huracán is an Italian lothario you might actually live with on a daily basis, without feeling cheap or getting your heart broken. ■

TOP: LAMBORGHINI; BOTTOM: FORD

Price: US $26,220 **Power train:** 242-kW (325-hp) 2.7-L twin-turbo V-6 **Overall fuel economy:** 10.7 L/100 km (22 mpg)

Boost V-6 engines. The smaller of the two wrings 242 kilowatts (325 horsepower) and 508 newton meters (375 foot-pounds) of torque from just 2.7 liters of displacement—the smallest engine in the truck's 67-year history.

Remember when pickups got 14 miles per gallon (16.8 liters per 100 kilometers) on the highway—with a strong tail wind? The 2.7-L F-Series is EPA rated at 26 mpg (9.0 L/100 km) on the highway, a record for a full-size gasoline pickup. The larger V-6 EcoBoost, at 272 kW (365 hp) and 3.5 L, cranks the Ford to 97 kilometers per hour (60 miles per hour) in an insane 5.6 seconds, with an equally remarkable absence of wind and road noise.

LED headlamps are another first for a pickup, with five times the life of ordinary incandescents. A bird's-eye camera system provides 360-degree views around this tall, burly truck, and the Trailer Hitch Assist uses a camera and digital guides to help drivers back up toward a trailer. There's even an optional active parking system that lets the Ford parallel park itself. The driving experience is also impressive: The mostly aluminum F-Series does feel lighter and more agile than its steel rivals. Other full-size Ford trucks, including the Expedition and Lincoln Navigator SUVs, will switch to aluminum bodies over the next three years.

Consider this F-Series the first giant leap on a scramble to 2025, when U.S.-market trucks like the F-150 must achieve 7.8 L/100 km (30 mpg) in overall fuel economy. ■

Pagani Huayra

A millionaire's gull-winged dream

PAGANI

Horacio Pagani is a throwback to the era of Henry J. Kaiser and Preston Tucker, when a visionary could start a car company on little more than brilliance and good connections. And like Kaiser and Tucker, the Argentina-born former Lamborghini engineer builds cars like no one else. As an object lesson, consider the Huayra (WHY-ra), a 537-kilowatt (720-horsepower) gull-winged, carbon-fiber goddess that makes many supercars seem like mass-market commoners.

Price: US $1.5–$1.8 million **Power train:** 537-kW (720-hp) 6.0-L twin-turbo V-12 **Overall fuel economy:** 19.6 L/100 km (12 mpg)

The Huayra has Bugatti-like exclusivity and value: Barely 40 units a year are being built to order at a price that can top US $1.8 million. Nevertheless, and happily enough, Pagani let me test one at the Monticello Motor Club in upstate New York. Note to self: Do not run this Fabergé-rare beauty into the nearest wall.

With a naked, clear-coated carbon-fiber body and monocoque chassis, the Huayra weighs 1,400 kilograms (3,100 pounds). Uniquely, Pagani also weaves titanium into the carbon fiber to boost rigidity and spread kinetic energy in a crash. (And how's this for obsessive: Every titanium bolt in the car is etched with the Pagani name.) To achieve street legality in the United States, the Huayra had to pass 47 crash tests using a pair of carbon-fiber monocoques that Pagani says remained usable despite repeated destructive tests. And the engine? It's built by Mercedes-AMG to Pagani's specifications, including a dizzying 1,000 newton meters (738 foot-pounds) of torque, and it's mated to a seven-speed single-clutch automated manual transmission.

On the day of my test drive, I arrive at Monticello to find the two-tone Huayra lolling in the pit lane, drawing a crowd and looking as though it's going 300 kilometers per hour while standing still. Exterior mirrors curl from slender stalks, like the antennae of some superpowered insect. Stepping into the car is like going back in time: The inside is a gleaming metal-and-leather cocoon that feels like the cockpit of a fantasy Howard Hughes aircraft. On the track, the car has a visceral, almost retromechanical feel. But there's nothing old-fashioned about the acceleration, which—inspired by Pagani's heroes, including Leonardo da Vinci—mimics the feel of a jet hurtling down the runway. Minus the liftoff, of course: Active aerodynamic flaps at all four corners rise and fall to mimic the aileron of a plane, keeping the body flat even as *g*-forces threaten to peel the skin off my face. It sounds like a gasping beast, its turbos huffing and chuffing with every mash of the throttle. The car rockets to 100 km/h (62 miles per hour) in about 3.2 seconds and keeps cooking to about 370 km/h (230 mph).

Standing in the pits at Monticello, Pagani tells me that, contrary to conventional wisdom, it is possible to build a business by making cars that cost as much as a three-bedroom co-op in Manhattan.

"If it's your passion and your dream, you will achieve your goal," he says. ■

Toyota Mirai

Doubling down on hydrogen fuel cells

Is it a Mirai or a mirage? Skepticism is necessary with hydrogen-fuel-cell cars. Dazzled by the promise of tailpipes that emit only water vapor, a generous driving range, and fill-ups within minutes, people often overlook the downsides: hardly any places to refuel and very high costs for both fuel and vehicles. But like Jason in the *Halloween* movies, hydrogen is ba-aaack, with several automakers insisting that they're *this* close to putting hydrogen cars in showrooms.

The Mirai is the latest trial balloon, coming to California in late 2015 in small volumes at US $58,325 (or about $45,000 after tax incentives). Fuel-cell cars are basically electric cars that replace the battery with a fuel stack that generates electricity via a one-way chemical reaction fueled by hydrogen; when the hydrogen runs out, you pump more in. Pressurized hydrogen reacts with a catalyst (usually platinum); electrons are stripped from hydrogen molecules to power an electric motor; freed protons then recombine with oxygen in the stack, creating good old H_20, which helps cool the fuel stack before exiting the tailpipe.

Enter the Mirai, the fruit of 20 years of Toyota fuel-cell R&D. Mirai means "future" in Japanese, but hopefully fuel-cell cars won't all be as aggressively frumpy as this Toyota. Looks aside, the Mirai leaps forward in technology: Toyota says the fuel stack packs 3.1 kilowatts per liter, more than twice the power density at a mere 5 percent of the total cost of its 2008 prototype. That stack of 370 fuel cells weighs just 56 kilograms (123 pounds) and powers an AC electric motor with 113 kilowatts (151 horsepower) and 335 newton meters (247 foot-pounds) of torque—good for a reasonable 9.0-second run to 97 kilometers per hour (60 miles per hour).

The Mirai stuffs 5 kilograms of hydrogen into a pair of carbon-fiber-reinforced tanks at 69,000 kilopascals (10,000 pounds per square inch). That's the energy equivalent of 19 liters (5 gallons) of gasoline, enough to travel 483 kilometers (300 miles)—about three times the range of a typical electric vehicle. One neat bit is a DC outlet that allows you to tap up to 60 kilowatt-hours when the Mirai is stationary; with a full tank you can power a typical household for six days. Survivalists, take note.

In the United States, the company will offer a $499-per-month lease with $3,700 down. Toyota expects about 200 people in the United States, and 500 more worldwide, to early-adopt this hydrogen baby in 2015. Global production is pegged in the tens of thousands after 2020. To that end, Toyota and the state of California are backing construction of up to 48 hydrogen stations by the end of 2016, with a dozen more planned in U.S. northeastern states. It's a start. To spur development of the technology, Toyota announced it will allow royalty-free use of nearly 5,700 fuel-cell patents around the world. ■

Price: US $58,325 **Power train:** 113-kW (151-hp) hydrogen fuel cell with AC motor **Overall fuel economy** (est.): energy equivalent of 3.9 L/100 km (60 mpge)

TOYOTA

THE LAW THAT'S NOT A LAW

CONTINUED FROM PAGE 40

G.M.: The semiconductor technology has some unique characteristics that I don't see duplicated many other places. By making things smaller, everything gets better. The performance of devices improves; the amount of power dissipated decreases; the reliability increases as we put more stuff on a single chip. It's a marvelous deal.

I used to give talks about how other industries might have progressed. You know, had the auto industry made progress at the same rate [as silicon microelectronics], you would have gotten a million miles per gallon of fuel, had cars that could go several hundred thousand miles an hour. It'd be more expensive to park [one] downtown for the night than to buy a new Rolls-Royce. And one of the members of the audience pointed out, yeah, but it'd only be 2 inches long and a half-inch high; it wouldn't be much good for your commute.

R.C.: You've predicted the end of Moore's Law several times in the past. How long do you think it will continue?

G.M.: Well, I have never quite predicted the end of it. I've said I could never see more than the next couple of [chip] generations, and after that it looked like [we'd] hit some kind of wall. But those walls keep receding. I'm amazed at how creative the engineers have been in finding ways around what we thought were going to be pretty hard stops. Now we're getting to the point where it's more and more difficult, and some of the laws are quite fundamental. I remember we had Stephen Hawking, the famous cosmologist, in Silicon Valley one time. He gave a talk, and afterward he was asked what he saw as the limits to the integrated circuit technology.

Now this is not his field, but he came up with two things: the finite velocity of light and the atomic nature of materials. And I think he's right. We're very close to the atomic limitation now. We take advantage of all the speed we can get, but the velocity of light limits performance. These are fundamentals I don't see how we [will] ever get around. And in the next couple of generations, we're right up against them.

R.C.: What happens then, once you've reached those limits?

G.M.: Well, things change when we get to that point. No longer can we depend on making things smaller and higher density. But we'll be able to make several billion transistors on an integrated circuit at that time. And the room this allows for creativity is phenomenal. Now there are other technologies that are proposed to extend beyond what we can do with silicon. Some of the things coming out of nanotechnology may have a role to play, and materials like graphene, a single layer of carbon hexagons, are very interesting. I'm not close enough to predict that any of them is going to be successful, but they have a tough competitor. [Multiple] billion transistors on a silicon chip is hard to beat.

R.C.: So do you think the kind of progress we expect from chips will change?

G.M.: Some things will change. We won't have the rate of progress that we've had over the last few decades. I think that's inevitable with any technology; it eventually saturates out. I guess I see Moore's

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

Law dying here in the next decade or so, but that's not surprising.

R.C.: Do you think the way we consume electronics will change as Moore's Law comes to a close?

G.M.: I don't think it's likely to change much. As long as the new products offer incremental capability, I think they will replace the older ones pretty rapidly. When we run out of ideas of what to add, then people may decide they don't need a new one every year, hang on to the same piece of equipment for three, four, five years. That'll slow down the industry quite a bit. But I think it's inevitable that something like that occur.

R.C.: There're the fundamental physical limits—the atomic scale, the speed of light—and then there's also the cost associated with fabricating smaller and smaller transistors. Which do you think we'll hit first? Is it going to be the cost or the fundamental physical limits? I guess they're tied together.

G.M.: They really are, yeah. Making things smaller is increasingly expensive. Fabs to operate on the newest technology nodes are absurd. It's hard to think Intel started with [US] $3 million total capital. Now you can't buy one tool, you can't even install one tool for that much, I don't think. The machines have gotten a lot more expensive and complex. On the other hand, their productivity in terms of transistors out per unit time has increased dramatically. So we can still afford to build a few fabs to utilize the modern technology.

We've had a lot of companies decide it was too expensive to move to the next generation already. There are only a few of us in the world that are investing in state-of-the-art fab facilities today. And I don't see that number changing much over the next generation or two.

R.C.: Your initial prediction was largely based on the idea that the cost of each component on a chip was decreasing. So is that going to be the thing that decides it in the end? It's an economic law, so it'll have an economic demise?

G.M.: I think it's going to be a technological demise rather than an economic one. People will continue to squeeze cost out of the products for quite a while after they can't make them any smaller. I'm sure that's happening already.

R.C.: I told a few people that I was going to meet you today, and I asked them what questions I should ask. Some just sort of laughed and said, "Can you ask him how we get out of this mess?" Because they're all struggling with these technological issues.

G.M.: Whoo. Well, you could always retire and move to Hawaii.

R.C.: I think they're trying to get to that point.

G.M.: Yeah, well, it's the nature of the business. There aren't many easy businesses, and this certainly isn't one of them.

This interview has been edited for length and clarity.

MOORE'S LAW IS DYING

CONTINUED FROM PAGE 45

for open-hardware practices. To see why this is the case, let's revisit the comparison I made earlier between Moore's Law exponential growth and a linear (75 percent, noncompounding) rate of technology improvement. But this time, consider a new Moore's Law scenario, with a doubling time of 36 months instead of 18 months.

