IEEE Spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
7.12
SPECTRUM.IEEE.ORG

APES WITH APPS

CAN APES REALLY TALK OR ARE THEY JUST MIMICKING HUMANS? RESEARCHERS USING A NEW TABLET COMPUTER APP AIM TO SETTLE THE DEBATE

STREAMING THE OLYMPICS

THE BIGGEST BROADCAST OPERATION IN SPORTS HISTORY WILL FEATURE LIVE, HIGH-DEFINITION COVERAGE OF EVERY EVENT

COMMENTARY

ALTERNATIVE ENERGY IS A MIRAGE

VACLAV SMIL ARGUES THAT TECHNICAL DIFFICULTIES AND HIGH COSTS WILL MARGINALIZE SOLAR, WIND, AND BIOMASS FOR YEARS TO COME

IEEE

volume 49 number 7 international

ieee spectrum

Winner of the 2012 National Magazine Award for General Excellence

7.12

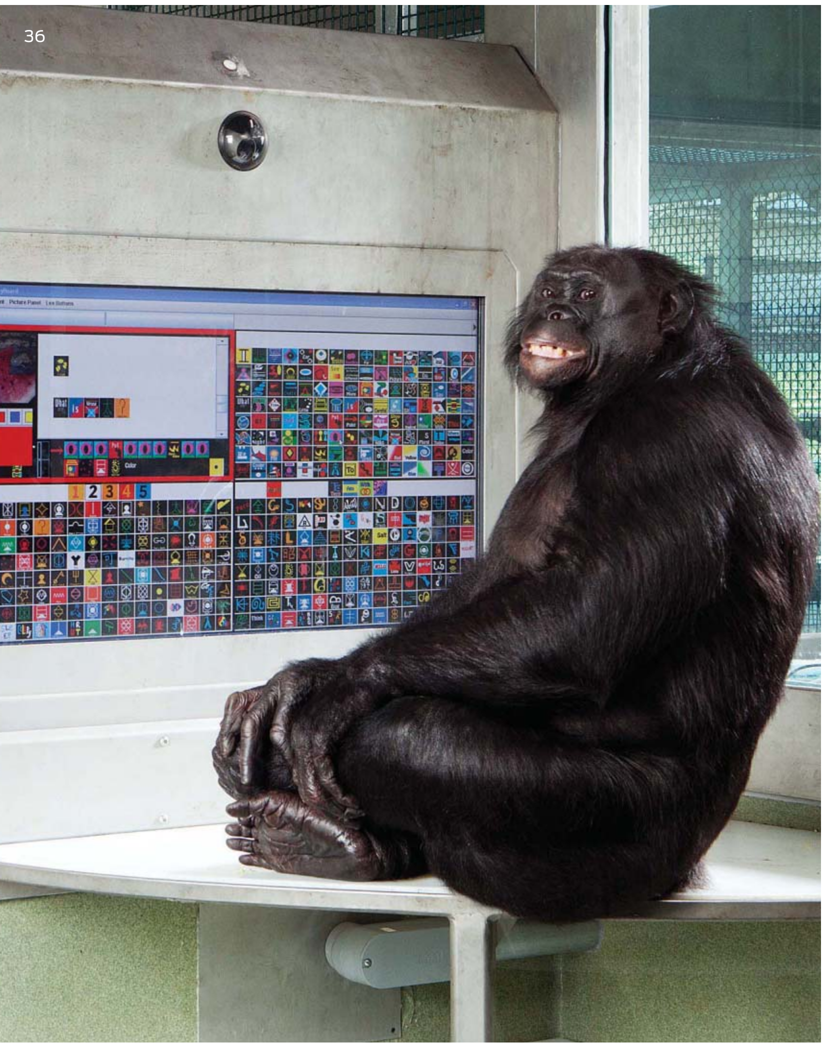
36

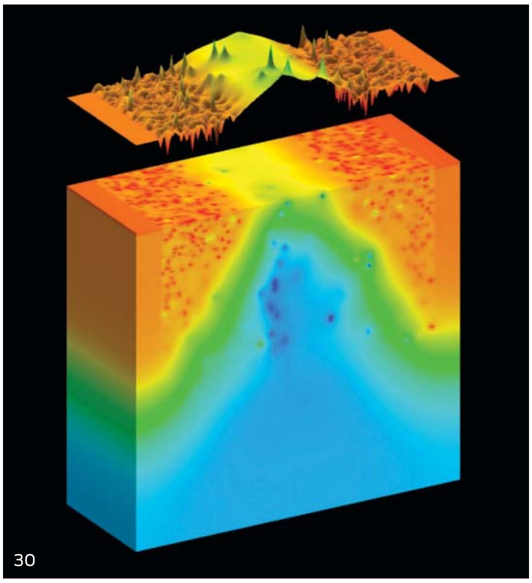
30


44

COVER: GREGG SEGAL; THIS PAGE, CLOCKWISE FROM LEFT: GREGG SEGAL; GOLD STANDARD SIMULATIONS; DAN PAGE

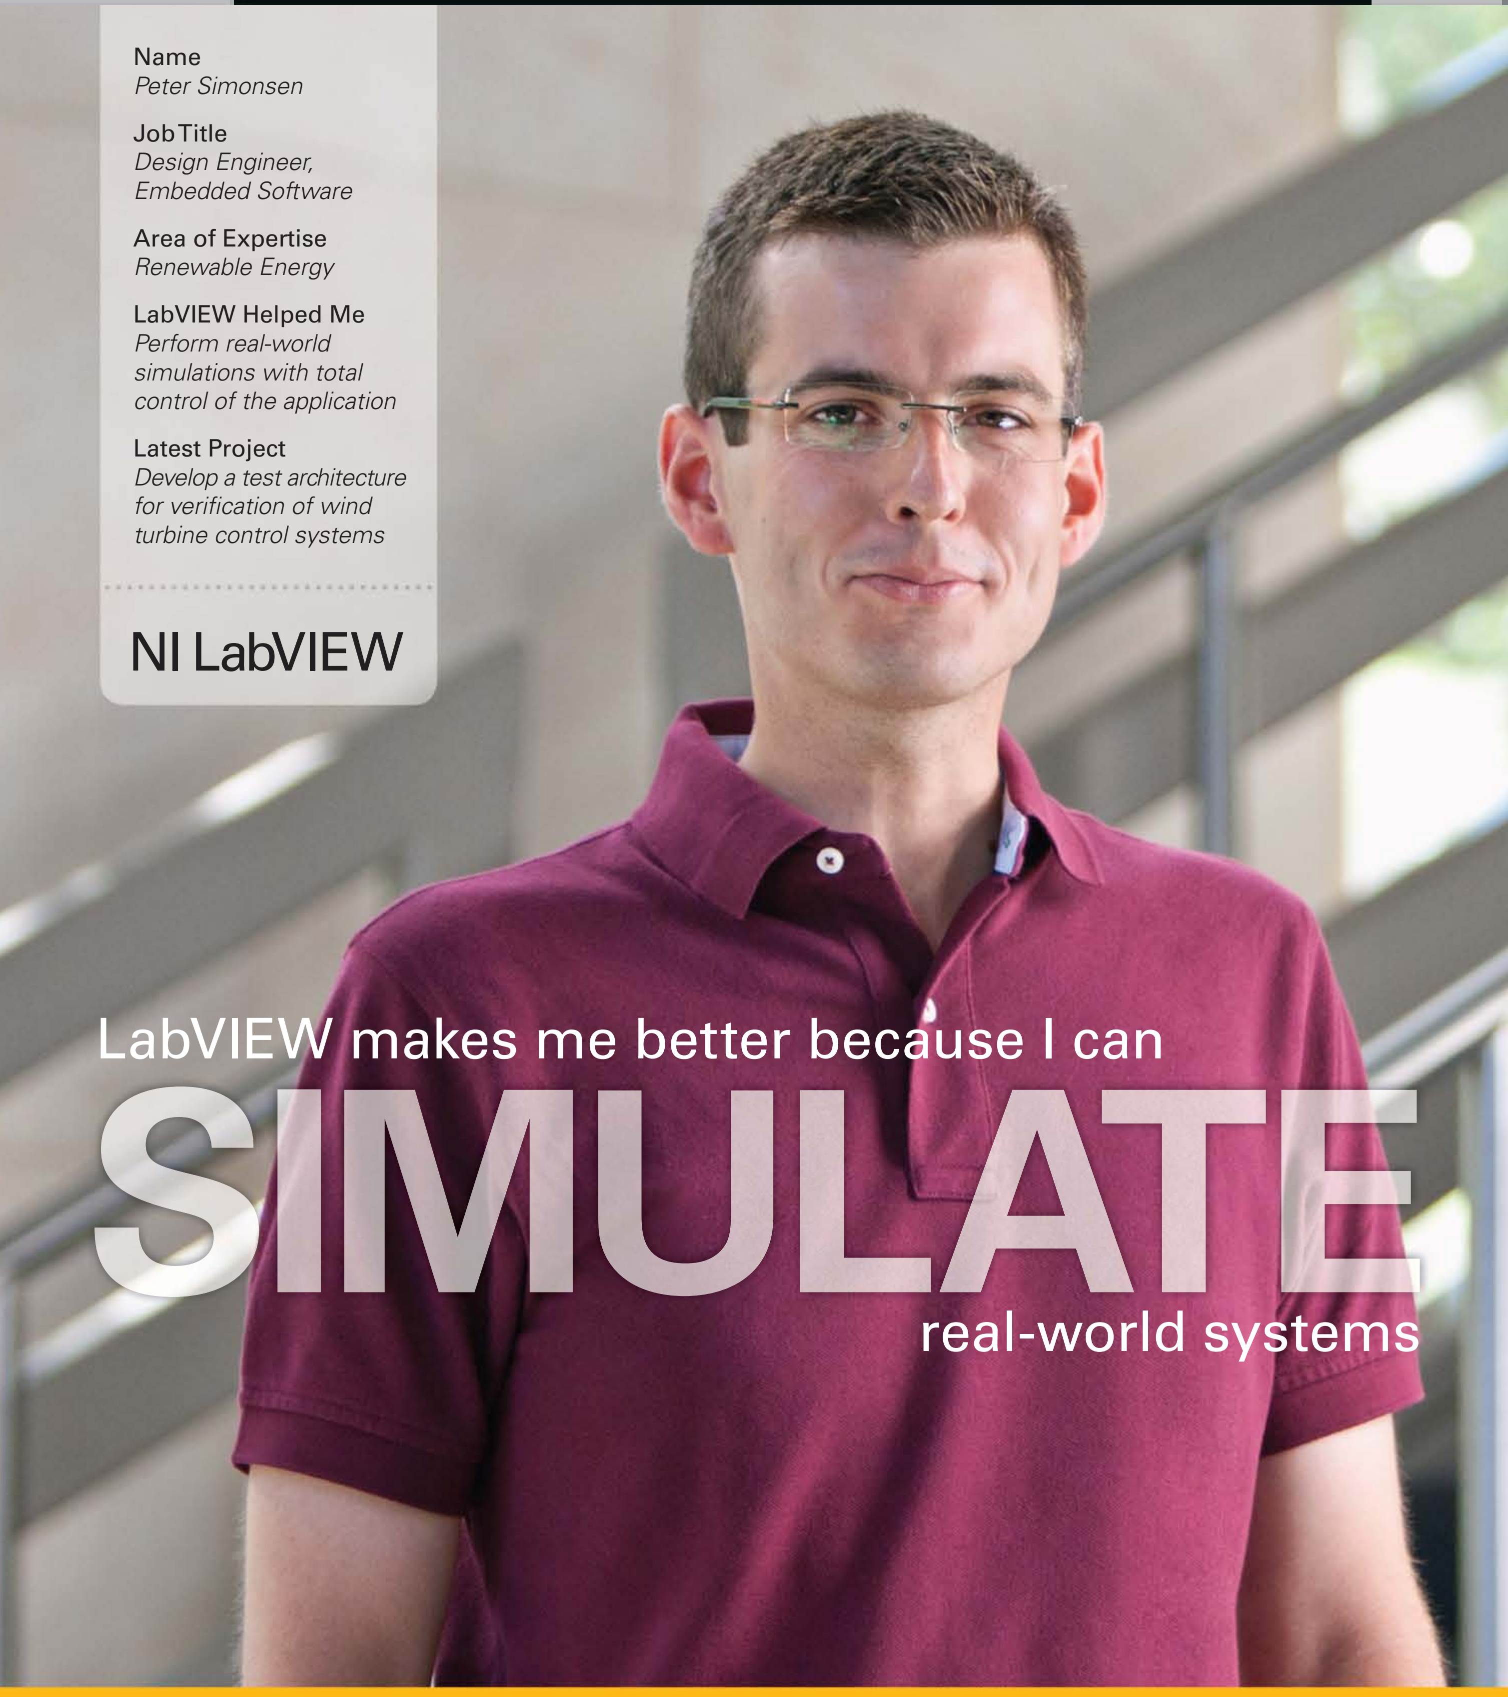
Name
Peter Simonsen
Job Title
Design Engineer, Embedded Software
Area of Expertise
Renewable Energy
LabVIEW Helped Me
Perform real-world simulations with total control of the application
Latest Project
Develop a test architecture for verification of wind turbine control systems
NI LabVIEW
LabVIEW makes me better because I can
SIMULATE
real-world systems

volume 49 number 7 international

ieee spectrum

7.12

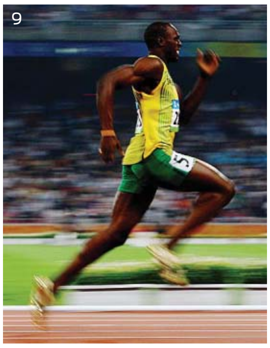

9

SPECTRUM.IEEE.ORG AVAILABLE 1 JULY

Robots and Human Evolution

What happens when you put successive generations of robotic tadpoles in a tank of water with a predatory robotic sea turtle? You find out why fish have the number of vertebrae they do and why they have two flippers instead of four.

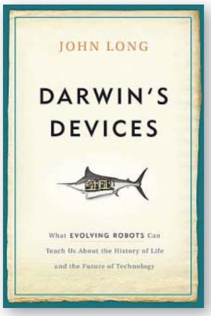

In a recent "Techwise Conversations," host Steven Cherry talked with biorobotician John Long about his new book, *Darwin's Devices: What Evolving Robots Can Teach Us About the History of Life and the Future of Technology.*

THEINSTITUTE.IEEE.ORG AVAILABLE 9 JULY

LIFE SCIENCE NEWSLETTER DEBUTS
Want to learn more about implantable neural interfaces, medical equipment standards, or the role that machine learning plays in computational biology? Then check out the new monthly *IEEE Life Sciences Newsletter*. Launched in April, the electronic publication aims to promote collaboration between researchers and practitioners in the fields of life sciences, engineering, and computing.

VOLUNTEER IN THE SPOTLIGHT
IEEE Foundation president Leah Jamieson talks about her role in creating the public visibility initiative, forming the New Technology Directions Committee, and more.

KEEPING MEDICAL DEVICES HACK FREE
Two IEEE Fellows are part of a research group determined to protect wireless medical devices like pacemakers, insulin-delivery systems, and brain implants from getting hacked. The group has developed a prototype firewall to block would-be hackers from tampering with the devices and causing catastrophic damage.

ONLINE WEBINARS & RESOURCES AVAILABLE AT http://spectrum.ieee.org/webinar

10 July: Introducing COMSOL Multiphysics Version 4.3

26 July: Modeling Magnetrons and TWTs With Vorpal

Symbolic Computation Techniques for Advanced Mathematical Modeling

Modernist Circuit Topologies and Transistors for Smart Grid Power Electronics

Embedded Meets Mobility—M2M Considerations and Concepts

Achieving Medical Device Software Compliance Under IEC 62304

ANSYS With Apache: A Chip Package System (CPS) Solution

Smart Energy Technology and Application-Specific Solutions Training

Simulating Smart Materials

The Online Car: How Will Cars Change When Permanently Connected to the Internet?

The Smarter Grid—Engineering—The Path Forward

Transformer and Inductor Modeling With COMSOL Multiphysics

New product release library: http://spectrum.ieee.org/static/new-product-release-library

Master Bond white paper library: http://spectrum.ieee.org/static/masterbond-whitepaper-library

IBM white paper library: http://spectrum.ieee.org/static/ibm-whitepaper-library

National Instruments resource library: http://spectrum.ieee.org/static/national-instruments-resource-library

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2012 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. 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 American Business Media, the Magazine Publishers of America, 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.

CLOCKWISE FROM LEFT: JULIAN FINNEY/GETTY IMAGES; BASIC BOOKS; NICOLAS LORAN/ISTOCKPHOTO

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

back story

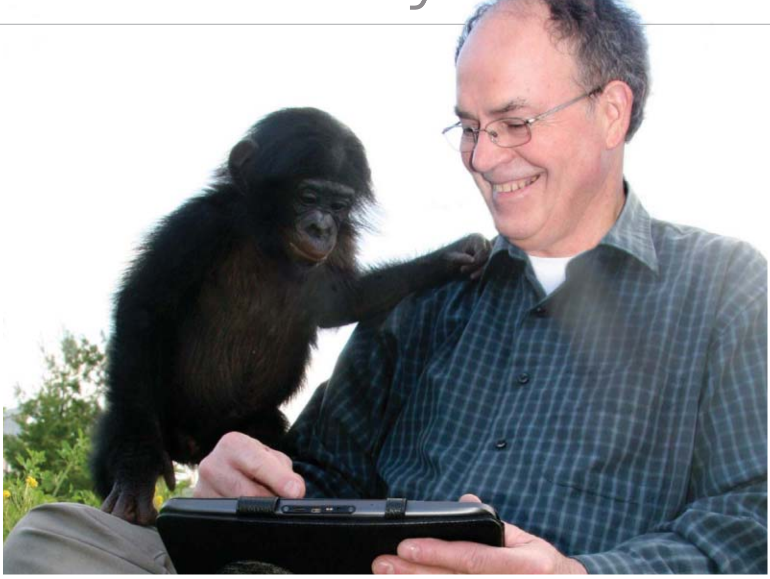

The Magic Word

How do you get an ape to pee? It's a challenge that Ken Schweller never thought he'd encounter when he embarked on a career in psychology and computer science 36 years ago. But Schweller, who spends part of his time at the Bonobo Hope Great Ape Trust Sanctuary in Des Moines, has one of the more unusual jobs in his field: He designs computer software to help the organization's linguistically adept bonobos communicate.

After a tour as an infantry sergeant in Vietnam, Schweller enrolled in graduate school at the University of Illinois at Urbana-Champaign. There, he became interested in the use of computers to study how children acquire language. The subject continues to fascinate him in his main job, as a professor of computer science and psychology at Buena Vista University, in Storm Lake, Iowa.

In 2004, shortly after the sanctuary opened, Schweller volunteered his services. His first task was to program *Pac-Man*–type mazes for the apes.

Initially, he says, the bonobos were wary of him: "Bonobos are matriarchal, so they give women more deference and respect. Men are often viewed as threatening." Eventually, though, they warmed to him, and Schweller became the organization's head programmer; last November he was named chairman of the board.

Probing the apes' language and cognitive abilities offers constant surprises, Schweller says. He mentions a recent incident in which a veterinarian needed a urine sample from a bonobo named Kanzi and wasn't sure how to go about getting it. A researcher told the vet, "Well, you could just ask him." Then she pointed to the lexigrams—the abstract symbols that the bonobos use to communicate—for *please* and *pee*. In short order, Kanzi walked to a corner of his enclosure, found a plastic cup, peed into it, and handed it to the amazed vet.

"I've learned never to underestimate these bonobos," Schweller says. "I'm sure the vet would agree." ❑

LIZ RUBERT-PUGH

CITING ARTICLES IN IEEE SPECTRUM

IEEE Spectrum publishes two editions. In the international edition, the abbreviation INT appears at the foot of each page. The North American edition is identified with the letters NA. 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, The Data is in *IEEE Spectrum*, Vol. 49, no. 7 (INT), July 2012, p. 52, or in *IEEE Spectrum*, Vol. 49, no. 7 (NA), July 2012, p. 60.

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

contributors

MIGUEL MIRANDA writes about the growing problem of semiconductor variability, in "When Every Atom Counts" [p. 30]. Miranda's interest in the issue was sparked about 10 years ago at the research firm Imec, in Belgium, where he joined a team that was trying to tackle various problems in the chip industry. He was surprised to find a big language divide between the design and manufacturing experts on the team. Miranda, who recently became a staff engineer at Qualcomm, says, "It was like taking a plunge in a cold pool."

MARC MOSKO and **VICTORIA BELLOTTI**, who wrote "Smart Conservation for the Lazy Consumer" [p. 24], are researchers at Xerox's Palo Alto Research Center (PARC). Mosko focuses on wireless data network access and routing; Bellotti helps clients build technology businesses around the world. When they're not at PARC, Bellotti moonlights as a jazz vocalist around the Bay Area, while Mosko continues his battle with local authorities to install a natural gas compressor in his 1890s home to fuel his CNG car.

ALFRED POOR is a senior member of the Society for Information Display. At *PC Magazine*, he was a contributing editor for more than 20 years and the publication's first lead analyst for business displays. TV channel choices once ranged from channels 2 to 13, but Poor says that today even "a control knob the size of a Studebaker hubcap" wouldn't suffice. In "Smart TV" [p. 20], he writes about the new world of TV controls that draw on speech and gesture recognition, data mining, and more.

GREGG SEGAL, a regular contributor to *IEEE Spectrum*, has loved being behind a camera since he was 11 years old. For this issue, he traveled to Iowa to photograph the bonobo apes with their apps [p. 36]. Kanzi, Panbanisha, and 2-year-old Teco are able to understand spoken English, respond using touch-screen icons, and even create words, such as *bread cheese tomato* for *pizza*. "It's just kind of jaw-dropping," says Segal.

VACLAV SMIL, a distinguished professor in the department of environment and geography at the University of Manitoba, in Canada, dates the interests that inform his article "A Skeptic Looks at Alternative Energy" [p. 44] to his student days at Prague's Carolinum University more than 50 years ago. "Contrary to the currently popular assertion of accelerating innovation," he says, "most technical improvements are evolutionary and take time to make a real difference. Consequently, I think little of claims of near-instant transformations of any complex system."

JEONG SUH, creator of the charts for "A Skeptic Looks at Alternative Energy" [p.44], is an illustrator and 3-D modeler for Bryan Christie Design. His favorite projects are architectural animations, such as the new World Trade Center rising from the ground he did for *Esquire*. He constructed the realistic 3-D buildings from scratch, detail by detail, based on actual specs. He also loves snowboarding.

