IEEE Spectrum

THE MAGAZINE OF TECHNOLOGY INSIDERS

7.10

HOW A LITTLE COMPANY IN POLAND MAKES THE SEMICONDUCTOR INDUSTRY'S HOT NEW ROCK

The Crystal That Will Change Everything

THE TROUBLE WITH MULTICORE MICROPROCESSORS

DECODING THE NEANDERTHAL GENOME

TAKING LEGO ONLINE

IEEE

U.S.A. $3.99
CANADA $4.99
DISPLAY UNTIL
3 AUGUST 2010
SPECTRUM.IEEE.ORG

volume 47 number 7 international

ieee spectrum 7.10

COVER:
PHOTO: ROBERT LASKA

ABOVE, CLOCKWISE FROM LEFT: HARRY CAMPBELL, ROBERT LASKA, DAN SAELINGER, HOLLY LINDEM; LEGO SCULPTURE BY NATHAN SAWAYA

COVER STORY

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

volume 47 number 7 international

ieee spectrum 7.10

SPECTRUM.IEEE.ORG

AVAILABLE 1 JULY

ANTARCTICA: LIFE ON THE ICE

When *IEEE Spectrum* Executive Editor Glenn Zorpette [left] traveled to Antarctica this past January, he had unrestricted access to U.S. bases and facilities, where he interviewed scientists, bakers, bureaucrats, technicians, disbursement clerks, environmentalists, survival trainers, helicopter pilots, painters, a bartender, and a masseuse. He captured the bleating of 3500 Adélie penguins in the world's southernmost rookery, the gentle trickling of 4000-year-old water from a glacier, and the throaty roar of a drill boring a hole 2.5 kilometers deep in polar ice. He spent a night at the South Pole, interviewed field scientists in one of Antarctica's legendary ice-free Dry Valleys, and visited the century-old huts of explorers Robert Falcon Scott and Ernest Shackleton. His special hour-long podcast reveals a continent alive with intellectual ferment, charismatic creatures, and poignant human interest.

TOP: COURTESY OF GLENN ZORPETTE

RIGHT: JAPAN AGENCY FOR MARINE-EARTH SCIENCE AND TECHNOLOGY (JAMSTEC)

TECH INSIDER WEBINARS

Check out all webinars, including these below, at **spectrum.ieee.org/webinar**

ONLINE WEBINARS & RESOURCES

- Agility Development of Safety-Critical Systems *Sponsored by IBM*
- Capture the Concept—A First Look at Comsol Multiphysics 4.0 *Sponsored by Comsol*
- Top 10 Clean Energy Business Opportunities—Design, Prototype, and Deploy *Sponsored by National Instruments*
- Calibre DRC Automated Waiver Technology *Sponsored by Mentor Graphics*
- Design resource library by Maxim: http://spectrum.ieee.org/static/maxim-resource-library
- Design resource library by Texas Instruments: http://spectrum.ieee.org/static/ti-resource-library
- White papers: http://spectrum.ieee.org/whitepapers

IEEE.ORG/THEINSTITUTE

AVAILABLE 7 JULY

MEDICINE IS NOT ONE SIZE FITS ALL
Everyone is unique—medically speaking, that is. Tailoring health care to each patient used to depend on trial and error, family history, and doctors' intuition. Today it relies increasingly on technology. Learn about advances in individualized health care discussed at a meeting sponsored by the American Medical Association and the IEEE Engineering in Medicine and Biology Society.

EXPLORING THE OCEAN
Researchers who want to know more about the marine environment, offshore structures, and ocean vehicles will be able to learn about these and other topics at the Oceans 2010 conference, to be held 20 to 23 September in Seattle.

APPLYING SYSTEM ENGINEERING IN HAITI
Nearly a dozen IEEE members were involved in assisting the U.S. military with data collection in the aftermath of the earthquake that struck Haiti in January.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2010 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 W224-N3322 Duplainville Rd., Pewaukee, WI 53072-4195, 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.

back story

RANDI SILBERMAN KLETT

Legos for a Living

SO HERE'S a job for you: Every day, you go to a Manhattan penthouse and build stuff out of an essentially unlimited supply of Lego building blocks. You build a cello; a life-size Steven Colbert, complete with business suit; a stunningly realistic iPhone. You even pull an all-nighter to build a boy and his laptop for an *IEEE Spectrum* feature story ("From Bricks to Bits," in this issue).

That's life for Nathan Sawaya [above], whose studio looks like an 8-year-old's dream come true. It's full of objects instantly recognizable and yet alien as well. Equally impressive is the workroom, where along every wall, individual bricks fill plastic tubs stacked from floor to ceiling. Sawaya figures he's got about 1.5 million individual bricks. But in case a large commission comes up, he keeps another 3 million to 4 million in storage.

The New York City–based artist used to work in more conventional forms, but six years ago he switched full-time to Lego bricks. Sawaya builds these sculptures for such clients as Donald Trump, who wanted a 10-foot-high sculpture for a hotel in Dubai. He also does fine-art pieces for museums and galleries; "The Art of the Brick" shows are currently touring North America.

He's not a Lego employee, but as a "Lego Certified Professional"—one of only nine in the world—he has a special relationship with the company, which allows him to buy Lego bricks at a steep discount. "I order them by the pallet," he explains. "It's not like there are lots of clients who call and say, 'I need 10 000 bricks of this type.'" Even with his discount, Sawaya says he spends six figures a year on the bricks.

Lego bricks are pretty sturdy to begin with, but for extra durability Sawaya glues the pieces together as he builds. That means that when a brick ends up where it doesn't belong, Sawaya turns to a more traditional sculpting tool: his chisel. ❏

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 first Update page is in *IEEE Spectrum,* Vol. 47, no. 7 (INT), July 2010, p. 9, or in *IEEE Spectrum*, Vol. 47, no. 7 (NA), July 2010, p. 11.

ieee spectrum

EDITORIAL

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

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

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

SENIOR EDITORS Harry Goldstein (Online), h.goldstein@ieee.org; Jean Kumagai, j.kumagai@ieee.org; Samuel K. Moore (News), s.k.moore@ieee.org; Tekla S. Perry, t.perry@ieee.org; Philip E. Ross, p.ross@ieee.org; David Schneider, d.a.schneider@ieee.org

SENIOR ASSOCIATE EDITOR Steven Cherry (Resources), s.cherry@ieee.org

ASSOCIATE EDITORS Sally Adee, s.adee@ieee.org; Erico Guizzo, e.guizzo@ieee.org; Joshua J. Romero (Online), j.j.romero@ieee.org; Sandra Upson, s.upson@ieee.org

ASSISTANT EDITOR Willie D. Jones, w.jones@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org

COPY EDITOR Michele Kogon, m.kogon@ieee.org

EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

EXECUTIVE PRODUCER, SPECTRUM RADIO Sharon Basco

ASSISTANT PRODUCER, SPECTRUM RADIO Francesco Ferorelli, f.ferorelli@ieee.org

ADMINISTRATIVE ASSISTANTS Ramona Gordon, r.gordon@ieee.org; Nancy T. Hantman, n.hantman@ieee.org

IEEE SPECTRUM JOURNALISM INTERN Anne-Marie Corley, a.corley@ieee.org

CONTRIBUTING EDITORS John Blau, Robert N. Charette, Peter Fairley, David Kushner, Robert W. Lucky, Paul McFedries, Prachi Patel, Carl Selinger, Seema Singh, William Sweet, John Voelcker

ART & PRODUCTION

SENIOR ART DIRECTOR Mark Montgomery

ASSOCIATE ART DIRECTOR Michael Solita

ASSISTANT ART DIRECTOR Brandon Palacio

PHOTO EDITOR Randi Silberman Klett

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy

EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia

SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

WEB PRODUCTION COORDINATOR Jacqueline L. Parker

MULTIMEDIA PRODUCTION SPECIALIST Michael Spector

EDITORIAL ADVISORY BOARD

Susan Hassler, *Chair*; Marc T. Apter, Francine D. Berman, Jan Brown, Raffaello D'Andrea, Hiromichi Fujisawa, Kenneth Y. Goldberg, Susan Hackwood, Bin He, Erik Heijne, Charles H. House, Christopher J. James, Ronald G. Jensen, Ruby B. Lee, Tak Ming Mak, David A. Mindell, C. Mohan, Fritz Morgan, Andrew M. Odlyzko, Barry L. Shoop, Curtis A. Siller Jr., Larry L. Smarr, Harry L. Tredennick III, Sergio Verdú, William Weihl, Başak Yüksel

EDITORIAL CORRESPONDENCE

IEEE Spectrum, 3 Park Ave., 17th Floor, New York, NY 10016-5997
Attn: Editorial Dept. Tel: +1 212 419 7555 Fax: +1 212 419 7570
Bureau: Palo Alto, Calif.; Tekla S. Perry +1 650 328 7570
Responsibility for the substance of articles rests upon the authors, not IEEE or its members. Articles published do not represent official positions of IEEE. Letters to the editor may be excerpted for publication.