Were you to plot those two lines, you would find that linear improvement can go on for more than eight years before it gets overtaken by the 36-month Moore's Law curve. And there is a point at around year two or three after product launch when a company would have a substantially better product if it was incrementally optimizing its design all along rather than just waiting for Moore's Law to operate. In other words, there is a genuine market window for profiting from innovative improvements made at a pace that the engineers at small businesses can manage.

Another benefit you can expect as Moore's Law decelerates is a greater standardization of platforms. A decade ago, it would have been ridiculous to create a standard tablet or mobile-phone chassis with interchangeable components–the underlying technology was evolving too quickly. But this has now become a reasonable proposition. [See "Your Phone Will Go to Pieces," *Spectrum*, January 2015.]

The creation of stable, performance-competitive open platforms will be enabling for small businesses. Such companies can still choose to keep their designs closed, but by doing so they will be forced to create a proprietary infrastructure to support the development of their products and to build on them. Many of those companies will find that they are wasting needless amounts of time and energy working on generic hardware–time that they could have spent refining the parts of their designs that make their products special. So more and more companies will choose to work with open hardware, just as many businesses have already done with open software.

Another change I foresee is that field-programmable gate arrays (FPGAs) may find themselves performing respectably compared with their hard-wired CPU kin. One reason is simple: As transistors shrink to atomic scales, the flaws that arise during fabrication are bound to become more common. With a CPU, such flaws can easily ruin the whole device. But with an FPGA, you can arrange the physical layout of your circuits to avoid small-scale imperfections.

Another reason to expect a bright future for FPGAs comes from a consideration of the relative difficulty of configuring their circuits. That exercise invariably requires a hardware-description language such as Verilog. So it resembles software programming–although it's typically more challenging. But the only real alternative involves programming multicore CPUs to eke out better performance from the massive parallelism they offer, and that is already difficult, and it promises to get even harder as chips with more and more cores are released.

So in the future, programming multicore processors and configuring FPGAs could reach parity in terms of the effort required. Should that come to pass, many more gadgets will surely be built with FPGAs. And when open-hardware companies switch to using FPGAs instead of CPUs, they will (by the very definition of "open") share their hardware-description-language files, too. Others will then be free to reconfigure the circuitry, down to individual gates inside the FPGA. So the open-hardware movement could penetrate microelectronic design to a very deep level.

Another welcome change I see coming is a rise in repair culture as technology becomes less disposable and more permanent. Replacing worn-out computer parts five years from their purchase date won't seem so silly when the replacement part has virtually the same specifications and price as the old one. This change in the keep-or-throw-away calculus will create a demand for schematics and spare parts, which in turn will facilitate the growth of open-hardware ecosystems and businesses.

The deceleration of Moore's Law is already showing its effect in markets that are not particularly sensitive to performance. Consider, for example, the Arduino microcontroller platform.

[See "The Making of Arduino," *Spectrum*, November 2011.]

The performance demands of Arduino users (for education, physical computing, and simple embedded-control applications) have not grown appreciably, and thus the platform can be very stable. Indeed, many Arduino boards have used virtually the same hardware since 2005. This stability in turn has enabled the Arduino to grow deep roots in a thriving user base, one that enjoys open standards for hardware add-ons.

Another example is the *shanzhai* phenomenon in China. The *shanzhai*, often dismissed as "pirates," are typically small businesses that rely on blueprints shared within their community to build low-end mobile phones. The market for these phones is largely insensitive to absolute levels of performance and thus to improvements in CPU technology. So the *shanzhai* get to use essentially the same core chipset for many years without compromising the competitiveness of their final products. This stability in turn affords these small, agile innovators time to learn the platform thoroughly and to produce riff after riff on the same theme. You may fault them for flouting intellectual-property laws, but you have to admit that they often achieve astonishingly creative results on a shoestring budget.

The scene is set, I believe, for many other kinds of open-hardware ecosystems to blossom. The inevitable slowdown of Moore's Law may spell trouble for today's technology giants, but it also creates an opportunity for the fledgling open-hardware movement to grow into something that potentially could be very big.

Personally, I'm looking forward to the changes—including the return of artisan engineering, where elegance, optimization, and balance are valued over raw speed and feature creep. Even if Moore's Law ends soon or abruptly, electrical engineers and consumers alike should learn to stop worrying and prepare to love what's in store. ■

Material in this article originally appeared in the author's blog post "Why the Best Days of Open Hardware Are Yet to Come."

POST YOUR COMMENTS at http://spectrum.ieee.org/mooreslawend0415

The Department of Electrical and Computer Engineering at Carnegie Mellon University has Tenure Track Faculty Openings.

Exceptional candidates in all areas of Electrical and Computer Engineering will be given consideration.

The College of Engineering and the Department of Electrical and Computer Engineering (ECE) are ranked among the top programs in the United States both at the undergraduate and graduate levels. The ECE faculty have strong ties with many multidisciplinary centers and institutes such as the Institute for Complex Engineered Systems (graduated NSF ERV), CyLab, Software Engineering Institute, Human Computer Interaction Institute, the Robotics Institute and the Information Networking Institute. The department is home to several multi disciplinary centers and laboratories; the Data Storage Systems Center (graduated NSF ERC), the Center for Silicon System Implementation, the MicroElectroMechanical Systems Laboratory, the SRC Smart Grid Research Center, the Parallel Data Laboratory, and the Computer Architecture Lab at Carnegie Mellon. Faculty members in the Department of Electrical and Computer Engineering also have ties and collaborations with colleagues around the world through a number of formal research and educational programs. The department has extensive experimental and computing infrastructure. ECE currently maintains state-of-the-art nanofabrication facilities and will be significantly expanding these in both capability and scale within the next few years.

A moderate teaching load allows time for quality research and close involvement with students and other faculty. We expect candidates for a tenure-track position to have a strong commitment both to research and teaching at the undergraduate and graduate levels. The Department is primarily seeking individuals for faculty positions at the Assistant Professor rank; however, under unusual circumstances and commensurate with the qualifications of the individuals, appointments may be made at the rank of Associate Professor or Professor.

Applicants should submit the following application materials as one document in PDF format and in the following order:

1. Resume/CV
2. List of professional references including name and contact information (3-5)
3. Statement of research (2 pages max)
4. Teaching interests (2 pages max)
5. Copy of two research papers (journal or conference)

Apply at **http://cmuecefacultyapps.catsone.com/careers/**
Letters of Recommendation should be e-mailed directly to **facultyapps@ece.cmu.edu**.
(Please reference candidate name and position in the subject line.)

Evaluation of applications will begin immediately and continue throughout the academic year until positions are filled. Candidates are encouraged to apply early. Carnegie Mellon is an Equal Opportunity Employer.

Carnegie Mellon University
College of Engineering

Stony Brook University

Faculty Positions in the Institute for Advanced Computational Science

Applications are invited for four tenure-track faculty positions of any rank (including endowed chairs), in applied mathematics and computer science in the Institute for Advanced Computational Science (IACS) at Stony Brook University. Candidates wishing to apply should have a doctoral degree in Applied Mathematics or Computer Science, though a degree in related fields may be considered. Ten years of faculty or professional experience is required for a senior position along with a demonstrated record of publications and research funding. A demonstrated record of publications and a demonstrated potential for research funding is required for any junior faculty. The selected candidate is expected to participate in interdisciplinary program development within the Institute and to establish a research program with a solid funding base through both internal and external collaborations. Of specific interest is research in, for example, programming models, algorithms, or numerical representations that advance scientific productivity or broaden the benefit and impact of high-performance computing. The selected candidates will have access to world-class facilities including those at nearby Brookhaven National Laboratory.

The Institute for Advanced Computational Science (**http://iacs.stonybrook.edu/**) was established in 2012 with an endowment of $20M, including $10M from the Simons Foundation. The current ten faculty members will double in number over the next few years to span all aspects of computation with the intent of creating a vibrant multi-disciplinary program. IACS seeks to make sustained advances in the fundamental techniques of computation and in high-impact applications including engineering and the physical, life, and social sciences. Our integrated, multidisciplinary team of faculty, students, and staff overcome the limitations at the very core of how we compute, collectively take on challenges of otherwise overwhelming complexity and scale, and individually and jointly define new frontiers and opportunities for discovery through computation. In coordination with the Center for Scientific Computing at Brookhaven National Laboratory, our dynamic and diverse institute serves as an ideal training and proving ground for new generations of students and researchers, and provides computational leadership and resources across the SBU campus and State of New York.

The search will remain open until suitable candidates are found with the first round of applications due May 15, 2015. All candidates must submit the required documentation online through the link provided below. Please input a cover letter, your curriculum vitae, a research plan (max. 2 pages) which should also describe how graduate and undergraduate students participate, a one-page statement of your teaching philosophy, a publication list, your funding record, and three reference letters to: **https://iacs-hiring.cs.stonybrook.edu**.

Stony Brook University/SUNY is an equal opportunity, affirmative action employer.

Joint Institute of Engineering

Faculty Positions available in Electrical and Computer Engineering

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

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

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

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

SHUNDE INTERNATIONAL

Joint Research Institute

Research positions available in Electrical and Computer Engineering

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

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

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

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

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

SUN YAT-SEN UNIVERSITY Carnegie Mellon University

Faculty Positions in Robotics and in Computer Science

NAZARBAYEV UNIVERSITY

Nazarbayev University is seeking highly qualified full-time faculty members at Assistant, Associate and Full Professor ranks to join its rapidly growing programs in Robotics and in Computer Science at the School of Science and Technology. Successful candidates must have an earned Ph.D. degree from a reputable research university, a demonstrated ability for research, excellent English-language communication skills and a commitment to graduate and undergraduate teaching.

Launched in 2010 as the premier national university of Kazakhstan, NU's mandate is to promote the emergence of Astana as the research and educational center of Eurasia. The strategic development of this English-language university is based on the Western model via partnerships with top-ranked world universities. NU is committed to become a world-class research university and is making significant investments in research infrastructure, equipment and personnel.

Applications are particularly encouraged from candidates with research interests in biomechatronics, unmanned aerial vehicles, computer-aided design, manufacturing, advanced machining and industrial automation for the **Department of Robotics** and in intelligent systems, high-performance computing, sensor networks, bio-informatics and big-data for the **Department of Computer Science**. Exceptional candidates with research interests in other topics of Robotics and Computer Science are also encouraged to apply.

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

Applicants should send a detailed CV, teaching and research statements, and list of publications to **sst@nu.edu.kz**. Review of applications will begin immediately but full consideration will be given to applications submitted no later than **May 15th, 2015**. For more information, visit **http://sst.nu.edu.kz** .

TU WIEN TECHNISCHE UNIVERSITÄT WIEN Vienna University of Technology

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

FULL PROFESSOR (TENURED) for High Frequency Engineering (Hochfrequenz Technik)

TU Wien invites applications for a Full Professor position (tenured) at the Fakultät für Elektrotechnik & Informationstechnik **www.etit.tuwien.ac.at** within its five main research areas Photonics, Micro&Nanoelectronics, Telecommunications, Automation and Control Systems, Energy Technologies and Systems. The Institute of Electrodynamics, Microwave and Circuit Engineering **www.emce.tuwien.ac.at** is active in microwave engineering and circuit design and contributes substantially to all of the above areas. We are looking for candidates being able to conduct excellent research in fields related to the future challenges of high frequency engineering in the frequency range between 1 GHz - 1 THz. The successful applicant will enhance our excellence in research by leading a research group in one of the following areas:

- •Passive and active RF-components, devices and systems
- •Modelling, analysis and simulation of ultrahigh frequency electromagnetic processes
- •Novel measurement techniques, metrology and circuit design concepts
- •Metamaterial-based approaches
- •HF-based sensors, imaging and ranging systems

More information concerning this position is available at **http://etit.tuwien.ac.at/fakultaet/offene-stellen-open-positions/**

The application package should be sent to the Dean in digital form (a single pdf or on a CD-ROM) **dekan.etit@tuwien.ac.at**.

Micron Technology, Inc. is seeking the following positions for its semiconductor R&D facility in Boise, ID; its manufacturing facility in Manassas, VA; its sales and design facilities in Folsom and Milpitas, CA; and design facility in Longmont, CO.

The following Micron subsidiaries are also seeking positions: Micron Semiconductor Products, Inc., at its headquarters in Boise, ID, and sales facilities in Meridian, ID; Folsom and Milpitas, CA; and its design facility Micron Technology Texas, LLC in Allen, TX:

Electrical, Electronics, Communications, Chemical, Industrial, Mechanical, Materials, Computer System Analysts and Software Engineering; Physics, Materials Science, Engineering Manager and other related Engineering occupations. Marketing, Sales, Logisticians, Accounting and other related business positions.

