Haiku OS | CAMERA BRAVURA: The End of the SLR ? | Redefining the Kilogram

IEEE Spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
SPECTRUM.IEEE.ORG
5.12

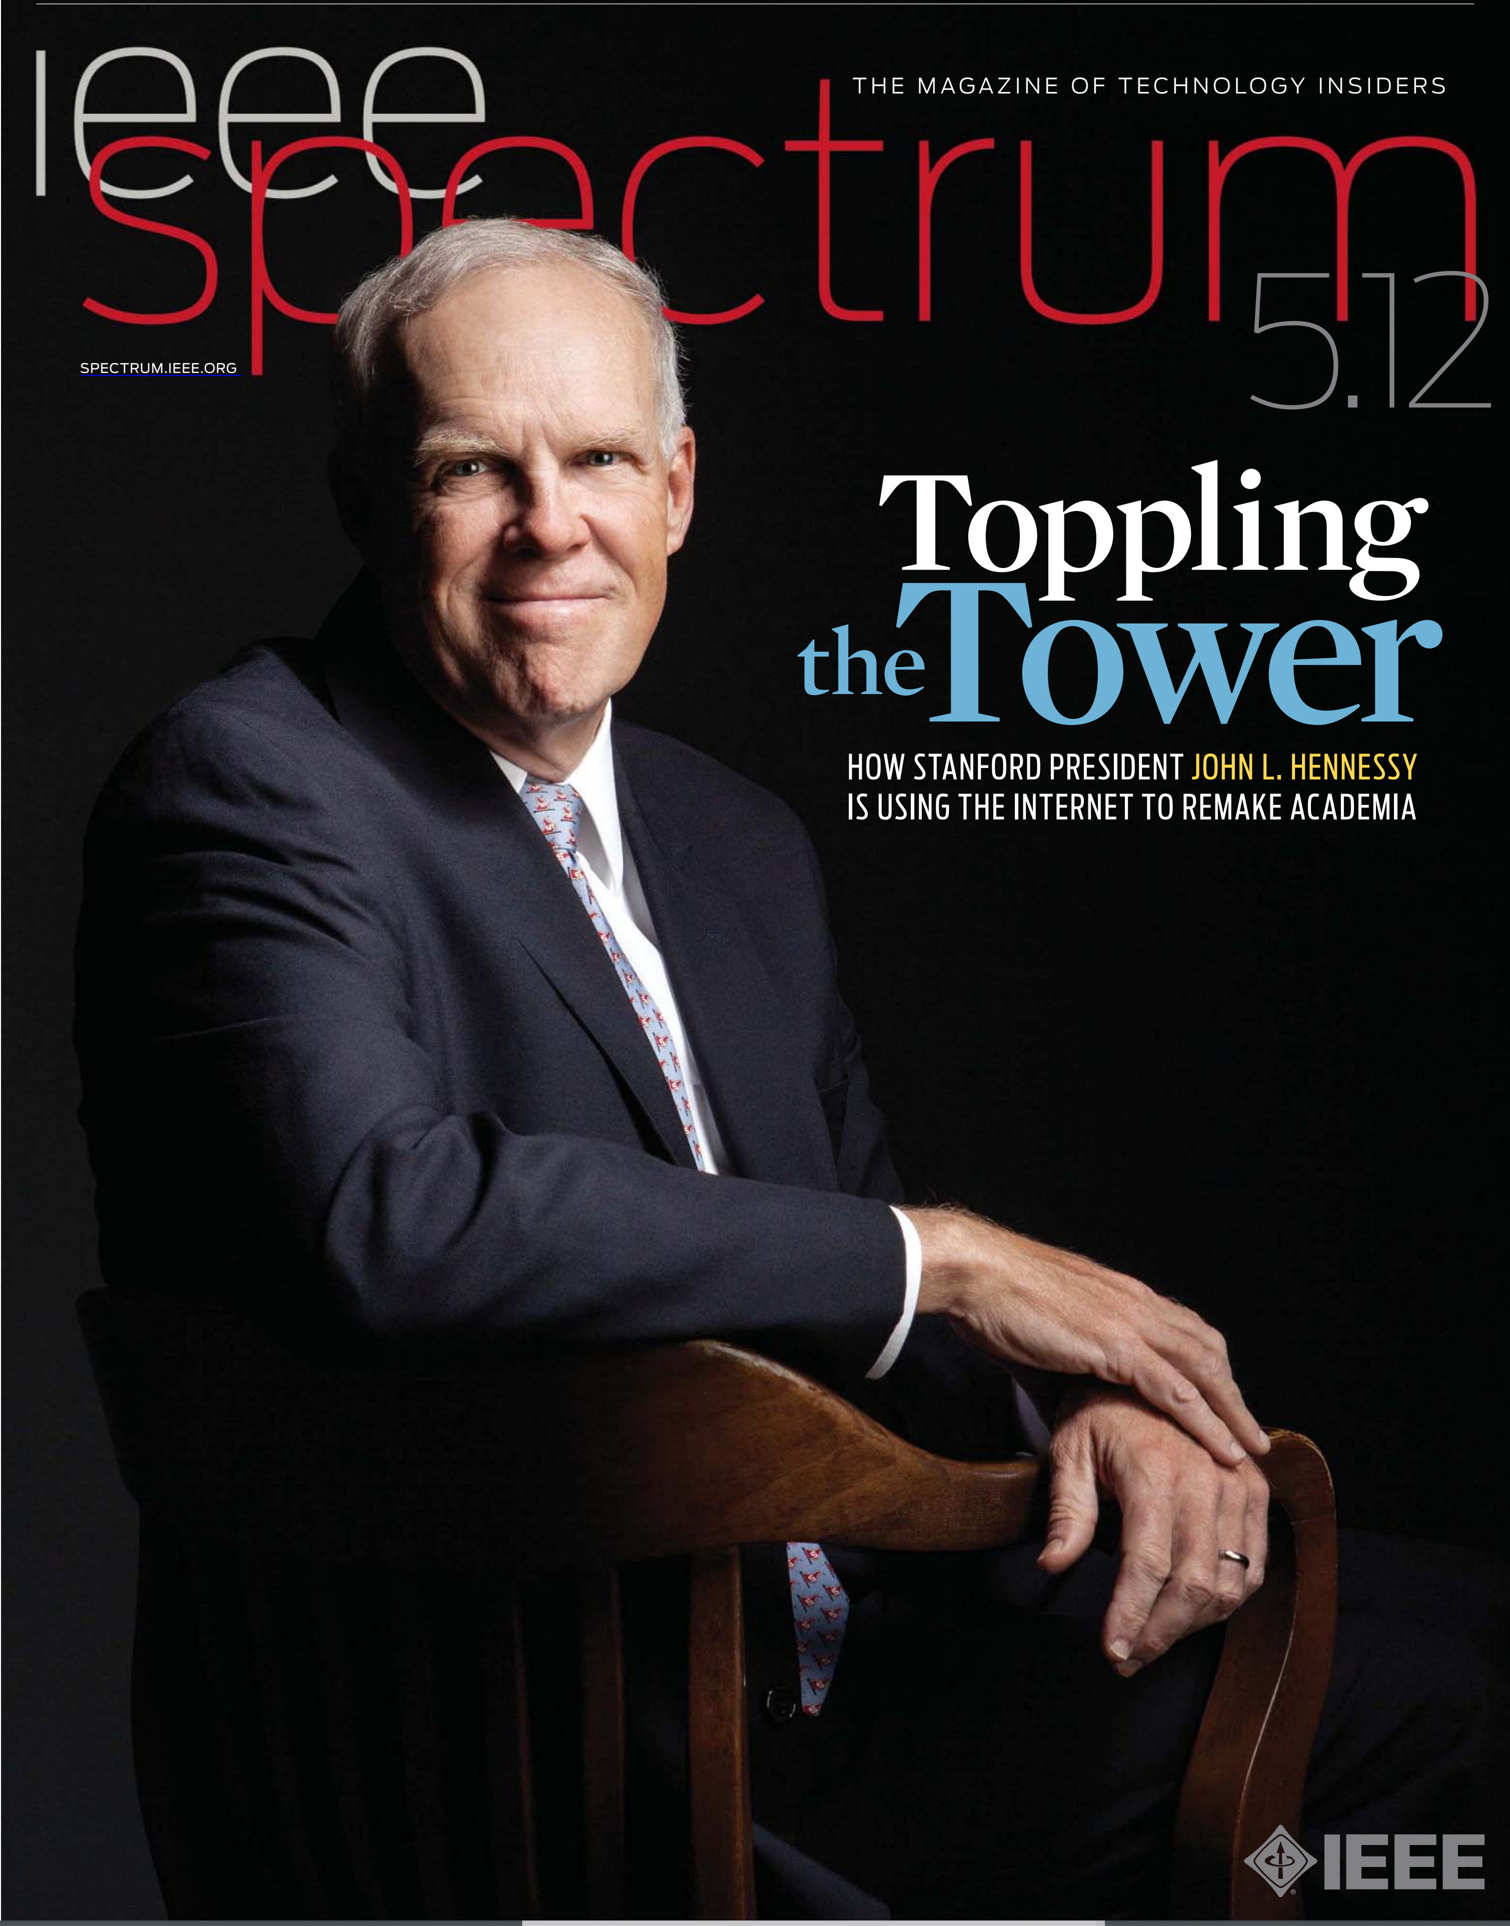

Toppling the Tower

HOW STANFORD PRESIDENT JOHN L. HENNESSY IS USING THE INTERNET TO REMAKE ACADEMIA

IEEE

HEAT EXCHANGER: Surface-to-surface radiation from the outside shell to the pipe walls transfers heat energy to a molten salt that then transports it to a turbine.

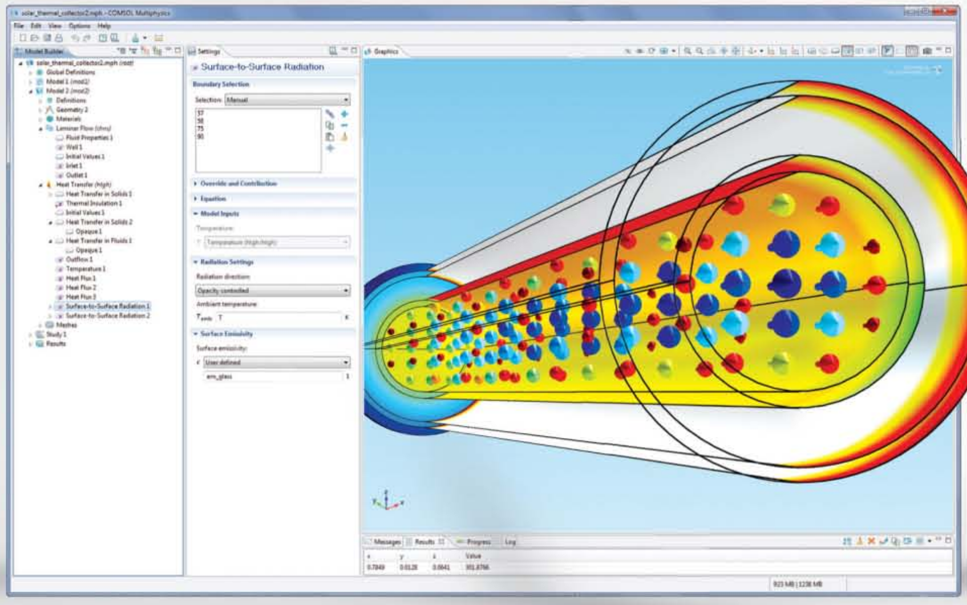

Verify and optimize your designs with COMSOL Multiphysics®

Multiphysics tools let you build simulations that accurately replicate the important characteristics of your designs. The key is the ability to include all physical effects that exist in the real world. Order a free CD containing 700 multiphysics presentations at www.comsol.com/confcd

Product Suite

COMSOL Multiphysics

FLUID
CFD Module
Microfluidics Module
Subsurface Flow Module

CHEMICAL
Chemical Reaction Engineering Module
Batteries & Fuel Cells Module
Electrodeposition Module

MECHANICAL
Heat Transfer Module
Structural Mechanics Module
Geomechanics Module
Acoustics Module

ELECTRICAL
AC/DC Module
RF Module
MEMS Module
Plasma Module

MULTIPURPOSE
Optimization Module
Material Library
Particle Tracing Module

INTERFACING
CAD Import Module
File Import for CATIA® V5
LiveLink™ for SolidWorks®
LiveLink™ for SpaceClaim®

LiveLink™ for Pro/ENGINEER®
LiveLink™ for Creo™ Parametric
LiveLink™ for Inventor®
LiveLink™ for AutoCAD®
LiveLink™ for MATLAB®

COMSOL

volume 49 number 5 international

ieee spectrum 5.12

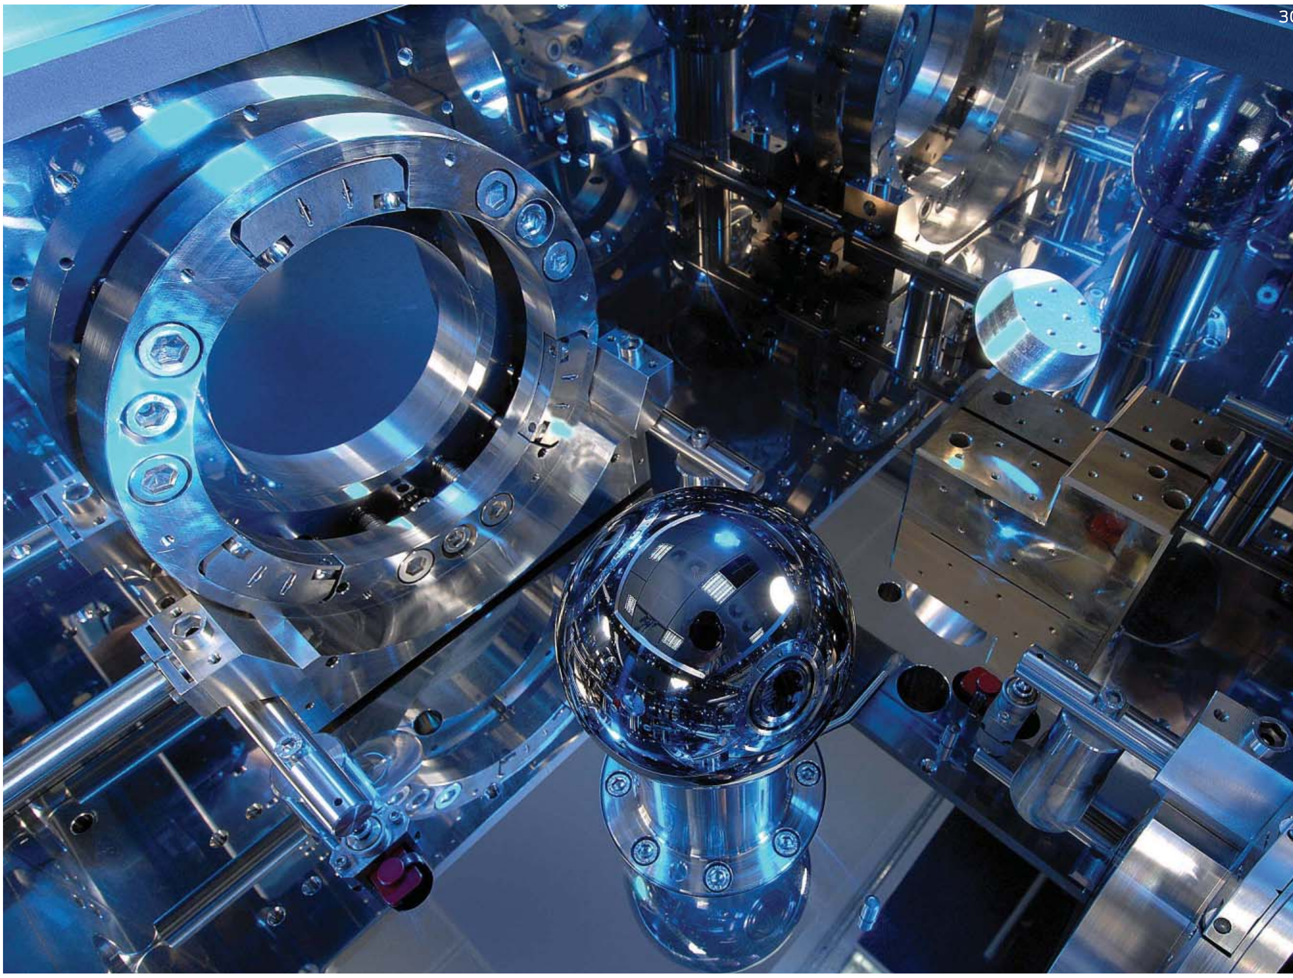

30

COVER: GABRIELA HASBUN; THIS PAGE, TOP: PHYSIKALISCH-TECHNISCHE BUNDESANSTALT; BOTTOM: GABRIELA HASBUN


Name
Dr. Christian Altenbach, Certified LabVIEW Associate Developer
Job Title
Research Ophthalmologist
Area of Expertise
Biophysics
LabVIEW Helped Me
Analyze and visualize data interactively using custom algorithms
Latest Project
Mapping molecular structure changes during activation of the light receptor, rhodopsin
NI LabVIEW
LabVIEW makes me better because I can
PROGRAM
the way I think

volume 49 number 5 international

ieee spectrum 5.12

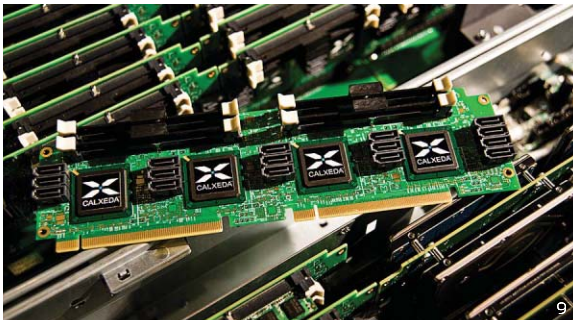

UPDATE

OPINION

DEPARTMENTS

SPECTRUM.IEEE.ORG AVAILABLE 1 MAY

Automotive Black Boxes

We've all heard of the flight-data recorder, or black box, on airplanes. Did you know that your car probably has a similar device? After an accident, this feature can provide accident analysts (and perhaps the police, courts, and insurers) with a data-rich record that details how fast you were traveling, when you hit the brakes, and whether you were wearing a seat belt. Is what you do behind the wheel no longer private? Read more at http://spectrum.ieee.org/carblackbox0512.

THEINSTITUTE.IEEE.ORG AVAILABLE 4 MAY

SPEED LIMIT FOR BIRDS AND DRONES
A team of MIT researchers—some of whom are IEEE members—has discovered the speed at which birds and aircraft will inevitably smash into something. The team found that given a certain density of obstacles, there exists a speed below which a bird—or any other flying object—has a fair chance of flying collision free. If it travels any faster than that, it is sure to crash.

A KNIGHT AMONG US
IEEE Fellow Christopher Snowden, vice chancellor of the University of Surrey, in Guildford, England, was awarded a knighthood this year for his contribution to higher education and his service to engineering.

DATA CONNECTIVITY IN PERU
Many health-care workers living in remote parts of developing countries do not enjoy the benefits of communications technologies that have become commonplace in other parts of the world. Data Connectivity for Remote Health Clinics, one of three branches of the IEEE Humanitarian Technology Challenge, seeks to change this.

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

CLOCKWISE FROM LEFT: CALXEDA; SHUJI KAJIYAMA/AP PHOTO; ISTOCKPHOTO

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.

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

back story

Lots of Lucky

It's 1957, and transistors still have the power to enthrall. Thus Robert W. Lucky, a 21-year-old EE fresh out of Purdue University, is the most fortunate guy in the world. He's got a summer job in the new products division of Westinghouse Electric Co., where he has access to all the transistors he wants, and where his chief responsibility is to design useful circuits with them.

"To get paid to do that was great," recalls Lucky. "It was the best summer job I ever had."

While still a student at Purdue, in West Lafayette, Ind., he had joined the Institute of Radio Engineers (IRE), one of IEEE's predecessors. Lucky, now an IEEE Life Fellow, has also been a hugely popular columnist for *IEEE Spectrum* since 1982. So he was the clear choice to write this month's Spectral Lines, commemorating the IRE's 100th anniversary.

Writing the essay gave him the opportunity to reflect on all that's changed in a century—and also in the 50-plus years spanned by his own remarkable career. In the photo, he's working on the oscillator circuit at Westinghouse that led to his very first patent.

ROBERT W. LUCKY

Lucky would go on to earn a Ph.D. from Purdue in 1961 and then join Bell Telephone Laboratories, which was still secure in its perch atop the U.S. corporate research edifice. At Bell Labs, he invented among other things the adaptive equalizer, which is still used in today's modems. He then moved into research management, eventually taking charge of the communications sciences research division. In 1992 he became corporate vice president in charge of research at Bellcore (now Telcordia Technologies) and remained there until retiring in 2002.

Readers of *Spectrum* may know Lucky best, however, for his poignant observations on engineering life and the profession. Though writing the Reflections column might seem effortless at this point, it's not. "Whenever a column is due, I go to a bookstore or a library, get a coffee, and sit down with my laptop," he says. "I have no idea what to write, and sometimes I think maybe this is the time I don't come up with anything. But in 30 years, that hasn't happened."

And, generous as always, he has advice for would-be engineer-columnists. "Over the years, people have asked me, How do you get to be a columnist?" he says. "I tell them: You have to be Lucky." ❏

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. 5 (INT), May 2012, p. 56, or in *IEEE Spectrum*, Vol. 49, no. 5 (NA), May 2012, p. 64.

contributors

MARK ANDERSON, an *IEEE Spectrum* contributing editor, wrote "This Is Your Brain on fMRI" [p. 21], about current attempts to "photograph thoughts," an idea Nikola Tesla proposed in 1933. Anderson's second book, *The Day the World Discovered the Sun: An Extraordinary Story of 18th Century Scientific Adventure & the Global Race to Track the Transit of Venus* (Da Capo), will be released this month.

SALLY WIENER GROTTA and **DANIEL GROTTA** are a husband-and-wife team who have been writing about digital imaging, photography, and printing since the early 1990s. For "Not All Pixels Are Created Equal" [p. 18], they compared four cameras, ranging from 8 to 40 megapixels. The most telling results were from the two cameras in the middle of the range, each with about 16 megapixels but with a price difference of over US $800. Daniel notes that a camera's price usually has a direct relationship to its image quality. "But if you can't tell the difference—or don't care—why pay the extra?" he says.

MARK HARRIS, a Seattle-based freelancer, used to be the reviews editor for *Digital Camera* magazine in the United Kingdom, so he was quite eager to write about the new "light-field" or "plenoptic" cameras, which allow photos to be refocused after they're taken ["Focusing on Everything," p. 42]. But Lytro's new light-field camera was a letdown, he says. "After just a handful of shots, it started to feel gimmicky." Harris expects the thrill to return when more sophisticated models emerge. "I can't wait for the first consumer Plenoptic 2.0 cameras, with higher resolutions and 3-D imaging."

GABRIELA HASBUN photographed John L. Hennessy for his IEEE Medal of Honor profile [p. 24]. Scheduling around the Stanford University president's sabbatical was difficult, but she says he was a joy to work with. Hasbun is currently expanding her series on San Francisco's Mission District, which she has been documenting through the faces and spaces of its stores and restaurants since 2003.

DAVID KUSHNER writes in this issue about "hacker-spaces" [p. 22], where do-it-yourself engineers meet to take things apart and put them back together in new ways. A *Spectrum* contributing editor, he has also written for *Wired*, *The New York Times Magazine*, and other publications. His latest book, *Jacked: The Outlaw Story of Grand Theft Auto* (Wiley), came out earlier this year.