ADVERTISING CORRESPONDENCE

IEEE Spectrum, 3 Park Ave., 17th Floor, New York, NY 10016-5997
Attn: Advertising Dept. +1 212 419 7760
The publisher reserves the right to reject any advertising.

contributors

HARRY CAMPBELL created an illustration for "The Trouble With Multicore" [p. 24] that depicts a multicore chip as a huge open office with just one occupant. "That guy at his desk represents the only working core. The others are out to lunch," he says. Campbell started out as a designer for Nickelodeon and Warner Brothers, where he did more humorous illustrations. His current graphic style works well for more conceptual work, like his political illustrations for *The New York Times*.

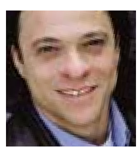

DAVID KUSHNER, like many adults his age, grew up playing with Lego bricks and now joins his kids on Lego projects. While researching the new computer game *Lego Universe*, he wondered, "How could you possibly do this online in a way that's authentic?" For "From Bricks to Bits" [p. 42], the *IEEE Spectrum* contributing editor traveled to the game developer's headquarters in Denver to learn what it took to solve that problem.

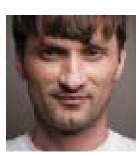

ROBERT LASKA shot the photographs for "The World's Best Gallium Nitride" [p. 36] in the midst of a natural disaster. Poland's flooding Vistula River nearly prevented the photo shoot from taking place. Normally a portraitist of Polish celebrities, Laska found that capturing flakes of gallium nitride was a whole new challenge. He says he never expected that a synthetically grown crystal, reflected in sunlight, could be so beautiful.

RICHARD P. MISLAN, who wrote "Cellphone Crime Solvers" [p. 30], is mourning the end of the TV series "Lost," a mystery that couldn't be solved by clues hidden in a cellphone. Mislan is an assistant professor for cyberforensics at Purdue University, a consultant to law enforcement, and director of the annual Mobile Forensics World conference. But he's probably best known for an appearance on a local news show a few years ago. "About once a week I still get a call from someone concerned that her spouse has installed spyware on her phone," Mislan says.

DAVID PATTERSON quite fittingly starts his article, "The Trouble With Multicore" [p. 24], with a sports analogy. Patterson, a professor of computer science at the University of California, Berkeley, who pioneered such now-common computer technologies (and acronyms) as RISC and RAID, is also an avid soccer player and weight lifter. "If programmers these days were treated like sports champions, maybe the U.S. programmers would do better," he says. "Or maybe not: After all, there's no steroid you can take to improve your coding."

RICHARD STEVENSON got to hold a 2-inch (51-millimeter) crystal of gallium nitride "worth as much as a BMW sports sedan" while touring the Warsaw factory of its maker, Ammono. The little company makes these jewels better than just about anyone else, which is why the manufacturers of blue lasers are beating a path to its door. Richardson, based in Wales, studied such compound semiconductors for his Ph.D. at the University of Cambridge. "The World's Best Gallium Nitride" [p. 36] is his third article on the subject for *Spectrum*.

CLOCKWISE FROM TOP LEFT: HARRY CAMPBELL; UNIVERSITY OF CALIFORNIA, BERKELEY; RANDI SILBERMAN KLETT; PURDUE UNIVERSITY; ROBERT LASKA; BETH KELLY

spectral lines

Radio Daze

AS A YOUNGSTER, I loved reading about home-brewed electronic projects. One, I recall distinctly, was "a tuna-can radio," a low-power transmitter that could be put in a package no bigger than 5 ounces of albacore.

I thought about that circuit again decades later when I came across something called a cantenna. Another mashup of radio and recycling, this homemade Wi-Fi antenna can supposedly boost your wireless range. A little research revealed other antenna configurations that promised even better results, including those of the FabFi collaboration in Afghanistan, which uses "common building materials and off-the-shelf electronics to transmit wireless Ethernet signals across distances of up to several miles," according to the Jalalabad FabLab's Web site.

With that for inspiration, I set about to rig my own high-gain antenna. My ambition was to access the Internet through my community's downtown wireless network, which is about 400 meters from my house.

It would be a fine DIY project, I thought, for this magazine's Hands On column. I was wrong. The reason: Federal Communications Commission regulations forbid what I intended to do—despite the fact that high-gain antennas designed to extend the range of Wi-Fi equipment are sold throughout the United States. You can buy these antennas from such ubiquitous U.S.-based retailers as Walmart and Office Depot, for example. This puzzling situation reveals some interesting things, both about the FCC's legal authority to police the airwaves and about its ability to keep up with the pace of technological advance. Let me explain.

One of the relevant regulations—Part 15, section 15.23(a)—seemed to give people like me a pass to hack such equipment together, so long as "good engineering practices" were used. The problem was that another section, 15.204(c), says, "[A transmitter] may be operated only with the antenna with which it is authorized." Confused? I was too. So I sought clarification from the FCC.

Rashmi Doshi, chief of the laboratory division of the FCC's Office of Engineering and Technology, patiently answered my questions, citing specific entries in the FCC's knowledge database. He explained that the regulation on home-built devices applied only to things that were built from scratch, not to modifications of FCC-certified radio equipment.

So the rules wouldn't, in fact, allow me to build and use the antenna I had envisioned.

But why, then, could I easily purchase high-gain add-on antennas for Wi-Fi equipment? Doshi explained: "We do not regulate antennas by themselves, as they do not generate RF signals. Accordingly, the sale of any antenna, including a high-gain antenna, is not illegal. However, it is illegal to use a high-gain antenna with a transmitter which it is not approved for." Captain Renault's famous line from *Casablanca*—"I'm shocked, *shocked*, to find that there's gambling going on here!"—echoed in my brain.

In fairness to the FCC, it may not have the legal authority to outlaw the sale of such antennas. What's more troubling to me is the FCC's interpretation of the regulation intended to deal with home-built equipment. That seems as antiquated as a tuna-can transmitter.

After all, these days radio hackers—like those involved with FabFi—are doing extraordinarily interesting things. But most of these experimenters are building on the shoulders, and circuits, of other do-it-yourselfers, open-source communities, and, yes, equipment manufacturers. Few people are still soldering transistors together one by one. Instead, they are modifying commercial equipment or building new gear using existing RF boards or modules. The definition of "scratch built" is thus quite murky.

Why not foster this creative energy by including all of it under the label "home built"? The danger, I suppose, is that it opens a door for an unscrupulous vendor to bypass FCC rules by selling boxes that anyone could plug together and call home built. Perhaps the FCC just needs to define a home-built device as something requiring skills beyond those of the typical consumer. In an era when policymakers struggle to figure out how to boost young people's technical know-how, the FCC should somehow recognize—and even embrace—the way today's tinkerers work.

—DAVID SCHNEIDER

COLIN ANDERSON/GETTY IMAGES

forum

LETTERS do not represent opinions of IEEE. Short, concise letters are preferred. They may be edited for space and clarity. To post your comments online, go to spectrum.ieee.org. Or write to Forum, *IEEE Spectrum*, 3 Park Ave., 17th Floor, New York, NY 10016-5997, U.S.A.; fax, +1 212 419 7570; e-mail, n.hantman@ieee.org.

CHARGE OF THE ELECTRIC BRIGADE

Where would we be today if automobile development had not been sidetracked by the gasoline-fueled internal-combustion-engine car? ["Top 10 Tech Cars," April]. Just compare the socioeconomic and environmental impacts of the infrastructure needed to support gasoline-powered vehicles versus that needed to support electric-powered ones. Certainly, if the same amount of money and time—100+ years—that were spent developing gasoline-fueled cars had been dedicated to electric-powered cars, many of the technical problems and end-product costs would not be factors today.

Robert C. Sponberg
IEEE Life Member
Columbia, Md.

STREET CRED

The May Back Story advises readers to "search on 'Tesla Street' in Google Maps" to find streets named after Nikola Tesla. If you search on "Telsa [*sic*] Terrace, Madison, WI," you'll find a would-be example right below Marconi Street. The electrical engineers in Madison have been trying to correct the misspelling, but those who live on the street oppose changing the name they have lived with for so long.

John G. Webster
IEEE Life Fellow
Madison, Wis.

WIRELESS POWER SURGES AHEAD

We at WiTricity Corp. read with great interest "Electrons Unplugged" [May] by David Schneider and were glad to see both the promise and challenges of midrange wireless power transfer systems explored in a thoughtful manner. His article focused on experimental results reported by MIT in the journal *Science* in 2007. The systems we build today, while based on the same science, are more developed than the system described in the *Science* article. Our current products have passed regulatory testing, and the plug-to-device efficiencies we achieve with our highly resonant systems exceed those of commercially available inductive proximity charging systems. We regret that because of the proprietary nature of our work, we were not able to provide more up-to-date performance data for Mr. Schneider's article.