Please submit your resume on the Web:
http://www.micron.com/jobs/search/

Resume and/or cover letter must reflect each requirement or it will be rejected. Upon hire, all applicants will be subject to drug testing/screening and background checks.

Note: some of these positions may require domestic and international travel for brief business purposes. Please read the full job description when applying online for such requirements.

EOE

The Department of Electrical and Computer Engineering at Carnegie Mellon University has research track faculty openings.

Exceptional candidates in all areas of Electrical and Computer Engineering will be given serious consideration.

The College of Engineering and the Department of Electrical and Computer Engineering (ECE) are ranked among the top programs in the United States both at the undergraduate and graduate levels. The ECE faculty have strong ties with many multidisciplinary centers and institutes such as the Institute for Complex Engineered Systems (graduated NSF ERC), CyLab, Software Engineering Institute, Human Computer Interaction Institute, the Robotics Institute, and the Information Networking Institute. The department is home to several multidisciplinary centers and laboratories; the Data Storage Systems Center (graduated NSF ERC), the Center for Silicon System Implementation, MicroElectroMechanical Systems Laboratory, SRC Smart Grid Research Center, the Parallel Data Laboratory, and the Computer Architecture Lab at Carnegie Mellon. Faculty members in the Department of Electrical and Computer Engineering also have ties and collaborations with colleagues around the world through a number of formal research and educational programs. The department has extensive experimental and computing infrastructure. ECE currently maintains state-of-the-art nanofabrication facilities and will be significantly expanding these in both capability and scale within the next few years. Research track faculty members are not required to teach and may therefore focus exclusively on developing their leadership within their area of research, developing research collaborations and supervising Ph.D. students. Occasionally, research track faculty members may teach if they and the department feel this would advance their programs. We expect candidates for a research track position to have a strong commitment to research. The department is primarily seeking individuals for faculty positions at the Assistant Research Professor rank. Applicants must have earned a doctorate in a related area by the date of appointment.

Applicants should submit the following application materials as one document in PDF format and in the following order:

1. Resume/CV
2. List of professional references including name and contact information (3-5)
3. Statement of research (2 pages max)
4. Teaching interests (2 pages max)
5. Copy of two research papers (journal or conference)

Apply at **http://cmuecefacultyapps.catsone.com/careers/**
Letters of Recommendation should be e-mailed directly to **facultyapps@ece.cmu.edu**.
(Please reference candidate name and position in the subject line.)

Evaluation of applications will begin immediately and continue throughout the academic year until positions are filled. Candidates are encouraged to apply early. Carnegie Mellon is an Equal Opportunity Employer.

Carnegie Mellon University
College of Engineering

---

*iCrea

Catalan Institution For Research and Advanced Studies

10 Senior Research Positions

Candidates for the areas of Experimental Sciences & Mathematics and Technology & Engineering with an outstanding research record and excellent leadership capabilities will be considered.

Further information and applications available at:

www.icrea.cat

Deadline: April 21, 2015

***ICREA** is a foundation sponsored by the Government of Catalonia *(Generalitat de Catalunya)*

Sr. Programmer Analyst (Regions - Birmingham, AL) Utilize minimum requirements to formulate and define specifications for complex operating software programming applications. Modify and maintain complex existing applications using engineering releases and utilities from programs.

Minimum requirements: Must have a Bachelor's degree or foreign degree equivalent in computer science, math, engineering, or related field and 5 years of experience in the job offered or 5 years of experience as a Sr. Programmer Analyst in a technical lead position. Experience may have been obtained concurrently and must include: (i) 5 years of business process automation experience in equity lending and direct lending; (ii) 5 years of Exigen application development with VisiFlow framework which must include each: Workflow, Exigen Queue Server, and Exigen integrator scripting; (iii) 3 years of ASP.net software development; (iii) 3 years of Web based front end development; (iv) 5 years of experience with D+H (Harland) Laserpro specific in consumer lending; (v) 3 years of experience with each: COM+, Entity Framework, MS Visual Studio and Team Foundation Server (TFS); 5 years of experience with Oracle Server, VBScript and VB6; and (vi) 5 years of experience using MS Host Integration Server.

Must have legal authority to work in U.S. Any suitable combination of education, training or experience is acceptable. Resume to Regions Financial Corporation, Attn: Jodie Everingham, 2090 Parkway Office Circle, Birmingham, AL 35244.

AND NOW THIS...

CARRY-ON CLOCK

TEMPORAL PRECISION IN A BOX

High-speed electronics require clocks that must be calibrated to precise standards. Once, the ultimate standard was based on astronomical observations. But in the mid-1950s, the first reliable atomic clocks were constructed and enthusiastically adopted: This ad for a portable atomic clock—weighing only 14 kilograms and with an 8-hour battery life—was published in April 1967, the same year the second was redefined in terms of the electronic transitions of a cesium atom. But atomic clocks didn't really get much more portable than this until 2011, following a breakthrough that squeezed an atomic clock down to the size of a computer chip—leading to, among other things, the world's first atomic pocket watch.

—STEPHEN CASS

PHOTOGRAPH BY Randi Klett


MathWorks®
Accelerating the pace of engineering and science
您会说
MATLAB吗？
Over one million people around the world speak MATLAB.
Engineers and scientists in every field from aerospace and semiconductors to biotech, financial services, and earth and ocean sciences use it to express their ideas.
Do you speak MATLAB?
Solar Image taken by the X-Ray Telescope: supplied courtesy of Smithsonian Astrophysical Observatory.
See related article at mathworks.com/ltc
Image Credit: SAO, NASA, JAXA/ISAS, NAOJ ©2011 The MathWorks, Inc.
MATLAB®
The language of technical computing

IEEE SPECTRUM

NEW CAREER GUIDE

YOUR PASSPORT TO TODAY'S GLOBAL JOB MARKET

PROGRAMMING THE NEXT GENERATION OF AUTONOMOUS ROBOTS

WELCOME BACK

TO THE SECOND EDITION OF THE IEEE SPECTRUM CAREER GUIDE!

Job hunting is serious business—and an ever-changing process. *IEEE Spectrum* wants to help new engineers navigate these waters and land their dream jobs. In 2014, *Spectrum* offered a digital publication full of advice, tactics, and strategies on winning the ideal jobs and launching successful careers. The Career Guide was designed especially for newly graduated and about-to-graduate engineers. This year's edition is full of ideas about tailoring your job search to fit your interests and skills. We hope you find it useful. Tell us how your search goes. Here's to your perfect career match.

—Mike Buryk
Director of IEEE Media Recruitment Advertising

Cover photo: iRobot Corp

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

PROGRAMMING THE NEXT GENERATION OF AUTONOMOUS ROBOTS

By Nick Hawes

As a future where humans interact regularly with "service robots" draws ever closer, there is a growing need for a new generation of software engineers who have the skills to work with robotic technology. Service robots are robots that can act autonomously (i.e., control themselves to some degree) in order to provide some kind of service to a human. Examples of such technology include robot porters in hospitals or hotels, automated forklifts and logistics vehicles in factories, and driverless cars. The International Federation of Robotics predicts that more than 130,000 service robots will be sold between now and 2017, with an estimated value of US $18.9 billion.

Creating, testing, and maintaining software systems for autonomous robots is an incredibly challenging task because of their inherent complexity. Robot software must handle a broad range of temporal properties, from low-level control loops measured in tens of milliseconds (often talking to hardware), to long-term task plans covering hours and days of robot activity (based on more abstract representations of robot capabilities). A similar range must also be spanned (often in parallel) in terms of representations of the world and the entities within it, from signal values in motors and perceptual systems, to kinematic trajectories, to rooms, objects and people. There is also the issue of testing: unit tests may show you when some small part of your code fails, but being able to test and debug your code under real working conditions (i.e., on board a robot in a close-to-production environment) is an essential part of determining whether your product actually performs. As more and more companies create or rely on service robots, employees with the skills to master these problems will become increasingly in demand.

Opportunities for robotics software engineers are likely to appear at a wide range of businesses in the coming years. There will be openings in companies that are developing their own autonomous robot products for the private market, from cleaning manufacturers like Dyson and iRobot all the way up to automotive manufacturers such as Nissan and Volkswagen, which produce autonomous "driver assistance" technologies that can take over driving from a human in certain situations. While these markets will be large in the future, it is the industrial use of robots that will create more jobs in the near term, as companies look to automate various parts of their businesses. Opportunities here may come from the companies themselves as they fund their own development of robot technology, or from companies that supply and contract to industrial users. Interesting examples of companies already looking to explore robot solutions include online super-

US $18.9 BILLION The estimated value of 130,000 service robots to be sold between now and 2017

iRobot provides hands-on robotics experience to college-level students through its internship program. The company's STEM education program educates K-12 students about the exciting world of robots through mentoring, classroom visits and other events. Photos: iRobot Corp.

market chain Ocado, delivery firm DHL, and companies working with nuclear technology, such as Avera and the Nuclear Decommissioning Authority. In Germany, and now influencing other European companies, the Industry 4.0 agenda–which features a push for more intelligence and automation in manufacturing–is driving large companies such as Siemens, Bosch, and Hewlett-Packard to invest in autonomous systems.

Alongside these developments are a growing range of companies that develop autonomous robots for industrial use, and these will certainly need to hire more software engineers. Looking across Europe, there are technology providers such as Guidance Automation and MoTuM, which work with autonomous ground vehicles; Shadow Robot Company, which supplies dexterous hands; and SeeByte, which provides unmanned underwater vehicles. In the United States, companies such as Clearpath Robotics and Rethink Robotics are developing new, exciting, and affordable robotic technologies. There are also more traditional manufacturers of industrial equipment, such as Kuka and Caterpillar, which are increasingly looking to deliver more autonomy from their products. Finally, although often ignored due to its relatively uncool image, agriculture could very well become one of the major industrial users of autonomy in the near future.

Autonomous robotics is experiencing a substantial growth period. If you are looking to move into this field, keeping an eye on the trends and companies outlined above will put you ahead of the competition. Jobs in this field are often advertised on the robotics-worldwide mailing list. In Europe, groups such as EU Robotics and SPARC will be driving the field forward in the coming years, and their websites and events will also provide information on opportunities that could arise for future employment. To make sure you are able to take advantage of opportunities that come your way, training in the key technologies required for understanding and programming autonomous robots is also essential. This can be done by working independently with open-source robotics tools (ROS) and simulators (MORSE, Gazebo) or seeking a robotics master's degree such as the one offered by the School of Computer Science at the University of Birmingham, in England. ■

Dr. Nick Hawes is a senior lecturer in the School of Computer Science at the University of Birmingham. His research is focused on applying techniques from artificial intelligence to allow robots to perform useful tasks in everyday environments (such as making your breakfast or supporting

Elevating Your Job Hunt

By Prachi Patel

4 Steps to Success:

- Discover your passion
- Stay on target
- Research the job, company
- Expand your network

While finishing her electrical and electronics engineering masters thesis at Columbia University in the spring of 2012, Natalia Baklitskaya saw a job posting for a hardware engineer at IBM. It sounded like a perfect real-world use of her technical skills. At the University's career fair, she made a beeline for IBM recruiters. "I told them why [it] was my dream job and how well it fit my skills," she says.

She landed an interview and then a job offer. A year later, they mentioned her excitement at recruiting sessions for other Columbia students. "It stood out to them that I was super-excited about the job and...they were excited in return," she says.

Many IEEE members stress that a few key steps can give you an edge in your engineering job hunt: discovering your passion, targeting your job hunting strategies, researching a job and the company, and expanding your networks.

Job websites can be black holes for resumes. To stand out, it is vital to connect with employers in person, Baklitskaya says. She has found that attending conferences and career fairs is the best way to do that. At the October 2013 Society of Women Engineers conference career fair, handing her resume and talking to Intel recruiters led to a job offer. "They texted me during the conference to come over and talk to them," she says.

For Rajlaxmi Chauhan, a Ph.D. student at the Indian Institute of Technology in Kharagpur, the true test of her passion came during her senior year in 2009. She had a job offer at the IT consulting and services powerhouse Wipro, one at the Indian Department of Atomic Energy, and a graduate school acceptance letter. In the end, her fascination with image processing and fondness for teaching won and she pursued higher studies.

Chauhan expects to earn her Ph.D. in early 2015 and is now getting academic job leads via job and professional networking websites, IEEE conferences and events, and through her professional and social networks.