RYAN LEAVENGOOD, author of "The Dawn of Haiku" [p. 36], is a computer consultant in Boynton Beach, Fla., who specializes in Ruby on Rails website development. He first became involved in the open-source effort to create the Haiku operating system in 2003 and is now the group's treasurer. No operating system is perfect, he notes, but "some are really imperfect. I figured any operating system I worked on couldn't be worse than what I was already using." When Leavengood isn't programming, he and his wife, Francesca, tend to their four cats, one St. Bernard, and a garden of native Florida plants.

CLOCKWISE FROM TOP LEFT: PENNY LEVERITT; GABRIELA HASBUN; BETH KELLY; FLASH PHOTO AND VIDEO; ADAM HEARN

spectral lines

TOGETHER: In 1915 the Institute of Radio Engineers held a joint meeting in San Francisco with the American Institute of Electrical Engineers [above]. In 1963 the two groups merged to form the IEEE.

PHOTO: IEEE

In Celebration of the IRE

In 1912, when the Institute of Radio Engineers was founded, my mother was 10 years old. She lived on a farm in a house with no electricity. She had never seen an automobile, and it would be quite a few years before she heard a radio station or saw an airplane fly through the sky above the farm. Yet she lived to see, on television, a man walk on the moon. What an age of achievement!

Radio itself wasn't too much older than my mother in 1912, and looking back, it seems remarkable that an institution like the IRE would have been founded around such a nascent technology. But those were propitious times for radio. Guglielmo Marconi's daughter Gioia once told me that her father considered 15 April 1912 the single most important day in the history of radio. That was the day the RMS *Titanic* sank, and the world's attention turned to the role of radio, in both the saving and the loss of lives in that disaster.

Surely the founders of the IRE had no idea of the miracles that would follow from the evolution of their experimental spark transmitters and crystal receivers. In the U.S. National Academy of Engineering's list of the 20 greatest engineering achievements of the 20th century, nearly half are directly related to electricity and electronics: electrification, electronics, radio and television, computers, telephone, Internet, imaging, household appliances, and laser and fiber optics. I have a strong sense of shared pride in contemplating that list. We engineers did those things. We changed the world.

I've been privileged to know and sometimes work with the pioneers who created those technologies. I feel as if I grew up with them, and that is largely because of the interactions enabled and encouraged by the IRE and its successor, the IEEE, which grew out of the merger of the IRE and the American Institute of Electrical Engineers in 1963. Before there was such a thing as a social network, the IEEE was our social network. Through its publications, conferences, and awards, we engineers became familiar with new technology and with new colleagues from around the world.

The world of radio engineers was a very small place in 1912. The IRE had just 46 charter members when it was founded; today the IEEE has more than 400 000, over half of whom reside outside the United States. As the IEEE has grown, it has branched out, as evident in the growing number of its societies and specializations. Despite this diversification, the center of our institution has held, and its purpose and activities are still fundamentally the same as those of the small society of radio engineers founded in 1912.

Can we expect the next 100 years to look like the last? Certainly, the last century saw some great inventions that spawned vast industries; the transistor and integrated circuit, the laser, and the Internet come immediately to mind. The one defining characteristic throughout was exponential change—a generalized Moore's Law of progress. I believe that the same pace of progress will continue. Whether electronics will play the same central role remains to be seen; perhaps biology will come to dominate. Regardless of what discipline prevails, however, the rigorous, analytical thought that characterizes engineering will still be necessary, and the supporting role of the IEEE will be invaluable—even if it should mean another merger. —Robert W. Lucky

IEEE Life Fellow Robert W. Lucky writes the Reflections column for IEEE Spectrum. *For more about him, see this issue's Back Story.*

Correction

The credit for the image of the Mariana Trench in "Voyage to the Bottom of the Sea" [March] should have read © 2008 The University of New Hampshire and its Center for Coastal and Ocean Mapping/Joint Hydrographic Center.

more online at spectrum.ieee.org

update

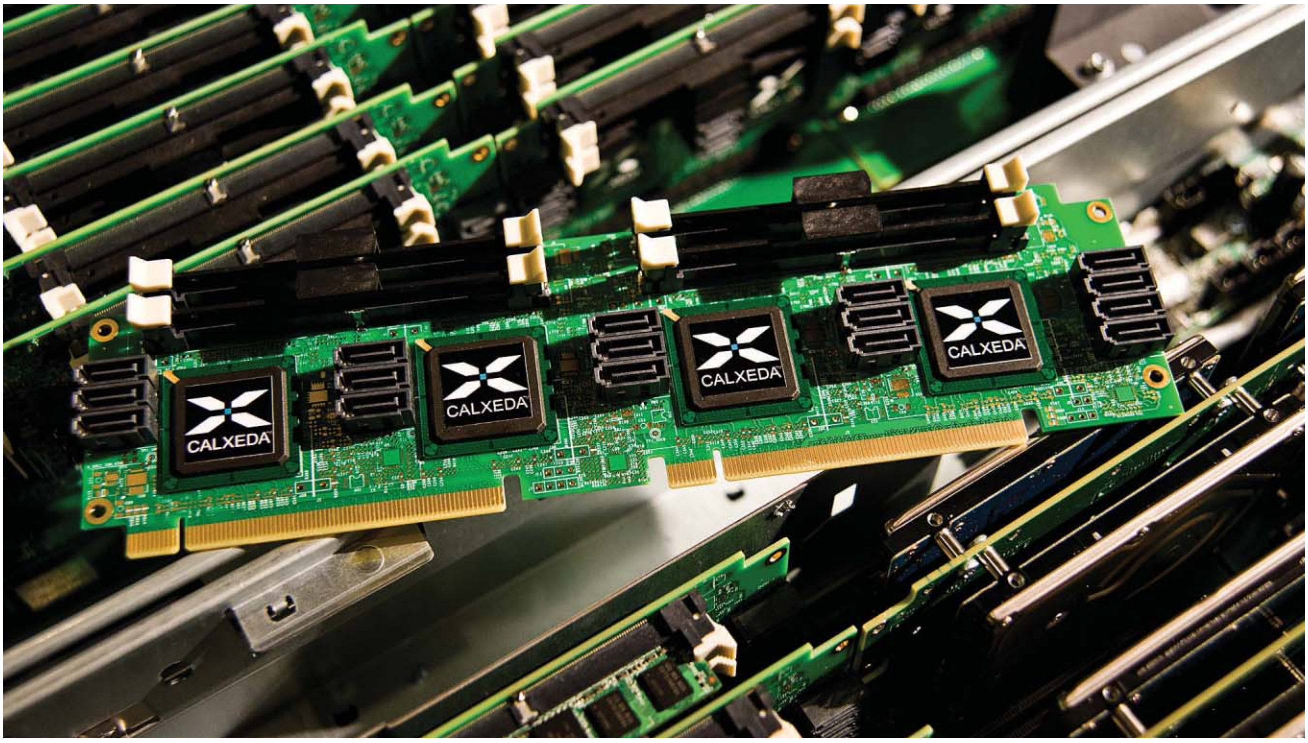

The High Stakes of Low Power

This is the year that ARM and Intel invade each other's turf

There are two giants in the computer processor industry. One is Intel, which builds most of the processors in today's PCs and servers. The other is ARM Holdings, in Cambridge, England, which thanks to its vast ecosystem of partners has established near-complete dominance of the market for the core logic inside smartphones and tablets.

But the demand for energy-efficient chips is reshaping the industry. As the PC market flattens, Intel aims to capture a sizable chunk of the rapidly growing mobile market, which rose to nearly half a billion smartphones in 2011. And chip designers in ARM's camp are eyeing a US $50 billion server market, fueled by the rise of social networking and cloud computing.

The coming months will see a number of volleys exchanged across the line that has traditionally divided the high-performance and low-power chip markets. One of the first will come from a small start-up in Austin, Texas, called Calxeda (pronounced cal-ZAY-dah). The fabless firm will begin shipping chips for servers based on 32-bit ARM mobile processor designs. They'll soon be joined by AppliedMicro, in Sunnyvale, Calif., which is working on an even speedier, 64-bit ARM-based chip. At the same time, Intel will leap into the mobile game; two companies—Lenovo and Motorola—plan to release phones based on Intel's low-power Atom processor by the end of this year.

Exactly how this competition shapes up will depend not on performance or power consumption but on the ratio between the two: performance per watt. And that metric is fueling a fiery debate over the fundamental differences between Intel's *x*86 chips and ARM's processors.

But the most obvious difference between the two may not actually be the important one, according to experts. ARM processors use reduced instruction set computing (RISC), while *x*86

PACKED IN: An HP system will take advantage of Calxeda's low-power ARM-based processors by cramming 288 of them into a single rack unit. *PHOTO: CALXEDA*

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

8 TRILLION ELECTRON VOLTS Energy of protons at CERN's Large Hadron Collider. The physics machine previously ran at 3.5 tera–electron volts.

update

processors rely on an older approach, retroactively dubbed complex instruction set computing (CISC).

Both RISC and CISC architectures govern the set of machine-level instructions, compiled from more complex code, that a chip can execute. CISC chips have a wider vocabulary—they can perform certain operations in one step that might require a series of commands on a RISC chip. But RISC chips can better handle speed-boosting tricks like allowing overlapping operations during each clock cycle.

As a result, over the years, Intel has incorporated decoders into its *x*86 chips to convert CISC instructions to RISC instructions to boost performance. This conversion process takes energy, but it's unclear whether this added step gives ARM an advantage when it comes to efficiency.

Instead, other differences between ARM and Intel chips may have more of a bearing on the coming competition. One key difference is microarchitecture—the particular way that processor resources such as cache and registers are distributed and instructions are scheduled. Today's high-performance processors, for example, are designed so instructions can be performed out of order. Every part of a computation is done as soon as possible to boost speed. Chips that employ this approach have built-in bookkeeping to make sure that the results are assembled in the right order at the end of the process.

Such tricks can have a big impact on efficiency and performance, says Benjamin C. Lee, an assistant professor of electrical and computer engineering at Duke University, in Durham, N.C. While a researcher at Microsoft, Lee studied how well the company's Bing Web search engine performed on Intel's out-of-order Xeon server chip and its in-order Atom netbook processor. Each core on the Atom chip could handle queries at half the rate of a Xeon chip core but required just 20 percent of the energy per request. However, Atom wasn't able to handle some of the more complex requests.

Microarchitecture will be a key battleground in any competition between Intel and ARM chips, says microprocessor industry analyst Linley Gwennap. Eking out even a slight improvement in performance can come at the expense of a large boost in power consumption. If ARM-based devices like AppliedMicro's 64-bit server chip are to compete with Intel chips in the server market, developers will have to accept similar diminishing returns, he says.

This logic isn't lost on Calxeda. The ARM licensee is pursuing applications like Web hosting that don't depend on raw performance. In many cases, "big data" companies like Facebook don't need speedy cores so much as they need a lot of servers that can handle simple tasks like fetching photos.

The processor in a dual-core version of the company's ARM-based server chip would consume about 1.5 watts of power, less than a tenth as much as a comparable Intel Xeon chip. Because the chips dissipate little heat, they can be densely packed. HP, the company's first client to reveal its plans, expects to stuff 288 Calxeda chips in a space that might otherwise be occupied by 8 Intel chips, says Karl Freund, Calxeda's vice president of marketing.

A good part of Calxeda's energy savings comes from innovations beyond the raw capability of the ARM cores, Freund says. Calxeda has integrated as much server infrastructure as possible—cores, cache, and sophisticated network switches—onto every single chip. Designing such a system-on-a-chip (SoC), which is a core technology in today's smartphones and tablets, cuts down on the power consumed by data lines in the computer and makes it easier to implement power management techniques. AppliedMicro is pursuing a similar integrated approach.

But Intel is unlikely to yield its server territory easily, says Mark Hung, a research director at Gartner Research in San Jose, Calif. "It's a very high-margin business for them," Hung says. "I expect Intel is going to be very protective of their turf."

At the same time, Intel has an added technological advantage: Its new 3-D transistors are the smallest, most energy-efficient around. Foundry giant Taiwan Semiconductor Manufacturing Co., which makes mobile chips for companies like Nvidia and Qualcomm, seems to be at least two years away from making a similar switch.

—Rachel Courtland

INTEL'S GOAL: Intel wants its processors powering smartphones. The company made this reference design to speed development.

PHOTO: INTEL

Soft Robots for Hard Problems

Squishy machines may move and manipulate objects in new ways

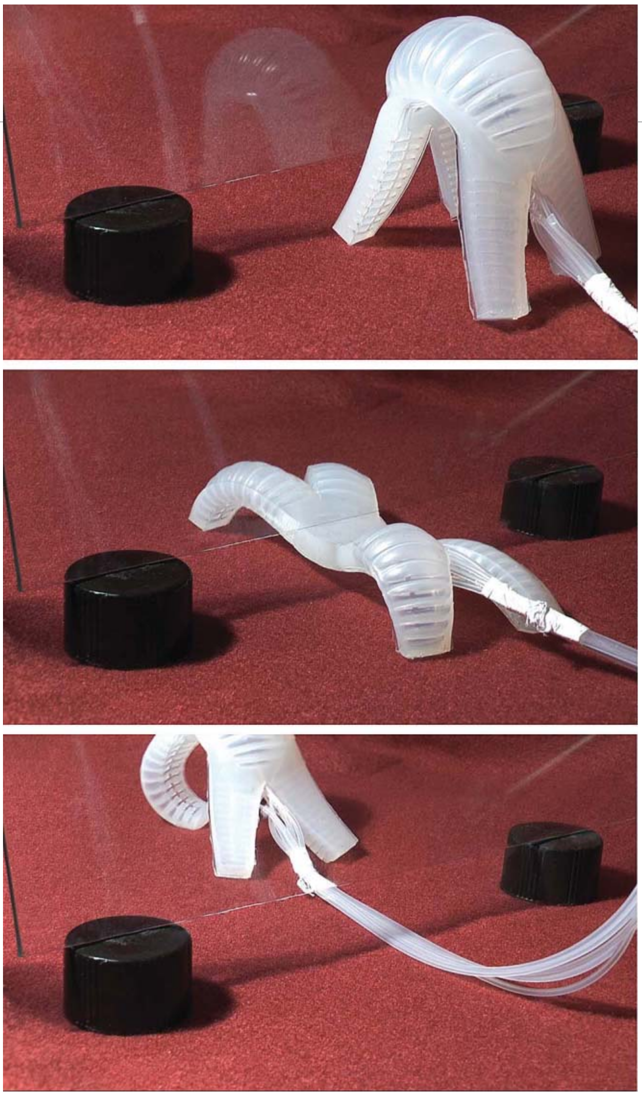

ROBOT LIMBO: A soft robot at Harvard University takes advantage of its rubbery nature to maneuver under an obstacle.
PHOTOS: GEORGE M. WHITESIDES/HARVARD UNIVERSITY

The word *robot* calls to mind images of C-3PO or the Terminator. But robots don't necessarily need to be gleamingly metallic and hard-edged; some might even be downright squishy. That at least is the vision of some robotics researchers, including Carmel Majidi, assistant professor of mechanical engineering at Carnegie Mellon University, in Pittsburgh, and head of the school's Soft Machines Lab.

"Nature is just full of examples of functionality without any rigid parts," says Majidi. Think of an octopus squeezing through a tight opening or a Venus flytrap snapping shut on an unwitting insect.

Soft robots could be built from various types of rubber or silicone. Adam Stokes, a postdoctoral fellow in George Whitesides's lab at Harvard University, has developed a soft rubber gripper that can pick up an uncooked egg or an anesthetized mouse without damaging them. The gripper, a small six-pointed star, is made of two types of rubber, one softer and more extensible than the other. Air is pumped into microchannels in the device through a small tube. "When you put air in, the softer rubber extends more and introduces curvature," Stokes explains. The star bends around the top of the egg, tight enough to lift it but with enough give so the egg doesn't break.

Majidi says the field of soft robotics is still fairly new and that researchers need to find alternatives to air pumps as a way to control the devices. The robots will also need ways to sense their own position. For that he's exploring the use of microfluidics, specifically liquid-filled microchannels inside a film of rubber. Something as simple as saltwater would render the channel conductive so that the device would become electronic. But there are other fluids that could work, such as Galinstan—an alloy of gallium, indium, and tin that's liquid at room temperature and a million times as conductive as saline, making it comparable to copper wire.

Because bending or stretching such a circuit changes the shape of the microfluidic channel, it also changes the circuit's conductivity and thus alters an electrical signal passing through it. "You get something that functions like a stretchable circuit," Majidi says. Such a device could act as a sensor that measures strain, pressure, or curvature. It could even be used as a stretchable antenna, as was recently demonstrated by a team led by Michael Dickey at North Carolina State University.

Soft-robot parts are relatively easy to build, the researchers say. The Harvard lab, for instance, uses a 3-D printer to make a mold. Stokes says that by selecting materials with the right mechanical properties, it should be possible to make devices in a wide range of sizes.

"There are lots of things you could do with this," says Stokes. Perhaps a robot could be a hybrid of hard parts, performing some of the functions robots perform now, plus soft parts for manipulating delicate objects.

The Whitesides lab also makes a four-limbed, air-powered robot that can crawl along a benchtop or bend itself to fit through a narrow opening. Stokes says they're studying the role that different gaits could play in letting robots move at different speeds.

The ways in which soft robots differ from their hard-bodied forebears give robotics researchers entirely new avenues to explore, many inspired by nature, Majidi says. "We can start building devices that are more multifunctional, more versatile," he says.

—Neil Savage

A version of this article appeared online in March.

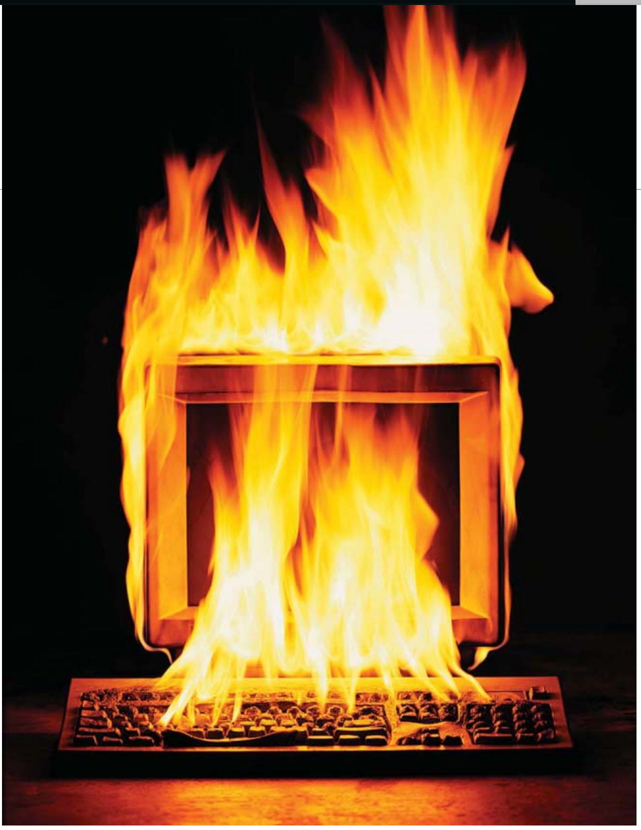

LAWRENCE LAWRY/GETTY IMAGES

update

Computing's Power Limit Demonstrated

Fifty-year-old principle is proved: Erasing information gives off heat

PHYSICISTS IN EUROPE have experimentally demonstrated, for the first time, that a theoretical principle limiting modern-day computing is real.

In 1961, Rolf Landauer posited that the act of erasing a bit of information gives off an amount of heat equal to $k_B T ln(2)$, which is the product of Boltzmann's constant, the temperature, and the natural logarithm of 2—a total of 3×10^{-21} joules at room temperature. Among other things, the theory has been used to address the famous problem of Maxwell's Demon, a thought experiment that suggested a minute monster could create energy for free by sorting particles by their speed, in apparent violation of the laws of thermodynamics. Charles Bennett applied Landauer's theory to the problem, arguing that the energy wasn't free because the demon had to erase a bit of information in its memory in order to sort each particle.

Though it sounds like a philosophical argument, the theory has real implications for computing. And the debate over its validity, some researchers claim, has influenced the direction of computing and semiconductor research.

Eric Lutz, who was at the University of Augsburg when the research was conducted, and a group of European scientists set out to demonstrate Landauer's limit by building his thought experiment in a real system. Landauer had imagined a memory consisting of a single particle. The particle can be in either of two wells of potential energy. If the particle is in the left-hand side, the bit is 0; if it's in the right-hand side, the bit is 1. The bit is erased by forcing it from an unknown starting point into the 1 well.

"It's the simplest memory you can think of," says Lutz. "It's completely generic."

Lutz and his colleagues built this system using a 2-micrometer-wide glass bead in water as the particle and created the potential wells using a modified version of a laboratory instrument called optical tweezers. The tweezers are a laser system that holds particles in place at the laser's focal point. The researchers made the two wells by alternating the focus between two points. Erasing the bit involves first manipulating the laser to lower the energy barrier between the wells, tilting the whole setup slightly to the right so the bead will fall in that direction, and then reestablishing the potential barrier.

At 3×10^{-21} J, the heat predicted to be released by this action is so small that no laboratory calorimeter could measure it. However, the measurement can be derived by closely following the bead's trajectory as the bit is erased. And that's what Lutz, now at Freie Universität Berlin, and his collaborators at École Normale Supérieure de Lyon, in France, and the University of Kaiserslautern, in Germany, were able to do. In the journal *Nature*, they reported that the heat released was exactly as Landauer predicted.

"This is beautiful experimental work," says IEEE Fellow Mark Lundstrom, a professor of electrical and computer engineering at Purdue University, in West Lafayette, Ind., and an expert on the limits of nanodevices.

The energy involved in computing today, Lundstrom points out, is hundreds of thousands of times the Landauer limit. Today's computers use CMOS electronics to implement what's called irreversible logic—you cannot run the process backward to reproduce the input. By its nature, this kind of logic destroys information with every cycle of the processor clock. In 2000, IEEE Medal of Honor recipient James Meindl calculated that the smallest amount of energy such a CMOS computer would need is limited by Landauer's principle. The new work from Europe proves it experimentally, showing that Landauer's limit is real for any irreversible operation.

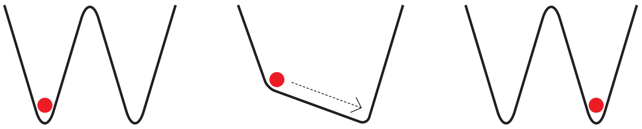

SIMPLE MEMORY: Erasing a bit involves lowering the energy barrier between two states and forcing the system into one state.

$169 BILLION Potential risk to the electronics supply chain from counterfeit parts, according to the market research firm IHS.

But irreversible logic isn't the only way to compute. An experimental form of computing called reversible logic doesn't necessitate destroying information. So, in principle, it could dodge Landauer's limit, its proponents say.

According to Gregory L. Snider, professor of electrical engineering at the University of Notre Dame, in Indiana, reversible logic does just that. In results to be published in the *Japanese Journal of Applied Physics*, he and his colleagues report an electronic reversible logic system that gives off less heat than Landauer's limit. "There is no limit if you're doing it reversibly," says Snider.

The need to get past this limit isn't urgent today, but extrapolating Moore's Law to that limit leads to some absurd ends. A chip built near the limit about a decade from now would throw off more energy per square centimeter than the surface of the sun, Snider estimates.

But others are skeptical of reversible computing's chances. Ralph K. Cavin, chief scientist at Semiconductor Research Corp., in Research Triangle Park, N.C., points out that electrons, unlike microscopic glass beads, can randomly tunnel through an energy barrier by virtue of their quantum nature. Any electronic computing scheme looking to dodge the Landauer limit would need to correct for that, and that correction would give off heat. In fact his colleagues christened the imaginary monster doing the job "Cavin's Demon."