CLOCKWISE FROM TOP LEFT: MIGUEL MIRANDA; DANI PALOMINO; FRANK MENTHA; CHRISTINE CHAN; COMPUTER SHOPPER

spectral lines

Solar for All? Not in My Lifetime

On a sunny afternoon in California, I drive past the shuttered factories of Solyndra Corp., the bankrupt maker of solar panels. A billboard-size "For Sale" sign sits prominently on the frontage road of Interstate 880, a main artery into Silicon Valley. Solyndra's empty factories are grim reminders of how—even in America's innovation heartland—life on the solar frontier can be nasty, brutish, and short.

My destination is Sand Hill Road, home to the world's largest concentration of venture capitalists. I meet a gang of them at a swanky bar where they are celebrating (a few days in advance) Facebook's extravagant initial public stock offering. While VCs who are invested in social media cheer their bonanza, those who had bet on solar energy ponder the difficulties of innovating in the energy field.

"We tried to apply the Silicon Valley model to solar," one VC explains. "Moore's Law would make everything go faster and cheaper. We'd become environmental heroes—and rich."

Demand for solar doubled last year, but technologists overpromised and underdelivered. The few solar champions that have avoided Solyndra's fate are firing staff, curtailing ambitions, and conserving cash. The solar panel, once seen as the "microprocessor" at the heart of an energy revolution, turns out to represent only a sliver of the actual cost of assembling and installing solar systems. Government subsidies to makers—Solyndra received US $535 million—and consumers contributed to a glut.

Then China's solar exporters, fueled by cheap capital, overwhelmed their American rivals by providing sturdy and reliable, if not particularly innovative, products, often at below the cost of raw materials, no less assembly and shipping.

"Forget solar," one of the VCs says. "Even the word energy scares me now."

The debacle on the solar frontier confirms anew that innovation in electricity is ridiculously hard—far more difficult than, say, creating Facebook or any number of photo-sharing programs. What ails solar is a dearth of revolutionary innovations—either present or in the pipeline—that would make solar generation, distribution, and storage far less expensive and far more compatible with the existing grid.

The need for more inventiveness is only part of what's holding back solar. Electricity is, after all, the product of a complex technological system that stretches back more than a hundred years. This system, while dependent on great gobs of (carbon-spewing, climate-altering) gas- and coal-powered plants, permits us to plug in and turn on, whenever we want, 24/7.

Electricity is also political. In the Great Depression, the federal government built a vast hydro network in the Appalachian region to bring electricity to some of America's poorest. In the 1950s and '60s, the government aggressively promoted nuclear-generated electricity. Today, reformers insist that modernizing electricity generation is so costly to undertake, and the benefits sufficiently difficult to capture, that ultimately private actors will underinvest in this area.

What's needed to correct the solar market failure is a powerful political movement that imposes a single set of standards for solar generation and delivery across the nation. The resulting economies of scale are impossible today because hundreds of local and national governments insist on differing rules and requirements for the technology.

This kind of regulatory standardization—which is distinct from the technological standardization that organizations like IEEE spearhead—is not going to happen. The United States is pluralistic: Let a thousand solar shoots bloom. We are not China, after all, a country led by a politburo of engineers who can impose a technological system on a billion people.

As with electric cars, solar as a major source of U.S. electricity is at least another generation—technologically and politically—from coming of age. This technology will not be ready for prime time during my lifetime (I'm 56). Solar will be a small player for the next 20 years: an expensive novelty that attracts enthusiasm, talent, and some investment but fails to deliver at scale.

—G. Pascal Zachary

G. Pascal Zachary is a professor of practice at the Consortium for Science, Policy and Outcomes at Arizona State University and a frequent contributor to IEEE Spectrum.

Correction

In "Top 10 Tech Cars 2012" [April] we wrote that the Mazda CX-5's direct fuel injector operated "under a pressure of 204 kilograms per square centimeter." We should have used the standard SI units for pressure, which would result in a measurement of either 20 megapascals or 20 meganewtons per square meter.

DAVID PAUL MORRIS/BLOOMBERG/GETTY IMAGES

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

update

more online at spectrum.ieee.org

The On-Demand Olympics

Engineers gear up for history's largest broadcasting operation

THE LAST time London hosted the Olympics was in 1948, the first year the games were broadcast to home televisions. Lauded at the time as a huge technological achievement, the operation would be considered quaint by today's standards.

As the sole broadcaster, the British Broadcasting Corp. shot the competitions with three then-state-of-the-art Emitron cameras wired to two control vans parked outside the famous Empire Stadium and its nearby pool (now Wembley Stadium and Arena). From the vans, engineers relayed pictures through coaxial cables strung along telephone poles to a central studio, where a producer chose the best program to pass on to the BBC's only transmission station, at Alexandra Palace. The broadcasts, totaling 68 hours 29 minutes, reached just 500 000 viewers—most within a 65-kilometer radius.

Sixty-four years later, broadcasting the Olympics from London couldn't be more different. By the start of the games this month, thousands of engineers, technicians, and producers from 148 broadcasting organizations will have descended on London in anticipation of the world's largest broadcast operation to date. Every event, for the first time, will be shot in high definition and streamed live over the Internet. Nearly 5 billion people will watch the games on home televisions, PCs, tablets, and smartphones.

Delivering such diverse and voluminous coverage is no small feat. The host broadcaster, Olympic Broadcasting Services (OBS), which is responsible for providing neutral coverage of all events, will do the bulk of the camera work. In total, OBS will shoot roughly 5600 hours of footage using more than 1000 HD cameras. The feeds will flow from venues across the United Kingdom over a vast network of fiber optic cables and converge at the International Broadcast Centre inside Olympic Park. "You're talking hundreds of gigabits per second coming into the IBC," says Tim Boden, a technical director at BT Global Services, the company that designed and operated the communications infrastructure for the London games. OBS production crews will then relay the data feeds to broadcasters, at the IBC and around the world, that have purchased the rights

EYES OF THE WORLD: If Usain Bolt breaks another record, you can watch it and most other Olympic events live on a smartphone.
PHOTO: JULIAN FINNEY/ GETTY IMAGES

update

Olympics Everywhere

How you will see the London Olympics

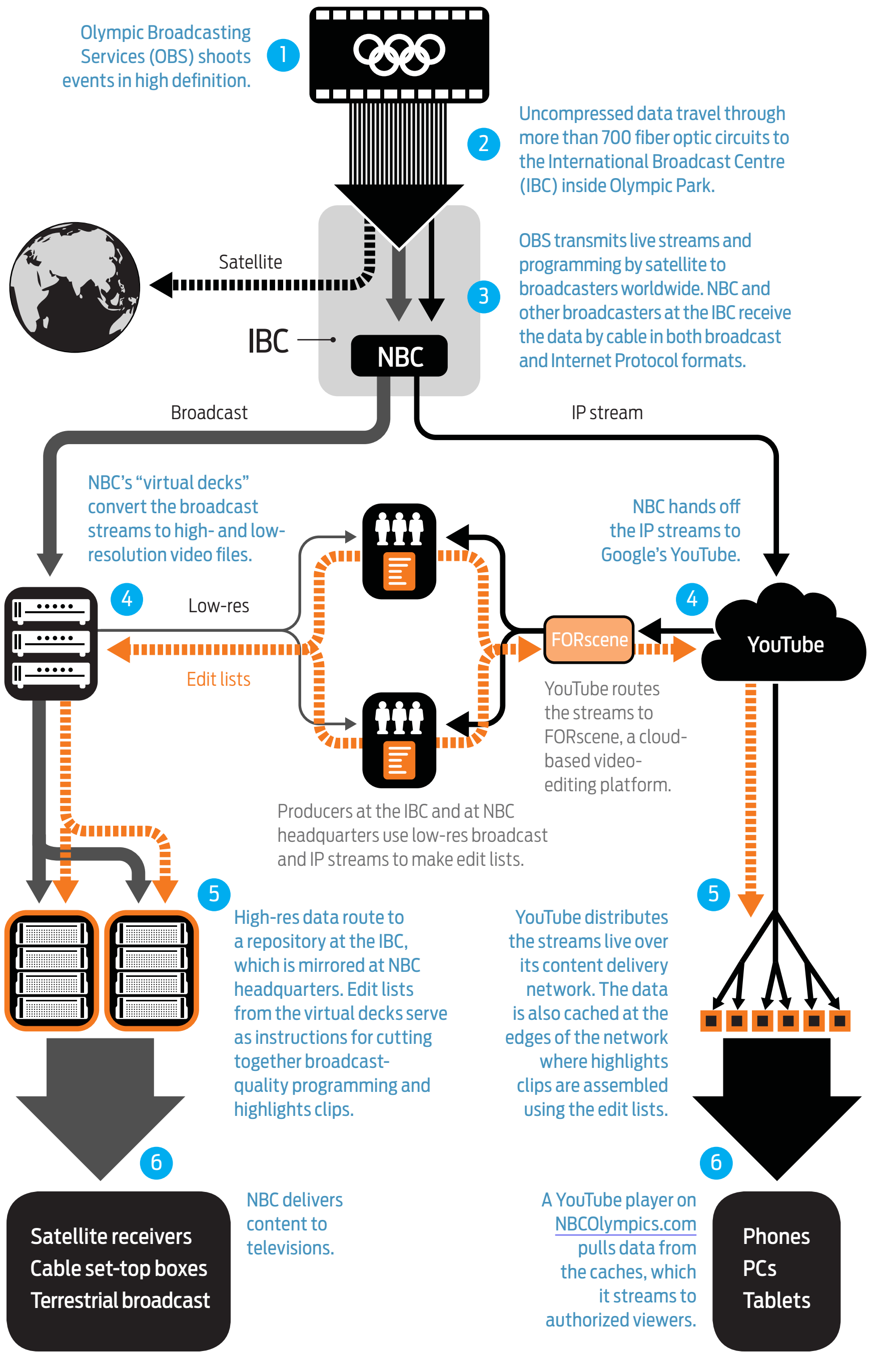

EMILY COOPER

to screen the games in their home countries [see "Olympics Everywhere"].

In addition to distributing Olympics coverage at an unprecedented scale, some broadcasters are debuting new production technologies. OBS, for example, will shoot some events, such as synchronized diving and gymnastics, in 3-D. And camera crews at the velodrome, basketball arena, and certain other venues will shoot in Super Hi-vision, a next-generation television format offering 16 times as many pixels as HD.

The U.S.-based National Broadcasting Co. (NBC) is also trying out new equipment. During past games, producers recorded incoming footage onto optical disks, which they manually loaded into edit decks to cut together programs. If multiple producers wanted to cut from the same footage, they had to record onto multiple disks. "In Beijing, we had hundreds of tape machines all recording the same thing," says Darryl Jefferson, NBC's post-production director. "It was extremely redundant and expensive." This year, NBC has replaced the physical decks with virtual ones. The new decks allow producers to view low-resolution versions of the streaming footage while high-resolution files are simultaneously written to a database. "Anyone can log in to the management system and start cutting in real time as the file is growing," Jefferson explains.

These Olympics are also the first time that NBC and the BBC will stream all Olympic events live on the Web. In partnership with Google, NBC will use the Internet company's YouTube player and content-delivery network to deliver Internet Protocol versions of its broadcast streams to viewers on computers and mobile devices.

"When you're watching the Olympics, it seems like a seamless and singular process," says Shujaat Ali, director of digital services at NBC. "But there's hundreds of pieces that have to work together. The breadth of what happens over 19 days is just incredible." —Ariel Bleicher

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

"By 2020, most computing and storage will be in the cloud, and we'll stop calling it the cloud."
GOOGLE ENGINEERING DIRECTOR PETER MAGNUSSON AT THE IEEE TECHNOLOGY TIME MACHINE 2012 SYMPOSIUM

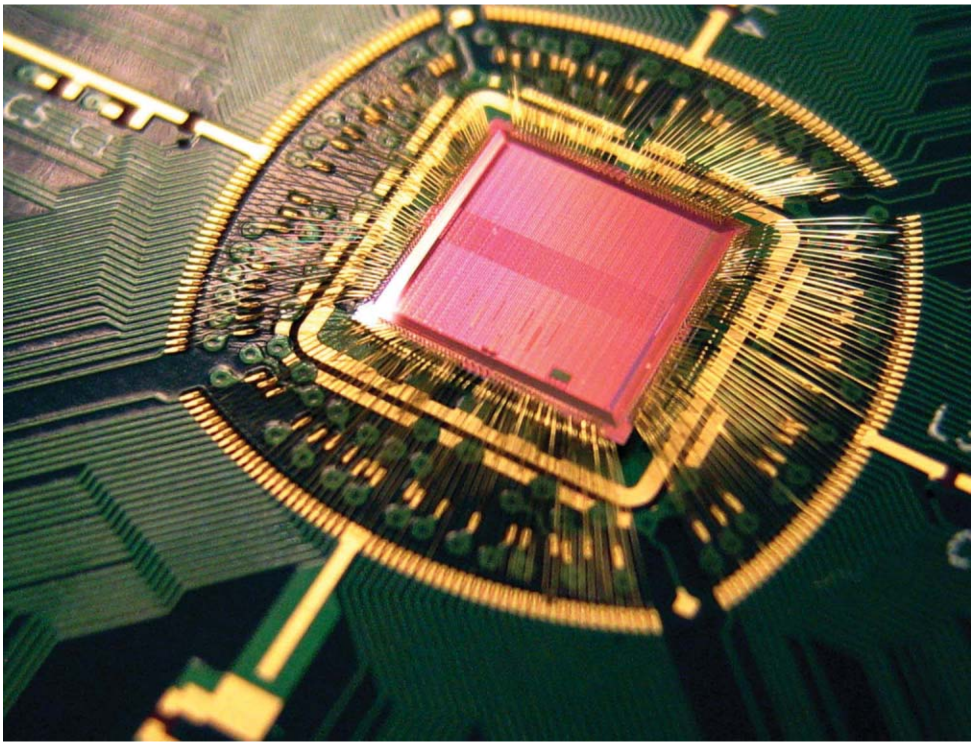

TICK TOCK: Cyclos Semiconductor tested its clock-energy recycling technology on an ARM processor.
PHOTO: CYCLOS

Good Timing

A method for recycling clock energy is making its way into high-performance chips

A LONG-STUDIED strategy for recycling the energy used to clock computer processors is now starting to make its way into commercial chips. The approach—dubbed resonant clocking—has been integrated into Advanced Micro Devices' new AMD A10-4600M processor. The chip, known by the code name Trinity, debuted in May and is now shipping in desktop and laptop computers.

Taking a cue from the pendulum, the technology uses inductors to periodically store and release the energy that would ordinarily be dissipated through the clock distribution network on a chip. According to AMD and Cyclos Semiconductor, which licensed the technology to the processor firm, the approach can cut down on the power needed to clock the chips by up to 24 percent. Alternately, the technology can be used to boost the top speed of the 4-gigahertz-capable chips by 100 megahertz without requiring more power.

Finding ways to save clock power, which can consume 30 percent of a high-performance chip's energy, is tricky. Chips need the steady heartbeat of a clock to drive computation. And the clock's signals must be distributed widely, to tens of thousands—or even hundreds of thousands—of locations.

Chipmakers have traditionally used "trees" to distribute clock signals. These employ many branching wires and networks of small circuits to pick up and propagate the clock signal. But trees don't work well for chips that are meant to operate at speeds of roughly 2 GHz and up, where slight variations in wire lengths, transistor properties, and power supply can significantly alter the time it takes for a clock signal to reach two different spots on a chip. To prevent timing-violation errors, chips are run at less than their maximum speed capabilities.

Designers of high-performance chips have largely solved this problem by adding a wire grid—often called a clock mesh—which shorts out the endpoints of the clock tree and thus minimizes timing variations. But this approach comes at the cost of power. Clock meshes are essentially big capacitors; injecting and pulling charge through all the added metal consumes a lot of energy.

About 10 years ago, researchers began eyeing a seemingly simple fix: Wire up a number of inductors to the mesh. The result is an LC oscillator that naturally bounces between two voltages as it alternately stores energy in the mesh and in the inductors. A chip that uses such an oscillator to make a "resonant clock" can rely on smaller, less power-hungry circuits to drive its clock distribution network. They'll be needed only to prime the oscillator and inject a little energy during each cycle to keep it running, like pushing a swing.

Cyclos, a start-up based in Berkeley, Calif., and Ann Arbor, Mich., has been in stealth mode since 2009, when the company reported a proof-of-principle test showing that it could use a single inductor, located off-chip, to recover as much as 85 percent of the energy in the clock distribution network of a 200-MHz ARM microcontroller chip.

Fully integrating the technology into a high-performance chip presented a few challenges. For one thing, says Cyclos cofounder Marios C. Papaefthymiou, an electrical engineering and computer science professor at the University of Michigan,

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

update

the team had to find a way to incorporate spiral loops of inductors into the already quite dense layers of interconnects that carry power and data around the chip.

"The way the power grids are designed, they're full of loops," says Papaefthymiou. "The inductors are close to power supply lines and close to ground lines, and the closer they get to these lines, the more they interfere with them." The team found a way to minimize interference by making the inductors a bit less ideal, winding them around in a spiral but also up and down, through the two top metal layers of the chip. The team's scheme for AMD's Piledriver CPU design—which forms the heart of the company's new Trinity chips—uses 92 of these 100-micrometer-wide inductors, spread out over each dual-core processor module. The basics of the design were reported earlier this year at the IEEE International Solid State Circuits Conference.

Although AMD saw clock power drop by a quarter, the potential power savings could be much greater, says Matthew Guthaus, an associate professor of computer engineering at the University of California, Santa Cruz. His team has shown that by customizing the size of inductors and their location on the chip, it may be possible to reduce clock power by as much as 90 percent.

Most of the savings would come from reducing the length of wiring between the inductors and the flip-flops at the end of the clock network, which are not distributed uniformly on the chip. Guthaus's team presented such a design approach in June at the Design Automation Conference in San Francisco.

Despite this potential, it's still unclear how far resonant clocking can be extended. Not all chips use clock meshes, and inductors can consume a lot of real estate. What's more, although a resonant clock doesn't have to be run directly at the resonant frequency of the LC circuit, its efficiency as an energy recycler goes down when the clock is run significantly faster or slower. That reduces the utility of one tried-and-true power-saving strategy—temporarily running chips at slower speed.

"Driving far from resonant frequency won't save power, and at some point [the circuit] won't work at all," says Phillip Restle, a member of the research staff at the IBM Thomas J. Watson Research Center in Yorktown Heights, N.Y. Cyclos has inserted a switch that allows the AMD chip to turn the resonant part of the clock on and off. But Restle, who performed some of the earliest work on resonant clocks in microprocessors, says the switch isn't a perfect fix because it adds to the power a chip consumes.

Still, Restle says this isn't the last we'll see of this clocking approach. "I consider this early days for resonant clocking," he says. "There are more rabbits to pull out of the hat."

—Rachel Courtland

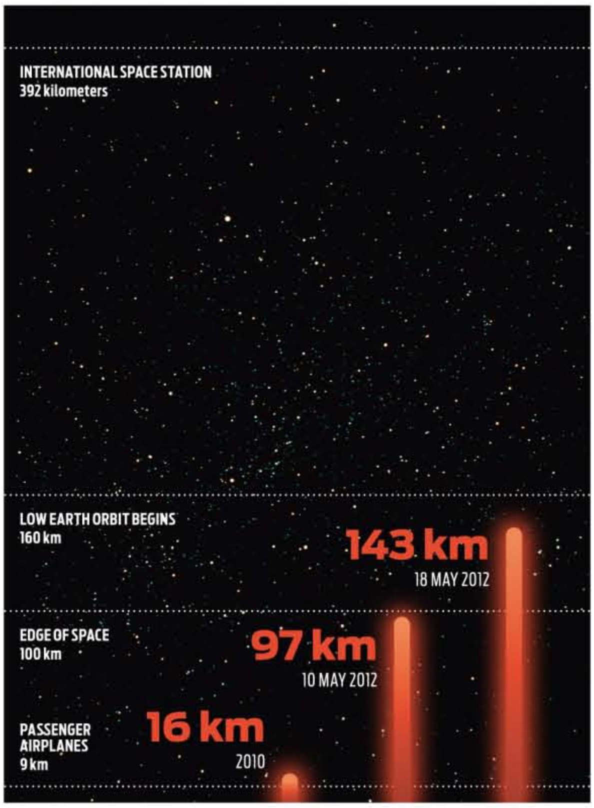

Beam Me Up

European scientists led by the University of Vienna's Anton Zeilinger say they've set a new record for the mysterious but potentially important phenomenon of quantum teleportation. The basis for a communications scheme that can't be hacked, quantum teleportation reproduces the quantum properties of one particle—a photon in this case—on another, distant particle. Scientists have been in a race to find out just how distant they can make these particles. Progress has been rapid, and the prize is nearly within reach: If you can establish a quantum connection between the ground and an orbiting satellite, you can, in principle, establish an impenetrable, globe-spanning network.

GETTY IMAGES

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

"This was more success than we had a reasonable right to expect."

ELON MUSK, FOLLOWING THE SAFE RETURN OF SPACEX'S DRAGON CAPSULE, THE FIRST COMMERCIAL SPACECRAFT TO DOCK WITH THE INTERNATIONAL SPACE STATION

Fixing Wind Power's Bat Problem

Turbines kill hundreds of thousands of bats each year, but new technology could drastically cut the toll

FACT ONE: Wind energy is booming around the world. Fact two: Wind turbines kill lots of bats. Fact three: Bats are worth billions to North American agriculture. Conclusion? Wind power has a bat problem. And with the Global Wind Energy Council predicting that worldwide wind capacity will double by 2016, the problem will only get worse. At a handful of U.S. sites, such concerns have already led planned wind facilities to be scaled back or put on hold.

This summer, engineers and bat biologists are coming together at a wind farm in Wisconsin to field-test a potential fix. They'll attach ultrasonic microphones to four or five turbine nacelles to record the squeaks and clicks bats emit for navigating and locating prey.

Based on the data collected there and at 40 other wind installations, software developers will create a predictive model that also factors in meteorological information like wind speed, temperature, and precipitation. The model will yield an indication of the risk to bats at the site at any given time. When the risk is high—meaning there are likely many bats present—the utility operator will be able to shut down the turbines and then bring them back up when the risk is low.

"This project is really focused on trying to reduce bat mortality at wind farms while at the same time maximizing electricity production," says John Goodrich-Mahoney, a senior project manager at the Electric Power Research Institute (EPRI), the sponsor of the field test.

Without question, he notes, these are bad times for bats. A fungal infection known as white-nose syndrome has killed at least 6 million bats in North America since it was discovered six years ago; some species may go extinct as a result. The economic impact of losing so many insect-eating animals is staggering: A study published in *Science* last year estimated that bat deaths could lead to annual agricultural losses in North America of more than US $3.7 billion.