Internships helped Liana Nicklaus hone her technical skills and identify her career calling. Nicklaus, a design engineer at ARM in Austin, Texas, graduated in May 2014 with a bachelors degree in electrical engineering from the University of Illinois in Urbana-Champaign. She found all her internships–at Wolfram Research, IBM, and Qualcomm–through the school's career fair. By junior year, she knew that she wanted to focus on logic design and computer architecture, thanks to an "incredibly fun and interesting" IBM internship doing digital design for a field-programmable gate array.

Nicklaus says she started looking for full-time jobs the minute she arrived on campus as a senior. With job offers from ARM, IBM, and Qualcomm in hand, she chose ARM. "It is a small, lean workforce, and to me that meant more responsibility and opportunity to do work that would have a big impact on the end product, which has turned out to be true," she says.

According to the spring 2014 Job Outlook Survey by the National Association of Colleges and Employers in Bethlehem, Pa. employers look for new graduates who are leaders, team-workers and good communicators. IEEE involvement can bolster all three. And they are valuable for extending professional networks.

"My IEEE connections always come in handy," says Dario Schor, a software engineer at Magellan Aerospace in Canada. Schor, who has a master's degree in computer engineering from the University of Manitoba, always dreamed about working in the aerospace sector. Early in his graduate studies, his advisor Witold Kinsner (IEEE Life Senior Member and president-elect of IEEE Canada) received a call from colleagues at Magellan Aerospace. They were looking for a graduate student to take on a short-term contract working on a spacecraft simulator for a satellite mission. After a few rounds of interviews, Schor landed the contract. "Over time, that developed into the full-time position I hold today," he says.

Meanwhile Baklitskaya, who served as president of the Columbia University IEEE student branch for one year, says that the experience balanced her technical side and rounded her out as an engineer.

Baklitskaya is now a design verification engineer at Qualcomm, a job she started in January 2014. That job was a result of her private social network. She gave her resume to a friend working at Qualcomm, who passed it along to a hiring manager.

Social connections also helped Josip Balen, a post-doctoral researcher at the University of J.J. Strossmayer in Osijek, Croatia who earned his Ph.D. in electrical engineering in October 2014. "I was lucky to get all four jobs I've had just by knowing people and by having good references," he says. "Networking is crucial for engineers even on a small scale." But his most important advice for job seekers: "Be strong-willed and do not give up. Brick walls are there for a reason." ■

Prachi Patel, who holds a master's degree in electrical engineering from Princeton University, is a contributing editor to *IEEE Spectrum* and a writer for *The Institute*. She is a freelance science and technology journalist based in Pittsburgh.

Career Paths: Technical or Management?

By Richard H. Spencer & Raymond E. Floyd

New engineers hiring into a company's engineering organization can expect to enter at a low level and then work their way up through a progression of increasing levels of responsibility, according to their ability and performance. Typical titles would be junior engineer or associate engineer. Over time these engineers will be given ever increasing difficult projects with which to demonstrate their understanding of the corporate requirements and their grasp of the technical nature of their work. At some point in time, they may discover what is known as the "dual ladder" in the corporate structure–the choice being whether to continue in the technical aspects of their career or move into the management side. While the choice may seem easy, one should consider several aspects of each path prior to making that important decision. What, then, are those aspects?

Technical Ladder

The successful engineers in the technical ladder career choice are those who like "dirt under their fingernails." These are individuals who like the feel of the equipment under their direct control: the design, the build, and the delivery of the final product. Technical-based engineers may operate as members of a team, but frequently find they are operating solo. Recognition by peers and management feeds their self-worth and is the reward they seek. If accomplishment is the major factor in measuring self-worth, the technical side of the business is probably the best fit. Unfortunately, many who fit this picture choose to move into management. When this is the choice, others may lose, as the individual no longer controls the project directly and the people working for him or her are not recognized properly. As a result, the corporation may lose a valuable resource–its people.

Management Ladder

Most corporations provide entry management positions staffed by individuals who have demonstrated a good technical grasp of their particular area. Unfortunately, sometimes that choice may be made for all the wrong reasons, and the new manager becomes a dismal failure due to the differences between a technical engineer and an engineering manager. Managers must be willing to help their people be successful. New managers must also recognize that their position is different from what it was. Before the move into management, others in the department were friends and buddies. Now they are subordinates to be evaluated, supported (including personal problems and judgments), and, in some cases, fired.

Most who enter management as a career choice are thinking of their opportunity to move from department manager to product/project manager and on up the corporate ladder to vice president, chief executive officer, and similar positions. Unfortunately, most forget that the management function is a pyramid structure, with a narrowing set of opportunities for advancement.

Dual Ladder Opportunities

The authors, early in their careers, discovered two things that were to shape their choices of career paths: They enjoyed the feel of equipment, and they enjoyed working directly with people. As a result, they both enjoyed working on projects for periods, moving into management for a period, and then returning to the technical side. The approach suggested by the authors is that it can be good for a career if the individual takes time to move from one ladder to the other. This allows skills to be renewed, removes personnel pressures, and offers a general opportunity to be concerned with yourself and your own feeling of accomplishment.

It is interesting to note the change in organizational structure with the advent of dot-com businesses. While some might identify these new ventures with software, that is not the limit! Some may offer a special app for your iPad or a new concept in garden sprayers. Within these organizations, there may be no apparent structure, not even an evident leader. However, there will still be some people, or person, who is the focal point for decisions and direction to the entire group. Whether in a pinstripe suit, white shirt, and sincere tie or a tank top and blue jean cutoffs, someone will become that leader (manager) and work with the techies to accomplish the "corporate" goals. Times change and management roles change, too, but in the end, you will still have a choice to be that leader or a follower doing your thing. ■

Raymond Floyd, BSEE, MSEE, and Ph.D. in Management. He spent twenty six years with IBM in a number of positions, retiring as a Senior Engineer. Richard Spencer, BSEE. He was with IBM thirty five years in a number of positions with the most time spent in Product Test. He retired as a Senior Engineer. Both authors are IEEE Life Senior members.

Biometrics Careers Are Poised for Growth

By John R. Platt

"On the cusp." That's how almost everyone I spoke with described the biometrics industry, which seems poised for massive growth over the next few years.

Biometrics–the ability to identify people by their physical characteristics or other traits–has been around for years, most notably in security applications. However Apple's addition of a fingerprint scanner to the iPhone 5 has kicked off a wave of new interest in biometric applications for consumer devices.

As a result, says Rawlson King, contributing editor of BiometricUpdate.com, "I think we're about to see an explosion in the industry." Art Stewart, vice president of the biometric products division at Synaptics, agrees: "We're in a massive period of growth."

It's not just consumer products: Plenty of brand-new applications for biometrics could also be rolled out over the next few years. "There's still a lot of room for innovation," says David Tunnell, CTO of NXT-ID.

A Truly Multidisciplinary Field

The term "biometrics" refers to a wide range of applications, including surveillance, voice verification, iris verification, facial recognition, and fingerprint recognition.

Developing the technologies to enable all of these applications requires people from a wide range of disciplines. "It can span everything from electrical engineering to computer science to mechanical and system engineering," Tunnell says. The field also needs people to write algorithms, come up with sensors, develop security protocols, worry about the legal ramifications, and understand the biology of the people who are being verified in the first place.

Creating technology that taps so many varied skills requires assembling a team capable of working together. "Collaboration is valued maybe a little more in this profession than in others," Tunnell says. "You need good sensor people. You need people who love algorithms. You need people who love the mechanicals and very strong software engineers, and then some hardware engineers for the integration. All the disciplines are important."

Getting In

Because the field employs so many disciplines, people can make the most of their passions. "You can carve out your own specialty," says Arun Ross, associate professor at Michigan State University and vice president of education for the IEEE Biometrics Council. "If you are interested in sensors, you can focus on developing biometric sensors. Others might be experts in image processing or computer vision." Still others might concentrate on the applications that interest them the most, for example, health care or homeland security.

The experts I spoke with don't expect many people to enter the field with specific "biometrics" degrees. Instead, Stewart says, "you're coming out of school with a discipline such as digital or analog design and using that tool set to create methodologies of imaging a fingerprint or scanning an iris or something like that."

King suspects that many people entering the field won't be biometrics experts on their first day. "The labor force won't be quite ready," he says. "This type of work will require a lot of on-the-job training." As a result, he says, "we can foresee somewhat of an employment gap within the sector," which may create opportunities for people to enter the field and grow within it.

What's Coming Next?

Although research into biometrics will continue to be conducted in academia, it will fall to industry to commercialize it, especially for consumer electronics, says Tunnell. King points out that some companies already have the technology in place but haven't quite figured out how to market it to a broad audience yet. "That's a big focus right now," he says.

Risks do exist for those considering the field. For one thing, consumers may not accept the idea of biometrics. "There's a perception that everyone's watching everybody else," Tunnell says. Fear of misuse of biometric data remains, something the industry will need to turn around.

In addition, Ross points out that biometrics have entered the consumer electronics market at a slower rate than expected. "A compelling case has not been made yet as to why we need to replace traditional authentication methods," he says. But King thinks that it won't take long before consumers adopt biometrics. "Then you'll really see an explosion in the industry."

Despite the risks, the experts I spoke with all had high expectations for both financial growth and employment in biometrics. "It's one of those areas where there will be a tremendous demand for jobs, especially in the future," says King. ■

John R. Platt covers technology careers for *Today's Engineer* and IEEE's *The Institute*. He also writes for *Scientific American, Consumer Reports*, and many other publications.

The Modern Engineer: New Skills, Opportunities, and Obligations

By Mark W. Werwath & Howard Wolfman

As companies of all sizes work to develop innovative products and services for the marketplace, many are turning to their more technically gifted employees—engineers and scientists—to step into cross-functional and leadership roles. To fill these roles, companies need employees who can think strategically, manage projects, lead teams, and think globally. Employees who are knowledgeable about intellectual property, regulatory strategy, and legal issues can facilitate better decision making throughout the innovation and design processes. With upfront thinking about these issues, it is much more likely that full advantage will be given to the entrepreneurial efforts of a client or firm.

But the broader involvement of engineers in management and commercialization activities creates new issues for the profession. These include ethical dilemmas that can arise when professionals with different areas of expertise come together to solve problems. These dilemmas include whistle-blowing, product safety, distribution and access, testing, compliance, environmental impact, due diligence, and regulatory influence.

These issues can be quite complex, and engineers may not always be sure how to handle them. That's why IEEE held its first conference that addresses such issues. The IEEE International Symposium on Ethics in Science, Technology, and Engineering was held in May 2014 at the Chicago Marriott O'Hare Hotel. It had panel discussions on topics such as nuclear energy, biomedical engineering, and responsible innovation in research. Workshops covered assessing outcomes; ethics, technology, and the engineer; and professional issues in the technical curriculum.

New Master's Degree Program

A related issue facing the engineering profession is an increased demand for those with broader knowledge and different skill sets beyond just technical skills. While engineers are always learning on the job, few have the opportunity to focus on deepening their knowledge and skills. This may explain why so many employers say they cannot find employees with the full skill sets they are looking for.

Academia has taken notice, and various universities have developed new programs designed to give engineers the skills that today's employers demand. One such program is from the law school at Northwestern University, in Chicago. The Master of Science in Law (MSL) is designed especially for engineers who want to learn more about intellectual property, regulation and regulatory analysis, and business/entrepreneurship. The program can be completed in one academic year by going full-time or could take up to four years on a part-time basis.

While the MSL is housed in the law school, it is not a lawyer-training program. Its sole focus is on training engineers and scientists. The MSL combines the most salient aspects of both law and business curricula into a practical, hands-on curriculum. Classes will be in areas such as intellectual property transactions, intellectual property valuation, written and oral communication, business formation, regulatory strategies, innovation policy and ethics, patent prosecution, information privacy and security, contracts and licensing, and entrepreneurial finance and venture capital.

Another challenge facing entrepreneurs and those innovating within large companies is determining when legal rules and administrative regulations are a constraint and when they are an opportunity. The MSL is designed to help its students assess these legal and regulatory situations and develop a comprehensive strategy for addressing them. Numerous situations demonstrate that these strategic decisions can make or break a business opportunity. These include trademark issues involving Apple's iPhone; a variety of business and legal issues involving car-sharing apps such as Uber and Lyft, as well as how these apps affect the business of traditional taxis; Tesla's battle with U.S. regulators surrounding the car dealership structure; patent infringement suits involving every major tech manufacturer; and DNA testing service 23andMe's legal problems with the U.S. Food and Drug Administration. Understanding the entire "innovation landscape" and its tremendous challenges and exciting opportunities is a necessary condition for success in the marketplace.