Katie Hall
Chief Technology Officer
WiTricity Corp.
Watertown, Mass.

FOR THE RECORD

It's frustrating to see articles about digital media that mention illegal but unprosecuted handling of media files ["Your Very Own Cloud," Tools & Toys, May]. For example, ripping your music CDs to MP3 and putting them on your iPod is technically illegal but not prosecuted as a crime—until the Recording Industry Association of America drags you into court and piles on extra charges for "commonplace" activities. How do you propose people get digital copies of feature films to put on the Pogoplug in your review? If you're ripping them from DVD or Blu-ray, you've severely violated the Digital Millennium Copyright Act by circumventing digital copyright protections on the disk, and even format switching is looked down on by the media companies. If consumers truly knew the scope of the laws concerning such commonplace activities—which should be perfectly legal under the old laws of personal possession and fair use—they would be up in arms.

Matthew Beckler
Graduate Student Member
Pittsburgh

Editor Steven Cherry responds: There is nothing inherently illegal about ripping a CD to an iPod, nor is it illegal to rip music purchased from iTunes or Amazon to a CD. There is also a wide variety of music and video content in the public domain or freely available under Creative Commons or other licensing schemes. If things were as dire as Beckler suggests, there would be almost nothing legal you could do with a TiVo, a Slingbox, or Apple TV. In a digital world, every recording and every viewing, on every server and every device, creates a new copy.

CORRECTION

In "The Saline Solution" [*IEEE Spectrum* Special Report: "Water vs. Energy," June], we mislabeled the orange arrows in the illustration "Spin Cycle: Better Pretreatment." The arrows appear to indicate centrifugal force but in fact were meant to illustrate the particles' separation from the rest of the water in the membrane-free desalination system.

update

more online at spectrum.ieee.org

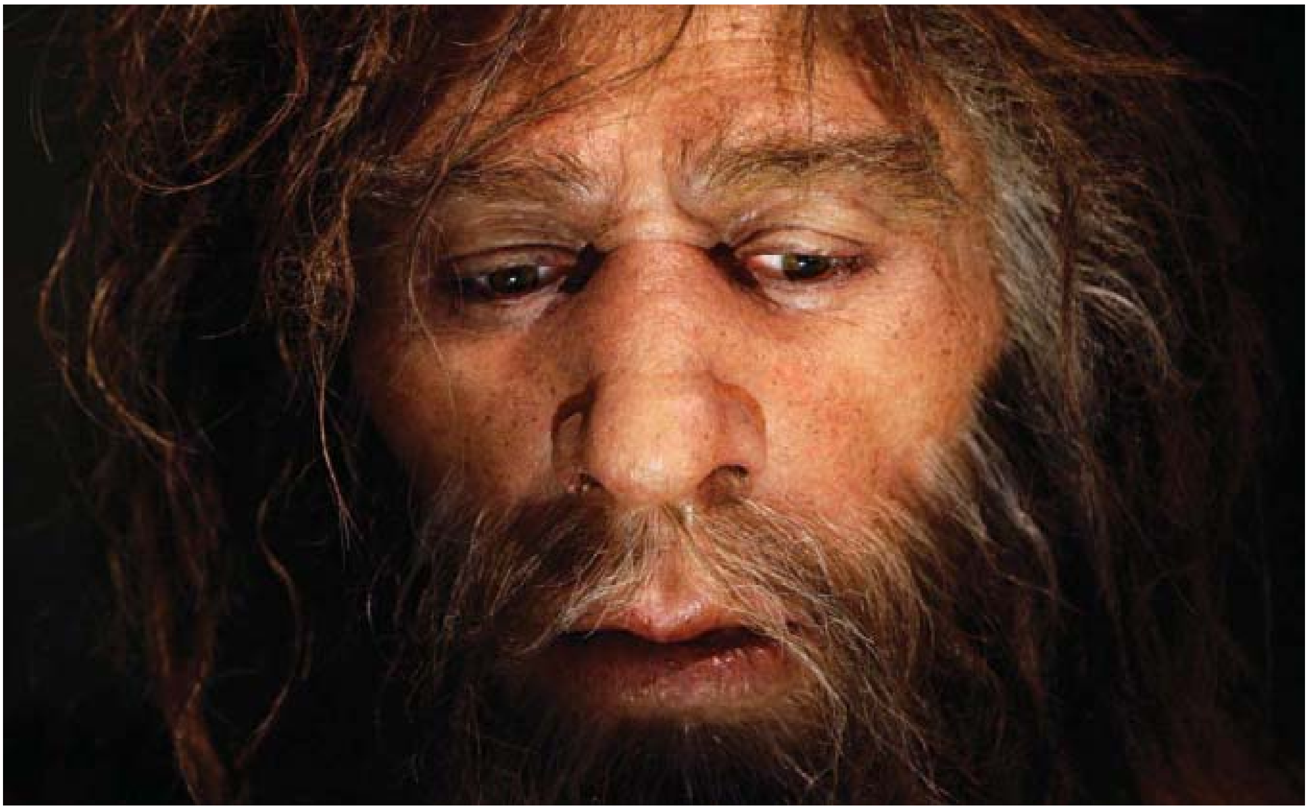

ANCIENT ANCESTOR: Computer code led to the discovery that there were more than a few human-Neanderthal assignations in the distant past. This handsome fellow is a reconstruction at the Neanderthal Museum in Krapina, Croatia.
PHOTO: NIKOLA SOLIC/REUTERS

Computing the Caveman

New software helped decode the Neanderthal genome

RECONSTRUCTING CAVEMEN from bits of fossil DNA is still the stuff of science fiction. But thanks to high-powered computing wizardry, we now have the blueprints you'd need to do it. An international team of scientists published the first draft of the Neanderthal genome in the journal *Science* on 7 May. The study showed, somewhat surprisingly, that early humans and Neanderthals interbred and that 1 to 4 percent of the DNA in modern Asians and Europeans comes from Neanderthals.

The bulk of the credit for decoding the Neanderthal goes to high-throughput sequencing technologies developed in the past five years, which turned bits of ancient DNA into millions of short strings of letters. But sequencing the Neanderthal genome would have been impossible without the sophisticated software that put all those millions of strings together in the right order.

Both the sequencing and computing of whole genomes have come a long way in the nine years since scientists first delivered the DNA blueprints for humans. Then, sequencing machines read 330 000 base molecules per day. (DNA contains a series of four different base molecules: adenine, cytosine, guanine, and thymine, or A, C, G, and T.) The latest machines made by Illumina, in San Diego, can sequence about 2 billion bases a day, amounting to hundreds of terabytes of data to store and analyze.

Assembling the human genome was like putting together a 100-piece jigsaw puzzle without any idea of what the picture should look like. The computers that were used to sequence the 3-billion-base human genome had to sort through about a hundred DNA fragments at a time, each made up of 2000 to 3000 bases. An assembly algorithm then compared these fragments, found overlaps, and strung together longer and longer stretches of the sequence.

The difference between now

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

update

and then, Neanderthal and human, is both scale and complexity. DNA fragments from fossil Neanderthal bones are very short, having degraded to an average length of only 50 bases, and there are millions of them, not hundreds. Some of these are chemically damaged in such a way that one base has mutated into another. And even worse, more than 96 percent of the DNA sequences scientists obtained came not from Neanderthals but from microbes that have contaminated the bones. To continue with the puzzle analogy, you now have millions of very small pieces, some with faded color, of which only a few thousand belong to your puzzle.

"The challenge is to find a needle in a haystack," says Janet Kelso, a bioinformatics researcher at the Max Planck Institute for Evolutionary Anthropology, in Leipzig, Germany, where much of the sequencing and data analysis was done.

DELICATE DRILLING: The Neanderthal genome was decoded from just 400 milligrams of bone powder.
PHOTO: FRANK VINKEN

Luckily, the completed human and chimpanzee genomes can serve as reference pictures for the Neanderthal puzzle—modern humans and Neanderthals share about 99 percent of their genome. "The computational analysis of this is not terribly difficult *in concept*," says Richard Green, a biomolecular engineering professor at the University of California, Santa Cruz, who led the computing efforts at Max Planck. "Take all of these sequences and align them against the human and chimp genome." That is, try to locate millions of 50-odd base sequences on the 3-billion-base human and chimp genomes.

This process, called mapping, might be conceptually simple, but in practice it's difficult enough to take days of computing time, even on a machine with hundreds of processor cores at work. Mapping is a common feature of modern genetics research, but it can fail with fossil DNA because of base mutations. So the researchers have developed their own mapping program, called Anfo Short Read Aligner/Mapper.