Cavin acknowledges that low-power computing is critical, but he doubts the Landauer limit can be surpassed. "It's a bit like perpetual motion: dreamed of… but not doable." —Samuel K. Moore

A version of this article appeared online in March.

SIGNAL RECEIVED: Scientists projected their results onto the Doge's Palace after successfully sending a message between two of Venice's islands using twisted radio waves. *PHOTO: FABRIZIO TAMBURINI*

A New Twist on Radio Waves

Using the angular momentum of light could make one radio channel into two, three, or more. But many wireless experts are skeptical

By now, you'd expect that communications engineers would have explored every trick in their century-old effort to cram more data into a limited number of frequencies.

But researchers in Italy and Sweden have shown that there is still uncharted territory. A little-explored quantum property, they claim, has the potential to boost the number of channels available in a single-frequency microwave link, perhaps as much as elevenfold.

In early March the researchers simultaneously transmitted two radio beams at exactly the same frequency between two of Venice's islands, a distance of 442 meters. The signals were received and decoded as clearly as if they'd been sent at two different frequencies. Ordinarily, coding schemes require either different frequencies to distinguish the signals or the division of a channel into time slots.

The new trick hinges on using a quantum state of photons called orbital angular momentum. A photon can carry angular momentum just as a rotating body does and can even transfer the momentum to small particles, causing them to rotate. The orbital angular momentum of photons has been intensively explored in the optical region of the electromagnetic spectrum. But its study in the radio-frequency region is quite new.

In theory, a photon can occupy any one of an infinite number of these quantum states, each associated with an integer value. These quantum states impart the radio beam with a distribution of phases as it travels through space that gives the beam the shape of fusilli pasta (a helix).

The researchers started with two off-the-shelf transmitters and receivers designed to operate at the Wi-Fi frequency of 2.414 gigahertz. "We chose this fre-

0.1 BITS PER SECOND Data rate of the first point-to-point communications link using neutrinos instead of photons. The ghostly particles passed undeterred through 240 meters of earth.

update

SLIGHT TWIST: Bending a dish antenna allowed researchers to project twisted radio waves.
PHOTO: FABRIZIO TAMBURINI

quency because equipment that can easily be controlled is available," says Fabrizio Tamburini, an astrophysicist at the University of Padova, in Italy. The setup also included two Yagi-Uda antennas (similar to old television aerials) for reception and a radio dish and a Yagi-Uda antenna for transmission, all commercial models. But the team modified the dish antenna by bending it into a somewhat helical shape.

The shape of the dish was important, because when an ordinary wave of radiation from the antenna horn struck it, what reflected off was a wave whose phase was shifted into the shape of a continuous helix, corresponding to an orbital angular momentum of 1. At the same time, a beam of exactly the same frequency was emitted with a Yagi-Uda antenna, which imparted no phase twist, corresponding to an orbital angular momentum of 0.

At the receiver, the team could easily separate the twisted radiation from the nontwisted beam by measuring the phase.

Tamburini stresses that the experiment was just a proof of principle and that real-world systems for data transmission will use phased-array antennas for transmission and reception instead of bent dishes. These consist of arrays of small antennas, each fed with a signal at a shifted phase to create the helical wave.

Tamburini and his colleagues plan to perform experiments with more transmission channels and to develop smart phased-array antennas that can generate radio waves with several orbital angular momentum states simultaneously. Tamburini says that satellite companies are interested in the technology and that he and his colleagues plan to start a spin-off company in collaboration with the university.

Convincing communication companies that twisted radio waves can add capacity where there was none before will require the simultaneous demonstration of three or more channels at the same frequency, each having different angular momentum, argues Michael Steer, an RF and microwave expert at North Carolina State University. "That would have been convincing; now the [researchers] are really asking us to trust them."

Indeed, some radio communication specialists are skeptical, saying that the technique will add no capacity. Ove Edfors and Anders J. Johansson, both at Lund University in Sweden, argued in the February 2012 issue of *IEEE Transactions on Antennas and Propagation* that at its heart, radio transmission using orbital angular momentum is no different from the multiple-input multiple-output (MIMO) communication technology in use today. MIMO, which involves transmitting and receiving on several antennas, increases data throughput and range without increasing power or using more bandwidth. Though the Lund group submitted their paper before the demonstration in Venice, Edfors says, "I still argue that this is traditional MIMO, but with a more esoteric antenna."

Stefano Maci, an IEEE Fellow and antenna expert at the University of Siena, agrees. "I have some doubt about the practical feasibility of actual systems based on radio vorticity. One should compare this system to a MIMO system," he says.

—Alexander Hellemans

62 PERCENT How much more radiation an obese man receives from a CT scan than a man of normal weight. Nuclear engineers at Rensselaer Polytechnic Institute, in Troy, N. Y., created software to help radiologists reduce the dose.

Snails in a Race for Biological Energy Harvesting

Blood-based fuel cells could power implanted electronics

BIOENGINEERS ARE getting better at replacing and enhancing body parts, but so far they've struggled to power implantable bionics without resorting to clunky batteries. Recently, researchers have turned to blood as a power source, because it carries energy in the form of electron-rich molecules like glucose and delivers it to all parts of the body. Chemist Evgeny Katz of Clarkson University, in Potsdam, N.Y., and his colleagues tested a new kind of fuel cell that, when implanted in *Neohelix albolabris* snails and immersed in the snails' blue, blood-like hemolymph, produced a small, steady supply of electricity over a period of months.

They started by coating the cell's anode with an enzyme that goes by the initialism PQQ-GDH. That enzyme pulls electrons from glucose molecules. The researchers then coated the cathode with a plant enzyme, laccase, which pushes electrons onto oxygen molecules. When immersed in hemolymph, the electrodes delivered a measurable current.

Choosing the right material for the electrodes is important because the intermediate enzymes usually interact little with electrodes, Katz explains. His team used buckypaper, a rough, 3-D structure built of tiny conducting carbon nanotubes. "Buckypaper electrodes are very practical," says biochemical engineer Sven Kerzenmacher at the University of Freiburg, in Germany, who was not involved with Katz's team but is also experimenting with buckypaper electrodes for biological energy harvesting. The nanomaterial exposes more of the electrode to electron-bearing molecules in the snail's hemolymph than other electrode materials would.

In the snails Katz studied, the electrodes extracted as much as 7.45 microwatts in a 20-minute burst and sustained 0.16 µW for up to an hour. Researchers at Case Western Reserve University, in Cleveland, recently implanted similar biofuel cells in cockroaches, extracting slightly less power. By comparison, a modern human pacemaker typically requires a steady supply of around 1 µW.

Other scientists are working on energy harvesters that could be powered by human blood instead of hemolymph. There are important differences: Human blood is better than snail hemolymph at distributing glucose, but most of its oxygen is trapped in hemoglobin. Humans also have strong immune systems that attack foreign bodies.

Biomedical engineer Philippe Cinquin of Joseph Fourier University in Grenoble, France, has tested his glucose biofuel cells for a period of 40 days in living rats, but in contrast with the Potsdam team's snail experiments, he had to build membranes to keep some harmful biochemicals away from his electrodes.

An initial application for human biofuel cells could be to power short-term implants, such as blood glucose sensors for diabetes, says Kerzenmacher. But once long-term, implantable electricity harvesters are available, they could enable tiny distributed devices to stimulate nerves and alleviate chronic pain, researchers say.

"There are a lot more biomedical devices we could build if we know we have energy for them," says Cinquin. —LUCAS LAURSEN

A version of this article appeared online in April.

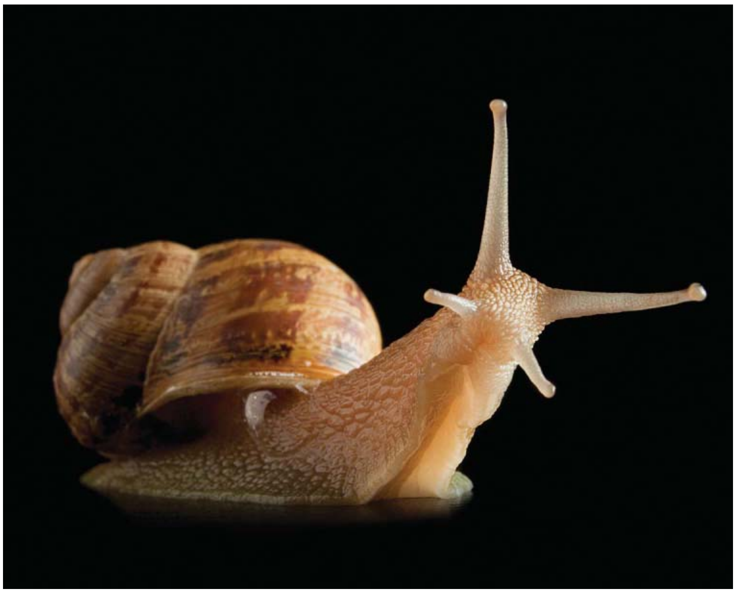

LEFT: MIKE KEMP/GETTY IMAGES; RIGHT: VIENNA UNIVERSITY OF TECHNOLOGY/REUTERS

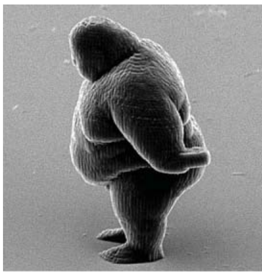

news brief

Small, Swift, but not Svelte
This grain-of-sand-size statue doesn't evoke speed, but in fact it has set a world record. Scientists at the Vienna University of Technology used a new 3-D laser-printing technique to construct the little fellow in just 4 minutes. Previously, the technique took hours or even days. The researchers hope to use it to make custom parts for medical procedures.

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

the big picture

A BIG TURN OFF

These aren't rabbit-ear antennas for a giant TV set. The unit is actually the world's first circuit breaker capable of switching on and off power lines that carry 1.2 megavolts. The circuit breaker was created by Siemens at the behest of the Indian government, which will test it at a transformer substation in Bina. India wants to replace the 800-kilovolt transmission lines that presently carry current from generating facilities in the country's mountainous regions to its largest cities. Because resistive power losses diminish as voltage is stepped up, a 1.2-MV transmission line can carry more than twice as much power as the 800-kV lines.

PHOTO: SIEMENS

tools & toys

Pentax Optio RZ18
(6.1 x 4.6 mm)

Pentax K-5 DSLR
(23.7 x 15.7 mm)

Pentax 645D
(sensor: 44 x 33 mm)

Samsung Galaxy S
(approximately 2.5 x 2.5 mm)

SIZE MATTERS: The size of a camera's lens is a good indication of the size of its image sensor.
PHOTO: SALLY WIENER GROTTA & DANIEL GROTTA

NOT ALL PIXELS ARE CREATED EQUAL

The size of a camera's sensor, not the pixel count, determines the quality of a photograph

THE METRIC most often used by camera manufacturers and marketers to tout their products has been pixel count. That's a shame, but it was probably inevitable—it's easy to measure, and consumers are used to the idea that more is better. However, the number of pixels is a measure of quantity, not quality.

To be sure, in the beginning—the 1990s—there was a great need for more pixels. But by 2000, pixel counts plateaued at 3.3 megapixels. At that number, the sensor was relatively cheap to produce, and the resulting images had just enough resolution for decent 8½- by 11-inch color enlargements from inexpensive inkjet photo printers. For a while, manufacturers competed not on size but with new features, such as geotagging, optical image stabilization, extended zoom ratio, and video. Then pixel escalation resumed with a vengeance. Nowadays, 12-, 14-, and 16-megapixel point-and-shoot digitals are the rule rather than the exception, and cellphones routinely offer 5-, 8-, and even 12-megapixel resolution.

What manufacturers didn't bother explaining was what a pixel is—and why we should care how many we have. Essentially, at the heart of every digital camera is an image sensor. The lens focuses photons reflected by the scene being photographed onto that image sensor. Etched into the image sensor's silicon are pixels (short for "picture elements")—technically, photoreceptor or photodiode sites. Each pixel is a single point that collects the electrons, which are then interpreted into information about color and light. As in the Postimpressionist style of pointillism, which used thousands of paint dots to create a work of art, the signals from pixels are processed into a recognizable image. The more pixels, the more information collected and the larger the photo.

But bigger pictures don't necessarily mean better pictures. In fact, pixel count alone cannot ensure a quality image. If it were otherwise, then a US $139 16-megapixel Nikon Coolpix S3300 point-and-shoot camera would produce pictures as good as those from a $5995 professional Nikon D4 digital single lens reflex (DSLR) with the same number of megapixels. Or, to take the comparison to an even greater extreme, the recently unveiled (and still unpriced) 41-megapixel Nokia 808 PureView smartphone (yes, Virginia, there is a 41-megapixel phone camera) would have image quality similar to that of a $17 500 40-megapixel Phase One 645DF camera. Some obvious differences account for the higher prices of the Nikon D4 and the Phase One 645DF: much faster performance, higher quality construction, more durable body, superior ergonomics, more precise settings and adjustments, interchangeable lenses, and so on. But those factors don't explain the most important advantage: vastly superior image quality.

If the number of pixels doesn't directly relate to image quality, what does? Actually, there are several factors that define and determine a digital camera's image quality: the physical size of the pixels and the image sensor, the filters (and usually the microlenses) bonded to the image sensor, the firmware that processes pixel data, and the camera lens.

But it all centers on the individual pixels.

Pixels on an image sensor are analogous to a bunch of red, green, and blue paint buckets placed side by side. (Red, green, and blue combine to create all colors.) The bigger the buckets, the more paint (electrons) they can capture.

Here's where it gets a little tricky, so it's best to explain by another analogy. Suppose you need to estimate how much rain falls onto a farm, and you have only a minute's worth of rainfall to do your measurements. Imagine that you spread 100 empty soup cans around the property, capped by funnels that are 10 centimeters in diameter. You might collect only a few hundred drops in each. Now suppose you could double the size of the funnel. The amount that can be collected increases exponentially. Calculating how much rain falls on the field by extrapolating the water collected with, say, 1000-cm funnels will yield vastly more accurate results. Here's why: If your raindrops are really photons, the signal-to-noise ratio is dominated by the fact that the noise is equal to the square root of the number of photons. Thus the more rain each soup can collects, the higher the signal-to-noise ratio.

Similarly, large pixels collect more electrons than small pixels, so the point when the pixel flips from a 0 to a 1 can be much more precise. That in turn means far less data processing. With small pixels, the camera's firmware is forced to extrapolate data.

What's more, there are technical complications that can affect the amount and quality of the data collected. For instance, on CMOS image sensors, electrons must first pass through a layer of metal wiring before being collected in the photoreceptor site. Some electrons can be deflected or absorbed and need to be re-created via image-processing firmware. While modern cameras have far better firmware than their predecessors, extrapolation is still at best an imperfect process. When the software guesses wrong, an image can have the wrong colors, lose detail, or produce visual artifacts like purple fringing.

It's axiomatic that at a given number of pixels, the larger the image sensor, the larger the individual pixels can be. Conversely, the smaller the image sensor, the smaller the pixels. What's more, cramming the same number of pixels into a much smaller area introduces other challenges that can generate technical errors, such as noise (electronic chaff) and blooming (light spilling off the sides of the pixels), which in turn lead to image degradation. Conversely, larger pixels and image sensors enhance image quality by producing a higher signal-to-noise ratio, greater dynamic range (the ability to capture details in the highlights and shadows), and truer colors. Higher-priced cameras have other advantages in addition to physically larger image sensors: better lenses with superior optics, better color filters and microlenses bonded onto the image sensor, and more advanced firmware.

To demonstrate all this, we grabbed four cameras from our studio. The first two were a $180 16-megapixel Pentax Optio RZ18 compact camera with a built-in 18X zoom lens and the $1000 16.3-megapixel Pentax K-5 DSLR with a 50-millimeter lens. (We used Pentax cameras for our tests because Sally already had them in the studio on a long-term loan; Pentax is an in-kind sponsor of Sally's *American Hands* fine art project.)

We set up a still-life scene that covered the full potential of dynamic range, including specular highlights and

COMPARE AND CONTRAST: Under magnification, even photos shot under ideal lighting conditions at 100 ISO show significant differences in how well the pixels define the details. Notice whether lines are well formed and if solid colors are truly solid or noisy. Less-expensive cameras boost contrast to make you think you're seeing sharpness.
PHOTOS: SALLY WIENER GROTTA & DANIEL GROTTA

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

tools & toys

lots of detail in highlights, midtones, and shadows. Since the two camera lenses have different focal lengths, we adjusted each camera so that the still life covered approximately the same percentage of the viewfinder (hence a similar number of megapixels were used to capture the different elements of the scene). We put the cameras on a rock-solid camera stand and used their self-timers, to ensure that the act of taking the pictures caused no motion blur. Then we took a series of photographs, first with studio lights at 100 and 400 ISO equivalencies and then without lights at 100, 400, and 1600 ISO. The cameras were set for comparable exposures.

Out of curiosity, we also took the same series of photographs with two other cameras: a $12 500 medium-format 40-megapixel Pentax 645D (with a 55-mm lens) and the 5-megapixel camera in Sally's $200 Samsung Galaxy S phone. (The phone has no ISO settings or tripod socket, so we carefully propped it on the camera stand, with and without lights but with no sensitivity adjustments.) The four image sensors cover quite a range of sizes, from 44 by 33 mm (for the Pentax 645D) to about 2.5 by 2.5 mm for the smartphone. (Samsung does not provide image-sensor size information, but it is quite tiny, about the size of a BB pellet.)

Under studio lights that gave ideal daylight illumination, each of the four cameras captured an attractive photograph when viewed on a computer screen at snapshot size. This did not surprise us—just about any camera can give you a decent picture in good light, if all you want to do is view it on a computer monitor. It's under magnification, when you want to print your photos, or if you are photographing under poor lighting that the differences among cameras become evident.

Even with studio lighting, of the two 16-megapixel cameras, the K-5 DSLR's larger image sensor (and better lens) captured more detail and less noise than the Optio's. But under low light, shot at 1600 ISO, the differences between the DSLR and the compact camera were extreme, with the Optio's shadow details completely lost. The 40-megapixel 645D did even better, even when we downsized the photos so they were the same resolution as those from the 16-megapixel Optio compact. The precision of capture with the larger image sensor and stellar optics of the 645D far outshone all the others and did so under all lighting conditions. Indeed, we're confident the same thing is true at the small end: The advantages of a comparatively larger image sensor, larger pixel size, and a better lens are so great that a compact camera would have done a job significantly superior to the phone's camera even at the same 5-megapixel size.

NIGHT AND DAY: Under low light conditions, with the ISO set to 1600, the differences among cameras becomes quite dramatic.
PHOTO: SALLY WIENER GROTTA & DANIEL GROTTA

So it makes a huge difference if your image sensor is the size of a BB pellet, a Chiclet, a postage stamp, or a tea cookie. Luckily, you don't have to research the Internet or white papers to find out the size of a camera's image sensor (a figure that not all manufacturers readily provide). Just look at the lens. Technically, an image sensor is a component in the camera's optical system. As such, it must match the lens. The diameter of the lens at the point where it is attached to the camera will give you an indication of the size of the image sensor.

For professional and fine art photographers who must have absolutely top quality photos and are likely to produce large prints, the difference in image quality offered by big image sensors (plus their matching superior lenses, better firmware and such) is critical. We've found that serious hobbyists are often just as particular as pros when it comes to getting the best possible picture their budgets can afford. For everyone else, whether or not the size of your pixels matters will depend on your plans to shoot in less-than-ideal lighting conditions or to print enlargements, or if you really care about details in your pictures. If all that's really important to you is that the baby is beaming, reach for the smartphone that's always in your purse or pocket.

—Sally Wiener Grotta & Daniel Grotta

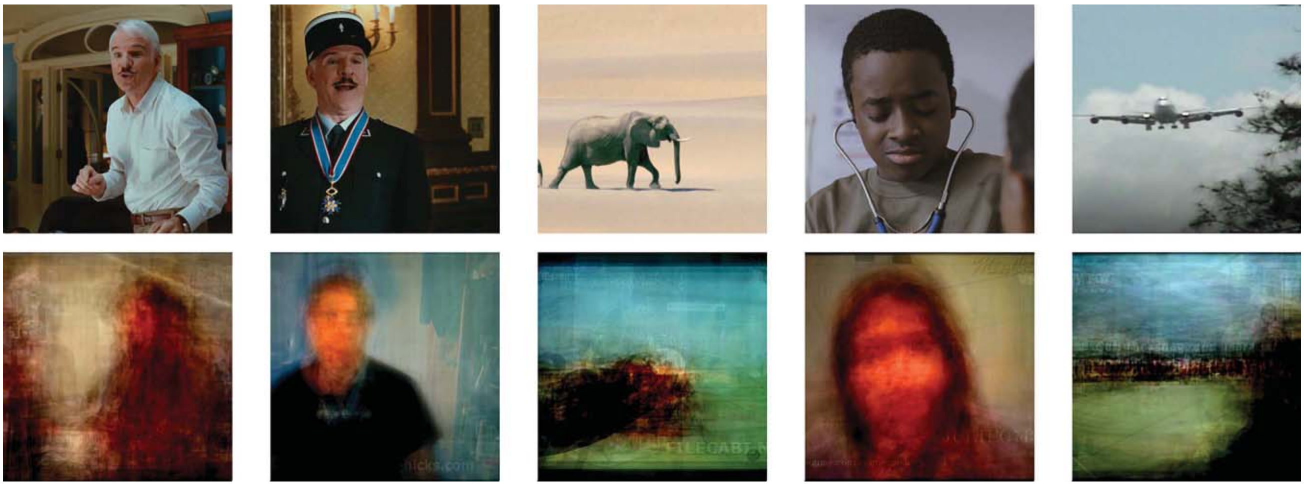

BRAIN RECONSTRUCTION: Segments of Hollywood movie trailers [top row] are shown to subjects while their brain activity is measured using fMRI. Then the images are reconstructed [bottom row] from the brain activity itself.
IMAGES: SHINJI NISHIMOTO

This Is Your Brain on fMRI

The science of mind reading is further along than you might think

At the end of last year, IBM predicted that by 2017 limited forms of mind reading would "no longer [be] science fiction." Along similar lines, though, in 1933 Nikola Tesla said he would soon be able to photograph people's thoughts.

Is IBM going to be equally wrong?

Maybe not. Surveying leading neurotech experts has turned up some support—albeit limited and carefully qualified—for the company's prediction. And oddly enough, one reason is that Tesla's prediction is—in very limited ways as well—coming true too.

Neuroscientists at the University of California, Berkeley, have used functional magnetic resonance imaging (fMRI) to produce rough representations of images as actually seen by a subject's visual cortex.

First, researchers created computer models of fMRI signals from each subject's visual cortex while the subject watched Hollywood movie trailers. Then the scientists recorded visual cortex activity during a second set of clips. They then simulated what the brain "saw" during the second set by running their models on some 5000 hours of online video clips to isolate snippets of images and motions that most closely matched the visual cortex's activity in each moment.

Lars Kai Hansen, professor of cognitive systems at the Technical University of Denmark, in Copenhagen, says his field has undergone dramatic change over the past 15 years. "I see no reason to doubt that...algorithms will allow us to speak 'brain' in much more detail than we can now," he says.

That said, no one expects cutting-edge scientific laboratories, let alone consumer technology, to be able to decode the brain's activities in any way that could enable actual mind reading. It's more that scientists have learned the first few phrases in a very complex language. Worse, the phrases rarely have a fixed meaning and don't combine to form sentences in a straightforward way.

Geraint Rees, director of the Institute of Cognitive Neuroscience at University College London, has been looking for the equivalent of a Rosetta Stone. "Is looking at red plus drinking a glass of wine equivalent to the sum of looking at red and, separately, drinking a glass of wine? Or is there some kind of nonlinear interaction? We don't know the rules of those combinations."