Ultimately, engineers who receive this training will bring more knowledge to their employers. These multitalented engineers will differentiate themselves in the marketplace and be well prepared for a variety of positions, including those in management.

Both the upcoming IEEE ethics symposium and the new MSL program at Northwestern Law School, which began accepting applications for its first class in September 2014, are responding to the changing role of the modern engineer and scientist in a global economy. ■

Mark Werwath, a certified Project Management Professional, has been an engineering project manager for 25 years. Werwath is currently the director of the Master of Engineering Management Program at Northwestern University's McCormick School of Engineering and Applied Sciences.

Howard Wolfman is an IEEE Life Senior member and the principal of Lumispec Consulting, specializing in energy-efficient lighting. He is also an adjunct professor at the University of Illinois at Chicago and a member of the IEEE Ethics Organizing Committee.

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

HOW TO GET HIRED IN THE AUTOMOTIVE INDUSTRY

By John R. Platt

The rapid development of such vehicular technologies as advanced safety, electrified power trains, autonomous driving, and advanced batteries means that automakers and their suppliers need lots of talented employees.

"There's a big demand for engineers knowledgeable about power electronics, as well as for those in disciplines—such as software and computer science—that have not traditionally been part of the automotive industry," says IEEE Fellow David Munson, dean of engineering at the University of Michigan, in Ann Arbor.

Munson sees the industry responding not only to new technologies but also to a growth in sales. More cars were sold in 2013 than in any year since 2007. As a result, the industry is hiring to rebuild its workforce after years of layoffs and hiring slowdowns. But amid reports of an engineering shortage, automakers report that it's tough finding qualified employees.

THE PATH LESS FOLLOWED

One reason is because few universities teach power electronics, one of the most important disciplines for building today's electronics-packed vehicles.

"We don't have enough trained people who can do the job," says IEEE Member Srabanti Chowdhury. She is an assistant professor in the School of Electrical, Computer and Energy Engineering at Arizona State University, in Tempe. She is also a principal investigator at the Next Generation Power Electronics National Manufacturing Innovation Institute, a project run by North Carolina State University, in Raleigh, to improve energy efficiencies in electric cars, among other applications. The U.S. Department of Energy chose the school in 2014 to lead the $140 million public-private program.

Chowdhury says the industry needs people "very comfortable handling high-voltage, high-current concepts in semiconductor devices." This may not seem as interesting to engineers as, say, nanoscale computer circuits for high-speed computer applications, which have attracted more attention and more students, but power electronics play a vital role in the newest vehicles, she continues: "The need for energy-efficient power conversion is making solid-state electronics play a vital role in power electronics." Chowdhury believes that engineers with a Ph.D. in power electronics devices, circuits, and systems will especially have a long-term advantage in the industry.

IEEE Fellow Ali Emadi concurs. "I strongly advise students or recent graduates to consider getting a master's degree or a Ph.D. in engineering focusing on automotive technologies because cars are getting so much more sophisticated and electrified." For students just getting started, Emadi suggests focusing on the fundamentals. "Get a really good background in electrical, mechanical, and computer engineering, and software, or pursue mechatronics, which is a combination of all four," he says.

JACK-OF-ALL-TRADES

With so many different types of technologies in today's vehicles, the ability to work in multidisciplinary teams is "nonnegotiable," according to Munson. "For electric cars, for example, we've got to have teams that know about communications, computers, electric power, and several other disciplines within electrical engineering." Students must also become more multidisciplinary themselves. Munson reports that some of his students interested in energy storage are taking chemical engineering courses and working on projects outside the classroom that involve new types of batteries.

Most car companies have online job boards and attend career fairs at engineering schools. "These companies are recruiting heavily at our university," Munson reports, saying their university's fall career fair attracts both automakers and auto suppliers. Many companies also recruit at events organized by professional organizations, such as the annual Society of Women Engineers conference.

Munson also suggests looking for jobs outside the big automakers themselves, as there are many other companies that serve the industry. "Some of these 'tier one' suppliers are gigantic companies that have very impressive research and development facilities in the state of Michigan," he says.

TRAINING

Even engineers who did not originally plan to work in the transportation sector can make the transition, adds Emadi. They can work on a part-time master's degree or enroll in short courses, such as the ones offered at the Educational Boot Camp of the IEEE Transportation Electrification Conference and Expo. The most important thing to learn, he says, is the language of the automotive industry, which may be different from what engineers are used to in other disciplines.

"Engineers might not realize just yet how big an opportunity we're talking about," Emadi says. The demand isn't just for the next few years, he says, but will continue for decades, expanding into aerospace, marine, buses, railroads, and trucks. "It's a big, big industry," he says. ■

John R. Platt covers technology careers for *Today's Engineer* and IEEE's *The Institute*. He also writes for *Scientific American*, *Consumer Reports,* and many other publications.

US$50 000 000 Cost of Google's program to encourage young girls to become coders.

Coding Programs Give Women Engineers a Leg Up in Their Careers

By Kathy Pretz

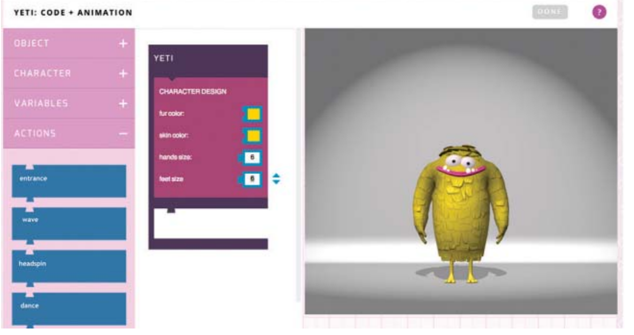

Coding games and animation opens doors for women looking to enter the booming programming job market.

High-tech companies are trying all sorts of ways to boost the number of women engineers. Google recently committed US $50 million over the next year to help encourage young girls to become coders with its Made with Code initiative. The project includes introductory coding projects like making a bracelet and printing it on a 3-D printer, and inspirational videos about women who are using code in their dream jobs. The initiative also includes partners such as the Girl Scouts of the USA, MIT Media Lab, and the National Center for Women and Information Technology. But Google isn't the first to encourage women and young girls to take up coding; other efforts have been under way for years.

IEEE Member Rupa Dachere launched the award-winning CodeChix, the first such program of its kind, in 2009. She used her own money to start this grassroots effort to educate, promote, and mentor women developers. Dachere is its executive director as well as a member of the technical staff at VMware, in San Francisco. Through chapters set up by local women developers, CodeChix strives to foster continuous learning, provide mentorship and guidance, and create a local community of women developers from industry and academia. It targets women engineers already in the field, recent female graduates, and those returning to work after a break. There are now 1,500 members. Chapters originally were formed in California's Bay Area but have expanded to Madison, Wis.; Milwaukee; Redmond, Wash.; and Seattle.

CodeChix held its first hands-on hardware and software workshop earlier this year to code Pi Doorbell, a home automation device that notifies homeowners of visitors. Other events include building "Star Trek"-like LED lights, various hackathons, and hands-on developer workshops. The chapters also have held technical talks and workshops that are open to the public. IEEE recently recognized Dachere's efforts with the 2013 IEEE Educational Activities Board's Meritorious Achievement Award in Informal Education for "establishing a non-profit organization that provides educational programs and mentoring for females in computer programming."

Square, the designer of credit card readers for smartphones and tablets, based in San Francisco, runs the High School Code Camp. This initiative teaches programming to female students to help prepare them for the advanced placement computer science exam. Square engineers developed the computer science curriculum and hold lectures for the students twice a week at the company's headquarters. The students also participate in hands-on lab projects; for this year's final project, student teams built Android mobile apps.

Square also runs a College Code Camp, which is a four-day immersion program that aims to inspire, educate, and empower the next generation of women in technology. The program uses leadership sessions, coding workshops, and a hackathon to build a stronger community around technically savvy women.

"We introduce the young women to Square engineers and our inspiring women executives, including chief financial officer Sarah Friar, business lead Françoise Brougher, and engineering lead Alyssa Henry," said Lindsay Wiese, the company's communications lead, who is involved with the camps. "By hosting and building Code Camp in-house, we are able to expose the young women to a wide variety of opportunities in engineering." ■

Kathy Pretz is the editor in chief of The Institute, *IEEE's member newspaper.*

Scaling Online Education to the World

By Rui Costa

Education is set for big changes, as the lowering costs of multimedia production lead many to spread ideas and knowledge in more creative and effective ways, while the growing reach of the Internet makes it ever easier to connect learners with teachers and mentors.

TED founder Richard Saul Wurman started recording TED conferences right from the very first talk, in 1984, by Nicholas Negroponte, founder of the MIT Media Lab. But only recently did technology emerge to make those talks available to everyone. Neat.

Khan Academy leverages the new tools of the Internet to reach learners all over the world. You can pause, go back, go forward, repeat, and engage with videos that teach you what before you could learn only in one-time-only classes or paper textbooks (often hard to consult).

Welcome to digitally enhanced education, where you can learn at your own rhythm, invent your own methods, and reinvent old ones. And this will not be the end of classes or schools or books or professors: rather, by complementing them, it will let them focus on what they do best.

The list of online educational initiatives is unending, but some concerns arise with all that diversity. Most initiatives adopt a one-size-fits-all model, creating content for a global audience all at the same time. Quality control seems to have low priority, as many projects favor "let's do something fast" over "let's do something good." But people gradually become demanding about what is offered to them, even if it's free; people already feel itchy when watching a video that was not recorded at high definition. A move-fast-breaking-things mind-set might work for disruptive tech start-ups, but it certainly does not when you want to help people learn and become independent.

Back in 2011, when we were student volunteers of IEEE at Técnico Lisboa, in Portugal, we started a pilot project with students and professors collaborating to create online videos (in Portuguese) about the most relevant (and difficult) content taught at our university.

After the May 2012 launch, the feedback was overwhelming: For the very first time, students felt as if their own professor (who writes and grades their exams) was right there with them, explaining exactly what they must know and exactly what they need to prepare for their examinations, laboratory experiments, and so on. Students credited us with helping them get inside a topic—and sometimes finish their degrees. Due to our innovative content, we even had views and feedback from other Portuguese-speaking countries, such as Brazil and Angola.

The innovation was that our topic-oriented videos (averaging 7 minutes) were developed locally, in the local language, and designed to truly suit the students' needs. Unlike the MOOCs (massive open online courses) that were becoming popular then, we aimed not to replace classes or classrooms but rather to provide an additional resource, based on multimedia and flavored with the requirements that students are used to in their local schools.

When we challenged others to join us, IEEE Academic was born: a distributed project that recruits students and professors from all over the world, provides a common infrastructure with very low costs (no need to waste time reinventing logistics that are the same everywhere), and empowers them to create content that addresses local needs—and often solves their problems, as a single video can help countless students over the years. We gather those contributions into a long-lasting library, freely available to everyone in the world.

Students can improve their learning. Professors can enhance their classes, flipping the classroom and freeing up time to engage with their students. Professionals can stay up to date with the latest developments or review something they studied many years ago, working actively toward their career development.

IEEE Academic, in other words, provides global, diversified, multilingual, and easy-to-access knowledge, for all ages and stages of the professional life.

This model has proved successful. Today we have volunteers all over the world. We aim not only to create high-quality content and put it online but also to take education to countries and locations that still do not have access to it. Education empowers people to advance, improve, and make better choices. Everyone deserves the same access to it. Despite all major investments in online education, there are still many people without access to educational content that is truly relevant to them. We want to reach them too.

By leveraging the efforts of a distributed international team of students, professors, and volunteers, IEEE Academic will combine the most innovative ideas and technologies to make education accessible to everyone, everywhere, for free. ■

Rui Costa is the founder and coordinator of IEEE Academic. He was the 2013 recipient of the Larry K. Wilson Regional Student Activities Award in IEEE Region 8, for the creation of IEEE-IST Academic. He currently works at Veniam, where he helps build the networking fabric for the Internet of Moving Things.