The devastating disease has in turn focused public attention on other threats to bats, including wind farms, says Susan Schumacher, a biologist with We Energies, the Wisconsin utility that will deploy the bat-detection system at its 90-turbine Glacier Hills Wind Park. The system, called Remote Bat Acoustic Technology, or ReBAT, was developed by a Gainesville, Fla., subsidiary of Normandeau Associates. In addition to analyzing the recordings in real time, it will also identify species by their calls, so that if an endangered species is present, the utility can take extra precautions.

MICHAEL DURHAM/MINDEN PICTURES/BAT CONSERVATION INTERNATIONAL

Other approaches are also being considered. Several years ago, a project sponsored by Bat Conservation International, in Austin, Texas, looked at whether ultrasonic "boom boxes" mounted on wind turbines could ward off bats. One hitch is that, much like the high-frequency bat calls themselves, the deterring signals don't travel far. Engineers are now readying the equipment for further field tests.

Another study, by Barry Nicholls and Paul Racey at Scotland's University of Aberdeen, considered radar as a bat deterrent; that work was sparked by Nicholls's observation that bats seem to avoid air-traffic control radars. The research stalled after Racey retired in 2009, but it's being picked up this summer by biologist Emma Stone at the University of Bristol, in England.

The only other proven way to reduce bat deaths is to curtail operations at night during migration season, by either shutting down the turbines altogether or raising their cut-in speed—the minimum wind speed at which the blades start turning. But both methods reduce the amount of electricity a facility can generate, by about 11 percent and 1 percent, respectively. EPRI's Goodrich-Mahoney says the Glacier Hills approach should cut generating capacity by a smaller amount, although it will depend on site-specific conditions.

Meanwhile, bat experts worry that present mortality rates might not be sustainable in the long run. But they can't know for sure because they don't have a good handle on how many bats there are. "If you ask 10 experts, you'll get 10 very different answers," says Crissy Sutter, a bat biologist with Normandeau Associates. "Getting some sort of population base is a critical next step."

—JEAN KUMAGAI

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

update

The Quest for 2-D Silicon

Silicene—the silicon analogue to graphene—could have amazing electronic abilities

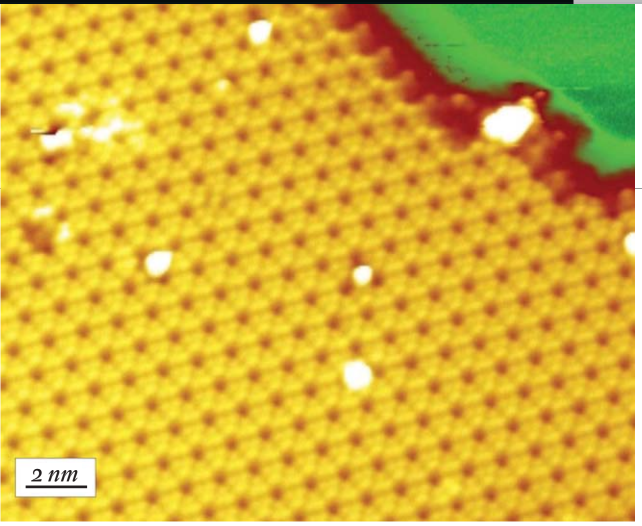

SILICON HONEYCOMB: Two-dimensional silicon forms hexagons on silver. IMAGE: THOMAS BRUHN

IN 2004 two researchers, Andre Geim and Konstantin Novoselev, at the University of Manchester, in England, announced the creation of graphene, a new, two-dimensional form of carbon with unparalleled electronic properties. A decade earlier, Kyozaburo Takeda and Kenji Shiraishi of NTT Basic Research Laboratories, in Atsugi, Japan, predicted that a similar structure of silicon atoms—silicene—should exist. Now, a group of scientists in Europe say they've finally managed to create a sample of the stuff. Experts expect that silicene will have some of graphene's amazing abilities—such as allowing electrons to speed through it as if they had no mass—but in an element more familiar to the semiconductor industry.

Geim and Novoselev had isolated graphene in an embarrassingly simple manner: They peeled it off graphite with Scotch tape. The creation of a silicene layer was much more difficult. "There is no equivalent to graphite where you could simply peel it off, and that was the problem," says Patrick Vogt, a physicist at the Technical University of Berlin. He led a team of researchers from France and Italy, who synthesized and investigated the properties of single layers of silicene and reported the results in April's *Physical Review Letters*.

The researchers created the monolayers by evaporating away silicon from a wafer in a vacuum and allowing it to carefully deposit on a silver surface. Silicon and silver atoms do not form chemical bonds, so the silicon atoms are forced to bond with each other and settle into a honeycomb structure, says Vogt. Unlike graphene, which is perfectly planar, the silicene is corrugated.

In 2010, other researchers claimed to have made silicene. A group led by Bernard Aufray at the National Center for Scientific Research–Interdisciplinary Nanoscience Center of Marseille (CINaM), in France, found a silicene-like honeycomb structure in their scanning tunneling microscopic (STM) images following a similar evaporation experiment.

But the interatomic distances they observed were viewed by some scientists to be too small to be ascribed to silicene. And there was another possible problem, according to Vogt. "The silver surface can mimic a honeycomb structure because of its interaction with the microscope," he says.

In order to prove that they were really dealing with silicene, Vogt, along with Guy LeLay from CINaM and Paolo De Padova from CNR-ISM in Rome, resorted to an additional experimental technique. They made STM observations of samples at different stages in the formation of the silicene layer. "We could see these silicene islands grow, and we could measure the height of the forming layer," he says.

Just like graphene, silicene is expected to be a good conductor because electrons can travel through it quickly, as if they were massless. However, critical measurements of the electronic properties are not yet possible. One of the problems is that the silver substrate on which silicene is grown is itself a good conductor. "This makes it very difficult to understand what comes from the silver and what comes from the silicene. If we had it on a semiconductor or an insulator, that would be much easier," says Vogt, who adds that aluminum nitride might be a possible substrate, an idea he plans to investigate.

For his part, Aufray intends to try to deposit silicene on sodium chloride surfaces, which have a crystalline structure close to that of silver. And his group claims in the April issue of *Journal of Physics: Condensed Matter* that its original work showed silicene as well. Aufray even argues that the layer obtained by Vogt and his team might not be a real silicon equivalent of graphene. "We think that silver also plays a role as a catalyst and could force the silicon atoms in specific positions," says Aufray.

So the jury is still out on whether anyone has truly created a silicon analogue to graphene. And if they have, would it be useful in electronics?

Walter de Heer, who studies graphene at Georgia Tech, thinks it is unlikely because silicene cannot exist on its own without a support. "It deforms very easily and turns into a drop of silicon," he says. "As compared to graphene, silicene is quite unstable."

However, some graphene researchers are open to the possibility. "Just because 3-D silicon is the most important material available today, we shouldn't jump to the conclusion that 2-D silicon will be equally important," says Aravind Vijayaraghavan, a graphene researcher at the University of Manchester. "On the other hand, it might turn out to be even better!"

—ALEXANDER HELLEMANS

the big picture

FOLDED FOR FLIGHT

Every child's fantasy has now been lived out by engineer Art Thompson and workers at the Pima Air & Space Museum in Tucson. On 21 March, they hooked up a 13.7-meter-long paper airplane to a helicopter, pulled the airplane to a height of 800 meters, and released it over the desert. The plane, dubbed Arturo's Desert Eagle, glided 1.5 kilometers before making a spectacular crash landing. The aircraft got its name (and the inspiration for its design) from one child in particular: 12-year-old Arturo Valdenegro of Tucson. Arturo beat out a bunch of other kids in a January contest held by the museum; his handcrafted paper airplane flew the farthest.

PHOTO: JOSHUA LOTT/REUTERS

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

hands on

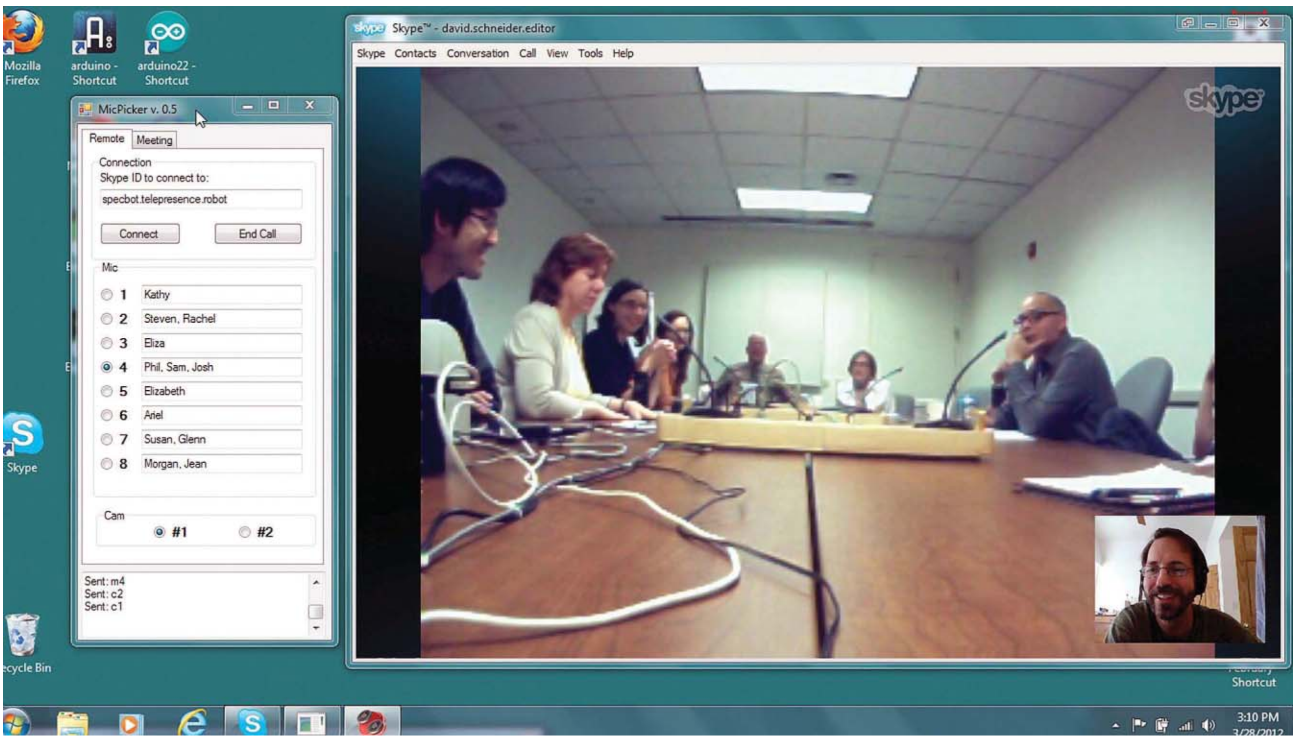

TELECONFERENCING ON THE CHEAP

With Skype and a little hacking, you can teleconference with a room full of people

A S A TELECOMMUTER, I love the Internet, for obvious reasons. Strangely, though, the teleconferencing system that *IEEE Spectrum* editors have long been using is just one step up from a speakerphone. It only provides audio—and crummy audio at that. As a consequence, every editorial meeting I "attend" is a painful experience: Imagine trying to pick out a single voice at a noisy restaurant, in the dark, while wearing earplugs.

A partial solution is to teleport into meetings using Skype. Skype's video capability gives you a window on what's going on at the other end, which of course is nice. But the real advantage is that it provides much higher bandwidth audio than you get over phone lines. That makes picking out voices in a crowd much easier.

Using Skype on a laptop to attend a meeting has some shortcomings, though. One is the webcam's field of view. Don't expect to take in much of the room this way. Another is following individual voices, which will vary in volume and quality with how far away the person is from the computer's microphone. Both difficulties could be overcome by connecting and controlling multiple microphones and webcams. And it turns out that with a modest bit of hacking, you can do just that.

The key is that Skype conveniently provides an application-programming interface (API). With it, another application can access many of the Skype client's functions, including its ability to select which microphone and webcam are active.

HEARING AID: Using Skype to teleconference with a room full of people works much better when you're able to adjust which microphone is live from the far end.
IMAGE: DAVID SCHNEIDER

My first attempt to use the Skype API to pull together a cheap-but-effective teleconferencing system used 10 separate webcams, each with its own microphone. I wrote a Windows program in C#, which I dubbed MicPicker, to send short text messages between two computers during a Skype session. Based on what message is received, this software selects which webcam and microphone the Skype client uses. So with this program running at both ends, a remote attendee can choose which webcam and mic will be live during a meeting, changing that selection on the fly to follow the conversation.

Buying 10 webcams can get pricey, though. I wanted the system to be inexpensive, but I also wanted the webcams to take in as much of the scene as possible. After hunting around, I discovered one that filled the bill: Creative's Live Cam Notebook Ultra (model VF0130). Its field of view is 85 degrees (most are 70 degrees or less), and I was able to purchase a bunch of them for US $10 each on eBay.

This particular model is no longer being manufactured, so buying them en masse, as I did, might be difficult. But it turns out that this is not a problem, for the simple reason that plugging 10 webcams into a computer doesn't work—at least it didn't with the computers I tried. Apparently webcams consume too much computational oomph—three or four seem to be the practical limit.

So I scaled the design down to two webcams, along with eight separate USB microphones. The microphones I purchased (Dynex USB Microphone DX-USBMIC) retail for $26, but they can be had for a lot less on eBay (I paid $8 to $12 each).

Of course, you'll rapidly run out of USB ports if you try hooking eight mics and two webcams into a laptop. But external USB hubs are cheap, too—especially if you get the same ones I did: just under $3 each on Amazon for a seven-port no-name USB hub. These units were probably selling for so little because they come with awkwardly wimpy power supplies. That was no worry to me, though, because I sent 5-volt power to all the USB peripherals using a separate supply (the NES-35-5 supply from Jameco Electronics for $23), which I housed in a small aluminum enclosure (CMC11949-R, also from Jameco, for $12).

To mount everything neatly and keep cable clutter to a minimum, I built four slender wooden platforms to support the microphones and webcams. In an attempt to give them something of a Danish modern look, I smoothly rounded the corners and covered them in white birch veneer. I mounted one microphone at each end of each platform and affixed

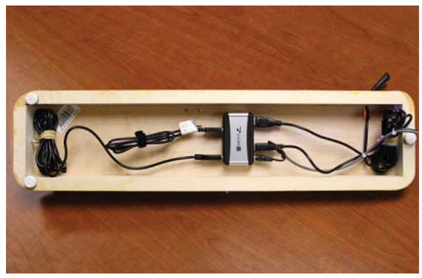

MICS EVERYWHERE: Eight USB microphones and two webcams plug into USB hubs installed within four wooden platforms [top two images]. Speakers, a 5-volt supply, and a netbook complete this inexpensive teleconferencing system [bottom two images].
PHOTOS: RANDI SILBERMAN KLETT

a webcam in the middle of two of them, putting those at the far ends of the conference table with the others spaced between. I stuck one USB hub in the hollow interior of each platform, daisy-chaining one hub to another so that ultimately there was only a single USB cable plugging into the computer, in this case a small netbook.

The webcam's field of vision is wide enough to allow two of them to cover the entire room. In fact, a single one shows me everything except the people who sit behind or alongside it. And with eight microphones, each attendee is always reasonably close to one of them. From those webcam views, it's easy to get a feel for who is sitting where and then switching to the best-situated mic—well, it's easy if you know which mic is which, and that turns out to be the biggest challenge.

With the help of a confederate on-site, I can figure things out easily enough. But if somebody rearranges my wooden boxes between meetings, I get a little lost as to which microphone is placed where. Fortunately, the mics are all sensitive enough, so it's not necessary to choose the optimal one. And trial-and-error switching from one to another soon reveals which mics work well enough to hear different people. (I've even added little text boxes to the microphone-selection radio buttons in the software so that I can easily note who is sitting near each mic.)

It's a hacker's teleconferencing solution, for sure, but it works surprisingly well—gobs better than the telephone-based unit my colleagues and I have long suffered with. And with Skype's new group-calling capability, other off-site editors can also enjoy the new system during staff meetings. My colleagues may gripe about always having to put up with my choices of which camera and microphone is active at any given moment—I have to be in control, because the MicPicker software requires a Windows machine, and the other editors use Macs. But if they want to direct the show, they're free to: They'll just have to study up on Skype's API for OS X and do their own hacking. —David Schneider

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

tools & toys

SMART TV

You won't need an Apple TV, Roku, or DVR box when your television runs its own software

MANY READERS are probably old enough to remember selecting a television channel by turning a knob. With just a dozen positions, it was the picture of simplicity.

Today, you could make a control knob the size of a Studebaker hubcap and it still wouldn't accommodate all our viewing options. Cable and satellite providers routinely offer hundreds of channels of programming and additional video-on-demand (VOD) choices.

In addition, your Blu-ray player or your Wii—or a media player such as a Roku or a Western Digital box—lets you watch movies and other shows from Netflix, Hulu Plus, Amazon Instant Video, and Vudu.

But wait, there's more. YouTube now gets one hour of new video uploaded to the site every second of every day. And that's not just Justin Bieber song videos and kittens doing cute things; YouTube offers full-length movies. How are you going to build a channel dial for all that?

Fortunately, televisions are smarter than ever (the newest Samsung sets boast dual-core processors). The first step, though, is controlling them.

The traditional five-button control (power on/off, channel up/down, volume up/down) just doesn't cut it when you need to navigate a complex electronic program guide (EPG), manage digital video recorder (DVR) functions, and stream content from the Internet. And designers know from bitter experience that just adding buttons does not get the job done.

One solution is to create a context-sensitive controller, which presents only the options that are meaningful at a given point. Making a selection results in a new control display from which you can make additional choices. Some controllers have achieved this using LCD panels lined with "soft" buttons that are reprogrammed depending on the image displayed on the screen. Some controllers now have touch screens, just like our tablets and smartphones.

Better, though, is to skip the remote control and just use those smartphones and tablets to directly control your TV. For example, Verizon Communications now has a unified smartphone app for its FiOS service that lets you manage your DVR settings, view EPG and VOD listings, and even control your set-top box over your wireless network.

Or just skip the controller entirely. The Microsoft Kinect and LG's Magic Wand remote control, to name just two, let you control your television experience by waving your hand or making some other gesture. To be sure, using gestures to type on an on-screen keyboard when searching for a particular show can be tedious. As a result, some TV manufacturers include a full keyboard. (Samsung is again in the lead; its latest models offer not only an optional keyboard but one with a touch-pad control surface.)

The next step is to eliminate even the gestures. LG, Samsung, and others are incorporating speech recognition technology so that viewers can just say what they want, and the television will interpret the command. (Note that you will often see this referred to as "voice recognition," a term best reserved for identifying a specific individual's voice, typically for biometric security applications.) Unfortunately, some systems use a microphone built in to the remote control, so don't think that eliminates the problem of "Where's the remote?"

But none of these interfaces make a dent in the bigger problem of finding the stuff you most want to watch.

An alphabetical or numerical list of all the scheduled and on-demand options is not very useful. So providers like iTunes organize items into genres, such as Action and Comedy, and other groupings, such as Top Movies.

The real power, however, comes when items are categorized across multiple axes. Perhaps you want to watch a recent action movie with, say, Bruce Willis or Nicolas Cage, one released later than 2005. This opens up the whole issue of metadata and the indexing of content. For any given movie or TV program, you might want to be able to search by a combination of title, cast member, year of release, content type, parental guidance rating,

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

DON'T TOUCH THAT DIAL: Samsung's latest TVs can be controlled by a smartphone. Other smart TVs can even be operated by gestures, doing away with extra devices entirely.
PHOTOS: LEFT, AHN YOUNG-JOON/AP PHOTO; RIGHT, THOMAS PETER/REUTERS

language, and more. This rapidly becomes a lot of data, especially for the many movies that fall into more than one category.

That's already a difficult task for the tens of thousands of titles in the Netflix streaming catalog, but what do you do about that hour of video being uploaded every second to YouTube? The data set becomes enormous.

Search is not a complete solution. Back in the old days, we learned all we needed to know from *TV Guide*'s fall preview edition, but that wouldn't get the job done today. We need an interface that also can make recommendations about movies and shows that it thinks we might enjoy.

There are three basic ways that systems are making these recommendations already. The simplest is based on popularity (for example, Top Shows or Most Viewed). There's at least a better-than-random chance that you'll like the titles that are getting the most views overall.

A more refined approach uses social media. If the system knows who your friends are and it knows what they're watching, then perhaps you will enjoy those titles as well. Many systems are experimenting with giving viewers the opportunity to post about what they are watching, with the idea that this may influence what their friends choose to watch and spark conversations about the show.

The most sophisticated approach tracks the choices that you make and then mines an entire database to match as many criteria as possible, finds other viewers with tastes similar to yours, and uses their rankings to suggest titles you might enjoy.

One complication is that the viewing preferences of a given household are not monolithic—the preferences of any two or more family members watching together will differ from their separate predilections. How would Netflix, say, know who's watching? To avoid the requirement of logging into the system each time, Samsung has built face recognition into some of its latest television models. Currently the feature controls only which apps are presented to the user, but it could be used for other functions as well.

Clearly this is a complex problem even for the well-organized virtual shelves of Netflix or Hulu Plus. The complexity grows exponentially when you try to create a universal guide to everything available in scheduled programming, VOD, and all the streaming options available on the anarchic Internet.

At this point, it becomes a Big Data problem, and the companies most likely to solve it are not the remote control makers, the television set makers, or even the broadcast networks themselves. Instead, the solution will most likely come from a company that already has developed expertise in handling enormous amounts of data on both information items and their users.

Amazon.com is one such company, as it already works hard to match shoppers with product choices from its wide-ranging catalog. This collection includes products that Amazon sells, as well as items that exist outside its own universe. The other obvious candidate is Google, which already knows a remarkable amount of information about its users and mines data as if it were a rare earth metal. Other companies could also play a role, such as Netflix, which continues to invest heavily in its recommendation systems.

Whatever the solution will turn out to be, it will not be based on a 12-position gang switch. Rather, it will draw from biometric and other sensors, social networks, and data mining. Your television will start to draw intelligence from the cloud, but it will also have to be pretty smart in its own right. —ALFRED POOR

Correction

In "The Poor Man's Reflow Soldering Oven" (Hands On, April) we stated that "when the [solid-state relay] is energized, it will supply a path to ground, allowing the elements to be powered." The relay should supply a path to the *neutral*, not to ground. A ground connection could result in the oven's metal chassis being charged to the mains voltage, which could mean injury or death.