Like conventional mappers, Anfo cuts up the reference genome (that of the humans and chimps) into small "words" and arranges them in an index. The 60-base sequences from the fossils are then broken into 16-base words. The algorithm looks for these words in the index, assigns scores based on similarity, and finds the best matches.

The key difference between Anfo and regular mappers is that the new program takes into account knowledge scientists have gained over years of sequencing fossil DNA. For example, scientists have found that degraded DNA fragments have a high rate of C to T mutations at the beginning of a fragment and G to A mutations at the end. So Anfo considers these factors in determining whether a base is correct.

With two dedicated Linux clusters totaling 500 processor cores, the program takes about a week to sort through the millions of fragments generated by the sequencing machine, decide which are Neanderthal and which are not, and then put them in order.

Anfo is a work in progress, with programmers adding upgrades regularly. But improvements to sequencing machines are doubling those devices' outputs about every year. So coders will have to race to keep up. —Prachi Patel

The Danger-Sensing Driver's Seat

A car seat pokes you in the ribs when an accident looms

Today's carmakers boast of their new vehicles' cocoonlike qualities, including the latest noise-reducing features and great shocks that smooth out those bumps on the road. Myriad warning alarms are designed to counteract the insular environment when there's trouble, but Yale University engineers are looking for a more subtle approach to signaling: vibrating motors and movable cams that send signals straight from the seat back to your back.

Such a system could be particularly useful for warning

7 Number of atoms in a working transistor built at the University of New South Wales, in Australia. The device is a step toward building solid-state quantum computers, researchers say.

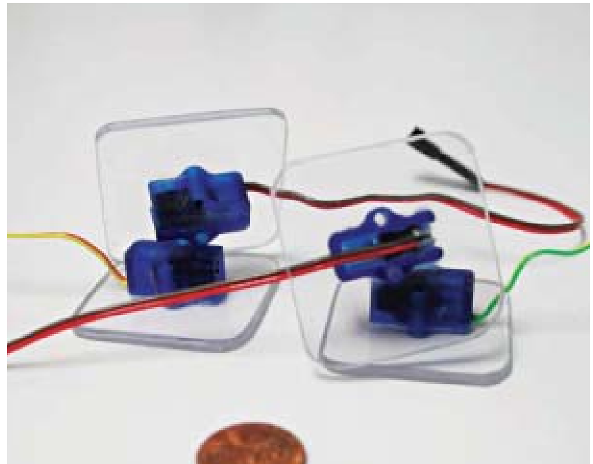

BUZZ, BUZZ: Vibrating motors warn of cars in a driver's blind spot.

PHOTOS: ARTHUR MCCLUNG III

you of cars in your blind spot—an area you can't see with either the side-view or rearview mirrors. The blind spot is blamed for sideswipe accidents that occur by the hundreds of thousands every year. Today's blind-spot monitors warn you by flashing a signal in the rearview mirror or on the car's dashboard.

But according to John Morrell, a former Segway engineer and now an assistant professor of mechanical engineering at Yale, a visual signal is the wrong one to send. The visual sense, he says, is already saturated in a car. On top of that, Morrell adds, an alert that appears in front of you about something that's behind you means you have to translate the information, costing crucial response time as well as mental effort that could distract you from other dangers.

Instead, Morrell's lab has designed a car seat that acts as a tactile interface between the driver and the environment. It uses vibrating motors as well as servos that press on the driver's back to signal the location of a following car.

The ultimate goal is a "renaissance user interface," Morrell says of the design, "using the mind and body to the fullest." The design was inspired by his work at Segway, where "every part of your body is coupled into the experience."

A demonstration setup includes a modified car seat, a steering wheel, foot pedals, and a computer running an open-source driving simulator called TORCS, short for "The Open Racing Car Simulator." Twenty cellphone-motor tactors are arranged in a rectangular array across the back of the seat. A car coming up directly behind the driver in the simulation activates the center vibrators, while a car to the right or the left will activate the same-side vibrators, giving the driver a directional cue. The closer a car gets, the more intense the vibration. Two cams on each side add to the signal strength by pressing on the driver's rib cage when a car shows up on either side.

In their preliminary study, the Yale engineers calculated the time a driver spent with another car in its blind spot during a "commuting exercise." Drivers were also given simultaneous math tasks to simulate distraction during driving. The results showed that the vibrotactile feedback improved drivers' performance over that attained by using the rearview mirror alone.

It's not the first time engineers have experimented with tactile signals to compensate for sensory overload during driving. Researchers at General Motors and TNO, a Dutch research organization, found that test subjects could actually pick up directional cues from a vibrating seat cushion while driving. Such a system could provide direct navigation information or even help warn of imminent collisions.

If systems like these turn out to be useful, car interiors will stay peaceful but may get more informative. —Anne-Marie Corley

news brief

Solar Cells on the Cheap

Researchers led by John Rogers of the University of Illinois at Urbana-Champaign have invented a cheaper way to build devices that use compound semiconductors. Their method involves growing stacks of thin films of semiconductor, peeling off the films one by one, and printing them onto cheaper substrates, such as silicon or glass. The trick could lead to cheaper solar cells and infrared cameras, says Rogers.

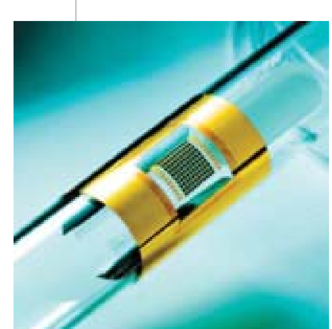

JOHN ROGERS/UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

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

update

WiMax for Smart Grids

The wireless network is unpopular with smartphones but could succeed with smart grids

Cellphone carriers are racing this year to implement 4G wireless networks so that future smartphones will have access to a fire hose of data. But the towers and infrastructure the carriers are putting in place may ultimately facilitate as many connections with appliances and electric meters as they do downloads of movies and music. The electricity grid, in other words, may be jumping from no G to 4G—and rather soon, according to analysts.

The reason, says Cecil Taylor of Blue Hawk Consultants, in Plano, Texas, is the array of data-rich applications expected in the next-generation electric grid: remote control of household electricity hogs like air conditioners, perhaps via a smartphone app; electricity price breaks at off-peak hours, leading to more automation for running appliances like dishwashers at night; and the ability to sell electricity generation (from solar panels) and storage (from a plug-in hybrid) back to the grid.

Once smart meters start appearing in the home, they'll need something approaching broadband speeds to handle those data-rich applications. According to Don Kintner of the Electric Power Research Institute, a city with a million smart meters that communicate just an hour a day will end up transmitting nearly 300 gigabytes of data per year.

"If you're a utility, you'd better get it right," says Taylor. "If you're going to start building a network that's going to last you 15 or 20 years, you're not going to want a slow, ISDN network. You're going to want a high-speed broadband network."

Kris Brown, smart-grid investment advisor at PricewaterhouseCoopers in Houston, says he thinks WiMax, a 4G IEEE standard, could be that "be-all, end-all" network, as he puts it. Pilot smart grids that use WiMax as part of their data communications have, in fact, already been rolled out in San Diego, Michigan, Texas, and parts of Australia.

This summer, 4G carrier Clearwire Corp. will be launching 3- to 6-megabit-per-second coverage with its cellular partner, Sprint Nextel, in at least 15 new markets in the United States. By summer's end, Clearwire's network, which uses WiMax, will be running in nearly 50 American cities. Later this year and into 2011, Verizon Wireless and AT&T will be launching their own competing 4G networks using a technology called LTE (for Long Term Evolution). Many cellphone carriers around the world are also expected to adopt LTE as their 4G standard—including T-Mobile, Vodafone, France Télécom, and Telecom Italia.

SMART-GRID SMACKDOWN

Broadband over Power Lines (BPL) vs. WiMax

	BPL	4G/WiMax
COST	Installation: US **$1000** per home (estimated)	Installation: **$440** per home (estimated)
COVERAGE	**Pervasive** (accessible to any home with power lines)	**Rolling out**; becoming widespread in 2011
NETWORK	Requires **new broadband infrastructure**, set up by power utilities (inexperienced broadband network operators)	Uses **existing broadband infrastructure**, set up by wireless carriers (experienced broadband network operators)

SOURCE: *Mind Commerce*

"There was a time when there was a great debate in communications whether WiMax or LTE would become the dominant wireless broadband technology," says Jeff Nichols, director of network and communication services for Sempra Energy (which is fielding a smart-grid test in San Diego). "That's pretty much played out now. The carriers are going to bet on LTE. And WiMax is finding new applications, including many in smart grids."

According to Taylor, LTE wireless networks will be choked with users doing high-bandwidth activities like video chats and mobile gaming, so those networks aren't going to look attractive to utility companies. Utilities will need guaranteed bandwidth at high and low traffic times. To the loser in the 4G broadband war, then, may go the spoils.