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

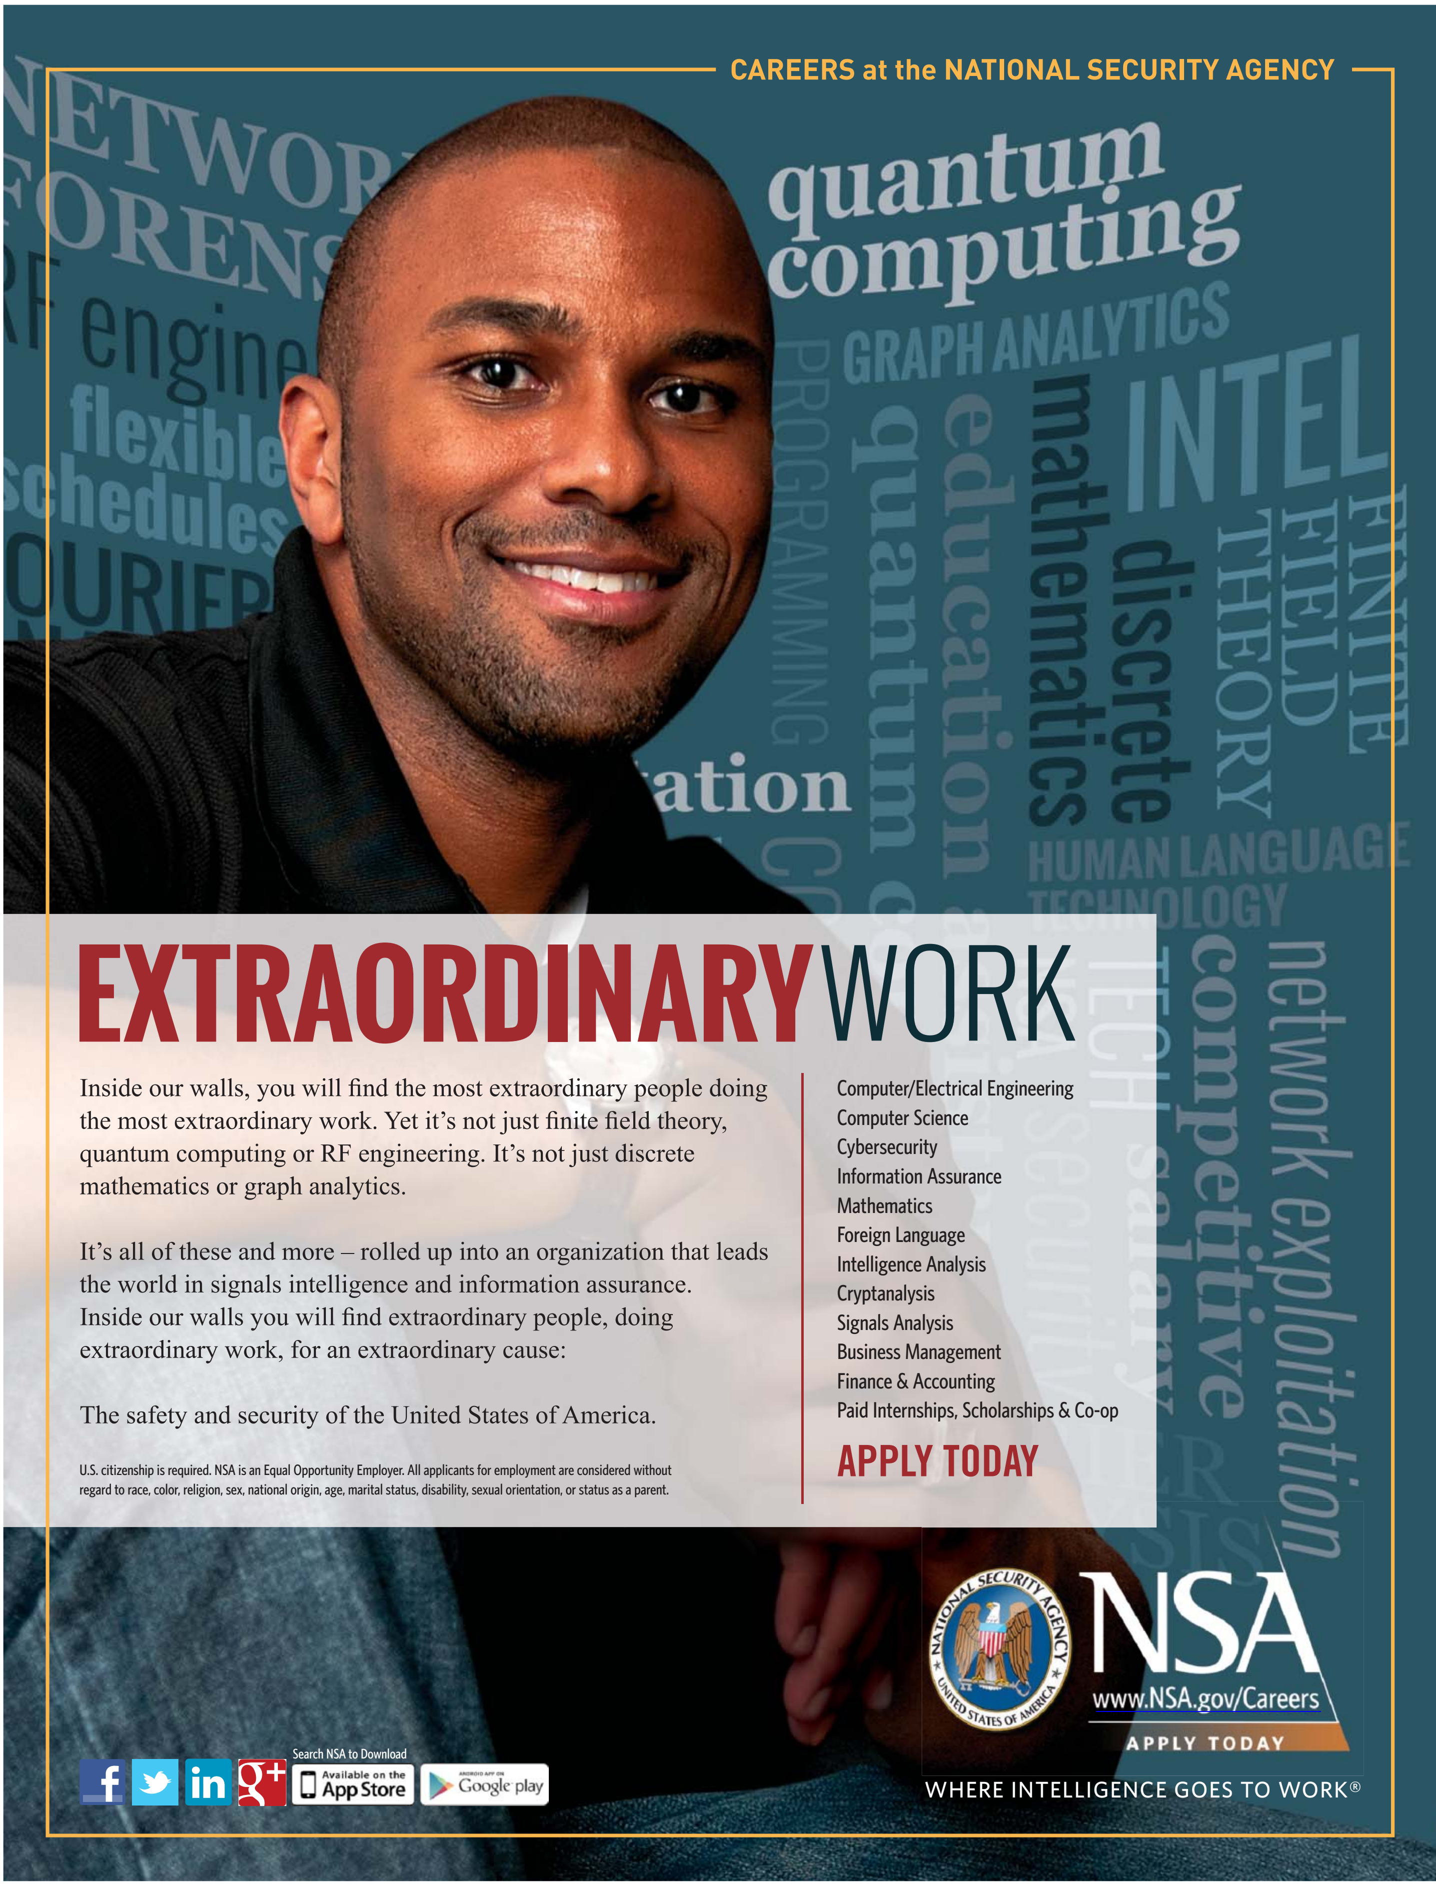
CAREERS at the NATIONAL SECURITY AGENCY
quantum computing
GRAPH ANALYTICS
INTEL
mathematics
discrete
FINITE FIELD THEORY
HUMAN LANGUAGE TECHNOLOGY
competitive
network exploitation
EXTRAORDINARY WORK
Inside our walls, you will find the most extraordinary people doing the most extraordinary work. Yet it's not just finite field theory, quantum computing or RF engineering. It's not just discrete mathematics or graph analytics.
It's all of these and more – rolled up into an organization that leads the world in signals intelligence and information assurance. Inside our walls you will find extraordinary people, doing extraordinary work, for an extraordinary cause:
The safety and security of the United States of America.
U.S. citizenship is required. NSA is an Equal Opportunity Employer. All applicants for employment are considered without regard to race, color, religion, sex, national origin, age, marital status, disability, sexual orientation, or status as a parent.
Computer/Electrical Engineering
Computer Science
Cybersecurity
Information Assurance
Mathematics
Foreign Language
Intelligence Analysis
Cryptanalysis
Signals Analysis
Business Management
Finance & Accounting
Paid Internships, Scholarships & Co-op
APPLY TODAY
NATIONAL SECURITY AGENCY UNITED STATES OF AMERICA
NSA
www.NSA.gov/Careers
APPLY TODAY
WHERE INTELLIGENCE GOES TO WORK®
Search NSA to Download
Available on the App Store
Google play

I CHANGE THE WORLD. I AM IEEE WOMEN IN ENGINEERING

By Nita Patel

Over the past few decades, women's participation in the workforce has increased significantly. Women are earning a larger share of university degrees and the demand for STEM educated professionals has grown more than four times the rate of the U.S. labor force as a whole. Nevertheless, women's representation in STEM occupations has remained relatively flat or is decreasing worldwide.

The reasons why more young girls are not pursuing STEM studies continue to be complex and nuanced; however, a few key themes–public stereotypes, inherent biases, and lack of role models–do emerge. Many organizations, including IEEE Women in Engineering, are focusing on how to inspire, engage, and advance more women in STEM professions. Developing one-on-one mentoring relationships, providing female role models, and increasing the awareness of inaccurate biases and stereotypes can make a difference. Although there are no magic elixirs, IEEE WIE has several programs aimed at changing this demographic.

ROLE MODELS

The lack of role models is a contributing factor to the underrepresentation of women in technology. IEEE WIE works to highlight female engineers through a visibility campaign called "I Change the World. I Am an Engineer." In addition, WIE facilitates online chats with WIE members, distributes posters spotlighting WIE members, provides a WIE app featuring more than IEEE WIE members, and highlights women in its biannual magazine and monthly newsletters.

IEEE WIE advocates for women in leadership roles, promotes and provides career advancement opportunities for women in the profession, educates the public about women in STEM, influences employers and academic institutions to create an inclusive environment, provides mentoring connections to individuals of all ages, honors women trail-blazers, and provides a nurturing, supportive environment for continued growth for women in engineering.

IEEE WIE administers the IEEE Student-Teacher and Research Engineer/Scientist (STAR) Program to mentor young women in junior high school and high school. Through this program, IEEE WIE has connected thousands of young girls each year to mentors within their communities. The 500+ WIE affinity groups around the world conduct STAR programs that engage professional women with other professional women or professional women with young girls. These activities might include classroom participation, humanitarian projects involving the pre-university community, competitions such as Code Hackathons or IEEEXtreme, hands on activities, local networking events, career coaching engagements, training of pre-university teachers, individual mentoring connections, public awareness activities, field trips, and technical meetings.

BIAS IN HIRING & ADVANCEMENT

Several research studies have shown that people create beliefs and stereotypes about gender that affect the way they perceive the behavior and attributes of men and women. For example, a 1999 MIT study on the status of women faculty in science showed women faculty increasingly marginalized as they progressed through their careers and were subject to disparities in salary, lab space, awards

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

and resources despite having equal professional accomplishments to their male colleagues. Other research highlights similar issues in an industrial, corporate environment.