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

geek life

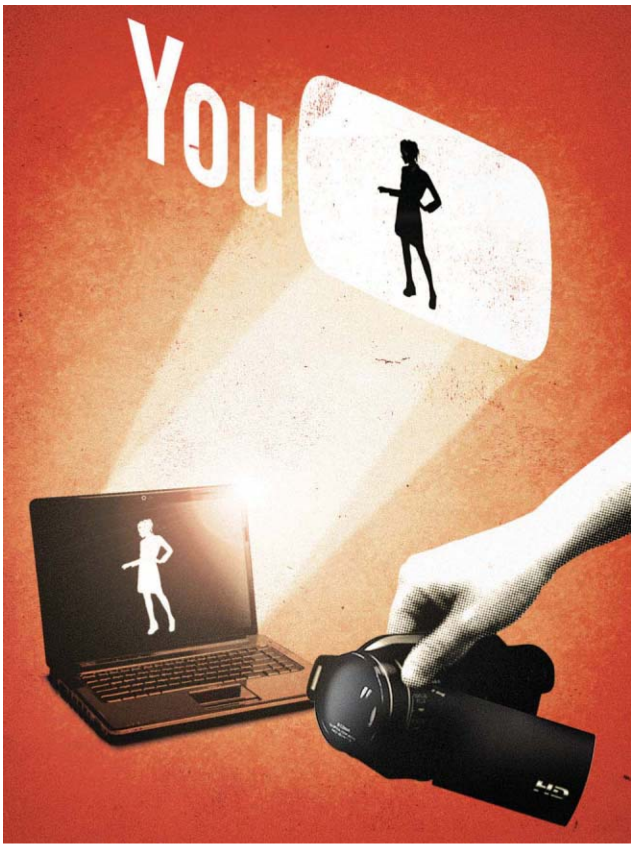

GLUEKIT

YOUTUBE'S YOUR FILM FESTIVAL

Google puts short films on the Web—and Sundance on notice

GETTING INTO a film festival today is like getting into college, only harder. Even Ivy League schools have acceptance rates higher than the 2 percent that's quoted on the Sundance Film Festival's website.

A new online festival of short films—the Your Film Festival, organized by Google's YouTube—may not prove less exclusive than Sundance. But it does drastically cut the cost of screening films. And anyone with a broadband connection can attend. Just go to http://www.youtube.com/yourfilmfestival to watch all the YFF-approved entries, and many rejected ones as well.

The general public was invited to submit entries for free, up to 15 minutes in length, by 31 March. Ridley Scott's London-based Scott Free Production Co., a partner in YFF, culled 50 semifinalists from a field of more than 15 000.

Public voting on the semifinalists, which began 1 June, will close 13 July. A permissive policy—one vote per film per IP address per day—will undoubtedly have entrants pestering friends and family to vote daily.

The top 15 vote-getting filmmakers will be flown to the Venice Film Festival in September, where their works will be screened. The top vote getter, to be announced at the festival, will win a US $500 000 grant to make a film with Scott Free. All 50 semifinalist films will be available on a special Your Film Festival on-demand in-flight channel on Emirates Airlines.

Although YouTube encouraged YFF entrants to post their submissions on the site as publicly viewable and searchable videos, it also recognized that many festivals won't touch a film that has already been released in any way. Thus many submissions to YFF came in either as private video uploads or as unlisted public uploads. (Unlisted entries on YouTube have no keywords, descriptions, or other indexed text and thus are effectively undiscoverable.)

"There is this traditional sense that people have had that there are gatekeepers to the film industry," says Nate Weinstein, entertainment marketing manager at YouTube. "And if you don't put your content through those gatekeepers, then you're squandering an opportunity."

Many filmmakers, though, are happy to have their works viewed. "They've really opened it up to everyone," says Thomas Price, an animator and filmmaker based in Nanango, Australia, who had previously submitted his short film *Paradise* to another festival, which rejected it. "It's a free platform. That's pretty incredible when you think about it."

A wordless story about a kind of stygian journey beyond death, *Paradise* is a visually arresting live-action short by any standard. In its first month on the YFF site—six weeks before YouTube announced its 50 semifinalists—the film had been viewed at least 6400 times—far more than any other first-round film festival submission would have garnered.

Then, too, it's vastly cheaper submitting to YFF than to traditional festivals. Reed Martin, author of *The Reel Truth: Everything You Didn't Know You Need to Know About Making an Independent Film* (Faber and Faber, 2009), says that transferring just a 10-minute HD video to 35-mm film can cost upward of $5000. Uploading the video to YouTube's servers is essentially free.

Martin says that YFF joins a larger continuum of online venues—such as CreateSpace, SnagFilms, and Vimeo—where emerging filmmakers can gain crucial exposure for minimal expense. "What's going to be the catalyst for seeing more of this stuff is somebody who cracks the code of making a TiVo for the Internet," Martin says. "Somewhere you can navigate all the stuff that's on offer and you can realize the dream of the audience of one."

—MARK ANDERSON

reflections BY ROBERT W. LUCKY

Curiosity

A RECENT conversation with a friend turned to the wonders of GPS. "How could we have ever lived without it?" he asked.

I agreed, and began to explain how GPS worked and how the critical clock correction was done. But I quickly saw a blank expression on my friend's face, and an averting of eyes.

"You're an engineer, and you care about things like that," he said dismissively.

The conversation left a strange aftertaste. Do you have to be an engineer to care about how things work? Do engineers have an innate sense of curiosity that is largely absent elsewhere?

Shortly thereafter, I brought up the question with two engineering friends. They looked at each other with an expression of collusion that excluded me. "No," said one, while the other nodded and said in agreement, "We don't have any special curiosity."

Nonetheless, I'm still curious about curiosity. This sense has always been a powerful stimulant for creativity and innovation. I have a mental image of Isaac Newton watching the apple fall. According to a contemporary biographer, Newton asked himself, "Why should that apple always descend perpendicularly to the ground?" Had no one ever been curious about this before? After all, apples had been falling since Eve.

Newton's curiosity led to his law of gravitation. But asking "Why?" is often a recursive exercise, like opening nested dolls. Two centuries later, Einstein asked himself why being in an accelerating elevator is similar to the effect of gravity. His curiosity led to the theory of general relativity. And even today physicists are curious about gravity. How does the apple know that Earth is pulling it downward—or perhaps I should ask, how does space get warped by mass?

Recently I read something that mentioned the Monty Hall paradox. Some years ago I had worked out an explanation in my mind, but now I had forgotten it. I lost sleep over it—I went to bed with my curiosity reverberating and wasn't satisfied until I worked it out again the next morning.

Still wondering if there was a curiosity meme that I could inflict on others, I tried the paradox on a dozen friends. The result of my small experiment was that no one else had any curiosity about it. However, I'm not giving up yet, and I'm about to try it on you, dear reader.

Years ago Monty Hall was the host of the TV game show "Let's Make a Deal." Contestants on his show could win a new car if they guessed which of three doors it was hidden behind. Behind the other two doors were goats. Presumably, they preferred the car to a goat—even if the show and their local ordinances allowed them to take a goat home.

Let's say you chose door No. 1. Instead of directly revealing whether you were right, Monty, who knew where the car was hidden, would open one of the other doors, say door No. 3, to reveal a goat. Thus the car was either hidden behind door No. 1 or door No. 2. So far, so good, but then Monty would offer you a puzzling alternative: "Do you now want to switch your choice?" he would ask.

Apparently, almost no one ever switched from the door they chose to the other remaining one. Among the dozen friends on whom I tried the experiment, no one switched. After all, it still appears an even bet, whether or not you switch. Moreover, the pressure is on, and if you switch and it turns out that the car was behind your original choice, you'll feel like a goat yourself.

However, it can be shown that switching doubles your probability of winning. This seems completely counterintuitive. Surely the probability of winning for either of the two remaining doors is the same. But it isn't. When I told my friends that they had missed the chance of doubling their chances of winning, they all denied that was the case and evinced no subsequent curiosity.

So I wonder. Are you curious about this paradox? ❑

GREG MABLY

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

tech life

BRAIN FREEZE: LG's smart refrigerators, shown here at the 2011 Consumer Electronics Show, can send an alert to your phone if you leave the door open.
PHOTO: DAVID BECKER/GETTY IMAGES

Smart Conservation for the Lazy Consumer

People aren't conserving energy for love or money—you have to trick them into it

If your electric company tells you to cut back your energy use or face the possibility of a blackout, you'll probably comply. You'll turn off unnecessary lights and appliances and use the air conditioner less. But the moment the crisis is averted, you'll quickly return to your old habits. That's because in spite of what people say about wanting to protect the environment and save money, they rarely limit their energy consumption to achieve these goals. Convenience always trumps conservation.

At one time, managers at electric utilities thought they could change consumer behavior by making electricity more expensive at certain times of day. That hasn't worked. The price premium needed to spur conservation seems to be higher than utilities will—or can—charge. That's the bad news. But there's also good news: Although price hikes may not work, clever tricks, it seems, will.

First, let's talk about the problem.

These days, it's not really about efficiency. Energy-efficient appliances have helped reduce overall electricity demand, but a big problem remains. Most people still get out of bed, leave for work, and return home at roughly the same time of day. And—no surprise—their schedules dictate their energy use, which tends to peak in the late afternoon on weekdays, from 3 to 6 p.m. or so, when offices are still open but people are also arriving home from work, turning on air conditioners, TVs, and ovens. The resulting surge in demand can double the amount of energy flowing through the power grid in a matter of minutes.

That surge means that power utilities have to keep an enormous amount of generating capacity in reserve. That's why they need "peaker plants." This extra generating capacity is expensive to build and maintain, so these plants are often old and inefficient ones that utilities would have otherwise shut down years ago. Peaker plants therefore cost a lot more to run than the utility's usual generating plants—as much as triple the cost. Consumers may not realize it, but they can basically blame peaker plants when they complain about high rates. And old plants are not necessarily the most environmentally friendly, either.

So consumers' pocketbooks and the environment would both benefit from leveling out people's electricity use. It should be easy: Two-thirds of residential electricity use in the United States (but somewhat less in Europe) comes from air conditioners, refrigerators, washers, dryers, and other electric appliances. People—or even better, automated controls—can adjust thermostats during the day and run washers and dryers in the middle of the night. Even small reductions in air conditioner use and delays in when appliances get switched on can help control the peaks in energy demand and reduce accompanying costs.

Utilities have for years dangled financial incentives in front of consumers to try to get them to trim or defer usage during peak hours. But taking advantage of those incentives requires awareness and active planning on the part of harried consumers. And that's just not happening. What's needed is automation that makes conservation simple, clever tricks to fool users into saving energy, and, perhaps, more impressive financial incentives, though it's clear that tricks and automation alone would go a long way.

Some residential consumers already let utilities shut off their air-conditioning systems when needed, in return for a credit on their bills. In California, for example, the Pacific Gas & Electric Company (PG&E) offers several such arrangements. The SmartAC program, with more than 100 000 participants, places a remote-controlled device on each home's air conditioner, allowing PG&E to reduce usage during peak summer hours. It cycles the A/C through 15 minutes of normal operation, then 15 minutes on blower only. Consumers receive a reward check—a pittance really, enough perhaps to buy one or two lattes a month.

In Canada about a year ago, the Ontario Energy Board made a well-publicized move to time-of-use pricing, with almost a two-to-one difference between peak and off-peak rates. Many other utilities offer such inducements to try to get consumers to shift their consumption away from peak hours. But these efforts have done little to alter usage patterns. Anthony Haines, Toronto Hydro Electric System's chief executive, believes that the savings are simply too modest. He thinks that peak electricity would have to cost 10 times as much as off-peak energy before utilities would move even 5 percent of their customers in the right direction.

In 2008, researchers from Xerox's Palo Alto Research Center went into homes in Northern California in an attempt to gauge the extent that money mattered to power consumers. Three investigators spent four months on this study, surveying 646 people. They reviewed energy bills, metered appliances, logged daily usage, and conducted interviews in 20 households.

ENERGY HOGS
A typical home uses over 30 percent of its energy for heating and air-conditioning, according to the U.S. Department of Energy. And televisions alone account for more than 10 percent of residential energy demand, according to the Natural Resources Defense Council.
ILLUSTRATION: GAVIN POTENZA

Basically, the study sought to quantify how much ordinary people would tolerate the inconvenience of avoiding, delaying, or reducing power consumption of various home appliances in response to higher electricity prices during certain times of day. The researchers put placards on people's appliances showing what it cost to use that appliance at various times. Electricity rates were cheapest early in the morning and overnight and most expensive in the afternoon and early evening. Participants then logged their usage.

To the researchers' surprise, most of the study participants showed no inclination to save money by modifying their choices of when to use energy. For the most part, it seems, Californians don't care at all about the cost of electricity.

Clearly, the current financial incentives are too small to motivate meaningful conservation. This is particularly true for devices that draw a relatively small amount of power—a laptop, for example. Indeed, in talking with these researchers, some study participants learned that computers were even cheaper to run than they had thought. And so they concluded that computers were not worth turning off to save energy. As one of the participants put it, "You're going to eat when you want to eat. You're going to watch TV when you want to watch TV. You know?" Strangely, several participants mentioned that they would turn off refrigerators or computers if they made too much noise.

tech life

Some also said they avoided using their ovens when the weather was hot. So while people were clearly willing to cut back on power usage to avoid minor discomfort, they were not inclined to do so for the purpose of saving money or conserving energy.

Several study participants did say they intended to reduce their energy consumption and briefly tried to do so. But the inconvenience or discomfort of coming home to a cold house, for example, or waiting for a computer to boot up typically outweighed any desire for cost savings, and they quickly fell back into their old habits.

The investigators who conducted this study noted that the user interfaces of the home appliances encouraged wastefulness, seemingly by design. For example, the default settings on clothes washers, dryers, and dishwashers aren't the most energy efficient. And while most appliances offer more-efficient cycles, users typically stick to default settings if they work well enough. The choice quickly becomes locked into a habit that lasts for the lifetime of the appliance. Users aren't likely to do comparison testing and will rarely bother to find out if a mode that uses less power (using cold water, perhaps, or a shorter cycle time on a clothes washer) does the job as well as the default setting.

This tendency to go with factory settings can, however, have an upside. Designers could simply tweak the user interfaces on appliances to make the energy-saving mode the default—in a sense tricking the user into conserving. Manufacturers could easily relabel dials so that the "normal" option would be the most efficient. If it's more trouble to hit a button for the high-energy cycle, the typical consumer is unlikely to do it.

Although reducing the amount of electricity that appliances draw in this way may be helpful, it won't go far enough. Far more of afternoon peak usage comes from HVAC (heating, ventilation, and air-conditioning) units; ovens, microwaves, and other cookers; and televisions. HVAC systems alone make up over 30 percent of a typical home's energy use, according to the U.S. Department of Energy. And televisions account for more than 10 percent of residential energy demand in the United States.

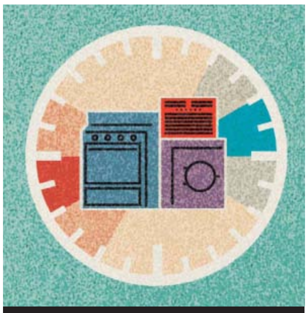

UPS AND DOWNS
Total hourly usage of electricity peaks at about 5 p.m. and drops to its lowest level at about 5 a.m., according to the Electric Reliability Council of Texas.
ILLUSTRATION: GAVIN POTENZA

The low-hanging fruit here are thermostat settings, which could be raised or lowered slightly for short periods of time, in response to real-time pricing or the use of other appliances in the home; temperatures might vary a little more than they do with a constant temperature setting but would remain comfortable.

But you can't expect consumers to manage such fine-grained control. After all, who wants to be constantly tweaking thermostats after a tiring day? More advanced, automated controls are clearly needed. And they are coming. In the next few years, new technologies are going to begin making their way into households, and some of these will dramatically change the domestic-energy picture.

These automated control systems will be connected to the Internet, of course, so they can follow moment-by-moment changes in the cost of electricity. But these systems won't just respond slavishly to real-time pricing data. They will also track what the people in the house are doing or figure out what they might want to do based on their habits. A system that turns off the TV when someone is engrossed in his favorite show is only going to do that once before the user disables it. These new smart-grid systems will have to be, well, smart.

Such technology is just now starting to hit the market. Toronto Hydro Corp.'s Peaksaver program, for example, attaches a smart switch to a home air conditioner that turns it on and off in response to the daily fluctuations in electricity demand. The off times are short, perhaps only 10 or 20 minutes, so the temperature of the home changes little.

Managing the electricity consumed outside of the HVAC system is, however, more complex. Other devices consume far less power individually (clothes washers, clothes dryers, and electric ovens) or have only one power setting when in use (TVs, desktop computers, radios, and speaker systems), or allow only minor adjustments (dishwashers and lighting). To have an impact, multiple appliances must be turned off, turned down, or delayed, at least temporarily. And that's the problem: The large family of appliances in a home represents a very unwieldy ensemble for a consumer to manage.

The solution will have to be far more sophisticated than anything we see on the market today. In these future systems, all these appliances would connect to a device that we call a domestic energy marketplace server. Such a server could turn appliances on and off, having them cooperate with one another to manage the household's overall energy use. Take the dishwasher. Right now, most dishwashers have a delay option: You can load your dishwasher now and set it to run, say, in an hour. So a particularly conscientious consumer

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

ENERGY BOSS: Intel's Home Dashboard shows how technology can help track and control when appliances use energy, but right now, it's just a concept product.
PHOTOS: INTEL

could—though most people won't—set the dishwasher to run after 8 p.m., when air conditioning use would likely diminish. This would help moderate energy demand a little, although not that much.

Here's a better idea. Suppose that dishwasher is integrated with a home's energy marketplace server. The server's software tracks rates and pricing trends, so the system understands that in the summer, the optimum time to run a dishwasher is likely after 8 or 9 p.m. So far, that's not a lot different from an ordinary delay switch, though the delay will be adjusted automatically, rather than by the user.

But say a power plant goes off-line somewhere on the local grid, causing prices to spike. Instead of starting that dishwasher at 8 p.m. as usual, your home energy server communicates with the utility and determines that there would be a considerable cost savings by delaying energy use, at which point it changes the start time. If the dishwasher had already started when the emergency occurred, the server would know whether the dishwasher could be reasonably interrupted. For example, it would stop the dishwasher if it was just heating water, but if it was in the midst of dispensing detergent, it would wait until after the next rinse cycle. Now add an electric oven to the mix. Again, the server would understand when that appliance could be shut down (during self-cleaning, say) and when it can't (while cooking dinner).

If each home in a neighborhood had such equipment in place, the energy marketplace servers in these homes could communicate among themselves to reduce demand within an entire neighborhood in response to pricing signals from the power provider. Having a whole neighborhood cooperate could mean big cost savings for individual homeowners because the neighborhood could work collectively—and still automatically—with the power provider to save energy in response to price incentives.

For the utility, this would not only reduce the costs of power generation, it would also reduce the cost of power distribution. That's because the midvoltage feed lines going into a neighborhood are sized to serve the worst case. All of those homes have air conditioners that cycle on and off at random times. It's possible, in that situation, that most of those air conditioners might coincidentally turn on at the same time, so the utility has to make sure that the lines and transformers that go into the neighborhood can handle that theoretical peak demand. But if energy marketplace servers in homes cooperated to coordinate the on-off cycling of air conditioners, those theoretical peaks—when every air conditioner turns on at the same time—just wouldn't happen, and the utility could build a distribution system for a much lower peak load.

Initially, an energy marketplace server would work directly with utilities, adjusting to the utilities' time-of-use pricing schedules. In the future, though, these servers could save homeowners a lot more money by working with

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

an energy aggregator. An energy aggregator would sign up people who agreed to reduce their energy use on demand at negotiated prices, and then pay consumers for that reduction. Then the aggregator could turn around and sell this demand reduction to a utility at essentially the same price it would cost the utility to buy the equivalent amount of energy on the spot market. Customers would likely make more through this type of spot-market negotiation than they would through standard time-of-use pricing, and aggregators could profit by setting a small differential between the price they charge the utility and the price they pay the customer.

We're not there yet, however. In the meantime, you'll likely see home energy-use controllers that are less sophisticated. But they'll be smart enough to run appliances on a cost-efficient schedule, given the household's electric rate plan, while taking into account user preferences. Some people really like getting their dishes done in time to empty the machine before bed, for example, whereas others don't care.

Several companies have taken the first steps at creating computerized energy dashboards with some limited appliance control. Such controllers adjust temperatures, lighting, and appliance settings based on predicted usage or user input. Some can also display the expected cost of running an appliance at a given time, so customers can choose an optimal schedule. This technology could encourage consumers to develop energy-saving habits.

For example, Intel has developed (but is not yet marketing) a home energy management system. And Control4, a home-automation company based in Salt Lake City, has an energy dashboard already on the market, the EC-100. Earlier this year, California's PG&E began distributing these dashboards to a test group of customers.

Further in the future, systems that get to know you could make adjustments you don't even notice, by monitoring your behavior and learning from it. For example, a domestic power-management system could use video, audio, and other types of sensors to detect what you're doing. By tracking your patterns, it could eventually learn to predict when certain savings may be acceptable. For instance, the system could learn to postpone doing laundry that's never retrieved from the washer until the morning anyway, to turn off heating and lighting in empty rooms that tend to stay unoccupied for long periods once exited, to turn TVs and music players on and off in different rooms as you move around, even to turn down lights when the TV is on and turn off the TV when you fall asleep watching the late show.

And we'd be saving energy without even knowing it—the best way of all.