Sempra Energy's Nichols says utilities don't particularly care what labels or standards go on the telecom technology they'll be putting into their next-generation smart-grid electric meters. "Utilities are agnostic about something like a radio technology," he says. "What they're looking for is something that works and is cost-effective. I do think that WiMax will continue to play a role. It's practical for now, and it'll still be practical five years from now. Beyond that, who knows?"

—Mark Anderson

670 GIGAHERTZ Fastest frequency for a monolithic integrated circuit. Northrup Grumman engineers say their IC could be key to making compact terahertz imagers and explosives detectors.

tech in sight

The Robot Baby Reality Matrix

Some robot babies look real. Some act real. A few do both. *By Erico Guizzo*

ADVANCED CAPABILITIES

NAO is a little French humanoid created by Aldebaran Robotics. It talks, tracks faces, and with 25 degrees of freedom, it can even perform Michael Jackson choreography.

NEXI is the size of a 3-year-old child and was conceived at MIT's Media Lab. With a 15-degrees-of-freedom face, this little social bot can show you when it's happy, sad, mad—or bored.

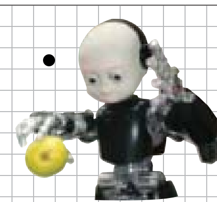

iCUB has multiple parents—11 European robotics labs. The size of a 3½-year-old, it's learning to walk, talk, and handle objects. It's probably the most advanced—and expensive—artificial child ever built.

DIEGO-SAN is the progeny of roboticists at the University of California, San Diego, and the Japanese firm Kokoro Co. They're teaching it to walk and hold objects—and they're also designing a smaller head.

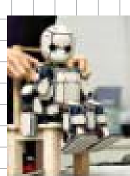

M3-NEONY is an Osaka University offspring. The robot has 90 tactile sensors, 22 motors, 2 cameras, a compact computer, and a fancy name. M3 stands for "man-made man."

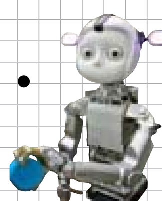

SIMON is the child of Georgia Tech researchers. This social robot has an expressive face, articulated torso, dexterous hands, and supercute ears.

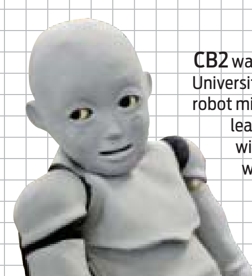

CB2 was born at Osaka University, in Japan. This robot mimics how infants learn by interacting with the world. And with 51 pneumatic actuators, it's powered by air.

ROBOTIC APPEARANCE

HUMAN APPEARANCE

ZENO is a small, cartoon-like humanoid designed at Hanson Robotics, in Richardson, Texas. It talks, understands speech, and can learn names and faces. It also has big green eyes that look as if they're ready to shoot lasers.

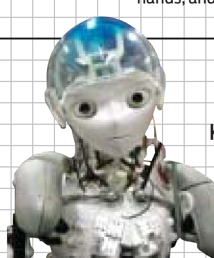

KOJIRO is a very special mechanical child built at the University of Tokyo. It has a flexible spine and a body that mimics our musculoskeletal system. Added bonus: You can control it with a PlayStation joystick.

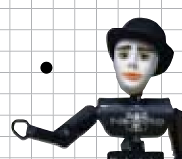

ROBOTINHO is the mechanical child of engineers at the University of Bonn, in Germany. They use the robot as a museum guide—and also to play robot soccer.

REPLIEE R1 is a copy of a real 4-year-old girl. Built at Osaka University, it has nine DC motors in its head, prosthetic eyeballs, and silicone skin. Opinions on how cute it is are mixed.

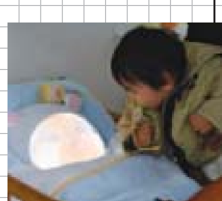

YOTARO is a baby robot that giggles, cries, and simulates a runny nose—but *not* a soiled diaper. Researchers at the University of Tsukuba, in Japan, designed it to show how rewarding babies can be. Really.

REALCARE BABY is a lifelike doll by Realityworks, in Eau Claire, Wis. You have to feed, burp, rock, and diaper the cyberinfant around the clock or it will cry—a real, recorded baby cry.

LIMITED CAPABILITIES

CLOCKWISE FROM TOP LEFT: ALDEBARAN ROBOTICS; THOMAS BREGARDIS/AFP/GETTY IMAGES; ROBOTCUB; KOKORO; ERICO GUIZZO; KOICHI KAMOSHIDA/GETTY IMAGES; UNIVERSITY OF BONN; ERICO GUIZZO; UNIVERSITY OF TSUKUBA; TONY GUTIERREZ/AP PHOTO; SANKEI/GETTY IMAGES; GEORGIA INSTITUTE OF TECHNOLOGY; YOSHIKAZU TSUNO/AFP/GETTY IMAGES update

Z-RAM Takes on DRAM

A new memory design might find a spot in the competitive DRAM market

A SWISS COMPANY, working with memory chipmaker Hynix Semiconductor, has introduced a design that it says will be a cheaper, lower-power replacement for the common computer memory known as dynamic RAM, or DRAM.

Innovative Silicon, in Lausanne, says it has redesigned its zero-capacitor RAM, or Z-RAM, so that it can be built on the same kind of wafers used for ordinary DRAM. This is a big advance for the company, because its previous devices required expensive specialty wafers. Z-RAM, unlike DRAM, doesn't require a capacitor, so the company estimates that the new design will be 25 to 30 percent cheaper as the critical features of memory drop below 40 nanometers over the next couple of years. What's more, the Z-RAM cells operate at voltages as low as 0.5 to 0.6 volt, in line with what future DRAM devices will require.

Innovative Silicon has licensed its Z-RAM technology to South Korea's Hynix, the No. 2 memory chip firm. Hynix built a test chip demonstrating the improvements and reported the results in June at the 2010 Symposium on VLSI Technology and Circuits, in Honolulu.

DRAM stores a bit of memory—a 1 or a 0—as charge on a capacitor. The problem is that DRAM makers have been cramming more memory onto less real estate, but the capacitors in a memory circuit can't be shrunk as fast as the transistors. The reason? A capacitor must be big enough to hold a recognizable charge. "Building the capacitor is becoming very challenging," says Jeff Lewis, senior vice president for marketing and business development at Innovative Silicon. "At 30 nm, people are pretty much running out of solutions to how to build them at all cost-effectively."

Instead of a capacitor, Z-RAM relies on what's called the floating-body effect. By building a transistor in a layer of silicon deposited atop silicon dioxide—an insulator—you electrically isolate the transistor. When current passes through the transistor, some electrons create electron-hole pairs. The transistor's drain lets the extra electrons leave, but it traps the holes, resulting in a net positive charge. That "floating" charge can be read as a 1. Increasing the voltage on the transistor gate empties the holes out through the source electrode, which can be read as a 0. With no capacitor needed, the bit cells only need to be the same size as the transistors, and those get smaller every year.

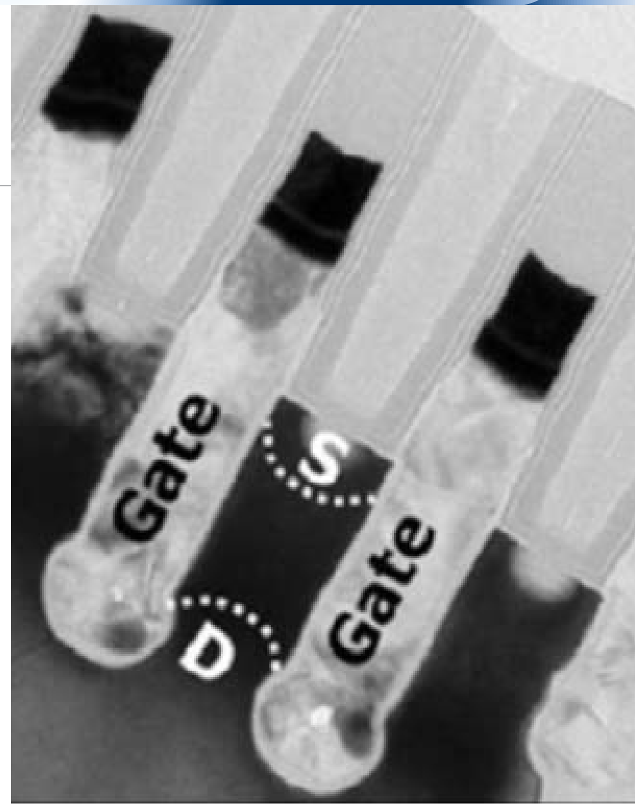

SIDEWAYS SOLUTION: Z-RAM works on ordinary silicon by turning the transistor vertically. *IMAGE: INNOVATIVE SILICON*

But building the device on a silicon-on-insulator wafer drove up the cost. Now the company has figured out a way to make memories on bulk silicon. Instead of building the transistor on top of an insulating layer, with the gate on top and the source and drain on the ends, Innovative Silicon aligns the transistor vertically so that the gates are on either side and the junctions are at the top and bottom. That provides the same isolation and the same floating-body effect. As a bonus, it fits in with the industry move toward a more 3-D transistor.