IEEE WIE provides a forum to not only highlight but also discuss these issues through a strong social media network, local affinity group meetings and international events. IEEE WIE holds professional skills development webinars and provides an online forum. The group also held the inaugural IEEE Women in Engineering International Leadership Conference (http://ieee-wie-ilc.org) in May in San Francisco, and plans to make this an annual event. Planning for the 2015 IEEE WIE ILC in the capital of Silicon Valley, San Jose, CA on 23-25 April already is under way. The conference theme is "Lead Beyond: Accelerating Innovative Women Who Change the World." With over 500 expected attendees and 50 incredible speakers, the conference will provide not only skills development workshops, but also sessions on work/life, stereotypes, and engaging men as advocates. ■

Nita Patel, (PE; IEEE Senior Member) is a Systems and Software Engineering Manager at L-3 Warrior Systems. Immediate past chair for Women in Engineering, she currently serves as chair for the 2015 IEEE WIE ILC and is on the Board of Governors for the Computer Society and Eta Kappa Nu.

Moving Forward

Stereotypes and biases can negatively influence young girls from pursuing careers in STEM fields and prevent women from remaining in STEM fields. Parents, teachers, counselors, engineers and executives can take positive, concrete actions to help increase young girls' interest in STEM and retain women in STEM.

A few ideas include:

- Encourage creativity and experimentation through hands-on activities.
- Talk with young children about the inspirational aspects of a STEM career.
- Be aware of the message you are portraying about science and engineering.
- Teach girls that intellectual skills are acquired, not inherent.
- Raise awareness of stereotypes and bias against women in STEM fields.
- Highlight the high job satisfaction among practicing engineers.
- Emphasize the importance of teamwork in engineering.
- Create clear, visible, and transparent advancement paths for senior leadership positions.

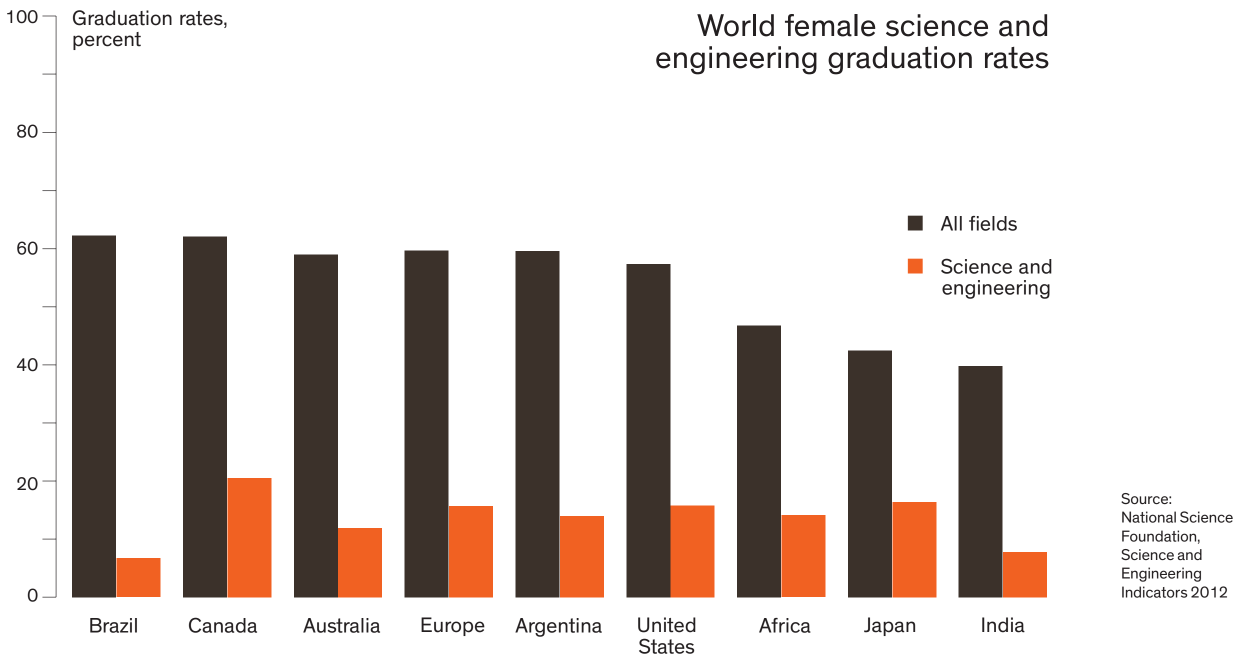

IEEE SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

WORLDWIDE EMPLOYMENT SURVEY

By Ron Schneiderman

In North America, Engineers Go With the Flow

Little has changed in the job outlook for engineering students and recent graduates in North America. If you have the needed technical skills, the grades, and some experience, you have a job. The National Association of Colleges and Employers' Job Outlook 2015 survey says employers are also placing a premium on people with problem-solving/analytic and communications skills. It also helps if you're willing to move. The U.S. Bureau of Labor Statistics (BLS) lists Texas as the fastest-growing state for high-tech professional-level jobs, with Florida ranking second. Austin, Texas, is a particularly fast-growing location for top tech talent. California-based Dropbox, the cloud-based storage provider, plans to hire 200 people in Austin by mid-2015–up from 40 in 2014. AT&T recently added an advanced development engineer to its staff in Austin as part of a plan to grow its team of experts in distributed systems. Florida-based SoftServe expects to almost triple its tech staff to 40 in 2015.

Software developers are in demand just about everywhere, but surveys indicate that the New York metro area has the highest demand for this discipline, with a four-year average annual growth rate of 35 percent. Bangalore-based Infosys is hiring at its technical facilities in Austin and San Francisco and says it plans to recruit up to 180 management and technology graduates from U.S. universities into its consulting practice in the United States. And don't discount offshore job opportunities. "We have people in the U.K. who are U.S. nationals," says Tim Fowler, head of the wireless division of U.K.-based Cambridge Consultants (CC), which designs products for other companies.

CC recently relocated its U.S. headquarters from Cambridge, Mass., to Boston, nearly doubling its space, to accommodate the growth of its U.S. staff from 32 to 100 over the next three years. Its shopping list includes EEs, senior embedded software engineers, senior human factors engineers, and program managers. CC also lists several internships, open mostly to university undergraduates, on its job openings website. "In general, the engineering recruitment market in North America and the U.K. has become much more competitive in the past few years," says Fowler. "There may be a point when some of our U.S. nationals decide they would rather be working in the U.S. We also have some non-U.S. nationals working in the Boston office. It's their expertise that our customers are buying, so being close to MIT is no accident."

Skill sets continue to be critical in hiring. Cisco Systems, which has cut thousands of jobs in the past two years, says it "continues to recruit engineers with skills in cloud, security, virtualization, analytics, and Internet of Things."

Most of the respondents to a survey of U.S. and Canadian electronics manufacturers by the IPC-Association Connecting Electronics Industries, a global trade group, complained that a lack of technical skills and experience in the talent pool in North America is making it difficult to fill job slots for process, test, and design engineers; quality control engineers; and engineering managers.

While applications to college computer science programs continue to soar, the BLS estimates that 1 million programming jobs will remain unfilled in the United States by 2020. Software engineers can pretty much write their own ticket and also have job choices beyond the tech community. Geico, the auto insurance company, is just one of many large non-tech companies that

Engineering Shortages In Europe Create New Opportunities

As in pretty much the rest of the world, Europe's technology companies are in need of specialized skills, and they're looking across the continent, and even globally, to find them.

The Reed Job Index suggests that the talent crunch has created a growing climate of competition for tech talent in Europe. "Companies are keen to snap up top tech and engineering talent," says James Reed, chairman of reed.co.uk. "Substantial skill gaps are emerging, and higher salary expectations are leading to more movement in the jobs market." The increase in engineering job vacancies is particularly apparent in the United Kingdom, with a posting of more than 500 job openings in Oxford and Cambridge.

U.K.-based Cambridge Consultants (CC) recently listed 30 job openings in the United Kingdom for software engineers, analytics engineers, DSP/control software engineers, and human factors engineers, among several other specific openings. The company, which designs products for other companies, also hires interns, mostly university undergraduates, through its website.

Is there a shortage of engineers in the United Kingdom? "Overall, I would say probably yes," says Tim Fowler, who heads CC's wireless division. To fill the

Ron Schneiderman covers technology and business management issues. His latest book, *"Modern Standardization: Case Studies at the Crossroads of Technology, Economics, and Politics,"* is being published under the Wiley/IEEE Press imprint.

are heavily recruiting computer science and computer engineering graduates. Goldman Sachs recently sponsored a well-attended recruiting event in New York City for recent and soon-to-be graduates with strong programming and analytical skills.

Job opportunities in Canada vary geographically. Microsoft says it plans to hire and train 400 software developers from around the world to work on mobile and cloud projects in Vancouver, B.C. Engineers Canada doesn't track tech company recruitment, but Kim Allen, the organization's chief executive officer, says Quebec, New Brunswick, and Newfoundland have the most openings for EEs. "We know that Canada is facing a potential shortage of engineers with more than 10 years of experience." Allen says that's mainly due to impending retirements and the specialized experience required by companies. Immigration is expected to be a major source for filling that gap. Engineers Canada projects that an additional 80,000 engineers of various disciplines will arrive in Canada by 2020.

As for the demand for IT skills, a survey by Teksystems indicates that IT professionals across the United States are receiving an average of 34 job solicitations a week, a significant increase from its previous surveys. Eighty-one percent of the IT specialists surveyed indicated that they're open to new job opportunities–even when happily employed. Careerbuilder said pretty much the same thing in a recent study: "It's not uncommon to see talented programmers and developers fielding multiple job offers from recruiters over the course of a year."

Foote Partners LLC, another company that tracks IT requirements and certifications, says cloud, big data, security, mobile, architecture, and application development top its list of most in-demand IT skills. ■

void, CC, like other tech companies, recruits and employs engineers from across the European Union (EU).

More European tech companies now offer formal internship programs. Most interns work during summer vacations, but others take these positions upon graduation. Google runs post-university graduate training programs in Dublin and Wroclaw, the largest city in western Poland. Some interns are paid, but some are not.

CC offers a program called Tech Scholars, which allows recent high school graduates to spend a year with the company before entering a university. CC works with the U.K.-based Engineering Development Trust, a developer of national education programs, to help find students for the program. (Other companies in the United Kingdom have similar student placement programs, but they each run these in their own way.) Tech Scholars work full time on real projects that have some commercial benefit. "It's not unusual that they end up with patented ideas," says Fowler. "The program gives them a better insight into what is valuable in the real world before they start working on their undergrad degree." They don't have to return to CC after they obtain an engineering or related degree, but several former CC Tech Scholars have rejoined the company after graduation and some have become company program directors.

Qualcomm has engineering openings around the world, but it has focused its recruiting in Europe on its chipset design center in Farnborough, Hampshire in the United Kingdom, and in Ireland. In Farnborough, the company expects to fill at least two 12-month internships beginning in July, one for its RF design team and the other in near-field communications (NFC) and software and systems. Most of the several positions open in Cork, Ireland, require design engineering experience.

Ireland has also become a hot spot for tech startups, with Silicon Valley-based venture capital firms reportedly starting to work more closely with Enterprise Ireland and other early-stage development firms in the country.

Germany continues to have more vacancies for engineers than qualified people to fill them. The German VDE-Association for Electrical, Electronic & Information Technologies, says EE graduates are among the most in demand, but many of its more than 7,000 student members graduating from local or regional universities aren't immediately available because they move on to full-time master's programs. Intel recently posted about 50 engineering job openings in Germany, including internships, mostly in design, test, and software development, and in France, with internships in software engineering.

The European Space Agency (ESA) posts up to 80 vacancies once a year for its Young Graduate Trainee program. The program, whose applicants must have a master's degree in engineering, is designed to give participants "an opportunity to gain valuable experience in the development and operation of space missions."

Additional job openings are expected for EEs and other tech professionals following the EU's approval of several public infrastructure projects over the next three years, including expanding and updating the 28-nation organization's communications systems. ■

WORLDWIDE EMPLOYMENT SURVEY

By Ron Schneiderman

Asia-Pacific: A Growing Hot Spot for Tech Jobs

Much of the job growth in the Asia-Pacific region is the result of local and global companies launching new facilities there.

ARM Ltd., based in the United Kingdom, is opening a fourth global design center dedicated to emerging device development in Hsinchu Science Park, Taiwan–its first CPU design center in Asia. This new facility will focus on the design, verification, and delivery of ARM's Cortex series of processors for wearable devices and embedded applications. ARM said that while it had expected the core engineering team to be in place by the end of 2014, the recruiting process would continue into 2015.