—Marc Mosko & Victoria Bellotti

COFFEE MONEY
Participants in the Pacific Gas & Electric Co.'s SmartAC program allow their air conditioners to be controlled remotely in return for enough cash to buy a latte or two a month.
ILLUSTRATION: GAVIN POTENZA

When Every Atom Counts

AS TRANSISTORS SHRINK, THE PROBLEM OF CHIP VARIABILITY GROWS

BY MIGUEL MIRANDA

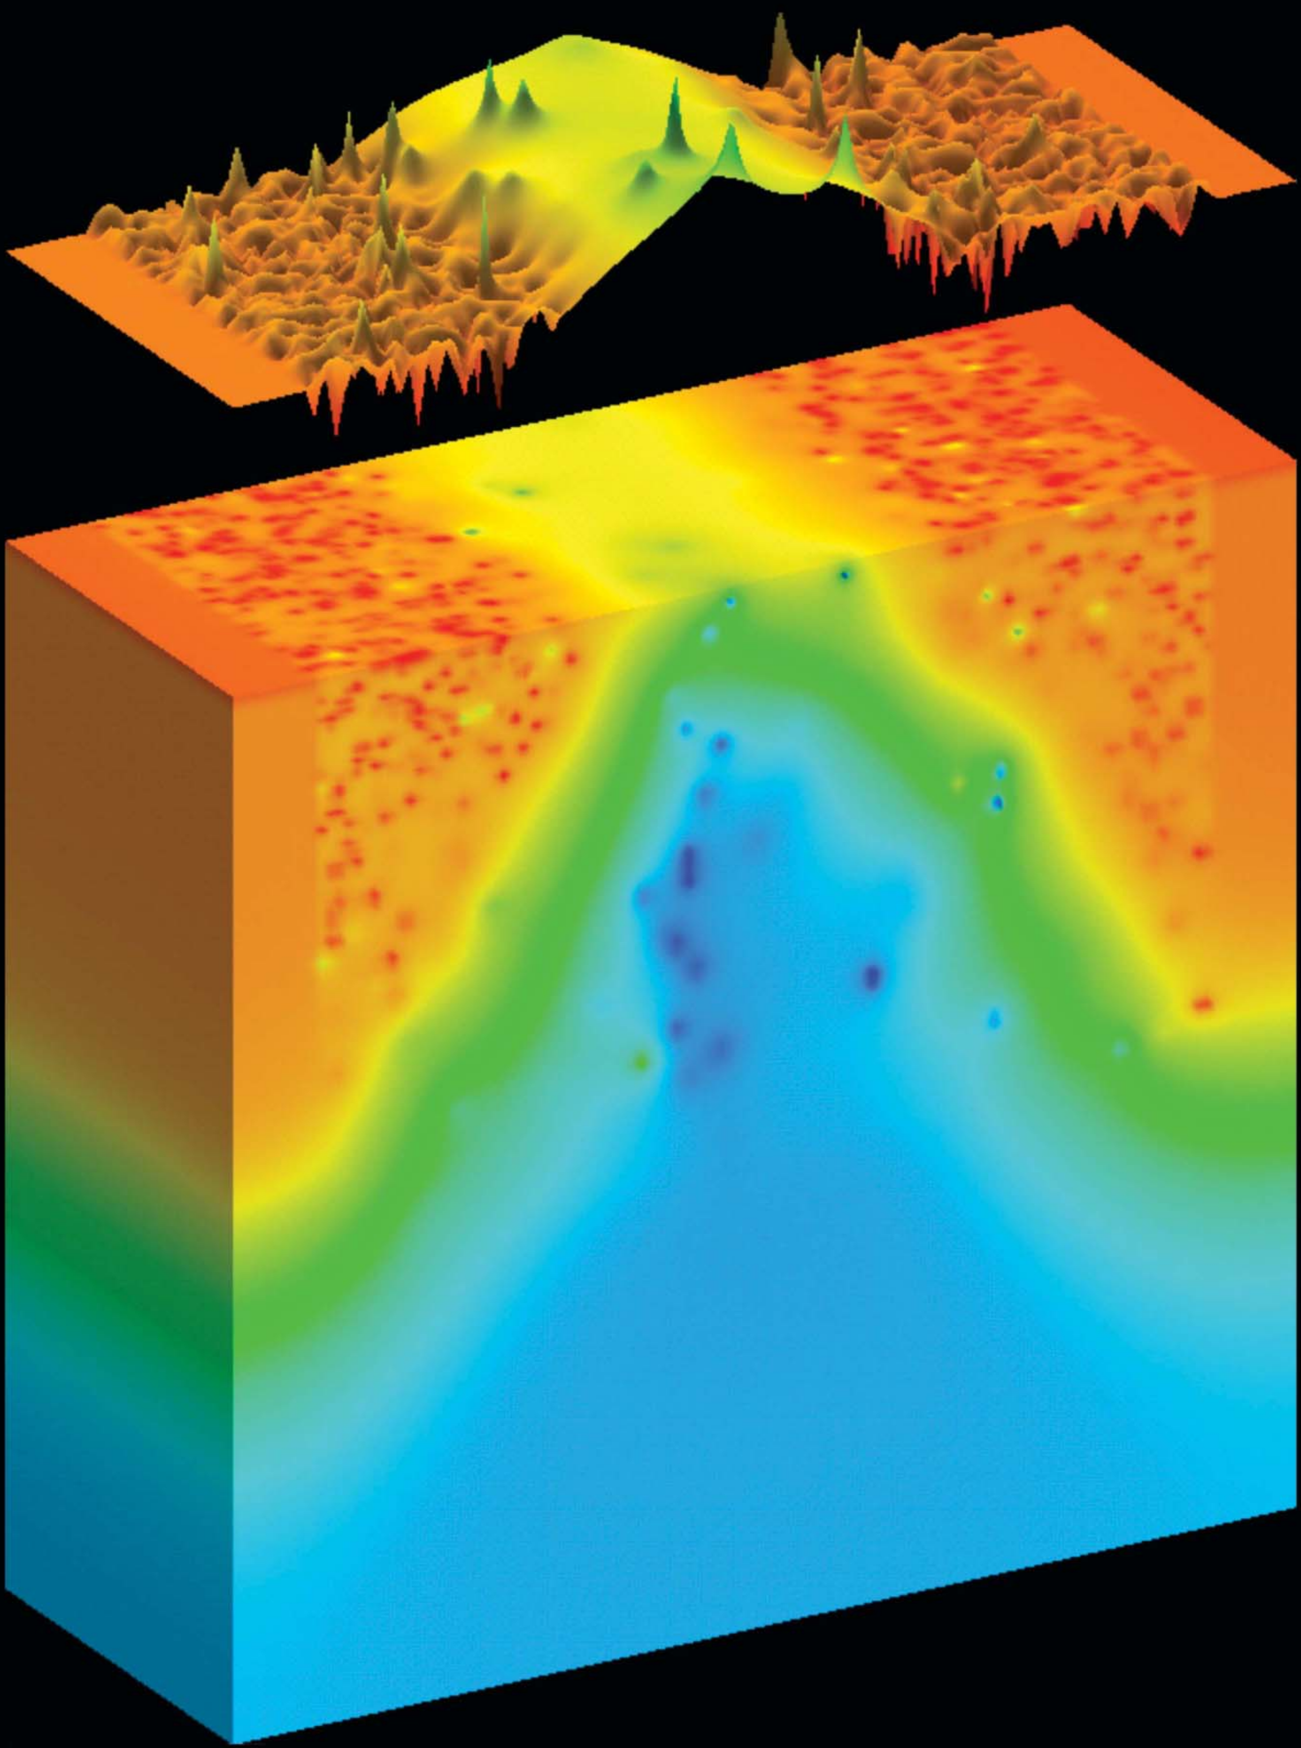

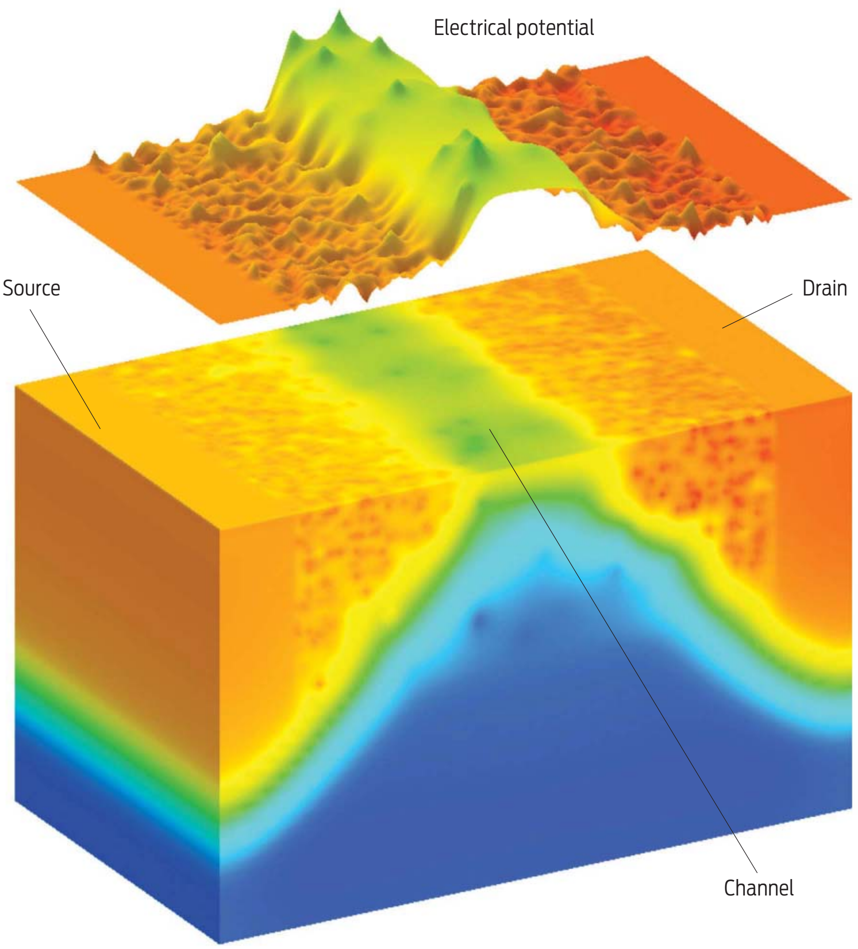

Today's chips are marvels of mass production. Tens of thousands of silicon wafers can move through a single fab in a month, each one carted from tool to tool, blasted by heat, bombarded by ions, immersed in vapor, coated with chemicals, hit with radiation, and exposed to acid. It can take months and hundreds of steps to transform a plain piece of silicon into an array of chips. But at the end of this elaborate assembly line, chipmakers finally get a pile of identical devices that perform just as their designers intended.

At least that's how it used to be. Then, about 10 years ago, chipmakers began to notice a problem: Even state-of-the-art manufacturing processes couldn't produce chips with consistent properties. Nowadays two transistors, fabricated a few dozen nanometers apart on the same piece of silicon, will not have the same electrical properties. It's one of the key barriers that the global chip industry—with sales of US $300 billion—must overcome to keep producing better, faster, cheaper, more energy-efficient chips.

The culprit is scaling. Chips have improved because their transistors and connecting wires have kept getting smaller, but now they're so small that random differences in the placement of an atom can have a big impact on electrical properties. Some batches vary so much that more than half will run 30 percent slower than intended or consume 10 times as much power as they should when on standby.

Some of these defective chips can be sold at a discount, but if they're for application-specific designs—say, for mobile phone communication or video encoding—they might find no better destination than the junkyard. And the defect rate will only get worse as transistors continue to shrink.

Chip variability is what the International Technology Roadmap for Semiconductors calls a "red brick" problem: one of a handful of important issues that lack any clear solution, forming a red brick wall that prevents forward progress.

But just because variability is here to stay doesn't mean we can't mitigate its effects. We could accomplish much by changing the way we design chips. This has traditionally been done by introducing a margin of error to account for the worst-case scenario. Now that ridiculously small defects have entered the mix, that approach no longer works, and chipmakers must overcompensate for the problem. The result is pessimistically designed chips that operate far slower and consume far more energy than they should.

Fortunately, a new family of design techniques promises to predict not only the worst-case scenario for a chip but also the likelihood that the scenario will happen. These approaches use statistical methods to make informed trade-offs between how fast the chips will run and how many good chips a given batch is likely to yield. Some makers of high-end microprocessors like IBM and foundries like the Taiwan Semiconductor Manufacturing Co. are already using some of these statistical techniques in their design flows. Although statistical tools are still far from being widely adopted, if we can push them along, these tools will help us make affordable chips that are as fast and efficient as those the semiconductor road map calls for—and perhaps then some.

VARIATION HAS BEEN PRESENT since the first integrated circuit was created. But its nature and scope—and the way it shows up in the behavior of digital circuits—have changed.

For a long time, most variation was global, caused by slight alterations in the manufacturing process. Such changes differentiate one chip from another or all the chips on one wafer from those on another. Global variations tend to affect the electrical characteristics of all transistors in

PREVIOUS PAGE: GOLD STANDARD SIMULATIONS; ABOVE, CLOCKWISE FROM LEFT: GOLD STANDARD SIMULATIONS (2); IMEC

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

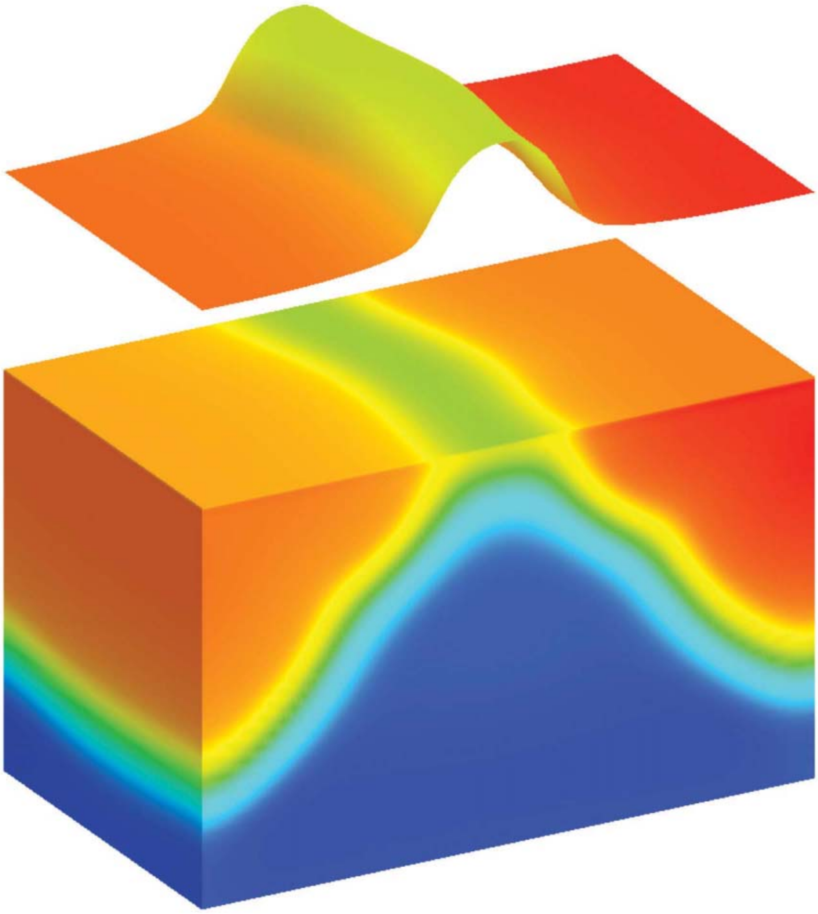

DOTS AND LINES: As transistors get smaller, random fluctuations in dopant location and concentration [left block] and the roughness of circuit features [right block] have a stronger impact on transistor properties. Both factors result in less than ideal electrical potential profiles [shown above each block] .

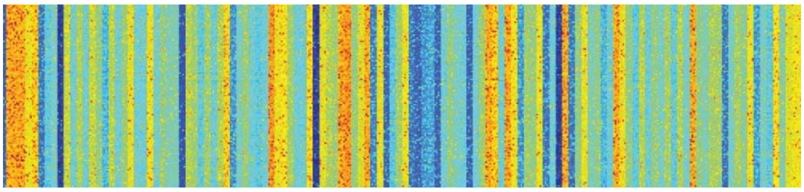

IMPERFECT MEMORY: This simulation of a 32-nanometer, 64-kilobit static RAM device illustrates the impact of local process variation. If there were no variation, all the cells in this graphic would be the same color.

the same, albeit unpredictable, way. For example, the light coming from a lithography tool that's used to print devices can be distorted by slight changes in optics from exposure to exposure, creating transistors that are slightly longer—and thus slower—than intended. Fluctuating environmental conditions can also create variation. If the temperature in a vapor deposition chamber drops too low, for instance, it can slow the growth of the insulating layer of oxide in a transistor. The resulting thin insulator can leave transistors a bit leakier than normal. But over the past few decades, chipmakers have been able to keep global variation under control by steadily improving manufacturing tools and manufacturing process control.

A second source of variation is often called local process variation or process variability, and it is proving far more difficult to address. It started appearing in digital circuits about 10 years ago, when chipmakers began producing transistors with channels less than 90 nanometers long, the span of a few hundred silicon atoms. At that scale, the electrical properties of a transistor begin to be affected by random sources of variation, such as the roughness of a transistor's edges or the granularity in the crystal of the metal electrode that turns a transistor on or off. Such variations have an independent effect on every transistor in any given integrated circuit. One transistor may end up being slower while its neighbor becomes speedier but also leaks more current.

One of the most dramatic sources of local process variation comes from dopants, the atoms of another material that are added to a silicon channel to speed up the switching of a transistor and, by extension, decrease the energy that switching consumes. Chipmakers typically add dopants by accelerating ions to high speeds and shooting them into a wafer. But this approach wasn't designed for work on an atomic scale; it's difficult to control how many atoms make it into a transistor and exactly where they fall. Transistor channels once contained tens of thousands of dopant atoms. Nowadays chipmakers produce transistors that can accommodate only a few hundred of them. And in that case, the absence of a single atom is much more noticeable and can alter how much voltage is needed to turn a transistor on or off by a few percent.

The random, uncorrelated nature of these variations poses a problem for circuit designers. Link up many such transistors in an integrated circuit, with its sensitive dependencies and timing requirements, and the variabilities can magnify one another: The resulting system may be even more randomly variable than its parts. Nor can you accommodate local variation by using hand-me-down tools developed to tackle global variation. We need a new approach.

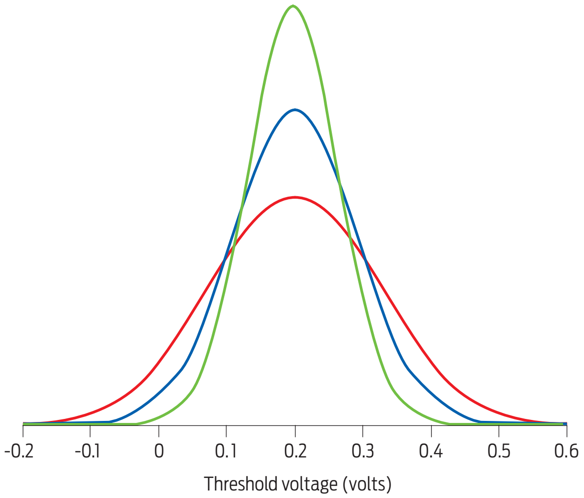

WIDENING VARIABILITY: The range of threshold voltages—those needed to turn a transistor on—is widening as circuit feature sizes shrink from 28 nanometers [green line] to 20 nm [blue] and 14 nm [red]. As feature sizes get smaller, an increasing fraction of transistors are made with threshold voltages at or below zero, which means they no longer operate as switches.

FIRST, THOUGH, LET'S TAKE A LOOK at the long-standing strategy for tackling global variations. This approach relies on models that outline extreme process deviations—all the possible manufacturing scenarios for a chip. Fabs and foundries tend to measure these parameters on the production floor, by deliberately fabricating test chips that are way out of spec in one direction or another. Engineers then measure the electrical properties of the transistors on these extreme chips. The resulting set of specs—which can be represented using a square—is called a corner model.

The four corners of the model can help engineers anticipate how any circuit will behave. In one corner, both the positively and negatively doped field-effect transistors—pFETs and nFETs, respectively—are fast. In the other three corners you find the three other pairings: slow pFETS with slow nFETS, fast pFETs with slow nFETs, and slow pFETs with fast nFETs. Chipmakers and designers assume there is an equal chance that a circuit will fall into any one of these four corners.

This model breaks down, however, when circuits suffer strongly from local process variation. With process variability, every single nFET or pFET may operate either slower or faster than average, and the performance of one device will be completely uncorrelated with that of its neighbor. For most circuits, there will be far more than four possible behaviors, and not all of them will be equally likely. Because the number of combinatorial possibilities effectively grows exponentially with the number of devices in the circuit, the problem would challenge even the largest computing clusters—called compute farms—that chipmakers use to characterize their circuit libraries.

Consider the simplest digital circuit, an inverter, built using one pFET and one nFET. A single inverter would have four possible behaviors and thus a clear worst-case behavior. But combine two inverters in an integrated circuit and you'll either double or quadruple the number of possible behaviors, depending on how they're connected. The problem only gets worse as you scale up. One step up from the inverter in complexity is the NAND gate, which contains two nFETs and two pFETS and could take on 16 possible behaviors. A very simple arithmetic logic unit, the core logic circuitry at the heart of a CPU, can easily contain hundreds of transistors and require the analysis of trillions of combinations, each of which would take a few minutes to calculate. Today's microprocessors easily contain more than a billion transistors, resulting in an effectively infinite number of combinations.

It turns out this exercise in combinatorics is more difficult for some sorts of circuits than for others. For a basic logic gate, it's actually fairly easy to at least predict the worst-case scenario of a circuit. That's because the pFETs and nFETs in logic are only on for a fraction of the time and never at the same time, making the calculations much easier. But that isn't the case for the flip-flop circuitry that's interleaved throughout a processor and used to clock computations or for static RAM (SRAM) memory cells, which now easily take up half the space on a smartphone chip.

EVEN IF IT WERE EASY to figure out the worst-case scenario, we wouldn't be able to estimate how many such bad chips we'll generally get in a batch. We can't even say whether the chance of producing such a chip is one in a billion or just one in a thousand. Yet you need some guidelines of this sort to push the performance of your design as far as possible without reducing the yield to uneconomic levels. In the chipmaking business, there is no more unforgiving measurement than yield. It is what determines your break-even point, the point below which you might as well close shop.

Fortunately, we can get guidelines for yield even before the manufacturing process starts, something that was never really possible with the corner model. Since around 2005, a new class of circuit analysis tools based on statistics has let engineers anticipate an unacceptably low yield and introduce modifications that will make the final manufactured IC more robust, such as increasing the size of the transistors that cause most of the delay and variation in a circuit.

So far, most of the effort in the digital realm has focused on revamping one of the core tools for logic design: static timing analysis. STA predicts the maximum operating frequency of an IC, which is determined by the longest time it takes a signal to propagate through logic gates. Nowadays, however, STA doesn't work so well, since each chip will have its own delay, and some will be able to run quite a bit faster than others.

A statistical version of the STA tools, called SSTA, can be used to calculate not only the theoretical operating frequency of a design but also the probability that the manufactured version will in fact operate correctly at that frequency. This turns out to be fairly easy to calculate, because the delay in logic is almost always a linear operation and the computational power needed to handle the analysis is not excessive.

But SSTA is only as good as the data you supply to it. The tools need timing information based on statistics describing basic transistor variation. You could create such statistics by simply summing the effect of each source of variation on the behavior of a circuit, but the problem with such "sensitivity analysis" is that most sources of variation can't be considered independently. If, for instance, a transistor's channel is longer than expected, then more voltage will be needed to turn the transistor on. As a result, voltage and channel length are linked. If you add the variation from each source separately, you'll effectively be counting the same effect multiple times, overestimating how much local process variation affects circuit timing.

DATA: ASEN ASENOV

The case is particularly bad for clocking circuitry: My colleagues and I calculated that sensitivity analysis overestimates timing delay, as well as two key measures of flip-flop performance, by a factor of two to three. Chipmakers who rely on this kind of analysis will end up overestimating how much they must widen transistors to increase current flow and thus the circuit speed. When the resulting processors are made, they'll run 10 to 20 percent faster than expected, but they'll also carry—and leak—more current than they should. For flip-flops, which are very active, switching every single clock cycle, the design choice will result in a lot of wasted power.