Moreover, says Chris Frith, emeritus professor at the Wellcome Trust Centre for Neuroimaging at University College London, no one has figured out how to read an individual neuron without opening up a skull and attaching a tiny wire to it. Instead, he says, fMRI reads the activity of clusters of hundreds of thousands of neurons. And even the best electroencephalogram (EEG) technology still can resolve only down to the level of hundreds of neurons.

That's enough to do some useful work, though.

In 2009, an Australian company called Emotiv Systems released a commercial EEG unit called the EPOC. A US $299 fitted wireless headset, the EPOC reads faint neural signals through the device's 14 scalp electrodes, which envelop the head like plastic tentacles. According to Emotiv's website, the user's computer—running Emotiv's proprietary signal

hands on

processing software—then translates Emotiv's data stream into meaningful information about activity in various brain regions.

Scores of videos on YouTube show EPOC users training their brains to control a video game, keyboard, robot, or wheelchair. But this is a long way from—to steal Tesla's term—thought photography. "Learning to control things using our brain activity is not the same as mind reading," Frith says. "I agree that control via a brain is no longer science fiction. But it will not lead to thought monitoring."

Ken Norman, principal investigator at Princeton's Computational Memory Lab, says people should be "superskeptical" about any claims of mind-reading technology.

"Brain data is noisy," Norman says. "The cognitive states we're trying to detect are complicated. And we're measuring all sorts of things in the brain that have nothing to do with the thought we want to decode."

He adds, though, that devices like the EPOC could still find some amazing applications.

We can measure whether someone thinks they're doing something right or wrong, how much effort they're expending, their focus of attention, or their familiarity with an object, Norman says. "Most of the cleverness is going to come from people coming to grips with the limits of the technology and then being really smart about what sorts of applications you can make work with that signal." —Mark Anderson

SMASH HIT: After stripping printers of useful parts, hackers let attendees release their technofrustrations at a Hacktory event. *PHOTO: THE HACKTORY*

SULTANS OF SOLDER

At hackerspaces like the Hacktory, geeks turn on, tune in, and hack out

It's a Tuesday night in Philadelphia—hacker time. A small group of do-it-yourself engineers and hipster geeks gather in a cluttered space downtown. One holds a Kinect, the motion-sensing controller for Microsoft's Xbox 360 system. But he's not playing a video game. He's about to take the Kinect apart to see if he can get it to work with a shoot-'em-up space game for which it wasn't intended.

It's a typical challenge at the Hacktory, a free-wheeling engineering clubhouse and just one of the many self-described "hackerspaces" popping up in cities around the world. The group takes its name from Andy Warhol's famous 1960s hangout, the Factory, hoping to bring that sort of imaginative spirit to the technically inclined—and reclaim the word *hacker*. "It's a way to take something apart and put it back together in a new way," says Hacktory director Georgia Guthrie. "It's a creative act, not a destructive act."

While many people think of hacker communities as existing online, there's also a desire to collaborate in person. "Physical hackerspaces sprang out of a need to have a sense of community and a place to hang out," says Jonathan Lassoff, president of Noisebridge, in San Francisco, which operates out of a former sewing shop and at almost 500 square meters is one of the larger hackerspaces around.

Hackerspaces.org, a hub for the collectives, lists dozens of groups, from Toylab in Argentina to Blind Security in Uganda. In addition to being a clearinghouse for information on the scene, the site holds a monthly call-in, which allows groups around the world to share ideas (the audio files are kept for future reference). Hackerspaces.org also organizes hackathon events every month. A recent example, The Playing Card Box Challenge, required participants to "create a hackerspace gift that will fit inside the box from a set of playing cards and mail it to another hackerspace."

Launched in 2007, the Hacktory was an outgrowth of the burgeoning "maker" scene—the do-it-yourself subculture that spawned its own glossy magazine, *Make*, and convention series, the Maker Faire. The Philly group sought to distinguish itself by focusing on creativity, letting the engineering sneak in by the side door. For example, a design workshop in which kids created jewelry tacitly introduced them to light-emitting diodes. "We've decided to cater more to art and integrate technology with art in any way possible," says Guthrie, an art historian who became interested in electronics after getting disenchanted with her museum work. "I thought, museums are old and stuffy," she says. "They don't let you touch anything."

The Hacktory offers a weekly open house. The club hosts special events, such as Nerf Gun Hacking, and concerts by musicians who perform on Nintendo Game Boy devices. One of the more popular activities was a sex toy hacking workshop, in which participants made their own devices using materials as diverse as plastic eggs and bicycle parts. Newcomers look for a place to build something themselves and end up sharing their knowledge and tools. "The network is the most important thing," says Andrew Davidson, a graphic designer who helps organize Hacktory events. "The people I've become friends with at the Hacktory are reliable resources for information."

Nonprofit Technology Resources, a Philadelphia-based group that offers computer hardware and training to low-income families in the city, provides the Hacktory with classroom space, financing from state grants, and management of those and other funds. Stanley R. Pokras, president and chief executive officer of the group, says the Hacktory fits in with his organization's mission to educate and inspire people to express themselves through technology. People understand technology by actually making things, he says.

The Hacktory has tried to overcome the stereotype of hacking as something only for postadolescent boys. "Most hackerspaces strike me as male nerds who are so superinvolved in their nerd-dom that they can't comprehend something like sewing," says Davidson. "We try to be a little more female friendly—and kid friendly."

One of the more popular events is the egg drop, which challenges participants to wrap an egg so that it won't break when it's dropped from a building. Another favorite happening is the printer smash, a stress-relieving way to get rid of old machines—and salvage the shattered parts for reuse. The Hacktory lets people have at the devices with the help of a sledgehammer. Participants don goggles and gleefully whack at the printers, shattering the parts into pieces in a parking lot near the Hacktory. The smash was a highlight of Philly Tech Week, a springtime celebration of technology and innovation throughout the city.

With interest in hackerspaces growing, there comes a new challenge: making money. Collectives are brainstorming possible solutions. Platoniq, a Barcelona-based group, commissioned a report on business models for hackerspaces that recommends charging for services and manufacturing third-party open hardware. While many spaces are loose and casual, some of the bigger ones are getting more organized, with boards, officers, and membership dues that can reach US $80 per month. Some use the spaces as freelance work environments as well.

The Hacktory is no exception, says Davidson: "The next step is to turn into a business and generate income." The group is beginning to offer mainstream classes, such as basic Web development and Python programming. It's also collaborating with hackerspaces in other cities. "We like being one of the organizations that stands for the original definition of the word *hack*," Guthrie says. "We want to show that being a hacker is cool."

—David Kushner

A version of this article appeared online in November 2011.

JOHN L. HENNESSY

RISK TAKER

STANFORD UNIVERSITY'S PRESIDENT PREDICTS THE DEATH OF THE LECTURE HALL AS UNIVERSITY EDUCATION MOVES ONLINE

BY TEKLA S. PERRY

GABRIELA HASBUN

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

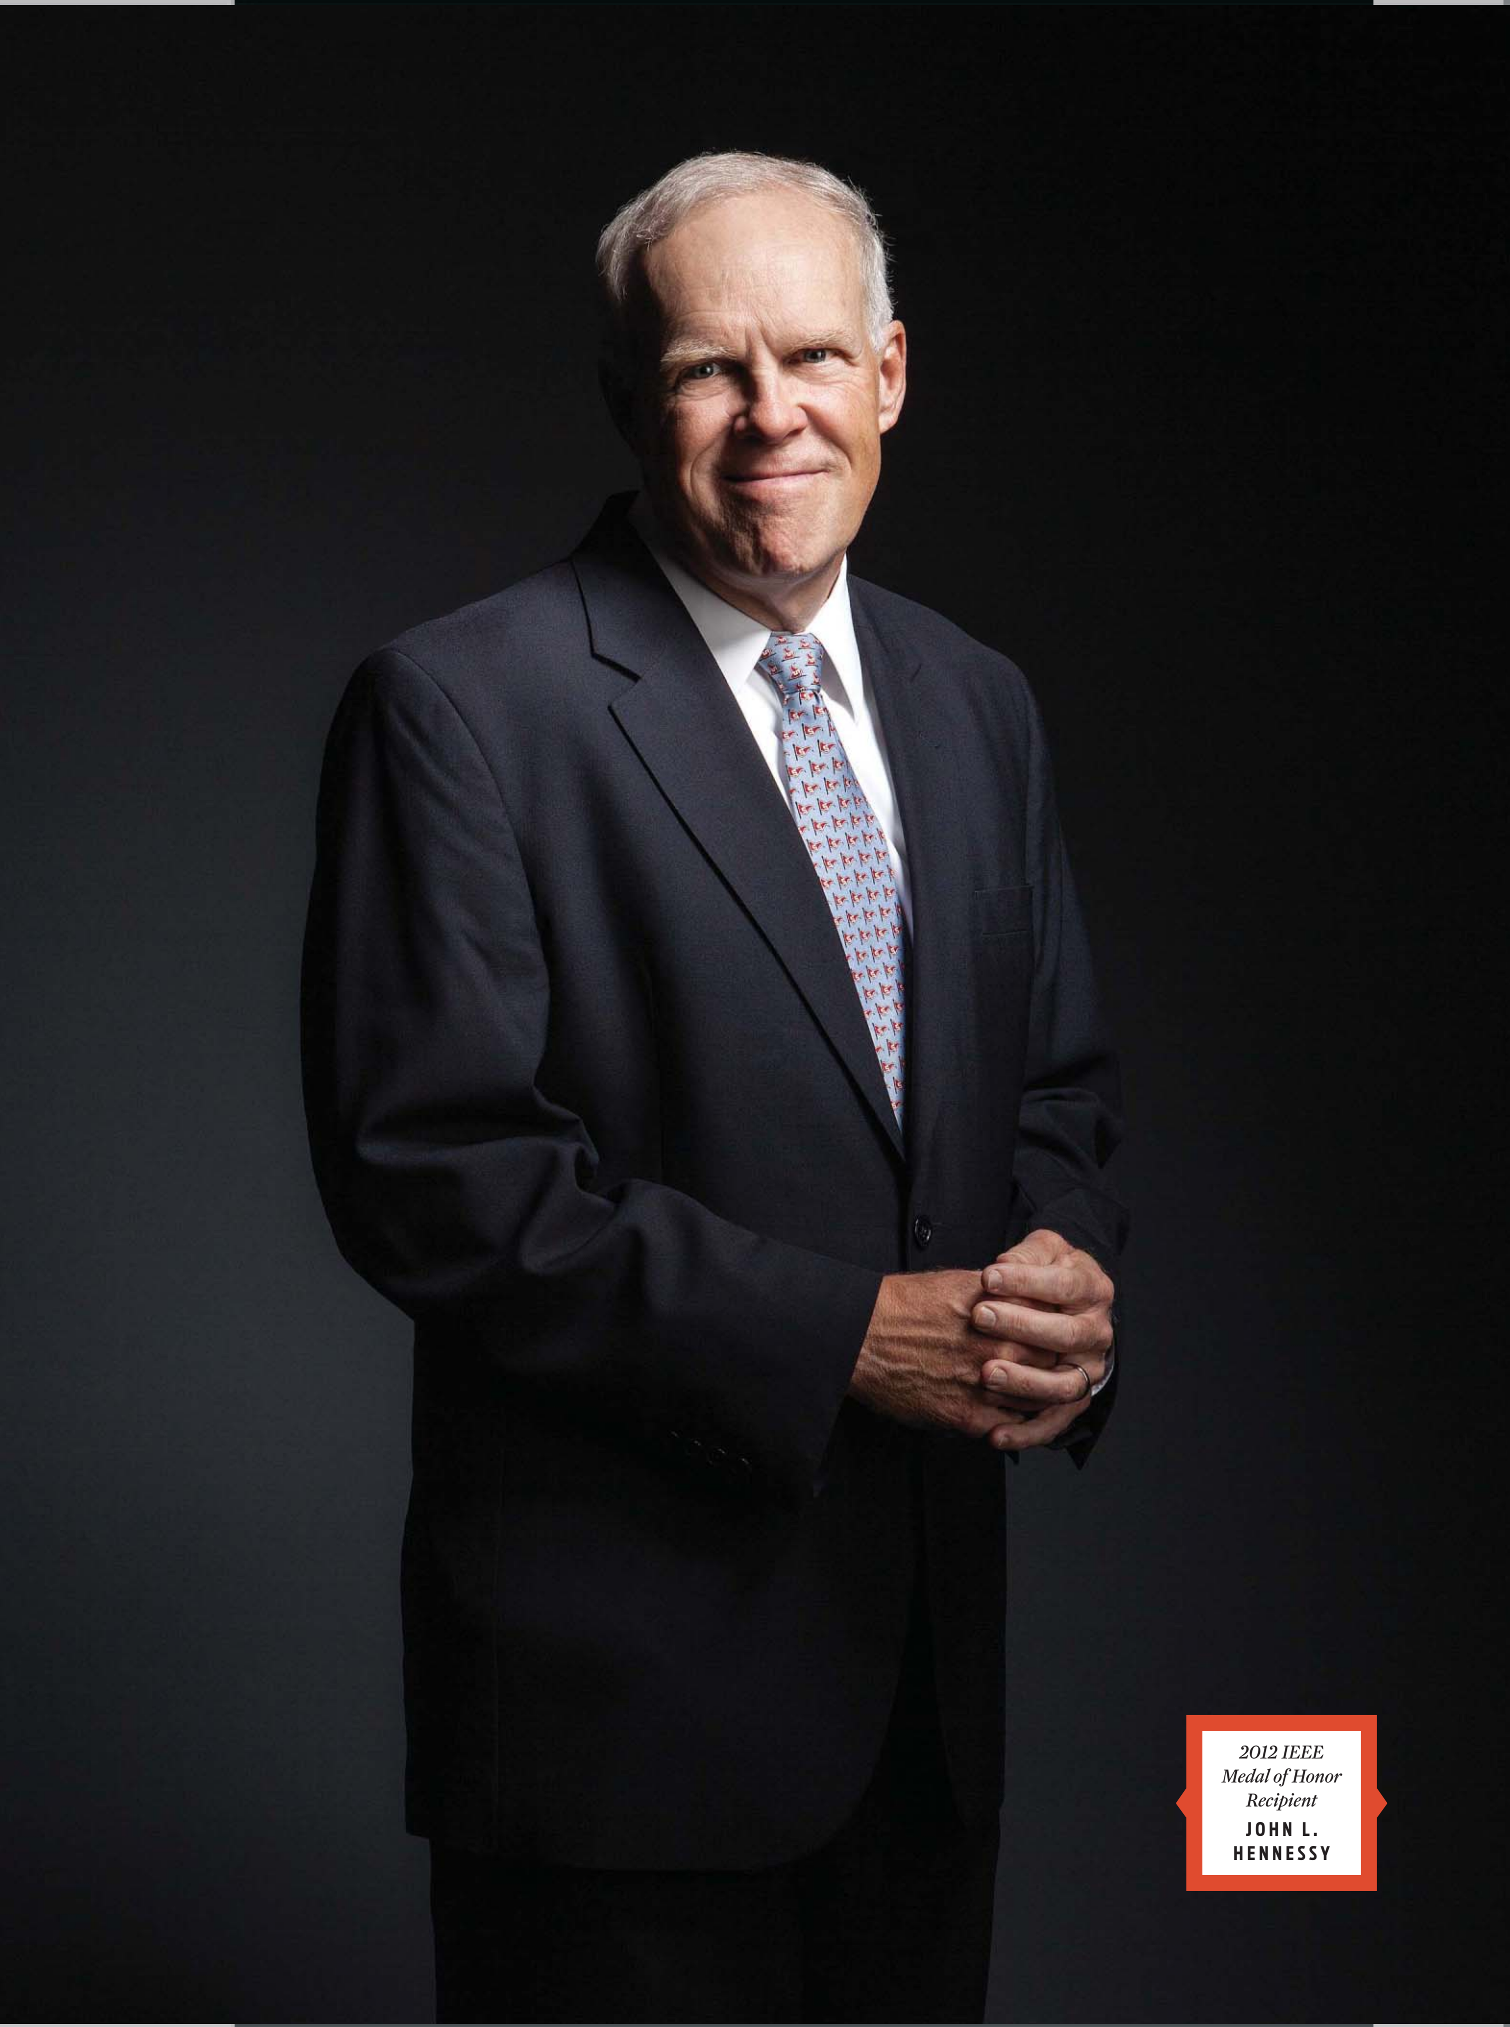

2012 IEEE Medal of Honor Recipient

JOHN L. HENNESSY

In the 1980s, John L. Hennessy, then a professor of electrical engineering at Stanford University, shook up the computer industry by taking the concepts of reduced instruction set computing (RISC) to the masses. Hennessy wrote papers, gave talks, designed chips, started companies, and even, literally, wrote the book (a textbook that's still used today). The RISC architecture, which focused on simpler, lower-cost microprocessors, was then thought to be an academic exercise with little practical use; today it plays a major role in the industry.

Hennessy, now president of Stanford, is once again designing, testing, and advocating a new architecture, this time in the field of university education. He first began rethinking research at universities and recently began reimagining university education itself.

For these efforts, in June Hennessy will receive the 2012 IEEE Medal of Honor "for pioneering the RISC processor architecture and for leadership in computer engineering and higher education."

IEEE Spectrum profiled Hennessy and his career as a computer architect and entrepreneur in "RISC Maker," [November 2002]. This year, we checked in on Hennessy's recent efforts to shake up higher education. Stanford has a long history in distance education, which in the 1990s moved from closed-circuit TV to Internet delivery. More recently, the university explored offering online courses to a much larger audience with a programming class for iPhone applications, first available in 2009, that has been downloaded more than one million times. Since then, Stanford has been developing and testing tools for producing, distributing, and enabling social networking for online courses. This past fall, more than 100 000 students around the world took three engineering classes—Machine Learning, Introduction to Artificial Intelligence, and Introduction to Databases. Hennessy says that's just the beginning. In fact, in his vision of the future, the lecture hall—those ubiquitous tiers of seats with fold-down writing arms, curving around a professor at a podium—will play a much smaller role.

But not everyone has applauded these educational experiments. One Stanford student who took the online machine learning class and blogged about it afterward said it didn't match the rigor of traditional courses. He also worried that a proliferation of online classes would reduce the value of a Stanford degree. Other critics take the opposite tack and wonder if the courses are too hard for the general public, noting that the vast majority of students who signed up for the online classes dropped out partway through.

Spectrum visited Hennessy at his office on Stanford's historic quad and asked him about his educational vision. Here's what he had to say:

If a video is simply a talking head, *it can be just as stupefying as sitting in a big classroom listening to a lecturer*

I'm a believer in online technology in education. I think we have learned enough about this to understand that it will be transformative. It's going to change the world, and it's going to change the way we think about education. Institutions like Stanford should be willing to fund the experiments, to try different things, to think about different models. We can do what other institutions would be strained financially to do, and they can learn from our experience.

It's going to filter down into high school, too, where we have an even more dramatic problem, considering the shortage of highly qualified high school teachers, particularly in science and math.

For us, it started a few years ago. One of my faculty colleagues, Daphne Kohler, said, "You know, I don't feel very useful when I stand in front of a classroom and give a set of lectures, 85 percent of which are the same as the year before. It's not very rewarding for the student, and sitting in large lecture halls is not the way students want to learn, particularly this generation." She pointed out that the large lecture hall is not a good learning environment, and it's not a good use of her time. And she was right. I agree that physical presence isn't all it's cracked up to be.

But it's not really about what I think. The students are rewriting the rules for us. That large lecture hall with nice banked seating and 300 people sitting with their attention focused on somebody standing in the front of the classroom is a model that lasted for many years, but the students have made it clear that that's not a model they find particularly attractive anymore.

Instead, this generation is completely comfortable watching a video online; for them, it's not markedly different than having a person up at the front of the classroom. They are happy using technology. They know how to hit the pause button; they know how to speed it up a little bit, to watch it 20 percent faster and make the process more efficient.

We started using technology to improve what was happening on campus. And we learned. We learned that if a video is simply a talking head with a set of PowerPoint slides, it can be just as stupefying as sitting in a big classroom listening to a lecturer. So what you need is more interactive learning. When Daphne Kohler started her experiment, she broke her lectures into chunks of 10 to 20 minutes and had a little miniquiz in between. And the miniquiz was able to identify that the student must have been asleep during that section or that there's one little thing to brush up on before moving on—hit this link and you'll go back to that place in the video.

We also learned that students like an online technology-mediated learning environment because it allows them to balance their lives. So if a student has a big project due for class XYZ this week, he may choose to fall behind one week of lectures in another course and then catch up on the weekend. Things like this can help with stress reduction, which is a big issue for many students.

And I think we learned that there is fun to be had with this online experiment. Last fall, a group of our faculty said, Well, why don't we just put these courses online, let anybody who wants take them. And who would have guessed you'd have more than 100 000 people sign up? Nobody. Courses like Machine Learning, Introduction to Artificial Intelligence, and Introduction to Databases are not of such obvious use as the iPhone programming course we put up a few years ago. That time, we could expect a lot of downloads because we knew people thought they'd soon be selling a widget in the App Store for a dollar a copy and making a lot of money. But Introduction to Artificial Intelligence isn't a skills-based course like that; it's a fundamental course. We were surprised how many students were willing to do the assignments and have them graded.

Of course, that creates a challenge, because it means that all assignments have to be graded automatically. There has been a lot of progress in grading technology. SAT writing exams [a standardized university admissions test], for example, are partially graded by computer right now. Each essay gets one grade from a computer and one from a person, with a second person looking at a test only if the scores differ significantly; logic courses have proof checking online. But at this point we don't know how far we can push grading technology and what is the right mixture of online and off-line grading.

We found that we were able to handle a lot of the Q&A through social networking. It's amazing that when you have 10 000 students in a class and a student puts up a question, the group quickly converges on the right answer: Several students put up an answer, other students come in and vote for what they think is the best answer, and there's a high probability that you'll converge to a pretty good answer in less than an hour. And if at the end of the two days there are three questions where the answers are not really clear, the instructor can come in and say, "This is the way you should think about this problem."

◆◆

Our experiments so far have raised questions we don't know the answers to. How much should we invest in production values? Take Sal Khan, of the Khan Academy online learning effort. His videos are pretty simple in terms of the production qualities; there's not a big investment. On the other hand, you go to anything that's done on TV, even something like "Myth Busters," and the investment in production is gigantic. Where on that spectrum do we really want to be?

One of the more interesting problems came up around problem sets. If you're teaching an on-campus course with 40 or 50 students or 100 students, you can probably live with a little ambiguity in the questions because there's usually a teaching assistant around. But if you're teaching a course online with 10 000 students and there's ambiguity in a homework question, it's a nightmare.

While we have experience using online technologies for teaching classes of up to a few hundred Stanford students, the experience of teaching a course with thousands or tens of thousands of students is very new. We are still very much in the early experiments. Teaching a smaller class of students who have been screened by our admissions procedures tends to ensure greater consistency in their abilities and allows more adaptation and feedback from instructors. Clearly, a class with 10 000 students allows almost no individual interaction and presents new learning challenges. There are also many experiences at Stanford that we do not see how we can re-create online, like small seminar courses and hands-on

Stanford Under Hennessy's Watch

Finances:
Stanford's endowment grew from US $8.9 billion in 2000 to $16.5 billion in 2011. Stanford received $7 billion in gifts during the fiscal years 2001 through 2011, including $1.1 billion in support of undergraduate education and $6.2 billion as part of the Stanford Challenge. These gifts fund more than 130 new endowed faculty appointments and more than 366 graduate fellowships and also provide $27 million in research seed grants.

Construction boom:
During the past 12 years, Stanford has completed more than 100 major construction projects, including academic buildings, housing, and athletic facilities. Notable is the new Science and Engineering Quad, a group of buildings that to date includes the Jerry Yang and Akiko Yamazaki Environment and Energy Building ($118.8 million), the Jen-Hsun Huang Engineering Center ($84.3 million), the Center for Nanoscale Science and Engineering ($77.5 million), and the Bioengineering and Chemical Engineering Building ($197.3 million). In 2011, the trustees approved more projects, including a $41.2 million computing facility and $438 million in improvements for campus energy systems.

project courses. For this reason, blended learning models may be very attractive.