Taiwan-based MediaTek, a fabless semiconductor company, following up on previously announced plans to diversity beyond mobile communications into connectivity and home entertainment sectors, has established an R&D center in Bengaluru, India. Akshay Aggarwal, general manager of MediaTek Bangalore, says he expects the staff at the new facility to grow from 100 technical professionals at the end of 2014 to more than 500 over the next few years.

Samsung Electronics, meanwhile, is investing $14.7 billion to build a new chip fabrication facility in Pyeongtaek, South Korea, south of Seoul. Construction begins this year on the plant, which is expected to be operational in 2017.

Boeing plans to open a cybersecurity analytics facility this year in Singapore as part of efforts to better serve customers in the Asia-Pacific region. The new facility will be located at Boeing's existing Singapore office. Boeing says it will hire local tech talent to operate the new, secure center once it is fully operational, and will train cybersecurity professionals locally. It did not provide details of how many people it will hire.

Anticipating a growing need for computer-aided engineering (CAE) services in new and expanding market sectors and Asia-Pacific, Altair Engineering, a wholly owned subsidiary of U.S.-based Altair, has located its new regional headquarters in Malaysia in Kuala Lumpur. Srirangam Srirangarajan, Altair's managing director of the southeastern Asian nations, Australia, and New Zealand, says Altair selected the location because of its strong talent pool and new business opportunities in the immediate area, mainly in aerospace and electronics.

Taiwan Semiconductor Manufacturing Co. says it plans to add at least 1,700 engineers to its tech staff. TSMC says it is looking for graduates with degrees in electronic/electrical engineering, mechanical engineering, optical science, physics, software engineering, chemistry, material science, chemical engineering, and industrial engineering.

With a renewed focus on skills, expansion, and startups across the industry, India, with strong support from the National Association of Software and Services Companies (Nasscom), has implemented several initiatives aimed at turning the country into a global hub for engineering and design and, as part of the process, technical talent development.

Wipro, India's third largest information technology (IT) company, has been building its tech staff and expanding its range of technology interests. Wipro says it will double its staff of 70 engineers in Muscat, Oman, into 2016 to meet increasing business requirements, and announced its first graduate hiring program in Australia for engineering and computer science graduates. Graduates hired through the program will be enrolled in a six-month training program conducted at Wipro's headquarters in India to further advance their technical skills.

Bangalore-based Infosys says it will hire up to 300 management and technology graduates–mostly computer science and software engineers–from U.S. universities, along with about 1,500 IT professionals during its current fiscal year.

Visa is opening a technology center in Bangalore in early 2015 and expects it to be fully staffed with more than 1,000 people by early 2017. "India is fast becoming a global technology epicenter with an incredible pool of technology talent," says Nitin Chandel, senior vice president of Visa's developer platform, based in India. Teams at the new tech center will focus on the development of key application programming interfaces (APIs) and software development kits.

Cisco Systems, which employs about 11,000 people in India, says it will create a new division in the country, mainly to work on eHealth and education programs but also to help India develop its smart cities and national broadband programs. As for staffing, Cisco has only said that it would move about 15 percent of its current group in India to the newly formed division.

Google and Microsoft are hiring in New Delhi and Mumbai. (Microsoft cofounder Bill Gates has invested in the Unitus Seed Fund, a venture capital group aimed at assisting entrepreneurs in India.) Google has changed from its usual practice of hiring through paper-based tests to coding competitions, called Google Jams, which have been used in other areas globally. ■

Effective Skype Job Interviews

By John R. Platt

Few things in life are more stressful than job interviews. Most interviews are still conducted in person, but today more and more are being conducted over Skype. This is a great way to meet with HR or a hiring manager without having to visit a company's office, but it can also add its own levels of stress.

Don't worry, though. With preparation, a Skype interview can be easier than an in-person interview—and possibly more effective.

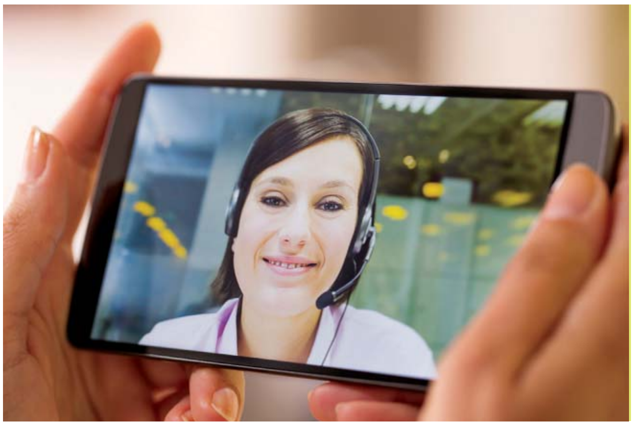

Before the Interview

Experts say you should spend as much time preparing for a Skype interview as you do for one in person. First, find a quiet spot. "If a candidate constantly has to speak to their children or yell at pets, it is not a good indication of someone who's focused and manages their time well," says Shilonda Downing, owner of Virtual Work Team LLC. If your house has the potential to be noisy, make sure that people know you need a good hour of privacy.

Take time to make sure the room lacks visual distractions. "I have interviewed people who have chosen their kids' bedroom for the interview," says Ian Jackson, managing partner of Enshored, an outsourcing company. "It does not help with that first impression if the interviewer can see Mickey Mouse behind you." Bare walls or an office-like setting are best, experts say.

Next, make sure that your camera is at eye level. This mimics the visuals both parties would experience in a real-world interview. Also, you should sit no more than an arm's length away from the webcam, says Bob Myhal, CEO of NextHire, and be centered on the screen.

Of course, all technology is fallible, so test everything out well before the interview. "The biggest challenge I see over and over is that candidates do not test their microphones before connecting," says software engineer Sid Savara. "It is a huge mistake."

If you're not a frequent Skype user, practice. "Skillful use of Skype—sharing documents, pointing to web addresses, etc.—gives you a great chance to show your technical prowess," says Denise Kalm, author of Career Savvy: Keeping and Transforming Your Job.

As with any interview, be prepared. Study the company and prepare questions in advance. Know your own story and be ready to talk about your experience.

Finally, even though you're calling from home, go ahead and dress for success. Make sure you look professional and put your best face forward. (Pants, however, are optional.)

During the Interview

Figuring out exactly when and where to look can be difficult, as eye contact is impossible on a Skype interview. To make up for that, Myhal suggests spending a good portion of your time looking directly at the camera.

Posture and body language also matter. "Sit forward in your chair so your energy is in your body," says Michelle Tillis Lederman, author of The 11 Laws of Likability. Use your hands: Gestures help bring you to life on the screen. "The animation will be projected into your voice and show your energy and enthusiasm," she says.

Be prepared for hiring managers to test you. "I have a short conversation using text messaging to see how candidates could deal with written English," says business coach Liz Scully. "It's easy to write well when given time, but texting quickly means you can see both their speed of reply, their accuracy and how they deal with misspellings."

Expect to show off your technical skills. "If you're a software engineer, your interview might consist of sharing your screen so the interviewer can see you writing code," says Todd Rhoad, managing director of BT Consulting, a career-consulting firm.

After the Interview

As with any interview, the process isn't over once the call ends. Send an email thank-you and follow up with questions. Show that you're interested in the job. "This may seem like an unnecessary formality, but hiring managers put a great deal of stock into this type of professionalism," says Myhal.

If the Skype interview goes well, you may be asked in for a second, in-person interview. That's a great opportunity to build on that relationship you established the first time and, hopefully, close the deal. ■

John R. Platt covers technology careers for *Today's Engineer* and IEEE's *The Institute.* He also writes for *Scientific American, Consumer Reports,* and many other publications.

IEEE Resume Lab Offers Specialized Tools for Job-Seeking

By Rory McCorkle

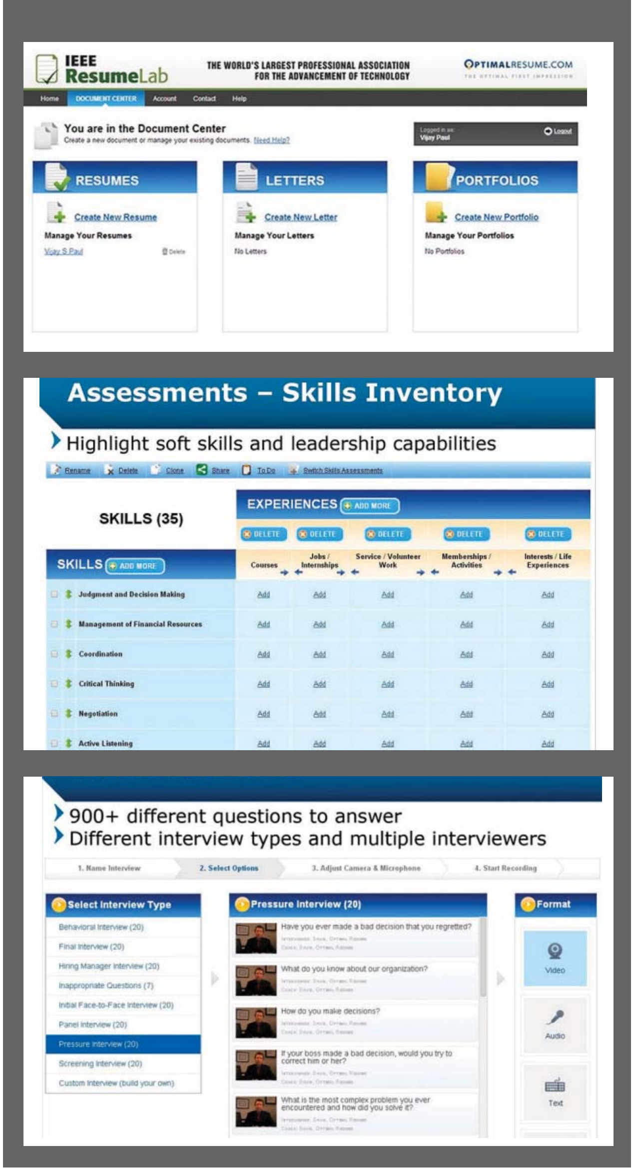

Resume Labs offers a comprehensive set of online tools to help you prepare successfully for your job hunt.

You can never underestimate the power of a positive first impression. And your résumé or curriculum vitae is a critical element of your first impression to a potential employer, whether you're applying for a job or an internship. This first impression can make or break your chances of getting your foot in the door at a potential employer. As a recent article on résumé writing remarked, "it's not the economy keeping people from jobs, it's themselves and their résumés."

IEEE has recognized this challenge and launched a new product to help members with their résumés and other aspects of the employment process. IEEE ResumeLab is an online service that allows IEEE members to develop a résumé or CV using specialized tools tailored for each step of the job-seeking process. This new product is added to the list of offerings that assist members as they find jobs and develop their careers.

IEEE ResumeLab isn't just about resumes, though. It has a series of modules that assist you through the employment process. For example, are you having trouble writing a cover letter? Or do you need help writing a follow-up communication after an interview? ResumeLab's Letters module can assist with templates across the employment process.

Have you had a chance to interview yet for a job or internship? If not, help prepare for this important step with ResumeLab's Mock Interview module. It has more than 900 questions loaded, so that you can answer a variety of questions or interview types. Even better, ResumeLab has built-in tips to help you with answering these sometimes tricky questions. Record your interview to share with a mentor or professor and get their feedback.

Other modules and their features include:

- **PORTFOLIOS**
  Do you have major papers or projects that you have completed during your university or college experience? Highlight your work by uploading and organizing this work to present to potential employers. You might also use this feature to provide writing samples.
- **SKILLS ASSESSMENT**
  Three types of skills assessments are embedded in ResumeLab. Use the Skills Inventory to highlight general skills and abilities, while demonstrating where you have learned and had experience with these skills. Alternatively, use the Proficiency List to assess your industry- or job-specific skills, along with your level of competency and experience.
- **VIDEO RESUMES**
  Record custom video messages for potential employers.

Share everything you create within ResumeLab online or via Facebook, LinkedIn, and other social media. This gives you the ability to share information not only with a potential employer but also with classmates, mentors, faculty, and others. Use it to get feedback on everything you develop in ResumeLab before you share with an employer.

Come to IEEE ResumeLab today and get started! ■

Dr. Rory McCorkle was formerly the Product Manager of Global Career Resources for IEEE. He now serves credentialing bodies as President of International Credentialing Associates (ICA), which researches, defines, and develops professional certifications for these organizations.