The vertical design, along with the new placement of the gates, also resulted in an operating voltage of 0.5 V, which the company says is 50 to 75 percent lower than any other floating-body memory device. The design therefore saves energy and makes the device compatible with existing power supplies, which keeps costs down, Lewis says. It also increases by a factor of 1000 the amount of time a cell can hold onto a bit amid the noise of other cells switching.

Innovative Silicon had originally envisioned using Z-RAM to replace static RAM, which is faster than DRAM and is integrated with the microprocessor. Because a Z-RAM cell is made up of only one transistor versus an SRAM's six, it can hold five times as many bits on the same area of silicon. Lewis says that the silicon-on-insulator market, necessary for the SRAM replacement design, never materialized, and that the need for a DRAM replacement has become more acute. Now, he says, "our focus is stand-alone DRAM only."

Lewis says the capabilities of Z-RAM match the demands outlined in the DRAM road map out to 2020. He expects Hynix to be making samples for testing in commercial devices within two years.

Other companies, such as Intel and Toshiba, are also pursuing floating-body memory. —NEIL SAVAGE

FLIGHT

CSM

CMP

EAT PERIOD
(1 HOUR)

CREW STATUS REPORT (SLEEP)
SELECT COMM NORMAL
LUNAR CONFIGURATION

COPY PAD
(SEE GET 114:20)

MSFN

0130 EDT
112:00

112:08

112:30

112:39

112:45

113:00

MSFN

CDR

FINAL PRE-EVA
DEPRESS CABIN
FINAL SYSTEMS

INITIAL EVA
EGRESS TO PLA
RELEASE MESA
DESCEND LADDE
REST/CHECK EM

ENVIRONMENTAL
ZATION
CHECK STABIL,

CONT SAMPLE CO
COLLECT AND S

MISSION	EDITION	DATE
APOLLO 11	FINAL	JULY 1, 1969

MSC Form 2189 (OT) (Nov 68)

FLIGHT PLANN

the big picture

ONE SMALL STEP

In April, Bonhams auction house held a space history sale in New York City timed to coincide with the 40th anniversary of Apollo 13. An Apollo 11 checklist page that details the countdown to the first footsteps on the moon—including instructions to "don gloves," "egress to platform," and "descend ladder"—sold for US $152 000. The sheet is framed with an Apollo 11 flight patch and boasts Neil Armstrong's signature, along with his historic first words from the moon's surface: "One small step for a man—one giant leap for mankind." The page was allegedly given as a gift to NASA press officer John McLeaish just days after the first moonwalkers splashed down in the Pacific Ocean. But the purchase is not without intrigue: Armstrong himself swears he never signed the quote for anyone.

PHOTO: BONHAMS

hands on

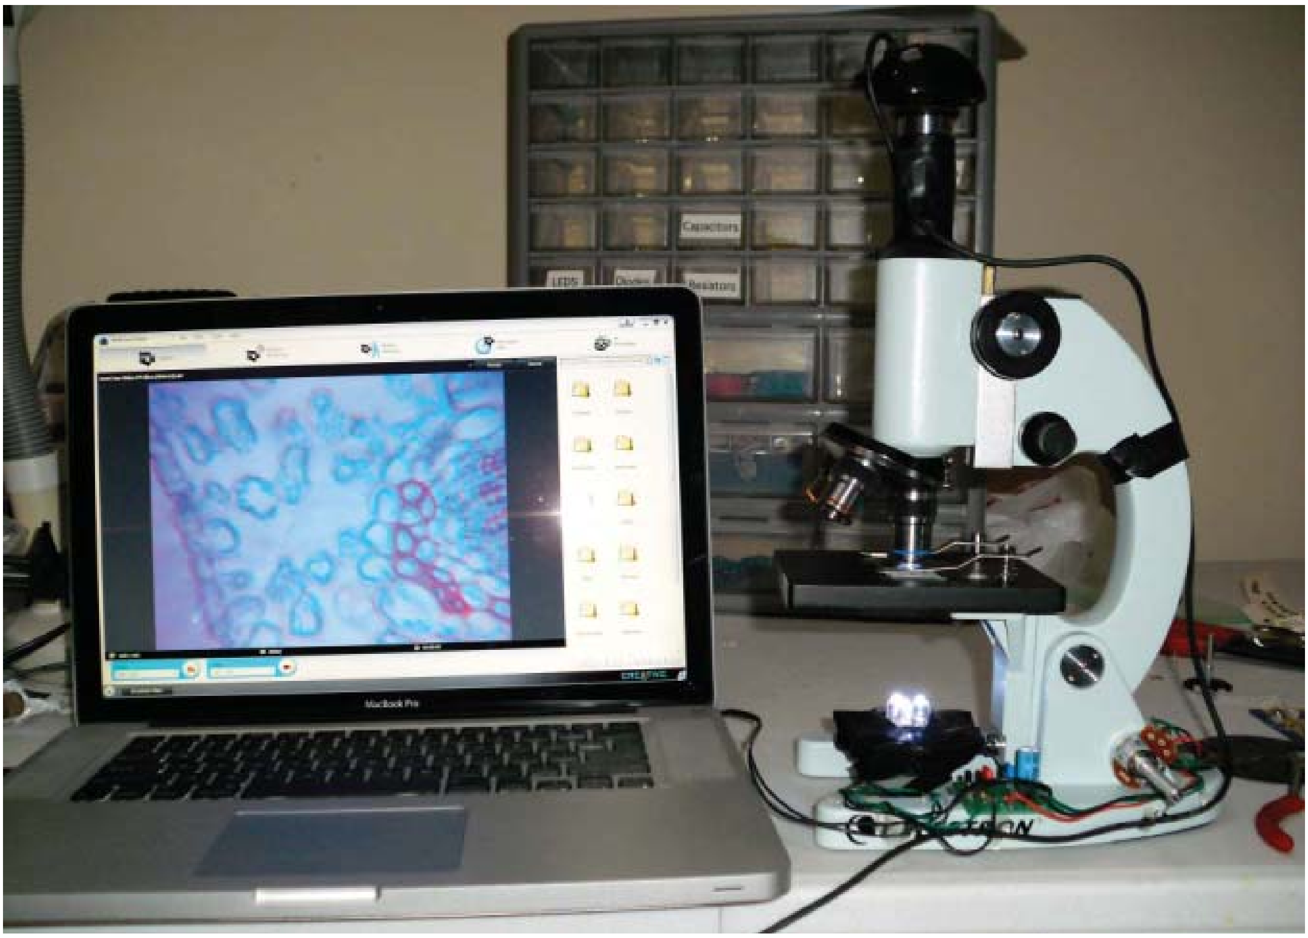

Think Big, See Small

How to build your own digital microscope

EVERYONE LOVES looking through a microscope. That got me wondering how hard it would be to convert one to capture images digitally. Sure, you can buy digital scopes ready made, but they tend to be either expensive or of poor optical quality. I wanted one that had decent optics and would also capture a decent-size image.

I started with a Celestron 44102 400x microscope (which cost me US $107 at Amazon) because it was highly rated and had lots of metal in the construction, making it a sturdy platform to attach things to. The only downside was the light source—a mirror. I wasn't about to depend on anything as primitive as the sun for my high-tech scope, but I left that issue for later.

The first step was to marry it to a digital camera. A webcam would conveniently transfer images to a computer via USB. I spent some time researching the optics and sought out some experts. Here's why: A microscope is designed to focus the image on the back of your eye when you look through the eyepiece. You can't just stick a webcam directly onto the eyepiece, because the webcam is also trying to bring an image into focus on the sensor inside the camera, resulting in a kind of optical double jeopardy. However, if you remove the lens from the webcam and hold it the right distance from the microscope eyepiece, you'll get a perfectly formed image.

I went back online to shop for a webcam, steering clear of the fancier ones. For one thing, I didn't need a microphone. Also, I wanted to avoid features like auto focus or face recognition because they might make it hard to get a sharp image. I ended up with a Creative Live Cam Video IM Ultra, a $50 model that shoots up to 1.3 megapixels—a lucky choice because it has a plastic ring that stands half an inch or so away from the body of the camera. After I removed the lens (a fairly painless procedure) and replaced the ring, there was just the right amount of space between the camera and the object to bring the image into perfect focus on the sensor. A few dabs of Krazy Glue attached the ring to the eyepiece, and I had a digital microscope. The first images I captured, using the stock slides included with the microscope, were amazing. Some of the stained plant cells in particular were extremely colorful and detailed.