Another technique is CPU-intensive Monte Carlo analysis. This approach uses loops, each of which selects values at random for basic transistor properties and "injects" them into a transistor model, thus simulating the impact of local process variations. Once this is done, an electrical simulation of the circuit is performed. The process is repeated hundreds or thousands of times, with different random values every time. As you might imagine, the problem with this approach is the computer time it consumes. If a logic library is composed of, say, 1000 building blocks or logic gates and every block takes 3 seconds on average to simulate, a Monte Carlo loop of 1000 iterations would blow up the characterization time from the single hour needed to simulate a corner in traditional corner analysis on a CPU to more than a month.

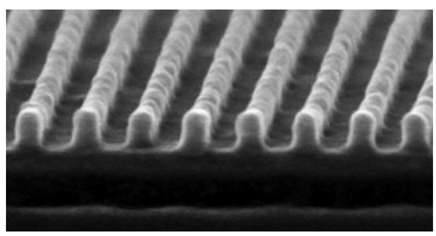

ROUGH GOING: These uneven rows of exposed photoresist, created using extreme ultraviolet light, will become 30-nanometer-long transistor gates.

IMEC

Memory also turns out to be especially difficult to simulate. Statistical models tend to use a single SRAM bit to model the behavior of the entire memory chip. But memory is only as fast as its slowest bit. Bits are read in groups—or bytes—and each group must wait for all its bits to return values. Because a memory chip is basically composed of a grab bag of bits with different properties, it is all too likely that at least one of the billion or so bits will have an extremely long delay time. Although there are tricks that can help make Monte Carlo simulations speedier, statistical tools that can optimize memory are still lacking.

MY COLLEAGUES AND I began to get a taste of how far we have to go—particularly in modeling memory—in 2008, when we began Reality, a €4.5 million, six-institution research project funded by the European Union. We aimed to create something completely new: a variability-sensitive, fully statistical model of a system-on-a-chip (SoC), the combination of memory and processor that forms the CPUs of most tablet computers and smartphones.

To accomplish this task, we combined transistor data provided by project member STMicroelectronics with statistical tools developed at the University of Glasgow and at Imec, in Belgium. This allowed us to model a representative subset of chip designer ARM Holdings' library of circuit building blocks and one of ARM's cores. Along the way, we devised a trick for quickly modeling memory variations in the ARM SoC by assigning random timing and power properties to components on the chip. The project produced more than a few surprises.

Our back-of-the-envelope calculations had led us to expect that the SRAM on the chip would cause only 10 to 20 percent of the chip-to-chip variation in signal-processing speed. It turned out to be closer to 70 percent. This variation has a big impact: If chipmakers use an SSTA model of the SoC that doesn't account for such variation in memory, they can overestimate the maximum operating frequency of a chip by as much as 10 percent. Such a mistake raises the probability of creating bad chips.

The Reality project has also helped identify what it would take to model the variation in an entire SoC. But although my team at Imec has transferred our memory tools for analyzing SRAM to a number of partners, including memory maker Samsung, these tools don't work by themselves. They require a larger modeling framework capable of accommodating statistical design, something that few commercial packages can do and none can do for an entire SoC. Memories can account for up to half the maximum timing delay between input and output on such a chip, which means that right now we're missing half of the problem and thus half of the solution. As it now stands, it's not clear whether such chip-level analysis tools will emerge by 2026, the most distant point now specified in the industry's road map. Chipmakers may simply opt to squeeze by with existing statistical tools for logic.

One project that's looking beyond the end of the current road map is called Terascale Reliable Adaptive Memory Systems (TRAMS). The €3.5 million effort, backed in part by the EU, aims to study the possible effects of variability on some of the switching architectures that might replace the traditional CMOS transistor. These include 3-D transistors built on an insulating layer, devices made with compound semiconductors, and logic based on carbon nanotubes. Some of these devices may be able to tune their electrical properties without dopants, and TRAMS will help determine whether this simplification will ease the variability problem or other sources of variation will crop up to take its place.

Most design efforts aim at avoiding variability. One notable exception is the multi-institution, U.S. National Science Foundation–funded Variability Expedition, which aims to develop ways to adapt to the problem. A promising way to do this is with circuit components that use proactive, on-chip hardware and software to monitor and dynamically adapt to imperfect chips [see "CPU, Heal Thyself," *IEEE Spectrum*, August 2009].

If all goes well, these research efforts will help identify the best way to handle variability. But it is clear that designers can't do the job alone. New devices and circuits will likely be crucial to reducing the impact of variability. And chipmakers may also get some relief from the problem in a few years' time, when they switch their manufacturing process from 300-millimeter-wide wafers to 450-mm versions. Fabs will then be able to produce more chips per hour and will therefore lose less money, proportionally, to out-of-spec chips.

In one form or another, variability will continue to plague us as we approach atom-scale switches. All we can do is make sure we have the tools we need to meet it head-on. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/variability0712

Apes With Apps

Using tablets and customized keyboards, bonobos can become great communicators

By Ken Schweller

Photography by Gregg Segal

CAN WE TALK?
Kanzi, a 31-year-old bonobo, can converse with humans by selecting "lexigram" symbols on his Motorola Xoom tablet.

PROGRAMMING LANGUAGE: When Kanzi presses a lexigram on the touch screen, the computer speaks the word and shows a corresponding picture.

Have you ever watched a toddler play with an iPhone?

Most likely, the child was completely captivated and surprisingly adept at manipulating the tiny icons. Two-year-old Teco is no different. Sitting with his Motorola Xoom tablet, he's rapt, his dark eyes fixed on the images, fingers pecking away at the touch screen. He can't speak, but with the aid of the tablet app I created for him, he's building a vocabulary that will likely total several thousand words. What's more, he'll be able to string those words together into simple sentences and ask questions, tell jokes, and carry on conversations.

Such talents wouldn't seem exceptional in a human child, but Teco is an ape—a bonobo, to be precise. To the uninitiated, bonobos look very much like chimpanzees, but they are in fact a separate species with distinct physical and behavioral traits. More collaborative and sociable than their chimp cousins, bonobos also seem to be more adept at learning human language. And they are endangered, found in the wild only in the Democratic Republic of the Congo. Recent estimates put the wild bonobo population at between 10 000 and 50 000. Fewer than 150 live in captivity. Along with the chimpanzee, they are our species' closest relatives.

For more than three decades, researchers have been working with a small group of bonobos, including Teco, to explore their amazing cognitive and linguistic abilities. Teco's father, Kanzi, is the group's most famous member: Anderson Cooper has interviewed him, and he's played piano with Paul McCartney and Peter Gabriel. Animal lovers worldwide have marveled at his ability to communicate by pointing to abstract symbols. He recognizes nearly 500 of these "lexigrams," which he uses to make requests, answer questions, and compose short sentences. The spoken words he understands number in the thousands.

Even so, many people question these abilities. Indeed, for more than a century scientists have debated whether apes could ever truly comprehend human language. Many researchers argue that language is the exclusive domain of humans, and several influential studies in the 1980s concluded that supposedly "talking" apes were merely demonstrating their capacity for imitation, with lots of unintentional cuing by the animals' handlers. Linguist Noam Chomsky has likewise argued that the human

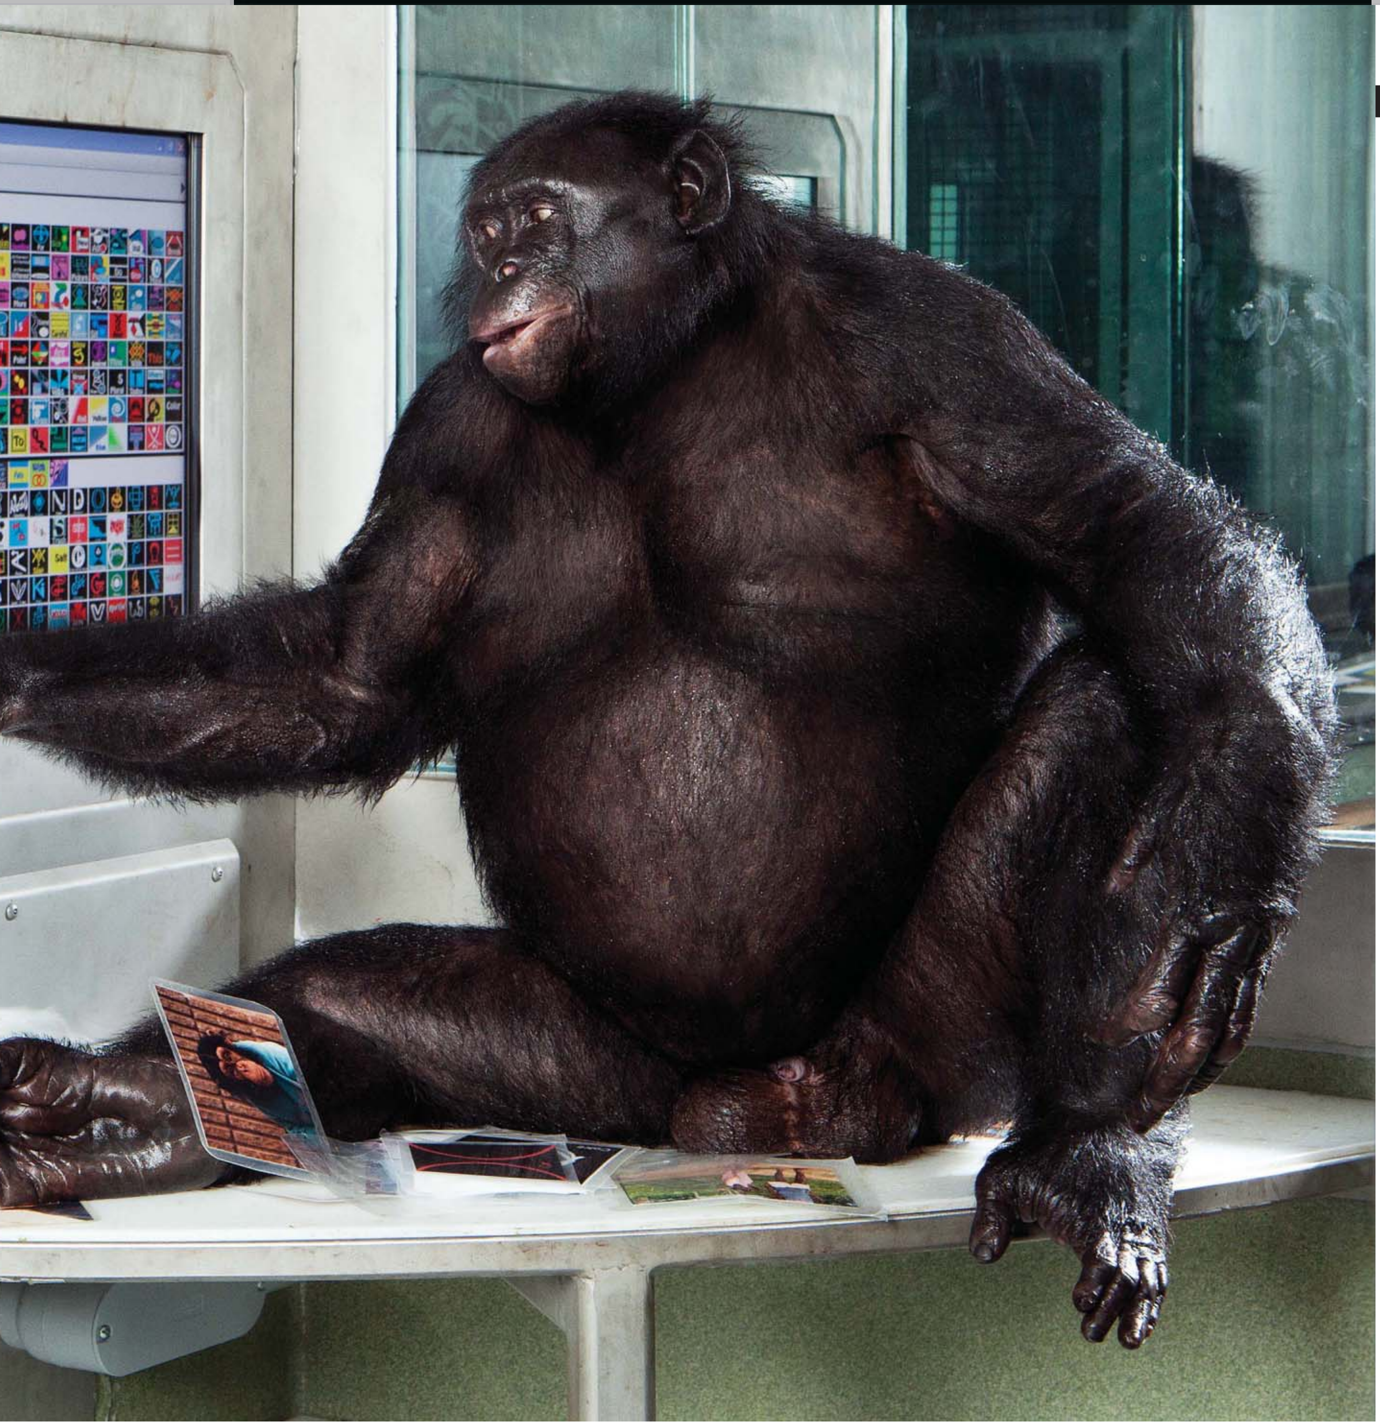

brain contains a species-specific "language acquisition device," which allows humans, and only humans, to acquire language.

But the bonobo research I've been involved with, led by primatologist Sue Savage-Rumbaugh at the Bonobo Hope Great Ape Trust Sanctuary, in Des Moines, strongly suggests otherwise. Today, the wide availability of touch screens, tablet computers, digital recording, and wireless networking is giving researchers the world over powerful new ways to study and unambiguously document ape communication. The results of these studies are in turn helping to spark a renaissance of technology-aided research into primate development and cognition and shedding light on the origins of culture, language, tools, and intelligence.

It's a typical workday, and Panbanisha, Kanzi's younger sister, is sitting before a 42-inch touch-screen display. She's doing a match-to-sample task: When she presses a green button in the middle of the screen, the computer's text-to-speech synthesizer says "apple," and then the lexigram for the word *apple* appears on screen, along with two other lexigrams, for *dog* and *sorry*. Panbanisha touches the one for *apple*. After five correct responses, she gets to pick a short video to watch. She selects one of her favorites: a clip from *Tarzan the Ape Man*.

Our research group is using tasks like this to measure the bonobos' vocabularies. We estimate that Panbanisha, like her

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

BABY TALK: Two-year-old Teco, shown with the author [left] and researcher Susannah Maisel, uses a simplified 25-lexigram app. His first lexigram was *grape*.

brother, understands several thousand words. These match-to-sample experiments are enabling us to determine the exact number and should also help dispel criticisms that the bonobos are simply displaying the "Clever Hans effect." Clever Hans was a horse that became renowned at the turn of the last century for solving arithmetic problems, telling time, and reading and understanding German. Later it was revealed that his trainer was subconsciously nodding whenever the horse tapped out the correct answer. Hans was indeed clever—clever at reading subtle cues from his trainer, not at adding and subtracting.

To avoid the Clever Hans effect, the bonobos aren't encouraged to use sign language, because it leaves too much open to interpretation. Instead they "talk" to us almost exclusively with lexigrams; the computer helps remove any ambiguity. One day, for instance, Savage-Rumbaugh was using the match-to-sample program to familiarize Panbanisha with new words. The computer's synthesized voice spoke the word "carrot" and then its screen displayed the lexigrams for *carrot*, *carry*, and *potato*. Panbanisha was about to hit the lexigram for *carrot*, but Savage-Rumbaugh, who'd misheard the word as "carry," told the ape she was mistaken. The ape, though, knew better and selected the *carrot* lexigram anyway.

For more free-form communication, the apes can use their lexigram program, which displays up to 600 symbols on screen [for examples, see the image "True Meaning"]. The bonobos can tap multiple keys to construct a sentence, and each sentence they write is time-stamped and recorded for further analysis.

In 1971 a primatologist named Duane Rumbaugh (Savage-Rumbaugh's ex-husband) came up with the idea of teaching language to apes by displaying abstract geometrical symbols on a computer screen. The first set of 120 symbols was then designed by Ernst von Glasersfeld, who also coined the word lexigram. Each symbol represented a noun, verb, adjective, or name. The lexigram lexicon was later expanded to 384 symbols, which were displayed on a keyboard. Researchers also used (and sometimes still use) a folding poster-board keyboard when greater mobility was required.

The latest version of the keyboard is created with software on a touch screen. These keyboards are easier to update and much less expensive to make than their hardware predecessors. Written in Java, the program will run on any reasonably up-to-date desktop or laptop. The keyboard software can also be wirelessly shared among several computers so that more than one researcher can communicate with a single bonobo. For easier translation, the researcher's keyboard displays the English word just below each lexigram. Or the researcher can type in a word or sentence in English, and the software does its best to translate it into a meaningful string of lexigrams. For example, there is no lexigram for *pizza*, so the program translates that word into the three-lexigram sequence for *bread cheese tomato*, a description the bonobos came up with themselves.

One of the newest improvements to the keyboard software allows new lexigrams to be defined on the fly. Unlike the original lexigrams, the newer symbols aren't limited to abstract shapes; they often include the English word, too, to make it easier for human users to understand. Here's an example: Let's say one of the bonobos has a toothache. There is no lexigram for *toothache*, so the researcher calls up a standard English keyboard on screen and types in "toothache." A second later, a new lexigram spelling the word in colorful letters appears on the bonobo's keyboard, and when the animal taps the key, the computer's flat, synthetic voice says "toothache." The researcher can then explain the link between what the ape is feeling and the new lexigram. Being able to add new lexigrams at will lets us reinforce new words, when the word and its meaning are fresh in the ape's mind. If the bonobo didn't have a toothache, it would be nearly impossible to explain the concept.

Another way the keyboard can be used is in picture mode. By pressing a lexigram key, the user can call up an image of the object, action, or concept that the lexigram describes.

I've spent many an hour watching the bonobos use the touchscreen keyboards, and they are incredibly good at operating them. Indeed, over the years, they've adapted to whatever new computer technology we've introduced. Before the touch screen, they used off-the-shelf gaming joysticks to control the computer cursor. The apes mastered that device in no time and soon graduated to using the joystick for its intended purpose: Kanzi is an expert at *Pac-Man*, while Panbanisha prefers "Sesame Street" video games.

I, ROBO-BONOBO: This squirt-gun-packing mobile robot, built by the author and his students, is designed to let the bonobos have fun with their visitors.
PHOTO: JAY BENEDICT/BUENA VISTA UNIVERSITY

Of course, studies of ape language didn't start with the advent of computers. Beginning in the early 1900s, several attempts were made to teach chimpanzees to speak human languages. Such efforts proved largely fruitless, however, because chimps cannot produce human vocal sounds.

Experimenters then turned to sign language. In the 1970s, for instance, the psychologists Beatrix and Allen Gardner reported in *Science* that they had raised a female chimp named Washoe to use and understand 85 signs. What's more, she could combine signs in novel and meaningful ways. For example, she referred to the refrigerator as "open food drink," even though the scientists themselves never used that phrase.

David Premack, a psychologist, and Ann Premack, a science writer, adopted another approach: They used colored plastic tokens to represent different words and concepts. In a series of experiments, a chimpanzee named Sarah learned to use the tokens to answer sometimes quite abstract questions about objects and their relations. For example, she was taught that a blue triangle token represented an apple; when she was later shown the token and asked what shape and color the object was, she replied that it was round and red, not triangular and blue.

Almost as soon as affordable computers came on the market in the 1960s, primatologists eagerly applied them to their work. As mentioned earlier, Duane Rumbaugh designed a computer-controlled lexigram keyboard for what came to be known as the Language Analogue, or LANA, Project. Lana is also the name of the first chimpanzee to use Rumbaugh's keyboard. It consisted of three aluminum panels, each about 30 by 60 centimeters, mounted side by side on the wall of Lana's cage. Each lexigram was embossed on a small square Lucite key, which was inserted into slots on the panel. The panel itself was connected by patch cords to a DEC PDP-8/E, an early minicomputer; its magnetic core memory could store a whopping 12 000 words. When a key was pressed, it glowed. The computer monitored the sequence of lexigrams Lana pressed and projected them above the keyboard.

Lana learned to use dozens of lexigrams. She could put these symbols in a grammatical sequence to generate sentences, sometimes quite complex ones. To request a treat from her trainer Tim, for instance, she might press the lexigrams for *Lana want Tim give M&M*.

TRUE MEANING: The 600 or so lexigrams that the bonobos use to communicate represent nouns, verbs, adjectives, and other parts of speech. The newer icons incorporate the spelled-out word, to make it easier for human users to understand. The latest version of the lexigram software also allows new symbols to be created on the fly. *IMAGE: BONOBO HOPE GREAT APE TRUST SANCTUARY*

Work with Kanzi and his fellow bonobos has taken our appreciation of ape language to a new level. Interestingly, Kanzi was never taught to use human language: He acquired it as children do, by being exposed to it. The process began when he was only 6 months old, while researchers were trying to teach lexigrams to his mother, Matata, a bonobo who had been raised in the wild. Baby Kanzi always accompanied Matata during her training sessions and so was in the perfect position to eavesdrop.

For two years, nobody suspected that Kanzi was paying even the slightest attention to the lexigram training, although he clearly liked the lights on the keyboard and the blinking projections above. It was only when Matata was taken away for a few weeks for breeding that researchers discovered how much Kanzi had picked up. After searching in vain for his mother, he spontaneously began using her keyboard to communicate with his caretakers. What is more, he understood the spoken words that the lexigrams represented, and he could locate their representations on the keyboard.

That event marked a paradigm shift in ape language studies. Previously, researchers had worked from a behaviorist psychology tradition, which held that mental events are products of reinforced training. So a scientist would show a chimp an apple, say "apple," and then make the sign for the word *apple*. If the chimp signed back with *apple*, he'd be rewarded with an apple. Kanzi showed us that bonobos don't really learn language that way; neither, of course, do people.

It now appears that 2-year-old Teco will equal or maybe even surpass his father and aunt in linguistic ability. Since birth, he has been totally immersed in human language—and to a much greater extent than Kanzi and Panbanisha have been. I built Teco his first toddler's keyboard app on the Motorola Xoom tablet with just a handful of oversize lexigrams: *grape*, *dog*, *go*, *cereal*, and *milk*. He's since graduated to a 25-lexigram keyboard.

At the age of 4 months, Teco recognized his first lexigram. Savage-Rumbaugh recalls that on the day it happened, the baby bonobo had been eating grapes. When the researcher told him she would give him a grape, he reached out his hand and touched the *grape* lexigram. She waited briefly and then asked him if he was ready for the grape. This time, he touched the *grape* symbol, but with his mouth instead of his hand.