It's important to keep in mind the difference between learning and credentialing. Universities—and high schools, for that matter—do both: They assist students in learning, and they provide a credential as evidence that a student has mastered certain material. I'm convinced that online learning will be widely used: The millions of downloads a day experienced by the Khan Academy and the broad interest in Stanford's online courses are clear evidence. But online credentialing is a new space and one that has much greater uncertainty. We have to consider a range of issues, from how credentials received online will be viewed by employers to how good online credentialing can be.

The biggest way this could go wrong *is if we assume that we can fully automate this, that we don't need any live instructors anymore*

◆◆◆

This can all add up to big changes in how we think about a university education. Could you imagine a situation where students do some of their degree work remotely and then do some on campus, particularly the small, experiential classes that don't transfer well to online?

This could be a way of keeping costs down, because you're living at home for a while or you're working part-time while you're doing the remote courses. So it could begin to address this issue of the cost of higher education. It's hard to predict how much you could reduce the cost of instruction, but it's important to understand that the cost of instructors is the No. 1 cost factor in education. If you could double the student-faculty ratio by reusing online material without reducing student learning, you could significantly reduce the cost of education.

◆◆◆◆

The biggest way this could go wrong is if we assume that we can fully automate this, if we think that we can throw students in front of terminals, that we don't need any live instructors anywhere, that students can be totally successful without ever talking to anybody. I think this is a model that will leave many students behind.

In a university, there is always a very small fraction of students who probably never need to come to class. They could just sit in their rooms, read the textbook, and they're capable enough, focused enough, disciplined enough, and driven enough that they could be successful. But that's a very small minority.

Likewise, there's a small minority of students who could watch everything online, never talk to anybody else, never engage with an instructor, never engage with teaching assistants, and learn just fine. But again, that's a very tiny minority.

The other potential problem is that we could see a lot of schlock out there. Because the Internet enables everybody to be a publisher, we could have a lot of things out there that are of lesser quality. Will people be able to distinguish quality in online education? This isn't resolved.

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/hennessy0512

We have to remember that the undergraduate experience is a whole experience. Part of that, particularly in the American higher education system, is the process of living and working with people of your own age. It involves learning how to work in teams, learning how to work with people with different backgrounds and different experiences. You certainly don't want to dismiss that part in a pursuit of doing it all online.

In spite of these hazards, I think online education and the role of technology in education are going to be transformative. And I'd like to think that what we've started here will not only continue at Stanford but that other universities will see it as a way to organize education and to play a larger role in the world.

To do this, however, universities have to be willing to change. Universities build on tradition and history, but they also have to be dynamic. And I think that struggle to balance those two opposing forces—to not become too attached to the past in such a way that you can't do something new, or to become too faddish in such a way that you lose your core values—is an ongoing challenge for all institutions.

But online education is going to happen; it's not going to wipe everything out, but it is going to happen. We have to embrace it. ❑

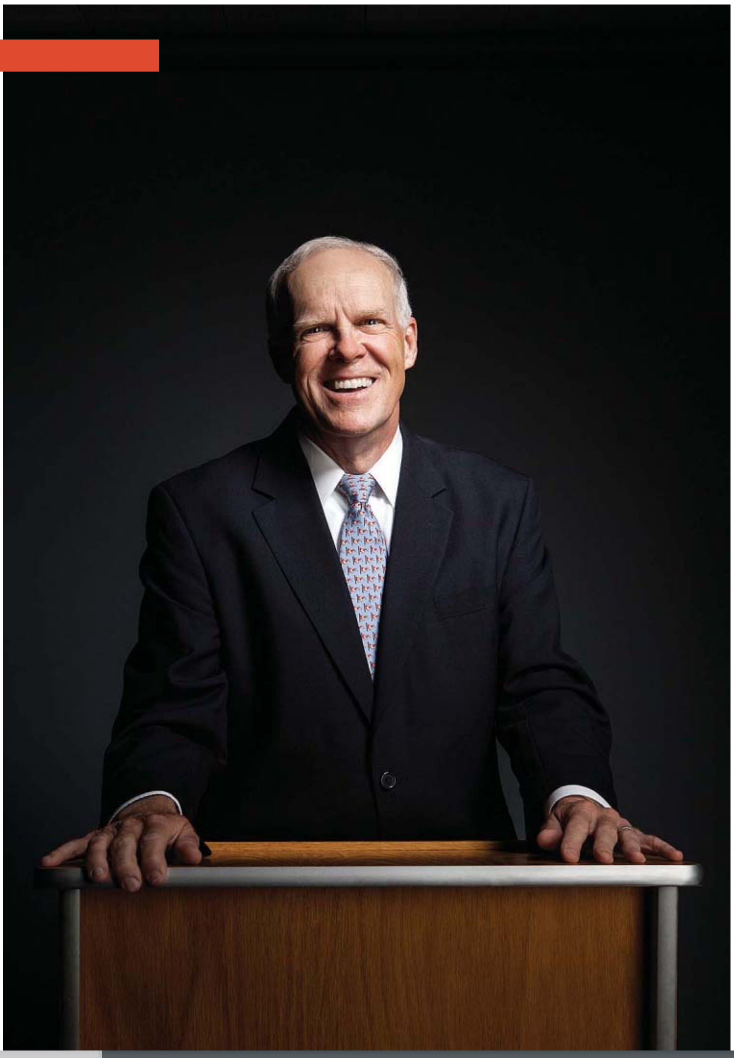

GABRIELA HASBUN


CST STUDIO SUITE 2012
Discover what happens...

Consider the kilogram

Two difficult experiments are poised to remake one of the world's most fundamental units

By Rachel Courtland

ROBERT RATHE/NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

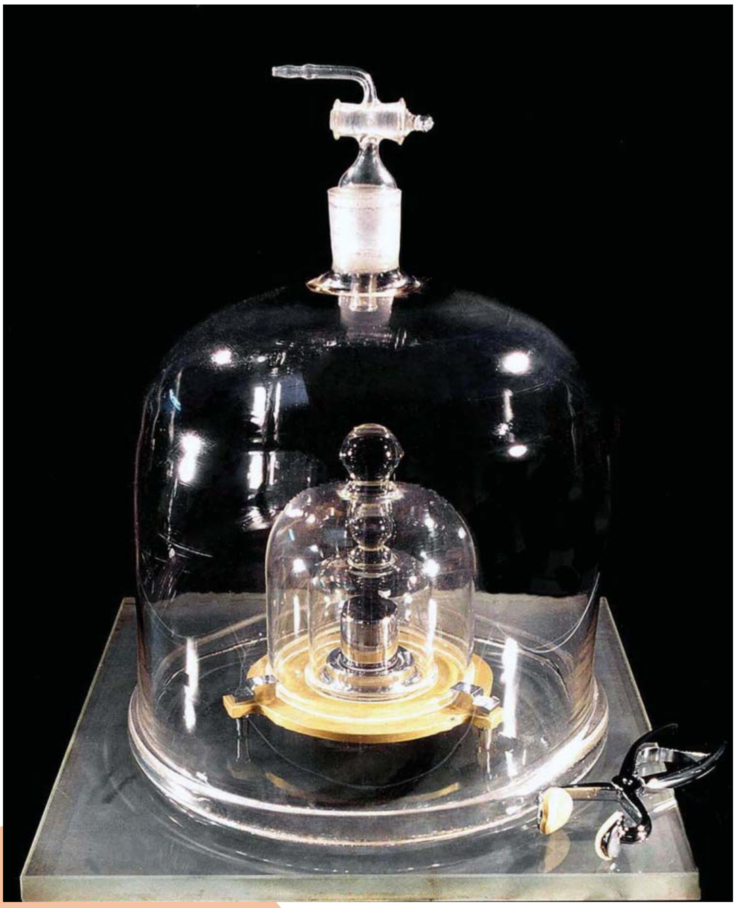

CYLINDERS AND SPHERES: The International Prototype of the Kilogram [left] is kept in a vault near Paris. One effort to redefine the kilogram measures ultrapristine silicon spheres. Two spheres are cut from a single crystal [top left] and polished [top right]; the spheres are then weighed with high precision against other objects in comparators like this one [bottom].

Once a year, three officials bearing three separate keys meet at the bottom of a stairwell at the International Bureau of Weights and Measures, in Sèvres, France. There they unlock a vault to check that a plum-size cylinder of platinum iridium alloy is exactly where it should be. Then they close the vault and leave the cylinder to sit alone, under three concentric bell jars, as it has for most of the past 125 years.

This lonely cylinder is the International Prototype of the Kilogram, known colloquially as Le Grand K, and it is the last remaining physical object to define a unit of measure. It's a quaint throwback to a time when people compared the ocean's depth to the span of a man's outstretched arms and the second to a tiny fraction of a year. Now we fix our rulers to the speed of light and our clocks to a spectral property of cesium. By thus linking measurement to fundamental and unchanging phenomena, scientists have paved the way for GPS satellites, gravity-wave detectors, and many other precision technologies that simply wouldn't have been possible before.

The trouble posed by the master kilogram is apparent in the many friction-filled steps by which it calibrates other masses. Once every few decades, a scientist plucks the cylinder from its perch with chamois-leather-padded pincers, rubs its surface with a cloth soaked in alcohol and ether, and steam-cleans it. Then he puts the prototype in a precise balance that compares it to the bureau's official copies, which are in turn compared to copies kept by member countries. And thus the prototype's mass trickles down to set the standard for the rest of the world.

The system has been far from seamless. When the cylinder was last removed from the vault in 1988, the bureau's metrologists were disappointed to discover that its mass and those of its official copies had drifted apart by as much as 70 micrograms since 1889. That discrepancy is tiny—comparable to the mass of a small grain of sugar—but it confirmed a troubling instability. All that metrologists can say is that the master kilogram seems to have lost as much as 50 μg over the course of a century relative to its siblings. But the actual drift could be up or down, and it might even be a lot more than 50 μg, because the prototype and its metallurgically identical copies could all be changing as an ensemble.

"It's a bit ridiculous in this day and age, because it's not just the mass that depends on the prototype. It's all energy, all force, all units that are linked in any way to the kilogram," retired metrologist Terry Quinn explains one gray afternoon at the bureau, about a week before an international committee was set to convene to decide the kilogram's fate.

A former director of the bureau (often referred to by its French acronym, BIPM), Quinn has been campaigning since the early 1990s to peg the kilogram to an unchanging aspect of nature. Such a change would be a boon to scientists who depend on stable units to perform long-term measurements, Quinn says. It could also have a big impact on electrical engineering, particularly for makers of precision multimeters and other basic tools.

When he began lobbying, Quinn says he thought it would take just a few years to come up with a better standard. But the effort to develop a precise way to measure mass with respect to a constant of nature has turned out to be phenomenally difficult.

CLOCKWISE FROM LEFT: BIPM; TILL TURSCHNER/FROM/LEIBNIZ-INSTITUTE FOR CRYSTAL GROWTH; CSIRO; BIPM

The October meeting marked a big turning point for the kilogram. Delegates from the bureau's then 55 member countries unanimously agreed on a tentative plan to base the kilogram on a fundamental constant of quantum mechanics. Three other core units—the ampere, the mole, and the kelvin—will likely change at the same time [see table, "Natural Units"].

This coup is largely the result, after decades of work, of steady strides in two challenging strategies for measuring mass. One approach attempts to pin down the exact electromagnetic force needed to balance the gravitational tug on an object. The other harnesses Cold War–era uranium enrichment technology and a host of experimental techniques to count the number of atoms in extremely round balls of ultrapristine silicon.

For years, the two approaches have produced starkly conflicting results. But over the past few months, metrologists have been excited to find glimmers of convergence, and the effort to pin down mass once and for all is beginning to pick up steam.

The scientists and intellectuals of the 18th century dreamed of a system of units based on something more fundamental and precise than the span of a hand or the mass of a seed. This notion solidified toward the dawn of the French Revolution, when Enlightenment thinkers shaped the beginnings of the modern metric system. They defined the meter as a small fraction of the circumference of the Earth. The kilogram they in turn derived from the meter, as the mass of a cubic decimeter of distilled water.

Linking units to the natural world was a lofty goal, but in the end, practicality demanded physical objects that could serve as references. The meter was a metal rod until the standard was retired in 1960, when the new measurement system—the International System of Units, or SI—pegged the meter to the wavelength of light spit out by atoms of krypton-86. The meter has since been redefined once again as the distance light moves through a vacuum in slightly more than three nanoseconds.

The kilogram's resistance to redefinition can in part be chalked up to politics. Getting an international consensus on a change to the SI takes time; one metrologist likened the process to steering a supertanker. But the big part of the delay comes from the sheer difficulty of measuring mass precisely with respect to unchanging standards of nature.

The challenge is most clearly seen in the effort to determine Planck's constant, the quantum-mechanical quantity that's set to become the new kilogram standard. Planck's constant governs the smallest amount of energy allowed by the laws of physics, and it appears in basic equations like $E=h\nu$, which relates the energy of a photon to the frequency of light. Because Planck's constant is linked to energy, it can also be linked to mass through Einstein's famous equivalence, $E=mc^2$.

The tool of choice for measuring Planck's constant is the watt balance. Rather than weigh one mass against another, such a balance weighs an object against the amount of electromagnetic force needed to keep the object in place. The electrical quantities measured by the balance can be linked back to Planck's constant using instruments that make precise quantum-electrical measurements.

A number of national labs are working on watt balances, but the largest one occupies a two-story, copper-foil-lined room in a remote building on the National Institute of Standards and Technology's sprawling, deer-studded campus northwest of Washington, D.C. NIST's project began in 1980, and the watt balance has been tweaked and overhauled so much that it's a matter of debate whether the instrument should be considered the third or fourth generation of the experiment. "There's a lot of Frankenstein here," says Jon Pratt, who took over management of the balance last year.

A mechanical engineer and former punk rocker with a soul patch, Pratt is a newcomer to the kilogram definition effort. He's been given the job of revamping and systematically studying the balance to see if it can be used to replicate the experiment's 2007 result, performed by Richard Steiner and others at NIST, which still stands as the most precise watt-balance measurement made to date.

During a December 2011 visit, I got a rare glimpse of the upstairs section of the balance. Ordinarily, the entire apparatus is housed in vacuum to eliminate error due to the buoyancy force of air, which can easily fluctuate with the temperature and pressure of the room. But the balance's vacuum chamber had been opened to allow technicians to work on a repair. The interior was a mess of wires, rods, and what looked like cantilevers. All the joints had been machined to be as loose as possible so that motions induced by gravity will always be oriented exactly downward. This became clear when, at one point, a guide accidentally nudged a component and set a few nearby parts jiggling as if they were Jello.

NIST's watt balance is the largest ever built, because its designers opted to build the magnet out of superconducting coil. But like all the world's watt balances, it works according to the principle outlined in 1975 by Bryan Kibble, then a physicist at the National Physical Laboratory in Teddington, England, who had been working on electrical balances

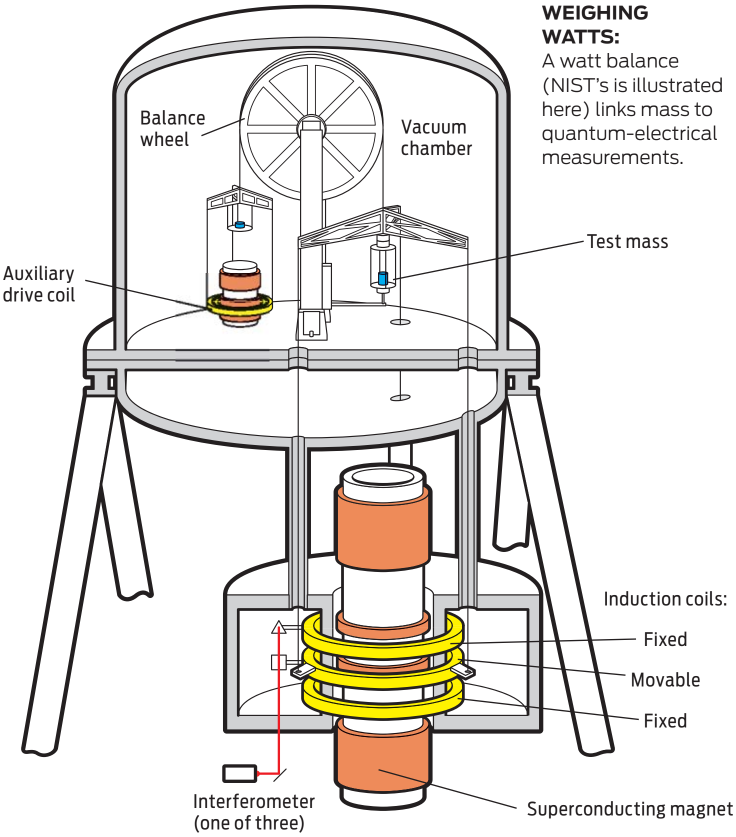

WEIGHING WATTS: A watt balance (NIST's is illustrated here) links mass to quantum-electrical measurements.

JAMES PROVOST

in order to make precision measurements of the ampere. A mass is placed on one part of the balance and a coil of conducting wire immersed in a magnetic field is used to balance the object's downward force. The experiment operates in two modes: In weighing mode, the coil remains static, and the machine is used to measure the current that must be run through the coil to enable it to maintain its position against the weight of the object. In the other mode, the current is switched off, and the coil is made to move up and down in the magnetic field, inducing a voltage across the coil.

Ideally, these two measurements can be combined to create a simple relation between mass, current, velocity, and voltage that does not depend in any way on the physical dimensions of the coil, which are difficult to pin down. The voltage and current in the coil can be measured with great precision by using calibration instruments that exploit the Josephson effect and the quantum Hall effect, quantum-mechanical effects that both depend on Planck's constant.

At the moment, watt balances like NIST's use objects of known mass, calibrated against the international prototype, to measure Planck's constant. In the future it'll be the other way around—the value of the Planck constant will be fixed, and watt balances will use this number to measure the mass of objects placed upon them.

That is, of course, if the physicists and engineers working on the experiments can make them accurate enough. The NIST group's 2007 measurement of the Planck constant pegged its value with an estimated precision of 36 parts per billion. In isolation, it was an impressive result that put the experimental uncertainty close to a target of 20 parts per billion. (The international prototype, however much it might be drifting, can be compared to other objects with about 10 times that precision.)

But results from another watt balance, now housed at a laboratory at the National Research Council of Canada, in Ottawa, tell a different tale. Last year, the Canadian team reported its own precise measurement of the Planck constant, which differs from the NIST group's by some 250 parts per billion. When the two conflicting results are plotted on the same graph, their error bars don't overlap.

Clearly, both experiments can't be right. But in some ways this discrepancy isn't all that surprising. Today's best watt balances are extraordinarily sensitive: Metrologists have found they can pick up earthquakes, magnetic interference from passing trains, the gravitational tug exerted by nearby snow cover, and even subtle shifts in ground level created when wind passes through nearby trees. Many of these sources of error are random and transient and can be eliminated by simply running the experiment longer. But there are still dozens of other quantities that must be carefully measured in order to get the right result. "The error budget is large," says Barry Wood, who manages the Canadian watt balance. "It's over 50 items that have to be calculated and assessed, and each of those is an experiment in itself."

Even the most careful work can overlook large sources of error. Before arriving in Ottawa in 2009, Canada's watt balance was built and run for nearly two decades across the Atlantic at the watt balance's birthplace—the National Physical Laboratory. Weeks before the instrument was set to be boxed up at the British lab, physicist Ian Robinson was performing last checks on the experiment when he discovered an overlooked issue: The angle of the balance beam changed when a mass was added, in part because some small metal joints, used to align the coil, had been compressing. The hard-to-detect change could offset measurements of Planck's constant by a significant amount. "That was one of those moments where the penny dropped, and it was not very pleasant," Robinson says. "But I was very, very glad that I thought of it."

The Canadian team has been working on a modification to fix the problem and is now partnering with the American team to resolve the discrepancy between the two results. But progress is slow. In February, representatives from both teams met at NIST to collaborate on just one of the corrections that must be made to the measurements: the exact value of gravitational force at the exact point where the mass sits in its pan. In the case of NIST's tall watt balance, which has a pan that sits a good 5 meters above ground level, the necessary correction can be dozens of times as big as the targeted precision of the machine. To resolve the two experiments, the watt-balance teams will likely have to run through dozens of other potential sources of error. All eyes are now on NIST, which is hoping to get a new result—the first in five years—later this year. Waiting in the wings are other watt-balance projects at the BIPM and national labs in China, France, New Zealand, and Switzerland.

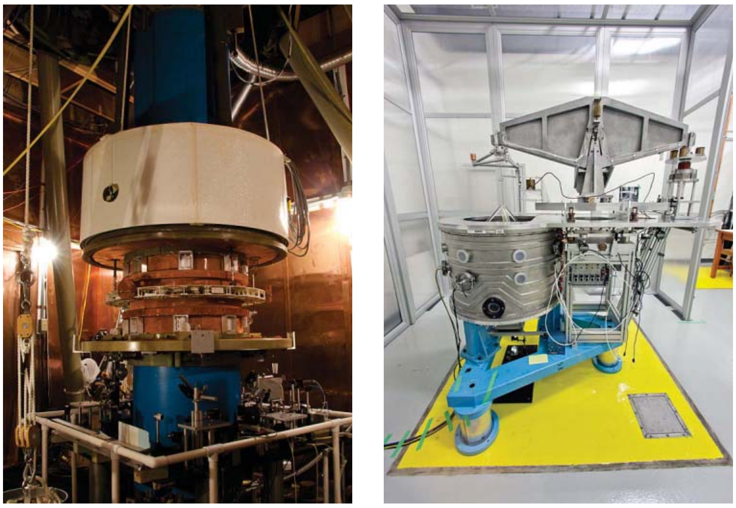

PURSUING PLANCK: The NIST watt-balance [left] and the NRC watt-balance [right] experiments are two leading efforts to measure Planck's constant.

LEFT: NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY; RIGHT: NATIONAL RESEARCH COUNCIL CANADA

Collaboration is new for the two leading watt-balance teams, which for years have had to operate on opposite sites of the Atlantic. But it's been a vital part of another effort to measure mass—by linking it not to Planck's constant but to the mass of an atom.

The effort centers around the Avogadro constant, a quantity often used in chemistry to link the masses of macroscopic objects to the mass of their molecular or atomic constituents. Avogadro's constant is currently defined as the number of atoms in a mole of a substance, and its value is identical to the number of carbon-12 atoms needed to make up exactly 12 grams. But Avogadro's constant can just as easily be used to work from any sort of atom all the way up to the kilogram.

Silicon is a natural candidate for such measurements. Thanks to the semiconductor industry, we can grow silicon crystals that are pretty much structurally perfect, making them ideal for measurements of basic crystal properties. But beating the uncertainty in Avogadro's constant down to a level that could enable it to replace the international prototype is a different story.

Natural silicon isn't pure. Although its dominant isotope is silicon-28, nearly 8 percent of it consists of the heavier isotopes silicon-29 and silicon-30. That slight contamination gives rise to a big uncertainty in silicon's molar mass. In the early 2000s, Peter Becker, who was spearheading efforts to measure Avogadro's constant at Germany's metrological institute, the Physikalisch-Technische Bundesanstalt, realized that natural silicon wouldn't work. But fortune intervened. A colleague from what was once East Germany suggested the team use a Russian uranium enrichment facility to purify some material. Becker raised €2 million to get the material. After about half a year and the efforts of roughly 250 gas centrifuges, the international team that had pitched in money was presented with 6 kg of 99.995 percent silicon-28.