As satisfying as the result was, it would have made for an awfully simple project. So, on to dealing with that light source issue!

As detailed in a previous article ["Board Certified," *IEEE Spectrum*, April 2010], I found a neat little circuit that would let me create a dimmable LED array, using a 555 timer, a transistor, a pair of capacitors, two diodes, and some resistors. I designed a printed circuit board and used a low-volume prototyping service to fabricate the actual PCB. A few weeks and $20 later, I had two perfect boards. I attached one to an array of 10-millimeter white LEDs and an old 12-volt power adapter I had lying around. I now had a finely controllable light source. By turning the 10K potentiometer, I had a nearly infinite variation of lighting levels available.

The result is a fully functional digital microscope—I can control the lighting to capture video or still images regardless of external conditions. And it's really nice to be able to sit back and look at a big image on my laptop rather than craning my neck over a lens.

For the truly lazy among us, it would be nice to have remote control over the position of the stage that holds the slide. That would let you move the image around with the click of a mouse, "bookmark" interesting features on a slide, or automatically capture high-resolution mosaic images. I decided to try to make a prototype.

With an infinite budget, I would use linear actuators or a rack-and-pinion gear system, but for my prototype I made a chain-driven system, using two 12-V stepper motors ($26 on eBay), one for each of the two drive chains. Next, I needed a controller board that could run the motors. From Peter Norberg Consulting, I got a USB board ($85) that can drive two motors and includes such nice features as the ability to use end-of-motion detection switches.


1


2

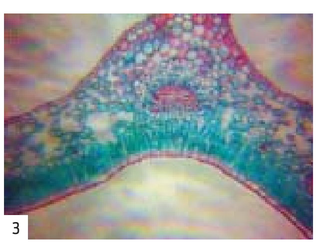
3


4

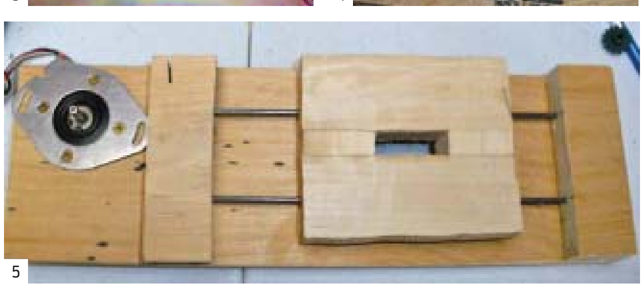
5

GOING DIGITAL: A fully digital microscope **[opposite page]** can look at a single blade of grass at 400x magnification on a computer screen, via an attached webcam. This page: **[1]** A close-up of the webcam, mounted to the eyepiece; **[2]** the LEDs and control board; **[3]** an image of plant cells, captured at 100x. The slide may be manipulated remotely via **[4]** reduction gearing and **[5]** a movable stage with a stepper motor installed. PHOTOS: JAMES TURNER

The last components were the gears and the chain connecting them, needed to gear down the motor, which had only eight steps per rotation. These components were surprisingly hard to find. I ended up with a pair of Vex Robotics sprocket-and-chain sets ($30 each). The nice thing about the Vex chain is that you can adjust the length easily, and it comes with a nice selection of gears.

I prototyped in wood rather than metal; it's cheaper and I have a lot more equipment for it. I built a stage for the slide, with nylon spacers glued to the bottom. Metal rods run through the spacers, letting the stage slide freely back and forth. A chain runs from the reduction gear coming off the motor, connecting to either end of the stage through another gear on the opposite side. It works like a charm, but to build a second gear-and-chain assembly that the first one would sit on, for the other axis, in the limited space available, would require purpose-built gearing and small motors as well as metal or plastic components. In other words, this project addition is not for the do-it-yourself faint of heart.

Meanwhile, I'm happy with my digital LED-lit scope. And a pair of fingers work just fine to move the slide around. —James Turner

Creative: *http://us.store.creative.com/index.html*
Peter Norberg Consulting: *http://www.stepperboard.com*
Vex Robotics: *http://www.vexrobotics.com*

DENIS POROY/AP PHOTO

THE KING OF COMIC-CON

In his spare time, software engineer John Rogers presides over one of the world's largest and strangest conventions

THE LAST thing you want to do is rankle a group of *Star Wars* stormtroopers. The annual geekfest known as Comic-Con International is not only a four-day exploration of the intersection of science fiction, comics, science, and pop culture; with 125 000 attendees, it's one of the largest conventions in the world, requiring heroic feats of engineering to keep the lines moving and the fans happy. Luckily, the guy at its helm is an engineer.

By day, John Rogers is a mild-mannered director of engineering at an undisclosed San Diego telecommunications company. But on the third week in July, he sheds his suit and tie and morphs into the president of Comic-Con. You'll probably spot him wandering the San Diego Convention Center, past a parade of Jokers, Supermen, and blue-skinned Na'vis, walkie-talkie in hand.

"My job at Comic-Con is all about managing crowds and people," Rogers says, but a list of his other responsibilities seems endless: avoiding empty seats in sold-out programs, arranging quick entrances and exits for panelists and audiences, pinpointing how many badge scanners and data-entry people are needed, and honoring fans who have lined up early. "If you're passionate enough to show up at dawn for a panel, then we have to work hard to make sure it happens for you," he says.

Rogers has done the honors since 1986. He devotes "more hours a month than I care to calculate" to managing the US $7 million budget, a staff of 20 full-time employees, and 3000 volunteers. He also oversees CCI's San Francisco–based sister conventions, WonderCon and Alternative Press Expo, which together attract another 40 000 attendees and 400 volunteers.

This year, Rogers is electronically streamlining the workflow by hiring a programmer to develop custom software that tracks where, when, and which kind of security and signage are needed throughout the convention center's 57 200 square meters (616 000 square feet). A database, fronted by an Excel spreadsheet, links the security shifts of the more than 300 daily guards with the 400-plus daytime and evening events. Another feature shows a detailed but changeable map indicating sign placement around artist tables, publishing-house booths, and movie studio kiosks. "As far as we know, nothing exists like this today," Rogers says.

"There's not a lot of technology or understanding of how you can use computers to make the process better," he adds. "If you look at the convention industry, you'll see people with binders that are 10 inches thick with explanations of where to place everything and position people."

Rogers began at Comic-Con as a projectionist when he was in high school. After graduating with a bachelor's in computer science from the University of San Diego in 1983, Rogers worked steadily in computer engineering—first for Sperry Corp., an electronics firm that became part of Unisys Corp., and then as a software developer at the now-defunct Simpact—before embarking on his current job in 1995.

Initially a comics-centric convention with science fiction and media components, Comic-Con has seen an increase in real-science programming, in part because Rogers loves sci-fi, programming director Eddie Ibrahim was formerly a biologist, and treasurer Mark Yturralde—a hardware engineer who helped develop the Comic-Con iPhone app—has a major jones for anything space related.

"One of my fondest memories was when we had people in to do Tesla coil demos and talk about the science behind it," says Rogers. "I have this really cool picture of me sitting on top of the coil at night, looking over the harbor, and holding two rods with lightning coming off of them." —SUSAN KARLIN

Comic-Con International
22–25 July 2010
http://www.comic-con.org

careers

WHERE THE JOBS ARE

The news is good but not great for engineers looking for work in 2010

LAST YEAR, pink slips were seemingly everywhere for engineers and computer scientists, as the likes of Texas Instruments and Microsoft laid off employees in record numbers. Things are better for engineers this year, but the signals are still mixed. Tech companies plan to hire at least as many electrical engineers as last year, but those already laid off are having a hard time finding jobs. And while new grads are getting fewer offers, they're doing better than their peers—according to a recent report, engineering degrees accounted for eight of the 10 highest paid degrees in the United States.

Though last year the average college graduate got a lower starting salary than the year before, computer-science majors saw an increase of 4.7 percent, to US $60 426, according to the latest salary report from the National Association of Colleges and Employers, in Bethlehem, Pa. Electrical engineers saw an increase of 3 percent, to $59 326.

"If we look at long-term unemployment data, there has been typically one peak in each decade," says IEEE-USA president Gordon Day. "We've had two—2003 and 2009—perhaps a sign of greater volatility in technology employment."

Hiring remains strong in aerospace, defense, and energy. There is also a shortage in the power and energy sector, Day says. "In areas like renewable energy and the smart grid, demand has increased faster than students can be educated."

Jobs for computer scientists and engineers will grow much faster than the average for other occupations over the next eight years, according to the U.S. Bureau of Labor Statistics (BLS), in part because of a boom in wireless technologies, electronic records, data processing, and information security. Other good employment bets include patent law firms, management consultancies, and financial companies.