What these and other incidents add up to is a rich picture of bonobo behavior. In the wild, the bonobos live in the rain forest of the Congo Basin, where, until the fairly recent arrival of logging crews, they had no natural predators and enjoyed an abundance of food. As a result, bonobos have gentle dispositions and are extremely sociable and collaborative. They are typically eager to interact with humans—and thus quite receptive to adopting human language and culture. What we now believe is that language, rather than being a uniquely human trait, is something other species can develop to varying degrees under the right circumstances—not to our level of sophistication but certainly to the point where we can communicate intelligently with them.

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/apes0712

If you're worried that the bonobos I work with are spending too much time staring at computer screens, rest assured: They have plenty of opportunity for more natural interactions and for just goofing around. One of the first projects I worked on was a 60-centimeter-tall mobile robot, dubbed Robo-Bonobo, which the apes can control using a wireless joystick. The bot, about the size of a small garbage can, is mounted on wheels and has an onboard camera and an animatronic chimp head that can be raised and lowered. The idea is to provide the bonobos with a safe way to interact with visitors and have some fun—the robot comes equipped with a squirt gun. Of course, there's also a scientific purpose: to study the apes' ability to solve problems that require them to take another observer's perspective.

My students have gotten very creative in developing new high-tech tools for this research. They've developed dozens of computer games, including puzzles and mazes, as well as a program called Keyboard Trainer, which helps people learn the hundreds of bonobo lexigrams. Students love working with the bonobos. On one visit, we watched Panbanisha tapping on a touch screen. Seeing us, she grabbed her poster-board keyboard and used the lexigrams to invite her visitors to have some juice. We all stopped working and sat down outside her glass enclosure to enjoy a drink together.

These days, much of my effort is directed toward finding better ways to document experiments and collect data. For example, it's helpful to know who exactly is using a particular touchscreen keyboard, whether human or bonobo. My colleagues and I are considering using thumbprint or retinal scanners, RFID bracelets, or face-recognition software to identify who's at each computer. We're also exploring ways to allow the apes to control their own environment—using their keyboards to open doors and windows, access vending machines, control cameras, and so on. And we're writing lots of apps for use with the bonobos' wireless and GPS-enabled tablets, which will allow us to collect data in the field. One app we're working on will let the apes alert security guards to any suspicious activity outside their enclosure, like the presence of wild dogs or other intruders.

Where might all this research lead? Beyond showing to what extent apes can use language, we anticipate that our studies will also shed light on people's development and use of language, the impact of culture on cognition, and the cognitive structures that must be in place for language to blossom. Our work also has implications for how to handle sentient animals such as apes in captivity. If apes can communicate with their keepers, for instance, they can ask for changes to their environment that would make it better suited to their needs. Such possibilities suggest that apes' cognitive abilities are a lot closer to humans' than anyone imagined just a short time ago. But maybe that shouldn't be so surprising, given the genetic similarities between the various great apes, a formal grouping that includes humans.

While the bonobo species still survives, we believe it's our obligation to learn as much as we can about these extraordinary animals. They are fascinating in their own right, and they are also a window into our not-too-distant evolutionary past. By studying them, we learn ultimately about ourselves. ❑

IT TAKES SEVERAL LIFETIMES TO PUT A NEW ENERGY SYSTEM INTO PLACE, AND WISHFUL THINKING CAN'T SPEED THINGS ALONG

A COMMENTARY BY VACLAV SMIL

A Skeptic Looks at Alternative Energy

In June 2004 the editor of an energy journal called to ask me to comment on a just-announced plan to build the world's largest photovoltaic electric generating plant. Where would it be, I asked—Arizona? Spain? North Africa? No, it was to be spread among three locations in rural Bavaria, southeast of Nuremberg.

I said there must be some mistake. I grew up not far from that place, just across the border with the Czech Republic, and I will never forget those seemingly endless days of summer spent inside while it rained incessantly. Bavaria is like Seattle in the United States or Sichuan province in China. You don't want to put a solar plant in Bavaria, but that is exactly where the Germans put it. The plant, with a peak output of 10 megawatts, went into operation in June 2005.

It happened for the best reason there is in politics: money. Welcome to the world of new renewable energies, where the subsidies rule—and consumers pay.

ILLUSTRATION: DAN PAGE

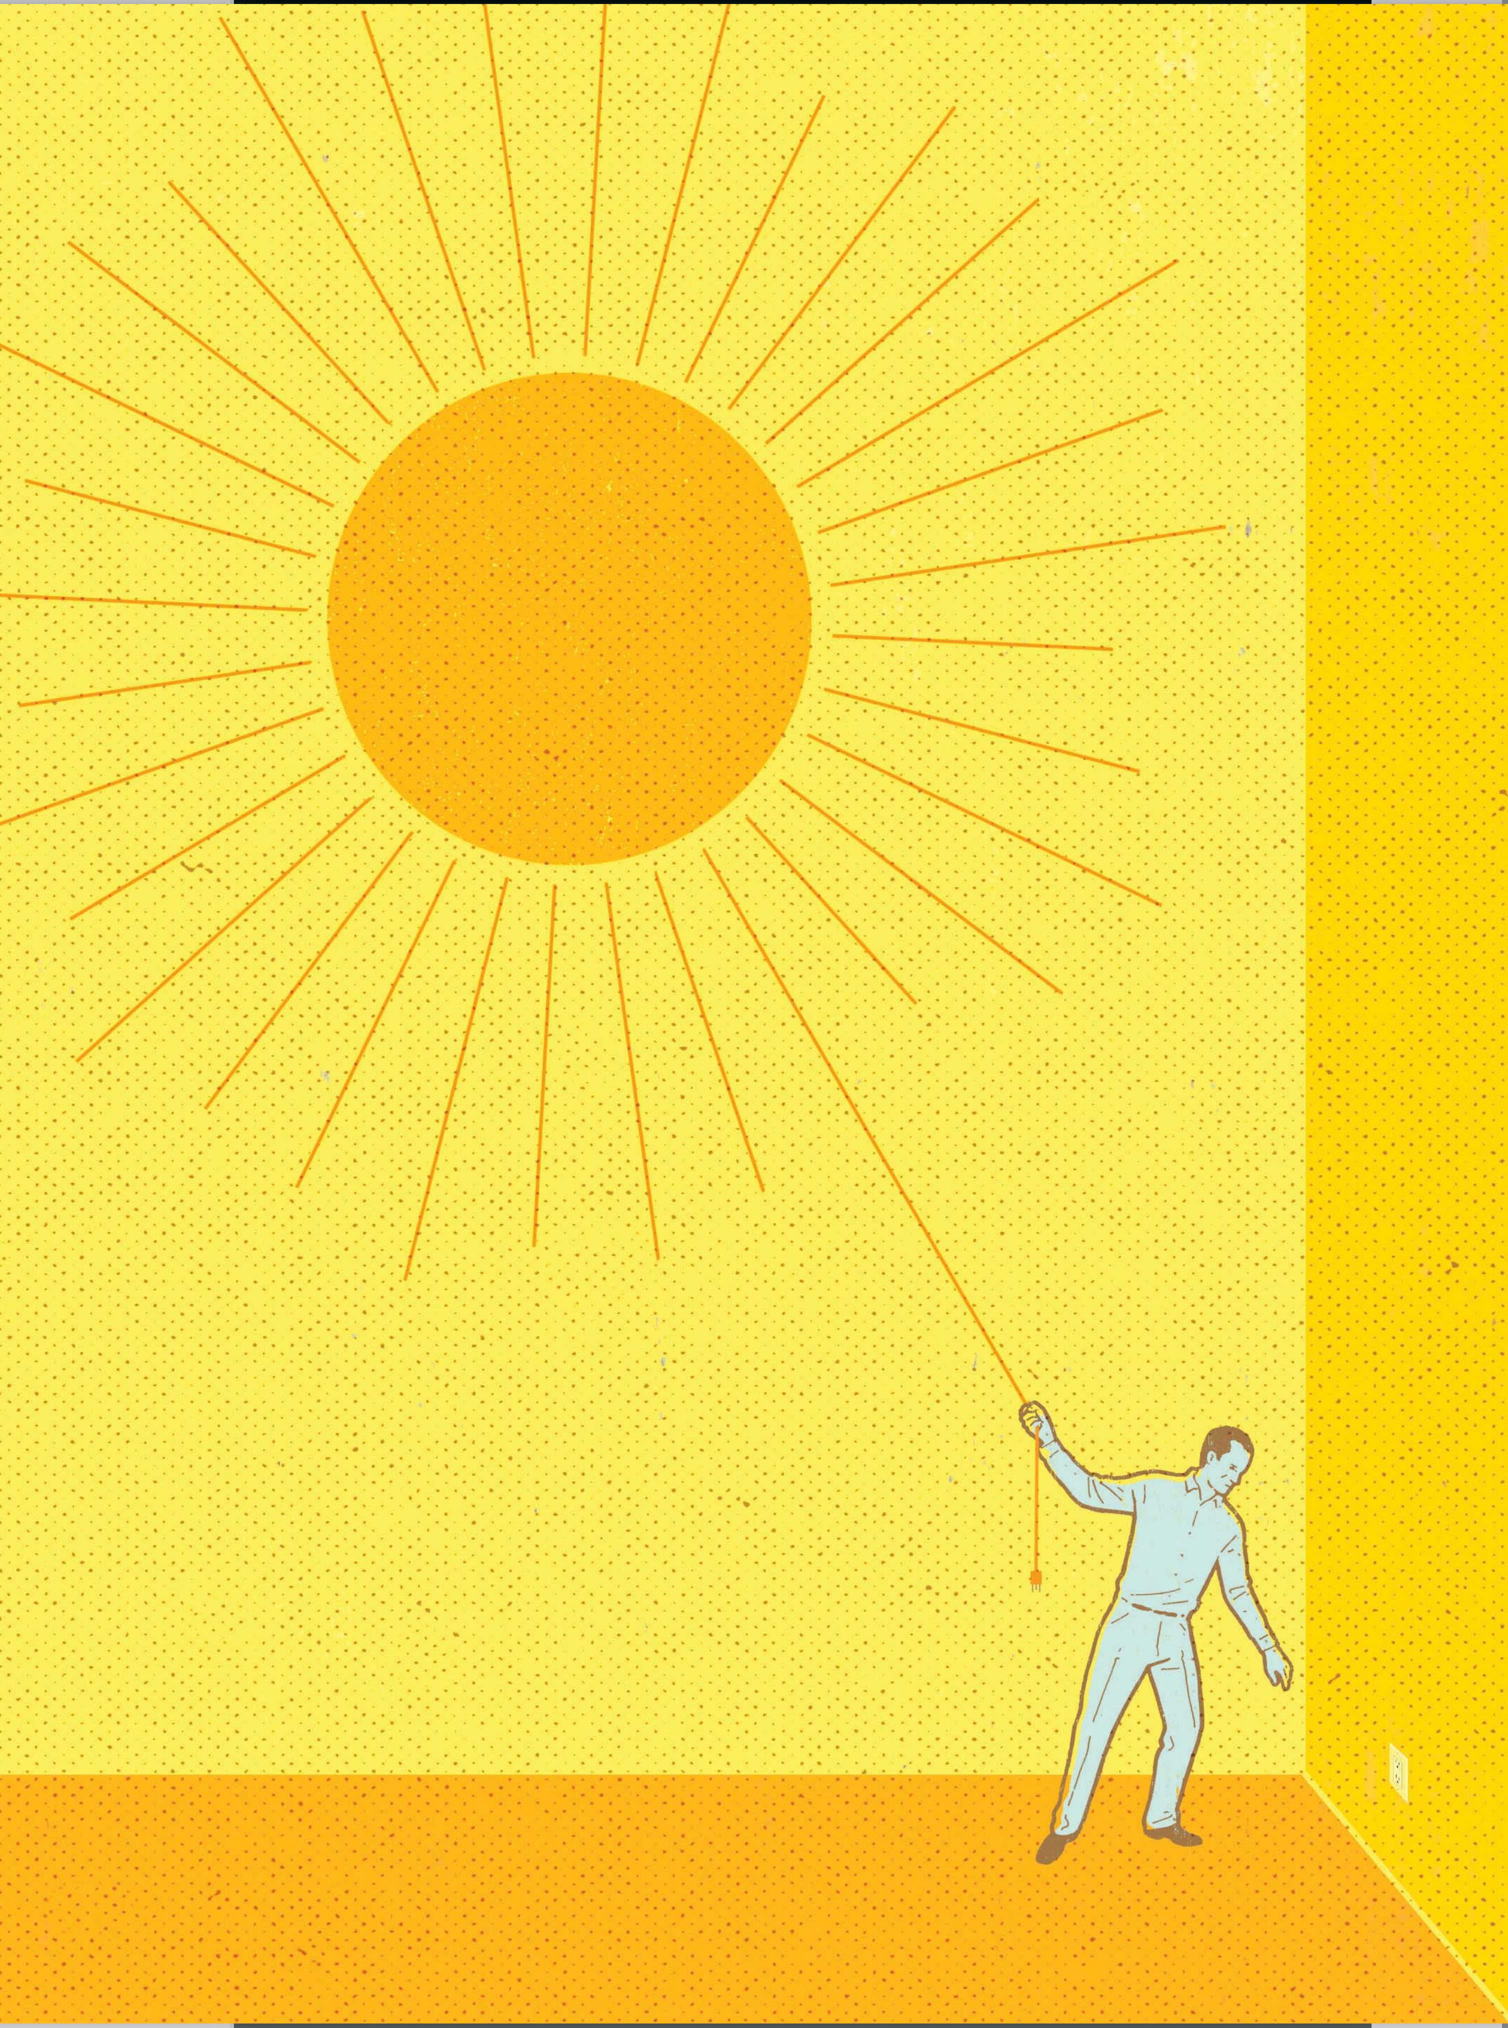

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

Without these subsidies, renewable energy plants other than hydroelectric and geothermal ones can't yet compete with conventional generators. There are several reasons, starting with relatively low capacity factors—the most electricity a plant can actually produce divided by what it would produce if it could be run full time. The capacity factor of a typical nuclear power plant is more than 90 percent; for a coal-fired generating plant it's about 65 to 70 percent. A photovoltaic installation can get close to 20 percent—in sunny Spain—and a wind turbine, well placed on dry land, from 25 to 30 percent. Put it offshore and it may even reach 40 percent. To convert to either of the latter two technologies, you must also figure in the need to string entirely new transmission lines to places where sun and wind abound, as well as the need to manage a more variable system load, due to the intermittent nature of the power.

All of these complications are well known, and all of them have been too lightly dismissed by alternative energy backers and the media. Most egregious of all is the boosters' failure to recognize the time it takes to convert to any new source of energy, no matter how compelling the arguments for it may be.

An example is the 2008 plan promoted by former vice president Al Gore, which called for replacing all fossil-fueled generation in the United States in just a decade. Another is Google's plan, announced in 2008 and abandoned in 2011, which envisaged cutting out coal generation by 2030. Trumping them all was a 2009 article in *Scientific American* by Mark Jacobson, a professor of civil engineering at Stanford University, and Mark Delucchi, a researcher in transportation studies at the University of California, Davis. They proposed converting the energy economy of the *entire world* to renewable sources by 2030.

History and a consideration of the technical requirements show that the problem is much greater than these advocates have supposed.

WHAT WAS THE GERMAN GOVERNMENT THINKING in 2004, when it offered a subsidy, known as a feed-in tariff, that guaranteed investors as much as €0.57 per kilowatt-hour for the next two decades of photovoltaic generation? At the time, the average price for electricity from other sources was about €0.20/kWh; by comparison, the average U.S. electricity price in 2004 was 7.6 cents, or about €0.06/kWh. With subsidies like that, it was no wonder that Bavaria Solarpark was just the beginning of a rush to build photovoltaic plants in Germany. By the end of 2011, Germany's PV installations had a capacity of nearly 25 gigawatts, which was more than a third of the global total. If you subsidize something enough, at first it can seem almost reasonable; only later does reality intervene. This past March, stung by the news that Germans were paying the second highest electricity rates in Europe, the German parliament voted to cut the various solar subsidies by up to 29 percent.

Such generous subsidies are by no means a German peculiarity. They have been the norm in the new world of renewable energies; only their targets differ. Spain also subsidized wind and PV generation before cutting its feed-in tariff for large installations by nearly 50 percent in 2010. China's benefits to its wind-turbine makers were so generous that the United States complained about them to the World Trade Organization in December 2010. In the United States the greatest beneficiary so far has been neither solar nor wind but biomass—specifically, corn used to produce ethanol.

According to the U.S. Government Accountability Office, the excise tax credit for ethanol production cost taxpayers US $6.1 billion in 2011. On top of that direct cost are three indirect ones: those related to soil erosion, the runoff of excess nitrate from fertilizers (which ends up in the Gulf of Mexico, where it creates dead zones in coastal waters), and the increased food costs that accrue when the world's largest exporter of grain diverts 40 percent of its corn to make ethanol. And topping all those off, the resulting fuel is used mostly in energy-inefficient vehicles.

YOU MIGHT ARGUE THAT subsidies aren't bad in themselves; indeed, there is a long history of using them to encourage new energy sources. The oil and gas industries have benefited from decades of tax relief designed to stimulate exploration. The nuclear industry has grown on the back of direct and enormous R&D support. In the United States it received almost 54 percent of all federal research funds between 1948 and 2007. In France it got the all-out support of the state electricity-generating company. Without that subsidy, the industry would never have managed to get its recent share of more than 75 percent of the French electricity market. We must therefore ask whether the subsidies for alternative energy can deliver what their promoters promise.

Make no mistake—they promise much. The most ardent supporters of solar, wind, and biomass argue that these sources can replace fossil fuels and create highly reliable, nonpolluting, carbon-free systems priced no higher than today's cheapest coal-fired electricity generation, all in just a few decades. That would be soon enough to prevent the rise of atmospheric carbon dioxide from its current level of 394 parts per million to more than 450 ppm—at which point, climatol-

DATA SOURCE: VACLAV SMIL

Energy Sources: They Grow Up So...Slowly

An energy technology takes a lifetime to mature. In the United States, for instance, it took coal 103 years to account for just 5 percent of the total energy consumed and an additional 26 years to reach 25 percent. Succeeding technologies hit the first benchmark sooner but the second one as late or even later: In the United States, nuclear power still hasn't gotten there.

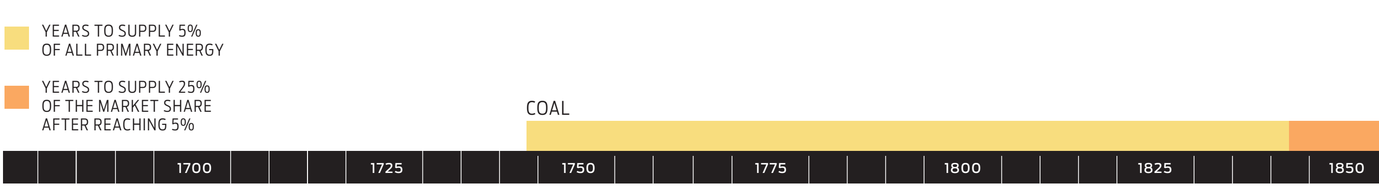

IEEE spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

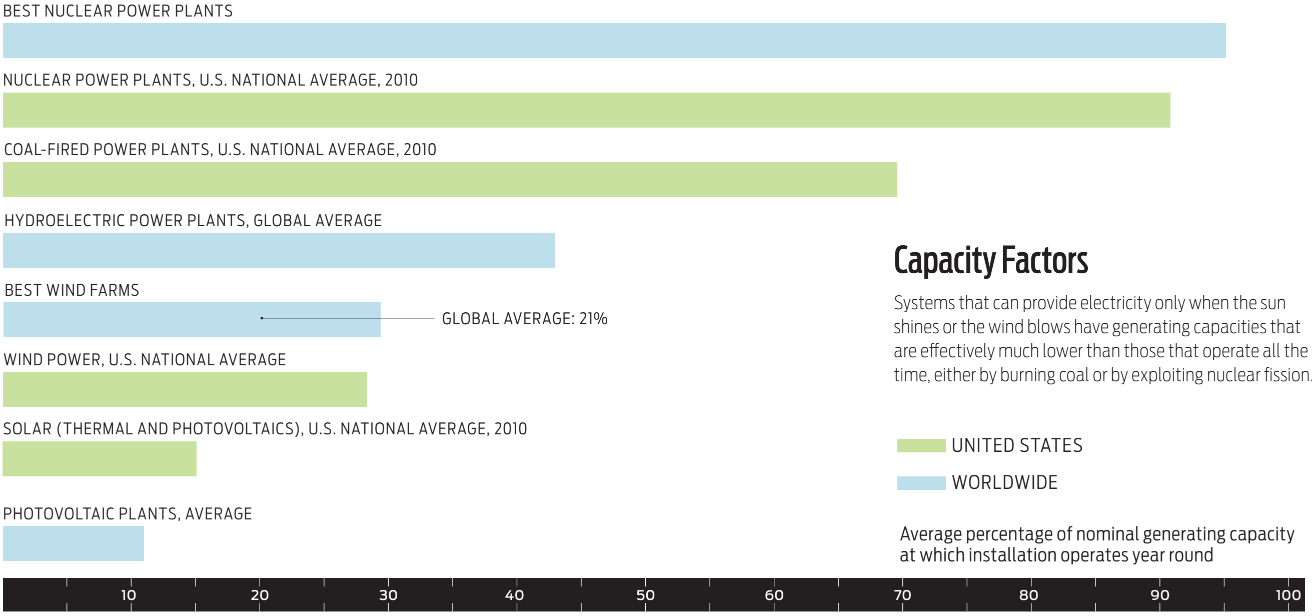

Capacity Factors

Systems that can provide electricity only when the sun shines or the wind blows have generating capacities that are effectively much lower than those that operate all the time, either by burning coal or by exploiting nuclear fission.

ogists estimate, the average global temperature will rise by 2 °C. I wish all these promises would come true, but I think instead I'll put my faith in clear-eyed technical assessments.

The matter of affordable costs is the hardest promise to assess, given the many assorted subsidies and the creative accounting techniques that have for years propped up alternative and renewable generation technologies. Both the European Wind Energy Association and the American Wind Energy Association claim that wind turbines already produce cheaper electricity than coal-fired power plants do, while the solar enthusiasts love to take the history of impressively declining prices for photovoltaic cells and project them forward to imply that we'll soon see installed costs that are amazingly low.

But other analyses refute the claims of cheap wind electricity, and still others take into account the fact that photovoltaic installations require not just cells but also frames, inverters, batteries, and labor. These associated expenses are not plummeting at all, and that is why the cost of electricity generated by residential solar systems in the United States has not changed dramatically since 2000. At that time the national mean was close to 40 U.S. cents per kilowatt-hour, while the latest Solarbuzz data for 2012 show 28.91 cents per kilowatt-hour in sunny climates and 63.60 cents per kilowatt-hour in cloudy ones. That's still far more expensive than using fossil fuels, which in the United States cost between 11 and 12 cents per kilowatt-hour in 2011. The age of mass-scale, decentralized photovoltaic generation is not here yet.