After a crystal was grown from that stock, it was cut into two spheres, which were polished to such extreme roundness that repeated measurements of a single parameter—the sphere's diameter—could be used to determine its volume. The spheres have since been toted from lab to lab in a transparent acrylic suitcase, always carried by hand onto airplanes because the objects are too precious to ship. Over the course of two years, an international team worked to pin down every property of the spheres. A reflective cavity was used to establish the diameter, and thus the volume, of each sphere. Teams at the BIPM and in Germany and Japan measured its mass against official kilogram copies. In Italy, a team used X-rays to work out the exact spacing between atoms in the crystal. In Switzerland, researchers characterized the chemical composition of the thin layer of quartz-like oxide that had grown around them. A team in Germany developed a novel dilution technique to work out the exact isotopic composition, which determines the molar mass of the spheres.

Last year, the Avogadro team published their first results, which pinpointed Avogadro's constant to a precision of 30 parts per billion. The precision was a fabulous improvement over the old results, but when it was converted to Planck's constant, the measurement was disappointingly off—it didn't agree with the American watt-balance measurements and had only a slight overlap with the Canadian watt-balance results. Since then, two teams of chemists in the United States and Canada have joined the effort to measure the molar mass of the silicon crystal. The measurement is difficult—the natural silicon in laboratory glassware or even in the dust in the air can throw off measurements, says Greg Turk, who is leading the effort at NIST. But early results from the Canadian team have pushed things closer to the results of the Canadian watt balance, suggesting that a convergence might just be under way.

The international committee in charge of mass has set some exacting standards for the kilogram experiments. At least two experiments with the watt balance, together with the Avogadro project, must demonstrate an accuracy equivalent to 50 parts per billion. One experiment must demonstrate an accuracy of at least 20 parts per billion. All the results must agree with one another.

So far, the NIST and Avogadro experiments meet the first requirement, but the second requirement is still out of reach, and there are unexplained differences between existing results. Some express bafflement over the stringency of the requirements. "Nobody talks about this, but the uncertainty in uncertainty is huge," says NIST's Pratt. "Sweating over 20 parts per billion or 36 parts per billion seems a bit pedantic."

The requirements are particularly nonsensical, says Quinn, given the fact that the real uncertainty in the mass of the international prototype is ignored. Despite its known drift with respect to its copies, the uncertainty in the mass of the kilogram is, by definition, zero. "The point I take and I emphasize is that our knowledge of the absolute mass of the kilogram is so poor that almost anything else is better," Quinn says. "I think that we could do [the redefinition] today and we would be better off."

One immediate benefit could be to electrical measurements. In the SI, the ampere is still impractically defined as the force between infinitely long conductors. Since 1990, those in need of precise voltage and resistance measurements have used a separate system of units, which employ what are now somewhat outdated values for several fundamental constants. Pegging the kilogram to Planck's constant—and the ampere to fundamental charge—will bring electrical units back into the fold.

The kilogram hunters are now eyeing 2014—the next General Conference on Weights and Measures—as the year when all their hard work will finally pay off. But it's far from guaranteed that the experiments will agree in time to meet such a deadline. Nonetheless, metrologists are getting ready.

Sometime soon, the keepers of the kilogram will take the cylinder out of its enclosure for the first time in more than 20 years. Once again, they will measure it against its copies, in the hope of establishing as close a link as possible between the present-day mass of the cylinder and the results of the experiments that would replace it. The retiring of the prototype will be slow and careful—which is only fitting for an object that has for so long supported the weight of the world. ❏

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/kilogram0512

Natural Units

Four of seven base units in the International System of Units could change as early as 2014

	Present	Proposed
A	**ampere (current)** The current required to produce a certain force between two infinitely long parallel conductors placed 1 meter apart in vacuum.	The ampere will be set by fixing the numerical value of *e*, the fundamental unit of charge.
K	**kelvin (temperature)** A certain fraction (1/273.16) of the thermodynamic temperature of the triple point of water.	The kelvin will be set by fixing the numerical value of *k*, the Boltzmann constant, which has units of joules per kelvin.
kg	**kilogram (mass)** The mass of the International Prototype of the Kilogram, kept in a vault in Sèvres, France.	The kilogram will be set by fixing the numerical value of *h*, the Planck constant, which has units that contain the kilogram.
mol	**mole (amount of matter)** The amount of a substance that contains as many elementary entities as there are atoms in 0.012 kilogram of carbon-12.	The mole will be set by fixing the numerical value of the Avogadro constant.

The Dawn of Haiku

How a volunteer crew brought A crack OS back

By **RYAN LEAVENGOOD**

It was the summer of 2001, and computer programmer Michael Phipps had a problem: His favorite operating system, BeOS, was about to go extinct. Having an emotional attachment to a piece of software may strike you as odd, but to Phipps and many others (including me), BeOS deserved it. It ran amazingly fast on the hardware of its day; it had a clean, intuitive user interface; and it offered a rich, fun, and modern programming environment. In short, we found it vastly superior to every other computer operating system available. But the company that had created BeOS couldn't cut it in the marketplace, and its assets, including BeOS, were being sold to a competitor.

Worried that under a new owner BeOS would die a slow, unsupported death, Phipps did the only logical thing he could think of: He decided to re-create BeOS completely from scratch, but as open-source code. An open-source system, he reasoned, isn't owned by any one company or person, and so it can't disappear just because a business goes belly-up or key developers leave.

Now if you've ever done any programming, you'll know that creating an operating system is a huge job. And expecting people to do that without paying them is a little nuts. But for the dozens of volunteer developers who have worked on Haiku, it has been a labor of love. In the 11 years since the project began, we've released three alpha versions of the software, and this month we expect to release the fourth and final alpha. After that we'll move to the beta stage, which we hope to get out by the end of the year, followed by the first official release, known as R1, in early 2013.

Even now, anybody can install and run the operating system on an Intel *x*86-based computer. Many of those who have done so comment that even the alpha releases of Haiku feel as stable as the final release of some other software. Indeed, of all the many alternative operating systems now in the works, Haiku is probably the best positioned to challenge the mainstream operating systems like Microsoft Windows and Mac OS. For both users and developers, the experience of running Haiku is incredibly consistent, and like BeOS, it is

ANDREAS TEICHMANN

TO BE: Software developers gathered recently in Düsseldorf, Germany, to code and collaborate on Haiku, an open-source effort that replicates—and improves on—an old operating system called BeOS.

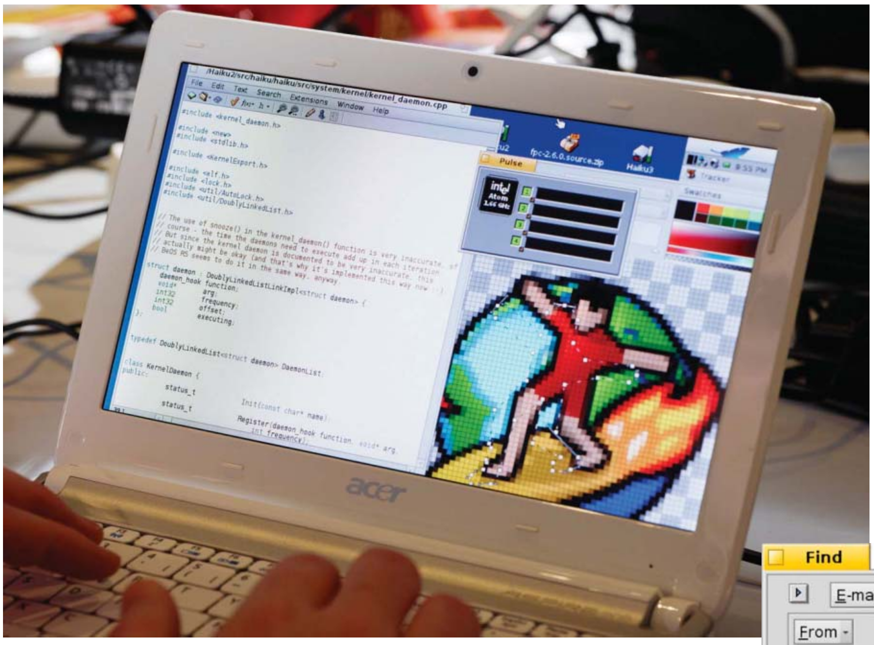

HAIKU CLOSE-UP: Though it's based on an older operating system, Haiku can run on present-day computers [left] as well as older machines that would otherwise be considered obsolete. Haiku's database-like file system organizes e-mail and other files according to their attributes [below], for easier searching. It also uses computer resources extremely efficiently and takes up very little storage space [middle left]. And its multithreaded design means that running many applications at once won't freeze up your machine [bottom].

IMAGES: TOP, ANDREAS TEICHMANN; OTHERS, HAIKU INC.

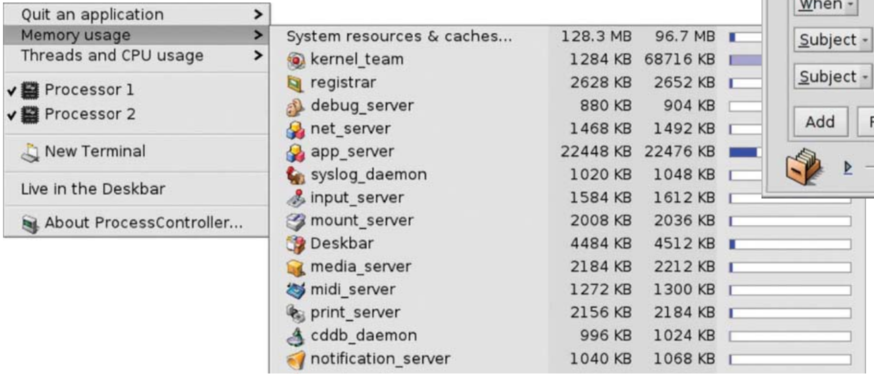

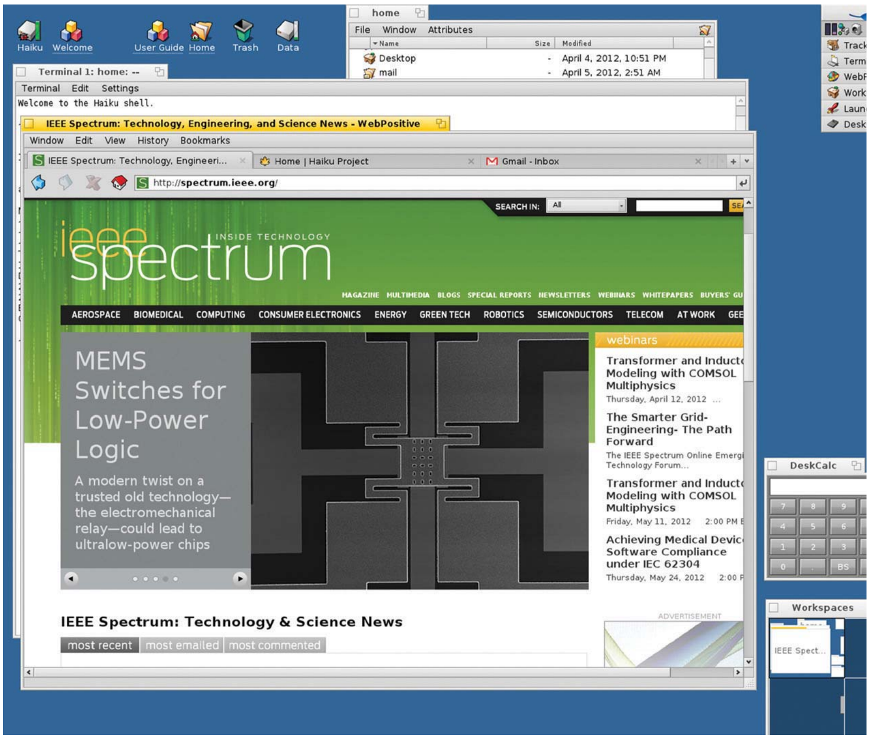

fast, responsive, and efficient. What's more, Haiku, unlike its more established competitors, is exceedingly good at tackling one of the toughest challenges of modern computing: multicore microprocessors. Let's take a look at why that is, how Haiku came to be, and whether the operating system running on your computer really performs as well as it should.

First, a little history. In 1991, a Frenchman named Jean-Louis Gassée and several other former Apple employees founded Be Inc. because they wanted to create a new kind of computer. In particular, they sought to escape the backward-compatibility trap they'd witnessed at Apple, where every new version of hardware and software had to take into account years of legacy systems, warts and all. The company's first product was a desktop computer called the BeBox. Finding no other operating system that met their needs, the Be engineers wrote their own.

Released in October 1995, the BeBox didn't last long. BeOS, on the other hand, quickly found a small yet loyal following, and it was soon running on

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

Intel *x*86-based PCs and Macintosh PowerPC clones. At one point Apple even considered BeOS as a replacement for its own operating system. The company eventually released a stripped-down version of BeOS for Internet appliances, but it wasn't enough. In 2001, Palm acquired Be for a reported US $11 million.

Even as the sale of Be was being finalized, Phipps launched the Haiku project, which back then was known as OpenBeOS. (The name changed a few years later; see box, "What's in a Name?")

From the start, Phipps decided that the project's singular focus would be to replicate the last official release of BeOS, known as R5. In hindsight, this was a stroke of genius. Open-source efforts, generally done for fun in people's free time, can sometimes fizzle out or get sidetracked without ever releasing anything tangible. Haiku has avoided this fate because the developers all believed in the goal of replicating R5.

Still, Phipps and those who joined him had their work cut out for them. An operating system is extremely complex, especially one as comprehensive as BeOS, whose various layers and applications had all been designed from the ground up to work together. Some of the most gifted engineers in Silicon Valley had developed the software over the course of more than 10 years. Re-creating such a system with an all-volunteer crew working in their free time was a crazy idea. But then, that's how Linux came to be, as well as the Free Software Foundation's GNU software, from the Gnu C Compiler to GnuCash, an accounting program.

Much like the layers of a cake, a computer system consists of the hardware, an operating system that manages the hardware, and the applications that run within the operating system, such as Web browsers, document editors, and fun things like games. The operating system is also the means by which programmers give instructions to the hardware.

Fortunately, BeOS had been written in a modular fashion, making it relatively straightforward to develop, test, and then replace each component of BeOS with its open-source equivalent. A few pieces of BeOS had already been released by Be as open source, such as the Tracker and Deskbar, which are equivalent to Windows' Explorer and Taskbar and to OS X's Finder and Dock. One of the first chunks of code that the volunteers tackled was the screen-saver kit, which has a very simple function but also a lot of moving parts. Among other things, it has to constantly monitor activity on the keyboard and mouse, load the screen-saver settings at the right times, and ask for a password when the screen saver is turned off. Once all the parts had been nailed down, Phipps replaced the BeOS screen saver with the open-source version, and much to everyone's surprise and delight, it worked. More important, it was a proof of concept that showed the developers they could replace each module in BeOS and know it would be fully compatible.

In 2008 we reached a milestone. Haiku, like any piece of software, is written in source code; to convert the source code into binary instructions that the computer can process, you need to compile it. The breakthrough was that we were finally able to compile the Haiku source code from within Haiku itself; programmers refer to that feat as "self-hosting." This step is critical in the development of any operating system, because without it, the OS will always be dependent on other systems.

Even as the development was getting under way, Phipps was working to create an organizational home for Haiku, eventually establishing Haiku Inc. as a nonprofit organization in 2003. In addition to overseeing the project and taking in donations, Haiku Inc. holds the copyrights to and the trademarks of the Haiku website, logo, and of course, the source code. All Haiku source code is now licensed under what's called the MIT License, which allows full and free use of the code by pretty much everyone, including private companies and other open-source projects.

Of course, the Haiku team also has to be mindful of not infringing on BeOS's patents, which are still in force. In general, we maintain a friendly relationship with the Japanese software company Access, which now owns the rights to BeOS; for example, the company allows us to post BeOS documentation on our website, and it has never challenged any of our efforts to replicate BeOS technology. Would Access or some other corporation that thinks it has a right to our code choose to sue a small open-source project like ours? Financially, it wouldn't make much sense, and it would also generate a lot of bad publicity. But we'll just have to wait and see. Given the current frenzied state of intellectual property litigation in the United States and elsewhere, no software project can be considered completely immune from legal troubles.

In 2007, Phipps announced he was leaving the project for personal reasons, and a new group of people took over Haiku Inc. Much as Phipps had anticipated at the outset, the loss of one leader hasn't doomed the project. I'm currently the treasurer and, along with the other executives of Haiku Inc., sit on its board of directors. We've kept the momentum going, and at their latest gathering in Düsseldorf, Germany, in early April, Haiku developers worked on the imminent alpha release.

At this point, we've replicated BeOS R5 so precisely that legacy BeOS applications, now more than 10 years old, can run on Haiku. When the project first started, many thought this sort of compatibility would be impossible. And in many key areas, Haiku has surpassed its predecessor. For example, Haiku supports more languages than BeOS did and is more amenable to internationalization in other ways as well. It can also handle modern video cards, newer processors, and wireless network access, and it supports more memory than BeOS could.

And in terms of programming, Haiku makes it much easier to design user interfaces for applications, because it has a built-in layout tool that lets you automatically place icons and other widgets on the screen. With BeOS, developers had to specify each minute detail of the layout, aligning buttons with check boxes by hand, for instance. These and other enhancements have allowed Haiku to remain relevant, despite the lightning-fast pace of innovation in computer hardware and software.

What's in a Name?

When the project to re-create BeOS began, it was known as OpenBeOS. But "BeOS" was trademarked, so a new name was needed. A naming contest held in 2002 drew a large number of entries, of which Haiku carried the day. The name harked back to an eccentricity in the wording of the BeOS Web browser's error messages: They were all rendered as haiku, the classic Japanese poetry form. So, for example, if a website could not be found, one error might be:

These three are certain:
Death, taxes, and site not found.
You, victim of one.

A list of BeOS error messages is available online at www.8325.org/haiku. There are plans to incorporate the messages into Haiku's Web browser. —*R.L.*

Even so, you may be wondering: With Windows, Mac OS X, hundreds of versions of Linux, and numerous mobile operating systems, does the world really need Haiku?

Yes, it does, for many reasons. Much as in nature, computer viruses flourish in a monoculture, and because so many of us use Windows computers, Windows viruses find plenty of victims. Further, if one operating system has a monopoly, its creators have little reason to improve their software. (Competition from the Mozilla Firefox browser is, after all, largely what prompted Microsoft to update Internet Explorer.) And diversity promotes interoperability because it compels software developers to create code that plays well with others. When there's very little competition, there's no incentive to do so.

But Haiku does more than just expand the gene pool of operating systems. One of the first things people notice about it is that it doesn't feel anything like Windows or OS X or Linux. It's unique. Linux, for instance, is based around a core—called a kernel—that was originally designed for use in servers and only later modified for desktop systems. As a consequence, the kernel sometimes gives short shrift to the user interface, which Linux users experience as annoying delays when their computers are doing especially taxing things, like burning a DVD or compiling code. Haiku's kernel has always been for a desktop system, and so it always gives priority to whatever is happening by way of its graphical user interface.

In between the kernel and the graphical user interface is what's known as the application programming interface, or API. An API is what application developers use to access other software systems, and the design of that API can affect both the developers and the eventual users of their software. Haiku has just one API. Linux, by contrast, has hundreds of APIs and about as many user interfaces, so you can't just switch seamlessly from one version of Linux to another. The version you use can greatly alter the appearance of your computer screen, the way programs boot and execute, and various other things, all of which make it quite difficult to develop software that consistently runs well on all Linux systems.

All of Haiku's components are designed to work together from the ground up; that includes its applications, like the media player and the Web browser. Its source code even uses a consistent style, something developers really appreciate because it allows them to quickly get up to speed. A developer writing code for Haiku can be confident it will run and act the same way on all Haiku installations. Even though Haiku is open source, there will always be only one official version. All this combines to provide a remarkably consistent user experience.

What really sets Haiku apart, though, is its efficiency and speed. On my three-year-old desktop computer, Haiku boots up from a cold start in 10 to 15 seconds, and when running it uses only about 170 megabytes of RAM. Compared with other operating systems, it's great at maximizing available computing resources. Haiku feels fast and responsive even when running on older systems that would otherwise be considered obsolete because they can no longer handle the inefficiency and bloat of other operating systems.

Much of Haiku's efficiency and speed is a direct result of its BeOS legacy. BeOS was designed from the start to make full use of threads, which in computing terms represent sequences of executing code. Much as fabric consists of many threads woven together, an operating system is made of threads that all have to share time on the processor. Generally, there is one thread for each application and one for the operating system's user interface as well. The problem is that the user interface in particular needs more than one thread. Those "application is not responding" messages in Windows and the "beach ball of death" in OS X—which most of us experience with some regularity and frustration—are a direct result of using just one thread for the user interface. When that thread gets too backed up with work or has another problem (like a slow network), the whole application interface locks up.

That didn't happen on BeOS, and it doesn't happen on Haiku. Where BeOS drove ahead of other operating systems of its time (and where Haiku is still ahead of contemporary operating systems) is that each individual application uses many threads. The core of the application itself has one thread, and each window the application creates has a thread. While this can make the code a little more complicated to write, the result is that applications almost never lag or lock up.

This highly threaded system means that Haiku can make great use of multiple processors or CPU cores, which are now ubiquitous. The Haiku kernel allows each thread to run on its own core or processor—even threads within the same application. Adding more threads for other tasks, such as retrieving data from a network, is also straightforward. Other operating systems make use of multiple cores only when running many applications at the same time or when a particular application has special multithreading code added to it. But this multithreaded code is difficult to write, and so most applications don't have it.

Haiku simplifies the process of writing multithreaded code by hiding most of the thread interaction so that the developer doesn't need to worry about it. A big part of what makes this work is the extensive use of message passing. Let's say Thread A and Thread B in an application both want to retrieve a piece of data from Thread C. Instead of directly accessing that data, Threads A and B *Continued on page 49*

6 Cool Things About the Haiku System

- It's free to use and designed to maximize user freedom.
- It's open source, with a very reasonable license. Companies can use it for commercial efforts with minimum fuss.
- It makes extensive use of threads, leveraging multicore processors and requiring minimal effort from developers.
- It's lean, fast, and efficient, providing new life for older machines.
- Its database-like file system makes searching and organizing a snap.
- Its design is consistent and integrated, from the source code to the kernel to the user interface.

◀ **THE BeBOX:** it didn't last, but its operating system lives on.

RENE GOLLENT

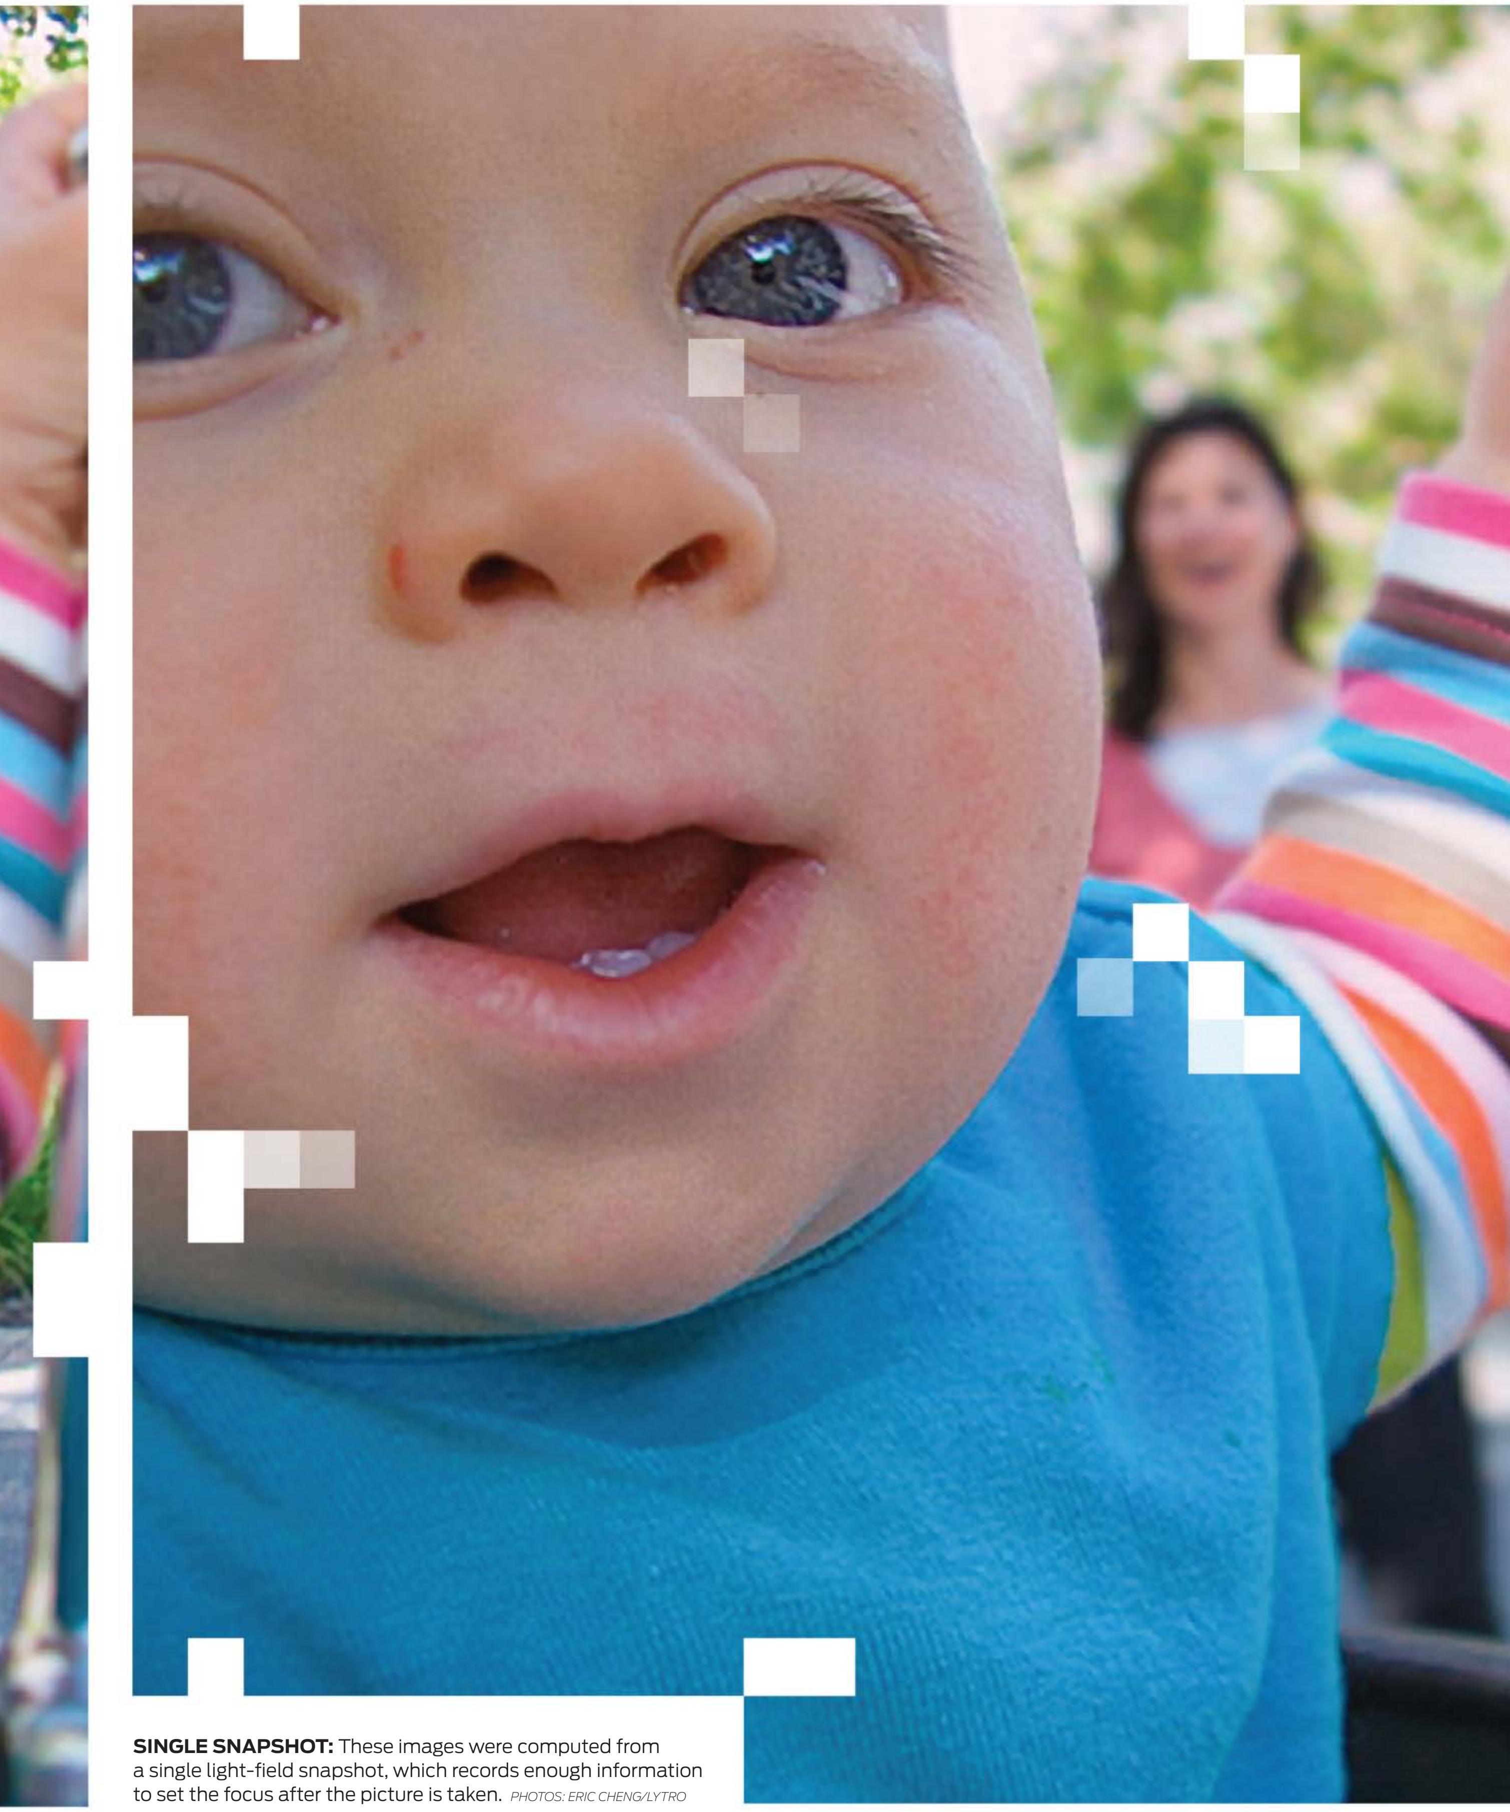

SINGLE SNAPSHOT: These images were computed from a single light-field snapshot, which records enough information to set the focus after the picture is taken. *PHOTOS: ERIC CHENG/LYTRO*

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

Focusing on Everything

Light-field cameras promise an imaging revolution

by Mark Harris

Leonardo da Vinci sketched out tanks, helicopters, and mechanical calculators centuries before the first examples were built. Now another of his flights of imagination has finally been realized—an imaging device capable of capturing every optical aspect of the scene before it.

Lytro, a Silicon Valley start-up, has just launched the world's first consumer light-field camera, which shoots pictures that can be focused long after they're captured, either on the camera itself or online. Lytro promises no more blurry subjects, and no shutter lag waiting for the camera's lens to focus. A software update to the camera, coming soon, will even let you produce 3-D images.

Light-field technology heralds one of the biggest changes to imaging since 1826, when Joseph-Nicéphore Niépce made the first permanent photograph of a scene from nature. A single light-field snapshot can provide photos where focus, exposure, and even depth of field are adjustable *after* the picture is taken. And that's just for starters. The next generation of light-field optical wizardry promises ultra-accurate facial-recognition systems, personalized 3-D televisions, and cameras that provide views of the world that are indistinguishable from what you'd see out a window.

But light-field cameras also demand serious computing power, challenge existing assumptions about resolution and image quality, and are forcing manufacturers to rethink standards and usability. Perhaps most important, these cameras require a fundamental shift in the way people think about the creative act of taking a photo.

In his manuscripts on painting, Leonardo wrote, "The air is full of an infinite number of radiant pyramids caused by the objects located in it. These pyramids intersect and interweave without interfering with each other....The semblance of a body is carried by them as a whole into all parts of the air, and each smallest part receives into itself the image that has been caused."

Nowadays, scientists and engineers prefer to think in terms of light rays rather than Leonardo's more poetic "radiant pyramids." But light-field photography is based precisely on his idea that the light arriving at any point—what he called the "smallest part" of the air—carries all the information necessary to reproduce any view that can be had from that position.

Doesn't an ordinary camera do that? Not at all. In a conventional digital camera, the light rays hitting each point on the image sensor combine. The sensor records the total intensity of the light rays landing on each point, or photosite, but in the process loses directional information about where the different rays came from. So the best a typical camera can provide is the familiar two-dimensional photograph, which has a fixed point of view and a focus determined entirely by how the lens was set when the photo was snapped.

Light-field photography is far more ambitious. Instead of merely recording the sum of all the light rays falling on each photosite, a light-field camera aims to measure the intensity and direction of every incoming ray. With that information, you can generate not just one but *every possible image* of whatever is within the camera's field of view at that moment. For example, a portrait photographer often adjusts the lens of the camera so that the subject's face is in focus, leaving what's behind purposefully blurry. Others might want to blur the face and make a tree in the background razor sharp. With light-field photography, you can attain either effect from the very same snapshot.

The information a light-field camera records is, mathematically speaking, part of something that optics specialists call the plenoptic function. This function describes the totality of light rays filling a given region of space at any one

TODAY'S OPTIONS: If you want to take light-field photos now, you have one consumer-grade option—the Lytro camera [left]—and a choice of three industrial-quality models from Raytrix, including the R11 [right].

PHOTOS: LEFT, LYTRO; RIGHT: RAYTRIX

moment. It's a function of five dimensions, because you need three (x, y, and z) to specify the position of each vantage point, plus two more (often denoted θ and φ) for the angle of every incoming ray.

When measuring light in a region that's free of any obstructions, you have to keep track of only four dimensions rather than five. Think about it: If you know that a ray isn't blocked, it's simple to follow where it goes. Record where it hits one plane (x and y) and the angle at which it hits (θ and φ) and you can work out where it came from and where it's headed. The same is true for any other ray hitting that plane at any angle. So with just the knowledge of the light crossing a single plane, you can calculate the position and direction of the rays filling the surrounding region, so long as there are no obstructions present. This four-dimensional function is called the light field (hence the term light-field camera).

All this has been known for many years. Back in 1908—the same year he won the Nobel Prize in physics for color photography—the French scientist Gabriel Lippmann invented something he called "the integral camera." His idea was to use an array of tiny lenses to project a scene onto a single sheet of film. The multiple views these lenses recorded could then be reconstituted into a 3-D image by viewing the processed film through an identical lens array. Three years later, Russian physicist P.P. Sokolov constructed the first integral camera using a pinhole array instead of the harder-to-fabricate lenses that Lippmann envisioned. Building the Lytro camera, however, required technologies that would not be realized for almost another century.

In place of film and pinholes, the Lytro camera uses a thin sheet containing thousands of microlenses, which are positioned between a main zoom lens and a standard 11-megapixel digital image sensor. The main lens focuses the subject onto the sheet of microlenses. Each microlens in turn focuses on the main lens, which from the perspective of a microlens is at optical infinity.

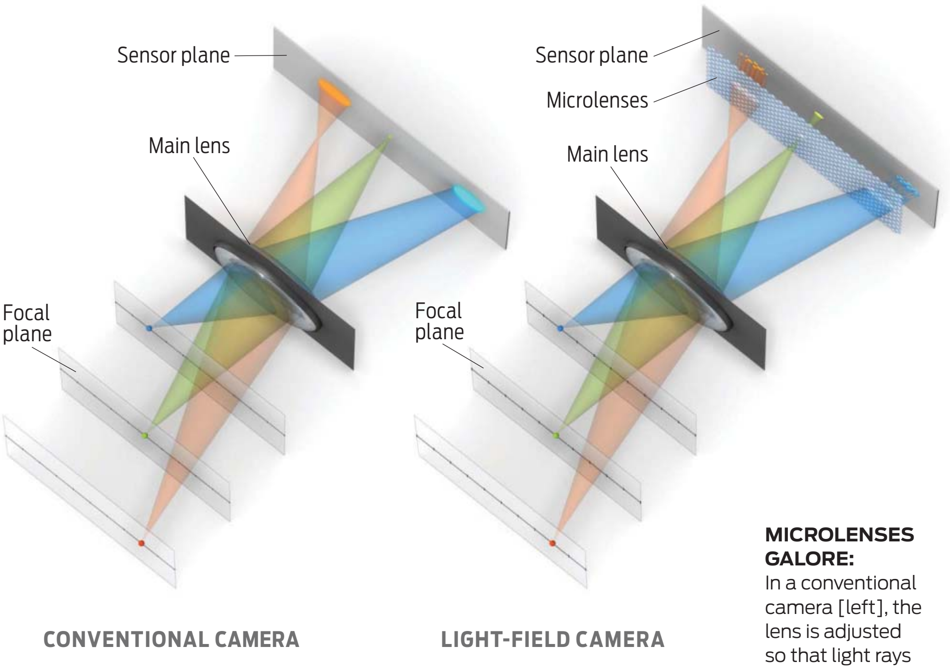

MICROLENSES GALORE: In a conventional camera [left], the lens is adjusted so that light rays emanating from one point on the focal plane [green] converge to one point on the sensor plane. Subjects nearer or farther will be out of focus [blue, orange]. In Lytro's light-field camera [right], the main lens focuses the light onto an array of microlenses, which in turn project it onto the sensor. This allows sharply focused images of subjects outside the focal plane to be digitally synthesized.

ILLUSTRATION: EMILY COOPER

It's not easy to visualize, but with the help of a diagram [see "Microlenses Galore"], you can see what's going on. Light rays arriving at the main lens from different angles are focused onto different microlenses. Each microlens in turn projects a tiny blurred image onto the sensor behind it. The light contributing to different parts of those tiny blurred images comes from rays that pass through different parts of the main lens. In this way the sensor records both the position of each ray as it passes through the main lens (the x and y) and its angle (θ and φ) in a single exposure. To understand why this lets you focus your picture later, think about what it means to focus a regular camera.

With a conventional camera, you have to adjust the focus so that all the light rays coming from one point on your subject converge at one point on the camera's sensor. Depending on whether the subject is near or far, you move the lens in or out to achieve proper focus. A light-field camera doesn't need to move its lens, because it can calculate the light field at any plane inside the camera. So it can generate images corresponding to various separations of lens and sensor, from close-ups to views of the distant horizon.

But capturing the full light field has a price. First, there's a vast increase in the amount of data the camera must acquire. Second, there is a significant loss in resolution of the final image, which is effectively limited to the number of microlenses rather than the resolution of the camera's sensor.

The transformation of the recorded four-dimensional light field into two-dimensional pictures requires computationally intensive Fourier-transform and ray-tracing algorithms. These in turn depend on access to processors that should be powerful, compact, and if the camera is to be a mass-market product, inexpensive.

The Lytro camera, which costs just US $400 or $500 (depending on the memory capacity), indeed packs some meaty processing hardware inside. And yet it has only three controls—an on-off switch, a shutter button, and a zoom for its eight-power main lens. There are no other settings to adjust, no lighting options to consider, and definitely no manual controls. When you shoot a picture, there's no shutter lag, although the camera takes around 5 seconds to show you an image. During that interval, the device computes what would have been recorded on several different virtual cameras with the focus set at various depths. When you then

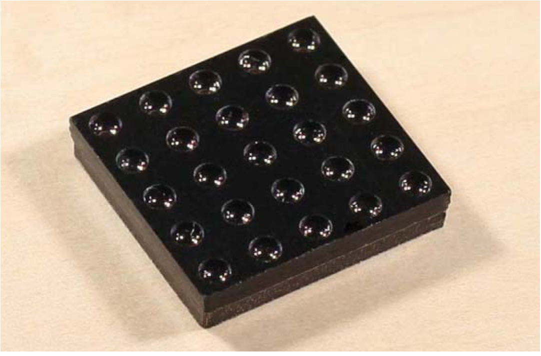

SMARTPHONE CAM: Pelican Imaging hopes to bring light-field photography to smartphones with a bug-eye-like camera array [top]. Images could then be focused at different depths, as shown in the center and bottom photos, after they are taken.
PHOTOS: PELICAN IMAGING

tap on the touch screen, the focus appears to shift to where you want it in just a fraction of a second.

"It's our ethos to have very sophisticated technology and algorithms and to [package] them in a way that's very easy to use," says Kurt Akeley, Lytro's chief technology officer. "Simpler is better." Images can also be viewed in desktop software (Mac only) or uploaded to Lytro's servers and shared on websites like Facebook, where viewers can click to refocus on whatever part of the image they select. But that's only a sample of what's possible with a light-field camera, and not everyone agrees with Lytro's keep-it-simple strategy.

"Plenoptics is about way more than refocusing images," says Winston Hendrickson, vice president of engineering for digital imaging at Adobe Systems, in San Jose, Calif. "In capturing all the spatial and angular information about a scene, you can do things like motion parallax, changing the perspective, and detecting objects: Single-lens stereo becomes easy." A motion-parallax function would allow the user to achieve a sense of depth by varying the vantage point slightly (this is the reason some birds keep bobbing their heads back and forth). "Lytro has a number of limitations, and I hope people see it as a work in progress," Hendrickson says.

Lytro promises that parallax and 3-D imaging will soon be added to its desktop and online offerings. One limitation that the company can't improve with software alone, however, is image resolution. At just 1.2 megapixels, Lytro's images are dwarfed by the 10- to 16-megapixel photos people have come to expect from their digital cameras and even, increasingly, from their smartphones.

Lytro argues that limited resolution isn't the drawback it appears to be. "There's inertia in the marketplace to think about quality as it relates to resolution," says Charles Chi, executive chairman at Lytro. "But when you think about how consumers use and share images today, it's with cellphones and computers whose screens have less than 2 megapixels of resolution."

Perhaps. But they also enjoy shooting and sharing video clips, another feature missing from the current Lytro camera. "If you don't have video, it's not going to get adopted," insists Paul Gallagher of Pelican Imaging Corp., a start-up based just a few doors from Lytro in Mountain View, Calif. Pelican is betting that light-field photography will find a natural home in smartphones, which can leverage their powerful application processors for the necessary computations.

"The generation of chips coming out the first half of this year should be more than adequate for light-field photography," says Gallagher, who is Pelican's vice president of business development and marketing. Pelican's system eschews a single, space-hogging main lens in favor of a bug-eye-like array of 16 to 25 microcameras, carefully aligned with a traditional (but higher resolution) image sensor. The device will be capable of both 3-D and video imaging. Pelican is hoping it will appear in smartphones before the end of 2013.

While Lytro and Pelican test the consumer market, a German firm has for two years been quietly shipping light-field cameras for industrial use. Raytrix is targeting applications in research, microscopy, and optical inspection in manufacturing. Its cameras use yet another kind of optical setup, one championed by Todor Georgiev, an optics researcher at San Diego–based Qualcomm, who has over 50 patents in plenoptics. Georgiev has dubbed the new configuration Plenoptic 2.0.

"Light-field cameras are happening now because of graphics processing units—GPUs," says Georgiev. "Each GPU has hundreds of processors packed into one little device, working in parallel. They are the supercomputers of today."

With the Raytrix system, as with the Lytro camera, there is a main lens, an array of microlenses behind it, and then the image sensor behind the array. Here, however, the microlenses are focused not at infinity (as in the Lytro camera) but on the image formed by the main lens in a plane that's some distance in front of the sensor. The main lens and each microlens act like the two lenses in a refracting telescope, creating an array of tiny, sharp, inverted images on the sensor. With this optical arrangement, the number of microlenses no longer limits the effective resolution of the final image. Indeed, the images produced can theoretically approach the full resolution of the sensor, although this has been tricky to achieve in practice.

With a glass array of around 20 000 microlenses in its €20 000 (about $26 500) R11 camera, for example, Raytrix manages to produce 2.7-megapixel still images from an 11-megapixel sensor and video at up to six frames per second. Unlike with Lytro, though, the number-crunching hardware needed for Plenoptic 2.0 cannot fit into a stylish anodized case. A Raytrix camera must be linked through a gigabit Ethernet cable to a PC that contains a high-end Nvidia GeForce GTX 580 graphics card, which itself costs more than the entire Lytro camera.

The Raytrix cameras are sophisticated enough to deliver on some of the more ambitious opportunities that light-field photography offers. For instance, the 20 000 microlenses in its R11 are not identical. Instead a mixture of three different focal lengths is used. That design sacrifices some lateral resolution but provides greater depth of field—the range through which everything in an image is sharp—up to six times what you'd get with a conventional camera using the same main lens. That's useful for optical inspection of components using macrophotography, gives striking all-in-focus results with telephoto pictures, and extends the distance over which Raytrix can calculate 3-D views.

And Georgiev isn't done testing the waters of Plenoptic 2.0. He is now experimenting with microlenses that have different aperture sizes to give his cameras higher dynamic range, enabling them to simultaneously capture detail in both the darkest shadows and the brightest highlights. He also envisions adding more color filters, to make single-exposure multispectral imaging possible, or adding polarizing filters, which would attenuate reflections to whatever degree the user wants, allowing crystal clear photos through glass or water.

Creating groundbreaking technology is one thing; getting consumers and companies to buy it is another. Both Lytro and Pelican Imaging hope to harness viral marketing—the wow factor that gets people talking about their products. But they are up against a $40 billion imaging industry that remains fixated on high pixel counts, powerful zoom lenses, and an ever-growing variety of flashy features. For most camera companies, and for most consumers, light-field imaging isn't on the agenda yet. "There's definitely a lot of education that needs to be done," says Lytro's Chi.

Like Apple, Lytro emphasizes style, design, and ease of use. Also like Apple, Lytro favors a closed vertical ecosystem, from capture to playback. That closed ecosystem, however, prevents other companies from working with its light-field files. "It was the inability of getting plenoptic data from Lytro that led us to doing our own research," says Adobe's Hendrickson. "We will introduce light-field editing in Photoshop when the time is right, but it's not ready for a broad scale right now." Georgiev shares this frustration. "If you buy a Lytro camera and you want to do, say, 3-D video, you can't, unless Lytro provides the software. I'm afraid Lytro may be shooting themselves in the foot," he says. "At first it looks like they're gaining competitive advantage, but it's actually closing the door for collaboration and progress." Lytro's position is that it wants to develop the technology to a level of stability before releasing its file format for wider use.

In this fast-moving field, though, such stability may be a long way off. Raytrix, for example, is already developing specialized microlenses to boost resolution for 3-D facial recognition in security systems, and the company promises to have video cameras capable of recording high-definition (1080p) plenoptic movies at 30 frames per second by the end of the year. These new cameras could revolutionize 3-D television and film, which currently require expensive stereo cameras and often time-consuming postproduction work. Not only could 3-D plenoptic movies be refocused after they were shot, they could also provide personalized experiences for different viewers, such as customizable stereo separation to give more natural 3-D effects (and fewer headaches for viewers).