Then consider the question of scale. Wind power is more advanced commercially than solar power, but with about 47 gigawatts in the United States at the end of 2011 it still accounted for less than 4 percent of the net installed summer generating capacity in that country. And because the capacity factors of U.S. wind turbines are so low, wind supplied less than 3 percent of all the electricity generated there in 2011.

It took 30 years—since the launch of small, modern wind turbines in 1980—to reach even that modest percentage. By comparison, nuclear power had accounted for 20 percent of all U.S. generation within 30 years of its launch in 1957, and gas turbines achieved 10 percent three decades after they went into operation in the early 1960s.

Projections of wind-power generation into the future have been misleadingly optimistic, because they are all based on initial increases from a minuscule base. So what if total global wind turbine capacity rose sixfold between 2001 and 2011? Such high growth rates are typical of systems in early stages of development, particularly when—as in this case—the growth has been driven primarily by subsidies.

And a new factor has been changing the prospects for wind and solar: the arrival of abundant supplies of natural gas extracted by hydraulic fracturing, or fracking, from shales. Fracking is uncommon outside the United States and Canada

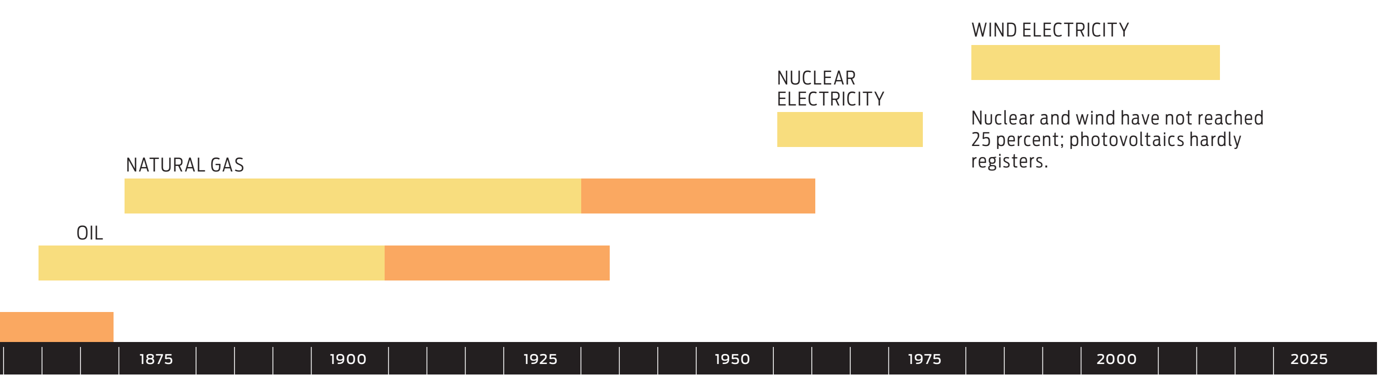

Nuclear and wind have not reached 25 percent; photovoltaics hardly registers.

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

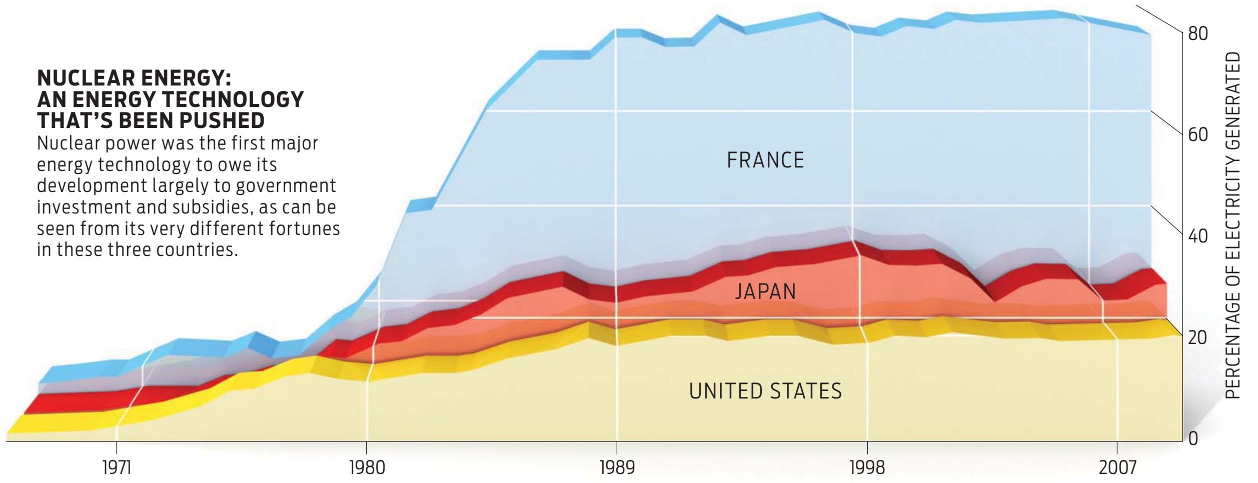

NUCLEAR ENERGY: AN ENERGY TECHNOLOGY THAT'S BEEN PUSHED Nuclear power was the first major energy technology to owe its development largely to government investment and subsidies, as can be seen from its very different fortunes in these three countries.

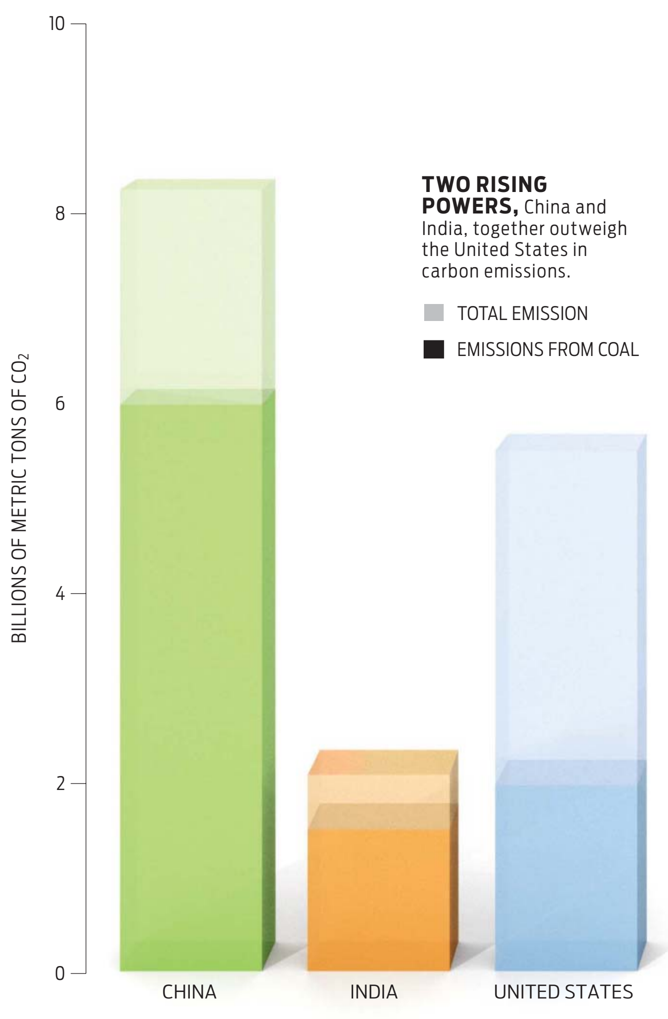

TWO RISING POWERS, China and India, together outweigh the United States in carbon emissions.

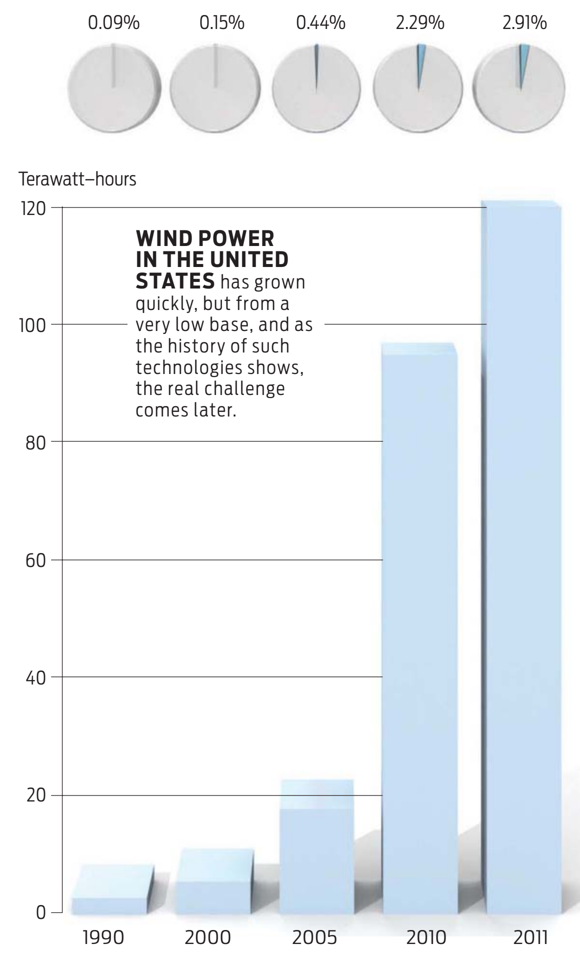

WIND POWER IN THE UNITED STATES has grown quickly, but from a very low base, and as the history of such technologies shows, the real challenge comes later.

INFOGRAPHICS: BRYAN CHRISTIE DESIGN; DATA SOURCE: VACLAV SMIL

at the moment, but it could be used in many countries in Europe, Asia, and Latin America, which also have large shale deposits. Some countries, such as France and Germany, have banned the technology for fear of possible environmental effects, but such concerns accompany all new energy technologies, even those touted for their environmental virtues. And natural gas can be used to generate electricity in particularly efficient ways. For example, combined-cycle gas plants exploit the heat leaving the gas turbine to produce steam and generate additional electricity using a steam turbine. What's more, gas turbine modules with up to 60 megawatts of capacity can be up and running within a month of delivery, and they can be conveniently sited so as to feed their output into existing transmission lines.

THE SITING OF MASSIVE WIND FARMS is also becoming increasingly contentious—many people don't like their look, object to their noise, or worry about their effect on migrating birds and bats [see "Fixing Wind Power's Bat Problem," in this issue]. This has become a problem even for some offshore projects. For example, a vast project off Martha's Vineyard island, in Massachusetts, which was supposed to be the first offshore wind farm in the United States, has been stalled for years because of local opposition. The intermittence of the wind makes it hard to estimate how much electricity can be generated in a few days' time, and the shortage of operating experience with large turbines introduces even greater uncertainty over the long term. We'll just have to wait to see how reliable they'll be over their supposed lifetimes of 20 to 30 years and how much repair and maintenance they will require.

And, of course, you can't use wind turbines unless you're prepared to hook them to the grid by building lots of additional high-voltage transmission lines, an expensive and typically legally challenging undertaking.

Assuming that any major wind farms in the United States would be built on the Great Plains, where there is sufficient wind and land, developers would need to construct many thousands of kilometers of transmission lines to connect those farms to the main markets for electricity on the coasts. Of course, the connection challenge is easier for small countries (particularly if they can rely on their neighbors), which is one reason why Denmark became a leader in wind power.

In the United States, the problem goes beyond building new lines; it is also necessary to add them to an existing grid that is already stressed and inadequate. The most recent *Report Card for American Infrastructure,* prepared with 2009 data by the American Society of Civil Engineers, gives the country's energy system a D+, largely because the grid is relatively old and its operations are repeatedly challenged by spikes of high summer demand. Raising that grade is more than a technical challenge, because improvements in infrastructure often face entrenched political opposition—the not-in-my-backyard syndrome.

As for Europe, there may be better interconnections, but it faces other problems in converting to wind and solar power. Its economic prospects are bleak, and that will limit its ability to invest massively in new technologies. Even Germany, the strongest European Union economy and a great proponent of new energies, has a difficult road ahead; it must find a replacement for its nuclear plants after having decided, following Japan's nuclear disaster in Fukushima, to phase them out. This is no small challenge at a time when Germany is cutting its subsidies for wind and solar power and its economy is close to recession.

Government intervention is needed because the odds are poor that any private program will be massive enough to speed the conversion to new sources of energy. But even governments in the rich countries are having trouble shoring up essential infrastructure, mainly because of mounting debts. Their causes include uncontained health-care costs, trade deficits, uncompetitive manufacturing, and tax-revenue shortfalls. At the same time, government subsidies to new energy technologies haven't delivered on an often-made promise: They haven't created many new, permanent, well-paid jobs either in the EU or the United States.

THE ULTIMATE JUSTIFICATION for alternative energy centers on its mitigation of global warming: Using wind, solar, and biomass sources of energy adds less greenhouse gas to the atmosphere. But because greenhouse gases have global effects, the efficacy of this substitution must be judged on a global scale. And then we have to face the fact that the Western world's wind and solar contributions to the reduction of carbon-dioxide emissions are being utterly swamped by the increased burning of coal in China and India.

The numbers are sobering. Between 2004 and 2009 the United States added about 28 GW of wind turbines. That's the equivalent of fewer than 10 GW of coal-fired capacity, given the very different load factors. During the same period China installed more than *30 times* as much new coal-fired capacity in large central plants, facilities that have an expected life of at least 30 years. In 2010 alone China's carbon-dioxide emissions increased by nearly 800 million metric tons, an equivalent of close to 15 percent of the U.S. total. In the same year the United States generated almost 95 terawatt-hours of electricity from wind, thus theoretically preventing the emission of only some 65 million tons of carbon dioxide. Furthermore, China is adding 200 GW of coal-fired plants by 2015, during which time the United States will add only about 30 GW of new wind capacity, equivalent to less than 15 GW of coal-fired generation. Of course, the rapid increase in the burning of Asian coal will eventually moderate, but even so, the concentration of carbon dioxide in the atmosphere cannot possibly stay below 450 ppm.

PERHAPS THE MOST MISUNDERSTOOD aspect of energy transitions is their speed. Substituting one form of energy for another takes a long time. U.S. nuclear generation began to deliver 10 percent of all electricity after 23 years of operation, and it took 38 years to reach a 20 percent share, which occurred in 1995. It has stayed around that mark ever since. Electricity generation by natural gas turbines took 45 years to reach 20 percent.

In 2025 modern wind turbines will have been around for some 30 years, and if by then they supply just 15 percent of the electricity in the United States, it will be a stunning success. And even the most optimistic projects for solar generation don't promise half that much. The quest for noncarbon sources of electricity is highly desirable, and eventually such sources will predominate. But this can happen only if planners have realistic expectations. The comparison to a giant oil tanker, uncomfortable as it is, fits perfectly: Turning it around takes lots of time.

And turning around the world's fossil-fuel-based energy system is a truly gargantuan task. That system now has an annual throughput of more than 7 billion metric tons of hard coal and lignite, about 4 billion metric tons of crude oil, and more than 3 trillion cubic meters of natural gas. This adds up to 14 trillion watts of power. And its infrastructure—coal mines, oil and gas fields, refineries, pipelines, trains, trucks, tankers, filling stations, power plants, transformers, transmission and distribution lines, and hundreds of millions of gasoline, kerosene, diesel, and fuel oil engines—constitutes the costliest and most extensive set of installations, networks, and machines that the world has ever built, one that has taken generations and tens of trillions of dollars to put in place.

It is impossible to displace this supersystem in a decade or two—or five, for that matter. Replacing it with an equally extensive and reliable alternative based on renewable energy flows is a task that will require decades of expensive commitment. It is the work of generations of engineers. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/altenergy0712

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

Aalto University

Aalto University School of Electrical Engineering provides advanced solutions developed in international scientific and technological projects and applications ranging from care of the elderly to space robotics.

SIX Tenure Track or Tenured PROFESSOR Positions

Aalto University School of Electrical Engineering seeks experts especially in the following fields:
- Electrical Engineering (power and energy), Automatic Control, Embedded Systems, Smart Living Environment;
- Radio Science and Engineering, Radio Astronomy, Space Science and Engineering, Optoelectronics;
- Communications and Networking.

The positions are open at all levels from assistant professor to full tenured professor.

Applicants must have a doctoral degree in a relevant field. For assistant and associate professor levels evaluation focuses on merits and potential for excellence, for full professors we look for demonstrated excellence in research and teaching. Relevant industrial experience is appreciated.

Application deadline is September 30, 2012. For further information and application details, please visit at:
www.elec.aalto.fi/tenuretrack

Aalto University is a new university with over a century of experience. Created from a high-profile merger between three leading universities in Finland – the Helsinki School of Economics, the Helsinki University of Technology and the University of Art and Design Helsinki – Aalto University opens up new possibilities for strong multi-disciplinary education and research. The university's ambitious goal is to rank among the top universities in the world in its areas of specialization. At Aalto University, there are 20 000 students and a staff of 5 000, of which 350 are professors.

aalto.fi

UCL

UCL Department of Electronic & Electrical Engineering

Lecturer/Senior Lecturer in Electronic Engineering

The appointment will be full time on the UCL Grade.
The salary range will be Grade 7: £35,707 - £38,744,
Grade 8: £39,818 - £46,972, Grade 9: £51,052 - £55,512 per annum, inclusive of London Allowance.

The Department wishes to make an appointment at Lecturer or Senior Lecturer level to provide intellectual leadership and to conduct a world-class research programme in her/his area of research as well as contribute to the teaching programme in the Department of EEE. For this appointment we particularly seek candidates with research experience in communications networks systems and circuits, whose research is either within or complementary to our existing programme.

The appointment is a full time permanent position.

For full details about the vacancy and how to apply online please go to http://www.ucl.ac.uk/hr/jobs/ and search for Reference Number 1250792. Please contact Vicky Coombes (email: v.coombes@ee.ucl.ac.uk) for queries regarding the application process. Please do not send CV's direct.

We particularly welcome female applicants and those from an ethnic minority, as they are under-represented within University College London at this level.

Closing Date: 27th July 2012 at 17.00.

UCL Taking Action for Equality

the data

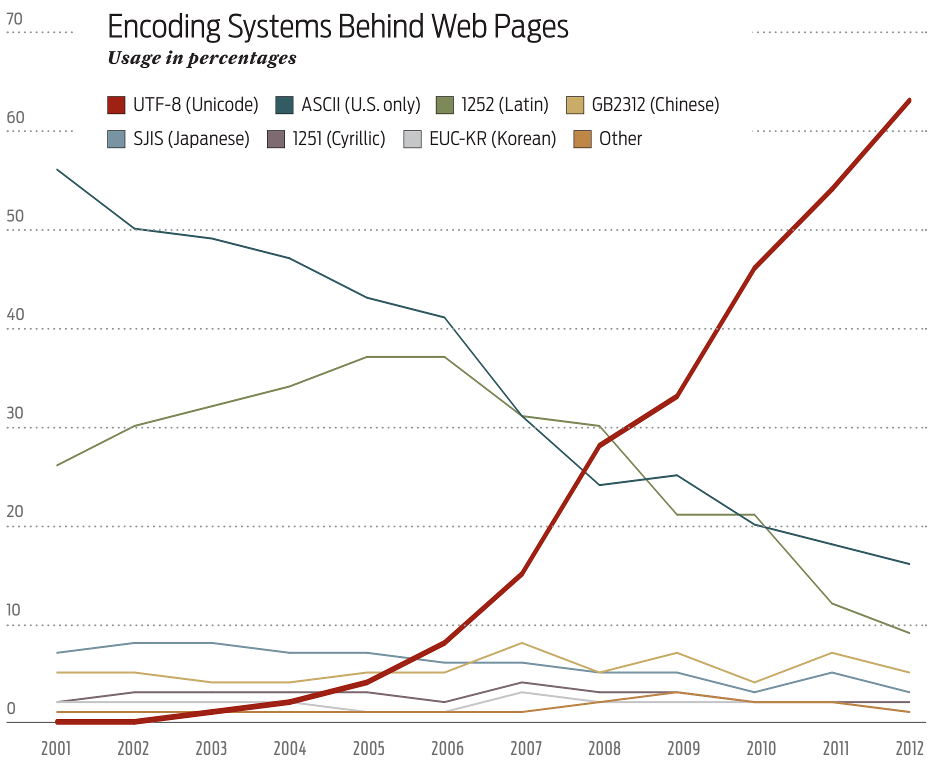

KEY: UTF-8: UCS Transformation Format-8; ASCII: American Standard Code for Information Interchange; 1252: Windows-1252, also known as CP-1252 and latin1 (Microsoft); GB2312: Guojia Biaozhun, which means "national standard" in Chinese; SJIS (or Shift Japanese International Standard): JIS X 0208 Appendix 1; 1251: Windows-1251; EUC-KR: Extended Unix Code-Korean

Source: Mark Davis

Will Unicode Soon Be the Universal Code?

A UNIVERSAL HUMAN LANGUAGE, though appealing in theory, has never gained much traction in real life. French, Chinese, and Arabic have served as lingua francas at one time or another, but almost no one is fluent in Esperanto, the global linguistic mash-up.

In computing, on the other hand, a universal way of encoding languages—that is, translating characters from any language into ones and zeros and vice versa—has been steadily growing since 2006. Unicode is now the encoding system of choice for over 60 percent of Web pages on the Internet, according to an analysis by Google.

The advantage of Unicode is that if everyone adopted it, it would eradicate the problem of *mojibake*, Japanese for "character transformation." Mojibake is the jumble that results when characters are encoded in one system but decoded in another. For example, ASCII, which predates Unicode but is now effectively a subset of it, cannot encode a curly apostrophe—but Unicode can. So when the contraction "that's" is written in Unicode and interpreted in ASCII, it comes out as "thatâ€™s."

Though Unicode has been around since 1991, the year the Web was born, it took a long time for the encoding system to become popular. Mark Davis, one of Unicode's creators, attributes the lag to two things: Most pages were produced in a single language in the early days of the Web, and the available development software tended to rely on different encoding systems, such as ASCII. But after 2006, "people were paying more and more attention to internationalization," says Davis, who is now the senior internationalization architect at Google. "And people started getting the tools to produce their websites in Unicode." *—Ritchie S. King*

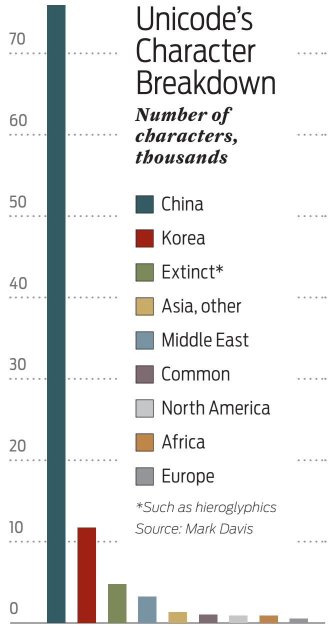

*Such as hieroglyphics
Source: Mark Davis


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