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/lightfieldcamera0512

Georgiev's dream is to build a large plenoptic camera that can capture multiple 3-D views and generate an image that's indistinguishable from what you get looking out a window. "Today, I can capture 10 views and interpolate the rest," says Georgiev. "But give me a 200-megapixel sensor and it would be much, much better."

While few photographers want or need the 40-megapixel sensors found in some of today's top-end models, tomorrow's plenoptic cameras will require all that resolution and more. "We can use as many pixels as we can get our hands on," says Lytro's Akeley. "There are huge opportunities to build cameras that capture many more rays and as a result produce sharper, bigger pictures, with greater features."

Ultimately, though, the main obstacle to the success of light-field imaging may hinge on the culture of photography rather than on its technology. Lytro's Chi concedes that many professional photographers feel threatened by technology that hands creative control to the viewer. Come the next generation of plenoptic cameras—which will allow you to adjust the picture's exposure, shift colors, remove reflections, and jump into 3-D at the click of a mouse—such reactions are bound to get even stronger.

How consumers will react is a big question mark too. Fixing a blurred image is one thing, but do we really want to spend hours tinkering with our pictures? Will light-field technology democratize photography or destroy it, splintering a form of art into billions of ultrapersonalized collections that mean little to anyone else?

Christopher Rauschenberg, owner of the Blue Sky photo gallery in Portland, Ore., offers a simple observation: "Viewers already decide what they see in a picture. Pick any artwork that speaks to you, think about what you see in it, and then ask someone else. They will see something different."

Beauty has always been in the eye of the beholder. But now, thanks to light-field photography, focus and perspective, among other attributes, will be too. ❑

The Dawn of Haiku

Continued from page 40

each pass a message, which is placed in a central message queue and dealt with in a specified order. Without that mechanism, the two threads would attempt to access Thread C simultaneously, creating a deadlock, in which they both just wait endlessly. With Haiku, messages can be passed within an application and also to other applications. Haiku's message-passing code thus manages all the complexities of threading so that the developer doesn't have to.

The main reason that BeOS used multiple threads was because the BeBox was conceived of as a media platform. One of BeOS's marketing monikers, in fact, was "the media OS." While that might have been a bit of hype, the system was very good at handling multimedia. In a typical product demo, a BeBox would play half a dozen different videos at the same time, with no lag or frame dropping, while the system still remained responsive. The efficiency of the system also allowed for low latency in audio processing. While Haiku still needs work in these areas, it is well ahead of many other systems.

Another advantage that Haiku has over other operating systems is that it makes extensive use of a database-like file system, which allows any file to have various attributes associated with it. These attributes can then be indexed and queried just as in a contemporary database. For example, each e-mail message in Haiku is stored along with attributes such as the subject, the sender's name and address, and the recipient's name and address. From the operating system's file manager, known as Tracker, you can search on any of these attributes. Attributes also let you extract the song information from MP3 files and then easily organize and search your music library through Tracker. Entries in Haiku's address book, known as People files, consist almost entirely of attributes.

Having your operating system organize your e-mail and address book offers a huge advantage: You're no longer tethered to just one program to manage your e-mail or your contacts. Anyone who has ever tried extracting old e-mails from a bloated Microsoft Outlook mail file will understand the beauty of this approach. Instead of having your e-mail locked into a proprietary format that's accessible from only one program, Haiku lets you use whatever program you want—you can even use multiple programs on the same set of e-mails. Similarly, Haiku's People files can be edited and managed by various programs and can even

Trying out Haiku On the Haiku website, you can download the latest alpha release or the more current nightly releases. For the best experience, we recommend at least a Pentium III with 256 megabytes of RAM and 2 gigabytes of storage space. If you are a developer looking for a cool project, come join us. Even if you just fix a minor bug, we would greatly appreciate your participation.

be edited from within Tracker. No other operating system has so successfully implemented a database-like file system.

For all its capabilities, Haiku does lack some key things that users have come to expect. For example, most operating systems today have interfaces that are hardware accelerated. That allows fancy graphics like window shadows, transparent window borders, and thumbnails of window contents to be handled in a fast, efficient way using video cards. Users of mainstream operating systems also have a huge array of applications to choose from. Haiku at the moment is limited to the few applications created specifically for it or left over from BeOS.

And like other open-source operating systems, Haiku suffers from a lack of hardware drivers. Modern video cards, printers, and other components are complicated—they're like miniature computers in their own right, and they require complex pieces of software, called drivers, to work. The driver acts as a bridge between the operating system and the hardware. But these days there are so many types and brands and models of hardware that it's hard for open-source developers to keep up.

Fortunately, Haiku can build off work already being done by the thousands of developers who work on Linux, FreeBSD, and other open-source operating systems. Haiku already makes use of FreeBSD's network- and wireless-card drivers through a special translation interface. Work is also under way to incorporate the new Gallium3D system, which looks to be the future for video-card drivers on Linux.

What's more, many of Haiku's strengths will likely never be matched by the leading operating systems, or at least not anytime soon. Linux applications aren't designed for multithreaded use to the same extent as Haiku's and thus can't take full advantage of modern multicore hardware. With its Snow Leopard release of OS X two years ago, Apple added a new technology, called Grand Central Dispatch, that makes it easier for applications to use multiple threads. But there's still just one thread for the application's user interface, and so the "beach ball of death" persists.

The latest release of Windows 7 is a big improvement over its predecessor, but it still has many issues, such as widespread interface inconsistency and just plain information overload. The control panel alone is enough to give a user a migraine. Windows is also huge: A bare-bones installation of Windows 7 Ultimate will consume 20 gigabytes of storage. Even if you include the dozens of free applications that come with Haiku, installing the operating system will cost you only about 700 MB—one-thirtieth the real estate required for Windows 7. And as for being a free and open system, neither Windows nor Mac OS will ever be either.

Ultimately, Haiku represents a different way of viewing your personal computer. If you think that software shouldn't be riddled with bugs and incompatibilities and inefficiencies, if you hate being forced to swap out your hardware and software every few years because "upgrades" have rendered them obsolete, and if you find that the idea of using an operating system that's fast, responsive, and simple is refreshingly novel and appealing, then maybe, just maybe, Haiku is for you. ❏

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/haiku0512

IEEE

Candidates in 2012 Election

THE IEEE BOARD OF DIRECTORS has received the names of the following candidates to be placed on this year's ballot. The candidates have been drawn from recommendations made by regional and divisional nominating committees. In addition, the names include candidates for positions in the IEEE Standards Association, IEEE Technical Activities, and IEEE-USA.

For more information on IEEE elections and candidates, please visit http://www.ieee.org/elections or e-mail corp-election@ieee.org.

Electronic ballot access and paper ballot packages will be mailed on or before 15 August to all IEEE members who are eligible to vote as of 30 June. To ensure all ballot packages are delivered to the proper mailing address, please visit http://www.ieee.org/go/my_account and update your member profile if necessary.

IEEE PRESIDENT-ELECT, 2013
J. Roberto B. de Marca
Tariq S. Durrani

DIVISION I DELEGATE-ELECT/ DIRECTOR-ELECT, 2013
Alfred E. Dunlop
Rakesh Kumar
Ellen J. Yoffa

DIVISION III DELEGATE-ELECT/ DIRECTOR-ELECT, 2013
Harvey A. Freeman
Byeong Gi Lee

DIVISION V DELEGATE-ELECT/ DIRECTOR-ELECT, 2013
James D. Isaak
Susan K. (Kathy) Land

DIVISION VII DELEGATE-ELECT/ DIRECTOR-ELECT, 2013
Wanda K. Reder
Noel N. Schulz

DIVISION IX DELEGATE-ELECT/ DIRECTOR-ELECT, 2013
Jerry C. Carroll
Jae Hong Lee
Marina Ruggieri

REGION 2 DELEGATE-ELECT/ DIRECTOR-ELECT, 2013–2014
Gerard J. Christman
Timothy P. Kurzweg
Emilio M. Salgueiro

REGION 4 DELEGATE-ELECT/ DIRECTOR-ELECT, 2013–2014
Robert C. Parro
Bernard T. Sander
Hamid Vakilzadian

REGION 6 DELEGATE-ELECT/ DIRECTOR-ELECT, 2013–2014
Thomas M. Coughlin
Paul J. Kostek

REGION 8 DELEGATE-ELECT/ DIRECTOR-ELECT, 2013–2014
Saurabh Sinha
Costas M. Stasopoulos

REGION 10 DELEGATE-ELECT/ DIRECTOR-ELECT, 2013–2014
Kukjin Chun
Ramakrishna Kappagantu

IEEE STANDARDS ASSOCIATION BOARD OF GOVERNORS MEMBER-AT-LARGE, 2013–2014
Wael W. Diab
Philip C. Wennblom

IEEE STANDARDS ASSOCIATION BOARD OF GOVERNORS MEMBER-AT-LARGE, 2013–2014
Paul Hearty
Robert S. Oshana

IEEE TECHNICAL ACTIVITIES VICE PRESIDENT-ELECT, 2013
Vincenzo Piuri
Jacek M. Zurada

IEEE-USA PRESIDENT-ELECT, 2013
Gary L. Blank
Clarence L. (Lee) Stogner

IEEE-USA MEMBER-AT-LARGE, 2013–2014
Dee S. Fultz
Ron T. Ogan

Universität Stuttgart

The Faculty of Computer Science, Electrical Engineering and Information Technology of the University of Stuttgart invites applications for the position of a

Professor (W3) for Telecommunications

and head of the institute for Telecommunications ("Nachrichtenübertragung") starting April 1, 2013.

The ideal candidate has profound knowledge in the area of wired, optical and wireless communication. The focus spans from coding theory and optimization of communication algorithms to aspects of the practical realization of communication devices and systems. The applicants are expected to have excellent scientific skills in at least two of the following topics and their methods:

- Optical and broadband communication systems
- Coding, communication algorithms and cryptographic methods
- Energy-efficient communication methods
- Co-operative wireless systems
- Adaptive communication systems
- Multiple-Input-Multiple-Output systems

The professor is expected to contribute actively to interdisciplinary research projects at the university level as well as in close collaboration with academic and industrial partners. A successful career in the private sector will be appreciated.

The professor contributes to undergraduate and graduate teaching in the area of telecommunications for the degrees "Electrical Engineering and Information Technology" (some of them in German language after about two years), the international Master program „Information Technology" (in English language) as well as for study programs of other faculties.

The requirements for employment listed in § 47 and § 50 Baden-Württemberg university law apply.

Written applications should be sent no later than **May 31, 2012** to the Dean (Dekan der Fakultät 5, Universität Stuttgart, Pfaffenwaldring 47, 70569 Stuttgart, Germany) or via e-mail to bewerbung-inue@f-iei.uni-stuttgart.de.

The University of Stuttgart has established a Dual Career Program to offer assistance to partners of those moving to Stuttgart. For more information please visit the web-page under: www.uni-stuttgart.de/dual-career/

The University of Stuttgart is an equal opportunity employer. Applications of women are strongly encouraged. Severely challenged persons will be given preference in case of equal qualifications.

COMPETENCE CENTRE | **Microelectronics**

AN ENTERPRISE IRELAND & IDA IRELAND INITIATIVE

Microelectronics Competence Centre Ireland (MCCI) is an Enterprise Ireland and IDA Ireland funded technology centre hosted at Tyndall National Institute in collaboration with the University of Limerick. Its mission is to carry out world-beating mixed-signal and RF circuit research in collaboration with its industry partners. For more information, please look at

www.mcci.ie

There are currently three open research positions within MCCI:

Analogue Mixed-Signal Principal Investigator The Analogue Mixed-Signal Principal Investigator will report directly to the MCCI Director and will be responsible for leading the MCCI core research team. They will grow the core research team with high calibre talent to create a centre of excellence in Digitally Assisted Analogue microelectronic circuit design. The Principal Investigator and their team will carry out innovative circuits research which will have impact internationally including publications in leading conferences such as ISSCC.

Postdoctoral Researcher in Analogue to Digital Converters A vacancy is available for postdoctoral researcher on the MCCI research project PSAR, to develop a 14bit 200MS/s pipeline SAR ADC on 40nm CMOS with digital background calibration. The successful candidate will work primarily on the key analogue circuits of the analogue-to-digital converter and contribute with other team members to the architecture and the design of calibration circuits. The successful candidate will report to the project principle investigator.

Postdoctoral Researcher High-Speed Analogue ASIC Design for Photonic Applications Applications are invited for a postdoctoral position on an MCCI project in high-speed ASIC design for photonic applications. The successful candidate's main task will involve the design of 56Gb/s, PAM-4 transceiver electronics for short-reach optical interconnect. In addition, he or she will contribute to the research programme targeting novel burst-mode electronics for long-reach optical access networks. The successful candidate will report to the project principle investigator.

For further information about any of the above positions, please contact the MCCI Director, Mark Barry (**mark.barry@mcci.ie**)

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

Max Planck Institute for Informatics

mpi max planck institut informatik

Federal Ministry of Education and Research

The Max Planck Institute for Informatics, as the coordinator of the Max Planck Center for Visual Computing and Communication (MPC-VCC), invites applications for

Junior Research Group Leaders

in the Max Planck Center for Visual Computing and Communication

The Max Planck Center for Visual Computing and Communications offers young scientists in information technology the opportunity to develop their own research program addressing important problems in areas such as

- image communication
- computer graphics
- geometric computing
- imaging systems
- computer vision
- human machine interface
- distributed multimedia architectures
- multimedia networking
- visual media security.

The center includes an outstanding group of faculty members at Stanford's Computer Science and Electrical Engineering Departments, the Max Planck Institute for Informatics, and Saarland University.

The program begins with a preparatory 1-2 year postdoc phase (**Phase P**) at the Max Planck Institute for Informatics, followed by a two-year appointment at Stanford University (**Phase I**) as a visiting assistant professor, and then a position at the Max Planck Institute for Informatics as a junior research group leader (**Phase II**). However, the program can be entered flexibly at each phase, commensurate with the experience of the applicant.

Applicants to the program must have completed an outstanding PhD. Exact duration of the preparatory postdoc phase is flexible, but we typically expect this to be about 1-2 years. Applicants who completed their PhD in Germany may enter Phase I of the program directly. Applicants for Phase II are expected to have completed a postdoc stay abroad and must have demonstrated their outstanding research potential and ability to successfully lead a research group.

Reviewing of applications will commence on **May 1, 2012**. The final deadline is **June 30, 2012**. Applicants should submit their CV, copies of their school and university reports, list of publications, reprints of five selected publications, names of references, a brief description of their previous research and a detailed description of the proposed research project (including possible opportunities for collaboration with existing research groups at Saarbrücken and Stanford) to:

Prof. Dr. Hans-Peter Seidel
Max Planck Institute for Informatics,
Campus E1 4, 66123 Saarbrücken, Germany
Email: **mpc-vcc@mpi-inf.mpg.de**

The Max Planck Center is an equal opportunity employer and women are encouraged to apply.

Additional information is available on the website **http://www.mpc-vcc.de**

Job Opening

professor KU Leuven (Belgium)

Efficient End Use of Electrical Energy

The group Electa of the KU Leuven, has an opening for a full-time academic staff member (any grade).

The KU Leuven is the oldest university of Belgium, and ranked within the top 100 universities in the world. The research unit ELECTA is a key partner in the field of electrical energy systems and smart grids, both locally and internationally, with both academic and industrial oriented projects. It received the mark 'excellent' during the recent research visitation (2011). The group is also co-responsible for the highly successful 'Master of Science in Engineering: Energy' at the KU Leuven. Electa is part of the EnergyVille research center.

Candidates should be excellent researchers and teachers in smart grids and cities.

More information: info26@esat.kuleuven.be or contact the professors at the research unit.

KATHOLIEKE UNIVERSITEIT LEUVEN
electa
Energy Ville

- www.esat.kuleuven.be/electa -

NUS National University of Singapore

DEAN, NUS SCHOOL OF COMPUTING

The National University of Singapore (NUS) invites applications for the position of Dean, School of Computing.

NUS is a multi-campus university of global standing, with distinctive strengths in education and research and an entrepreneurial dimension. It has an enrolment of 26,700 undergraduate and more than 10,500 graduate students from 100 countries. NUS has three Research Centres of Excellence (RCE) and 21 university-level research institutes and centres, and also enjoys a close association with 16 national-level research institutes and centres. Research activities are strategic and robust, and a 'no walls' collaborative culture forms the bedrock of NUS' research-intensive vibrancy. In addition, a spirit of entrepreneurship and innovation promotes creative enterprise university-wide. NUS plays an active role in international academic networks such as the International Alliance of Research Universities (IARU) and Association of Pacific Rim Universities (APRU). It is ranked amongst the best universities in the world, and is well-regarded for research and teaching in multiple disciplines. For more information, please visit: **www.nus.edu.sg**

Established in 1998, NUS School of Computing provides a stimulating environment that amalgamates the best of educational traditions, drawing its faculty members of 90 from leading universities around the world. It has a student population of about 2,300, out of which 500 are graduate students. It operates on the clear recognition that computer science fundamentals must play a critical role in many emerging technologies, and information systems and their effective management are of strategic importance to businesses, government agencies and charitable organizations. This operating philosophy drives the approach of the School to both teaching and research. For more information, please visit: **www.comp.nus.edu.sg**

The University seeks to appoint a distinguished scholar, with outstanding academic and administrative leadership as Dean of the School of Computing. The Dean will lead the School in all administrative and academic matters, including strategic planning, research and program development, funding and implementation of academic review process. The successful candidate is expected to be a dynamic leader with strong management and communication skills.

Please send applications, enquiries or nominations, with detailed curriculum vitae, by **15 June 2012** to:

Professor TAN Thiam Soon
Vice Provost (Education)
Chair, Dean Search Committee
National University of Singapore
University Hall, Lee Kong Chian Wing
#UHL-05-01D
21 Lower Kent Ridge Road
Singapore 119077
E-mail address: **pvobox7@nus.edu.sg**

Are You An Electrical Engineering Student Looking For An Internship?

Look No Further!

We Have Internship Opportunities for Students.

Visit us at:

www.AfterCollege.com/IEEE_Internships

IEEE Advancing Technology for Humanity
IEEE StudentJobSite

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

the data

E-mail Isn't Killing the Post Office

TRUE OR FALSE: The U.S. Postal Service is struggling financially because e-mail and the Internet have rendered it obsolete.

False.

Yes, this independent agency of the federal government has reported losses—billion-dollar losses—in recent years. But technology is not the cause.

Rather, in 2006, then president George W. Bush signed into law the Postal Accountability and Enhancement Act, which mandated that the USPS set aside 75 years of health-care costs for its retired employees over the ensuing 10 years.

"No other corporation or government agency that I know of has to prefund their retirees' health benefits [like this]," says Philip Rubio, assistant professor in the University Studies program at North Carolina A&T State University, in Greensboro. "We're talking about postal workers who haven't even been born yet."

Critics of the 2006 legislation say it's a poison pill, forced on an agency that has accommodated itself to change many times over—from railroads and telegraphs to airmail, AOL, and Amazon.com.

David Partenheimer, a spokesperson for the USPS, told *IEEE Spectrum* that the agency has asked Congress to restructure the health-care prepayment system and return the US $11 billion in overpayments it has already made to the Federal Employees Retirement System. "We receive no tax dollars for our operations, and that is something we don't want to change," he says.

—*Mark Anderson*

FIRST-CLASS AND STANDARD ("BULK") MAIL are the USPS's two cash cows. Since 1997, first-class mail has brought in a reliably steady US $33 billion to $38 billion per year. If there's a Facebook or an e-mail effect, it's a small net negative buried in bigger numbers. Standard mail's piece of the pie actually goes up from $12 billion to $15 billion from 1997 to 2001 and from $17 billion to $20 billion from 2002 to 2007. Packages have brought in a comparatively paltry $1–$2 billion annually throughout.

Despite the fact that the backbone of the USPS's business remains strong, in 2007 the **AGENCY'S ANNUAL LOSSES** go off the charts. By law it had to begin socking away the health-care costs of its retired workers through the year 2081. "The post office is really the victim of an invented crisis," says Rubio. "The Internet poses a challenge. But it also poses opportunities and possibilities. The real problem is the law Congress passed in 2006."

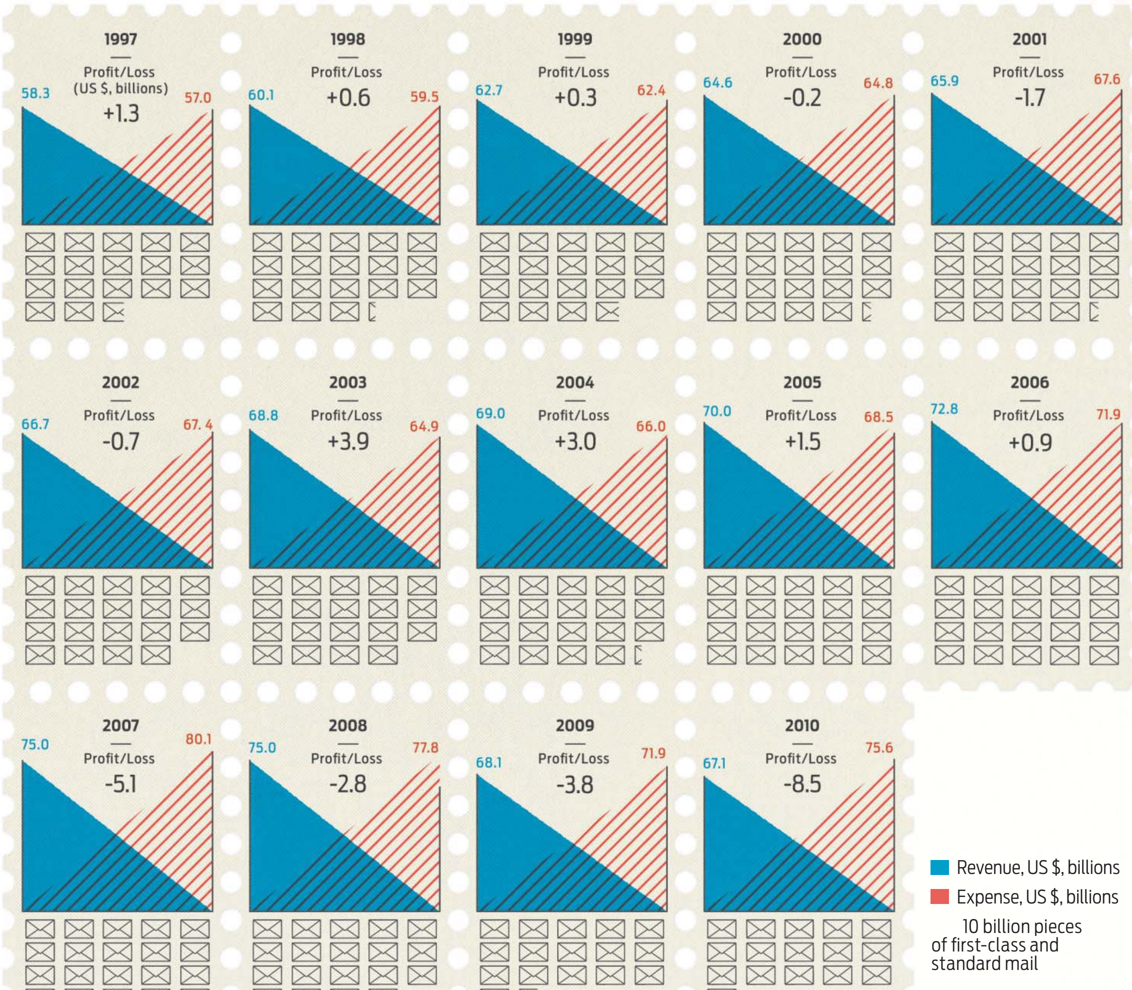

Source: USPS annual reports at http://about.usps.com/who-we-are/financials/welcome.htm

GAVIN POTENZA