STEPHEN WEBSTER/GETTY IMAGES

On the other hand, unemployment data published early this year by the BLS indicate things are not going well for laid-off electrical engineers. The number of EEs without jobs fell from 7.3 percent in the third quarter of 2009 to 5.2 percent in the fourth. Coupled with a 3 percent dip in the number of employed EEs, the data suggest that EEs who previously identified themselves as unemployed have either stopped looking for jobs or are switching to other fields.

In Europe, unemployment is about 3 percent for electrical and electronics engineers; overall rates of unemployment hover in the double digits. Hiring had slowed down along with the rest of the global economy but is slowly starting to pick up. Companies large and small are looking for a few new hires, especially those with specific skills in software and hardware design.

Reliable numbers for employment in Asia are hard to find, but indicators point to a healthy job market for EEs. The number of science and engineering doctoral degrees in China, India, Japan, South Korea, and Taiwan is increasing, as is R&D spending, which is expected to surpass that in the United States if it hasn't already, according to the National Science Foundation. U.S. and European firms have been trawling this well-stocked pool of engineering talent.

Employers are being selective about where they look for talent. With so many experienced professionals on the market, they have their pick of qualified candidates. "It makes more sense these days to hire someone with experience than someone new, and the cost difference might not be as big as it once was," says Suzanne Kahn, who represents nanotechnology company Nantero, of Woburn, Mass.

This has affected on-campus hiring. While recruiting has been vibrant at top engineering schools, other schools report a drop. At Ohio State University, recruiting for EEs and computer engineers is down 13 percent, and fewer students are getting offers. "Salaries continue to be good," says Douglas Williams, electrical and computer engineering professor at Georgia Tech. "But there are fewer offers per student. And the bottom half of the class is having a much harder time finding employment." —PRACHI PATEL

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

tools & toys

books

REVIEWED ONLINE THIS MONTH
Go to http://spectrum.ieee.org/reviews

Taming the Beloved Beast: How Medical Technology Costs Are Destroying Our Health Care System
By Daniel Callahan; Princeton University Press, 2009; 288 pp.; US $29.95; ISBN 978-1-4008-3094-7
Reviewed by Kenneth R. Foster

When the Lights Went Out
By David E. Nye; MIT Press, 2010; 304 pp.; $27.95; ISBN 978-0-262-01374-1
Reviewed by Mark Anderson

Early FM Radio: Incremental Technology in Twentieth-Century America
By Gary L. Frost; Johns Hopkins University Press, 2010; 208 pp.; $60.00; ISBN 978-0-8018-9440-4
Reviewed by Michael Riezenman

REVIEW: ALIENWARE M11X

The King Kong of gaming computers now comes in a Fay Wray–size package

Alienware M11x
US $800 and up
http://www.alienware.com

I DIDN'T KNOW Alienware made a netbook!"

"They don't."

That was a conversation I had several times at work after the arrival of my M11x, a tiny monster of a machine. Alienware, for those unfamiliar, is the prestige brand of Windows gaming PCs owned by Dell, and it never does anything by halves. If you want ultimate performance and can afford a top-of-the-line gaming laptop that costs as much as four or five "regular" PCs, Alienware can kit you out.

So it was with some bemusement that the gaming press greeted the news of the M11x, which looks for all the world like a netbook—11-inch screen, 3.25 centimeters thick, no DVD drive—but with an Alienware-ish price that starts at US $800 and climbs quickly. I was unenthusiastic as well; it's as if Jaguar had tried to make a subcompact car. Having owned a couple of netbooks, I've enjoyed their portability, but not the frustrating torpor of an Atom processor struggling to run Windows.

So the M11x didn't seem worthwhile—until I read the specs. Ultralow-voltage Core 2 Duo processor, Nvidia GT 335M video with 1 gigabyte of VRAM, up to 8 GB of RAM, HDMI video output...what is this thing? A visit to the Alienware booth at a trade show brought the answer.

It's a top-of-the-line notebook in the body of a netbook. Clearly, Alienware is sensitive about the image of its laptops, where 5 or 6 kilograms is pretty standard, and they set out to change that. What they came up with is a joy to use: a 2-kg yet solidly built machine that can stand with the lions and run with the gazelles. The company's over-the-top product design is there: alien-head power button, illuminated speakers, and keyboard with customizable colors. But I don't mind a little flash now and then. It puts me in the mood for gaming.

I haven't come close to filling the 500-GB, 7200-rpm hard drive, but I've installed *Crysis*, *Modern Warfare 2*, and *Mass Effect 2*, all of which can punish the average desktop machine. They play beautifully, maintaining high frame rates throughout. Game settings that rely on the video card, like texture resolution and post effects, can be maxed out, while features that tax the admittedly underpowered central processing unit—say, a huge number of scene objects—should be set somewhat lower.

The CPU was obviously chosen for low heat output and battery life: Using the low-power Intel graphics card (which is switchable on the fly with the Nvidia card), you can get more than 7 hours of basic work done on a single charge. Heavy gaming with Nvidia graphics will run 3 to 4 hours.

I've also run applications like Photoshop CS5 and ZBrush, powerful programs that demand good hardware, and the M11x delivers. Filters are fast, brush strokes smooth, and large image files can be edited with no noticeable performance degradation. Microsoft Office likewise runs extremely well.

It's not as elegant or thin as a MacBook Air or a Dell Adamo, but the M11x blows them away in performance in a similarly small package. And it has a glowing alien-head power button. What's not to love? Turns out Alienware does make a netbook—and then some.

—Harry Teasley

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

reflections

BY ROBERT W. LUCKY

Contemplating a Dead Computer

IT WAS my early morning ritual. Having pushed the button to boot my computer, I turned my attention to the waiting pleasure of steaming coffee. After a few delicious sips, I glanced at the monitor. Strangely, it was black. I pushed the button again. To my distress, this had no effect on the monitor or on the absolute quiet of my PC. I unplugged the PC, waited, and tried again, to no avail.

How could this be? Had the computer died in its sleep? It had worked flawlessly for a year, including the night before. Perhaps that first jolt of electricity in the morning had—like a too-strong cup of coffee—arrested its frantic little heartbeat.

I spent the next hour pushing the on button repeatedly in the ridiculous belief that the computer would spring to life. Then I went into troubleshooting mode. It wasn't hard to determine that the motherboard itself was the problem, and that issue, as they say, is above my pay grade. The computer was officially dead.

As I drove to the local computer store, I thought about how the world has changed in the last couple of decades. Once, I would have fixed that computer and bragged about my accomplishment. Once, I would have mourned my dead computer as a friend who had stood by my side. But that was no longer the case. As I was leaving the house, my wife had called out to me. "Buy one that will last," she had admonished. "None of them will," I replied.

I realized that I had come to think of computers and other electronic gadgets as disposable. Even if they don't break, they're good for only a couple of years anyway. Technological progress overruns them, and they get bogged down with the accumulation of junk. The sad reality is that they're not worth fixing. The cost of identifying and fixing the problem is more than that of a new system. After all, electronic circuitry gets ever less expensive, while the cost of skilled labor escalates.

BRAIN STAUFFER

That evening I was having dinner with some friends, and I mentioned what had happened. "But there aren't any moving parts," someone remarked. "What could go wrong?" It was a good question. Anywhere something whirs, like the hard drives, or where there is burning heat, like around the microprocessor, we expect bad things might happen. The good news is how remarkably robust electronic circuitry is. Hundreds of millions of transistors must do their jobs perfectly. I always consider it a miracle that computers work in the first place. But once they do, the circuits seldom fail.

I've always taken a great pride in the engineering achievements embodied in a computer. The sophistication and complexity of the software and the hardware—and in the network beyond—epitomize the wonderful progress we have made. But here I am at the store looking at new computers and thinking of how little I miss the one that has just quit working. I'm thinking instead that the real value was my data, and that for once, the hard drive had outlived its host. My precious data was still safe.

Ah, but I had jumped to this conclusion too quickly. First, my backup routines had been imperfect; I had not backed up frequently enough. After all, the computer had always worked, hadn't it? Worse, my computer had outsmarted itself. The manufacturer had cleverly used two hard drives, which appeared to the operating system as a single logical drive in a RAID array—brilliant!—but alas, striping the data across the drives in a way that only the dead motherboard fully understood. Not in a way that would be impossible to reconstruct, but once again the considerations were time and expense. And once again, it wasn't worth fixing.

Home again, I bid farewell to my deceased computer. The moment of truth was upon me. Could I put the machine out with the trash? Alas, no. I still harbor the hope that the next time I push the button, my computer will come alive. ❑

THE TROUBLE WITH MULTI-CORE

CHIPMAKERS ARE BUSY DESIGNING MICROPROCESSORS THAT MOST PROGRAMMERS CAN'T HANDLE

BY **DAVID PATTERSON**

HARRY CAMPBELL

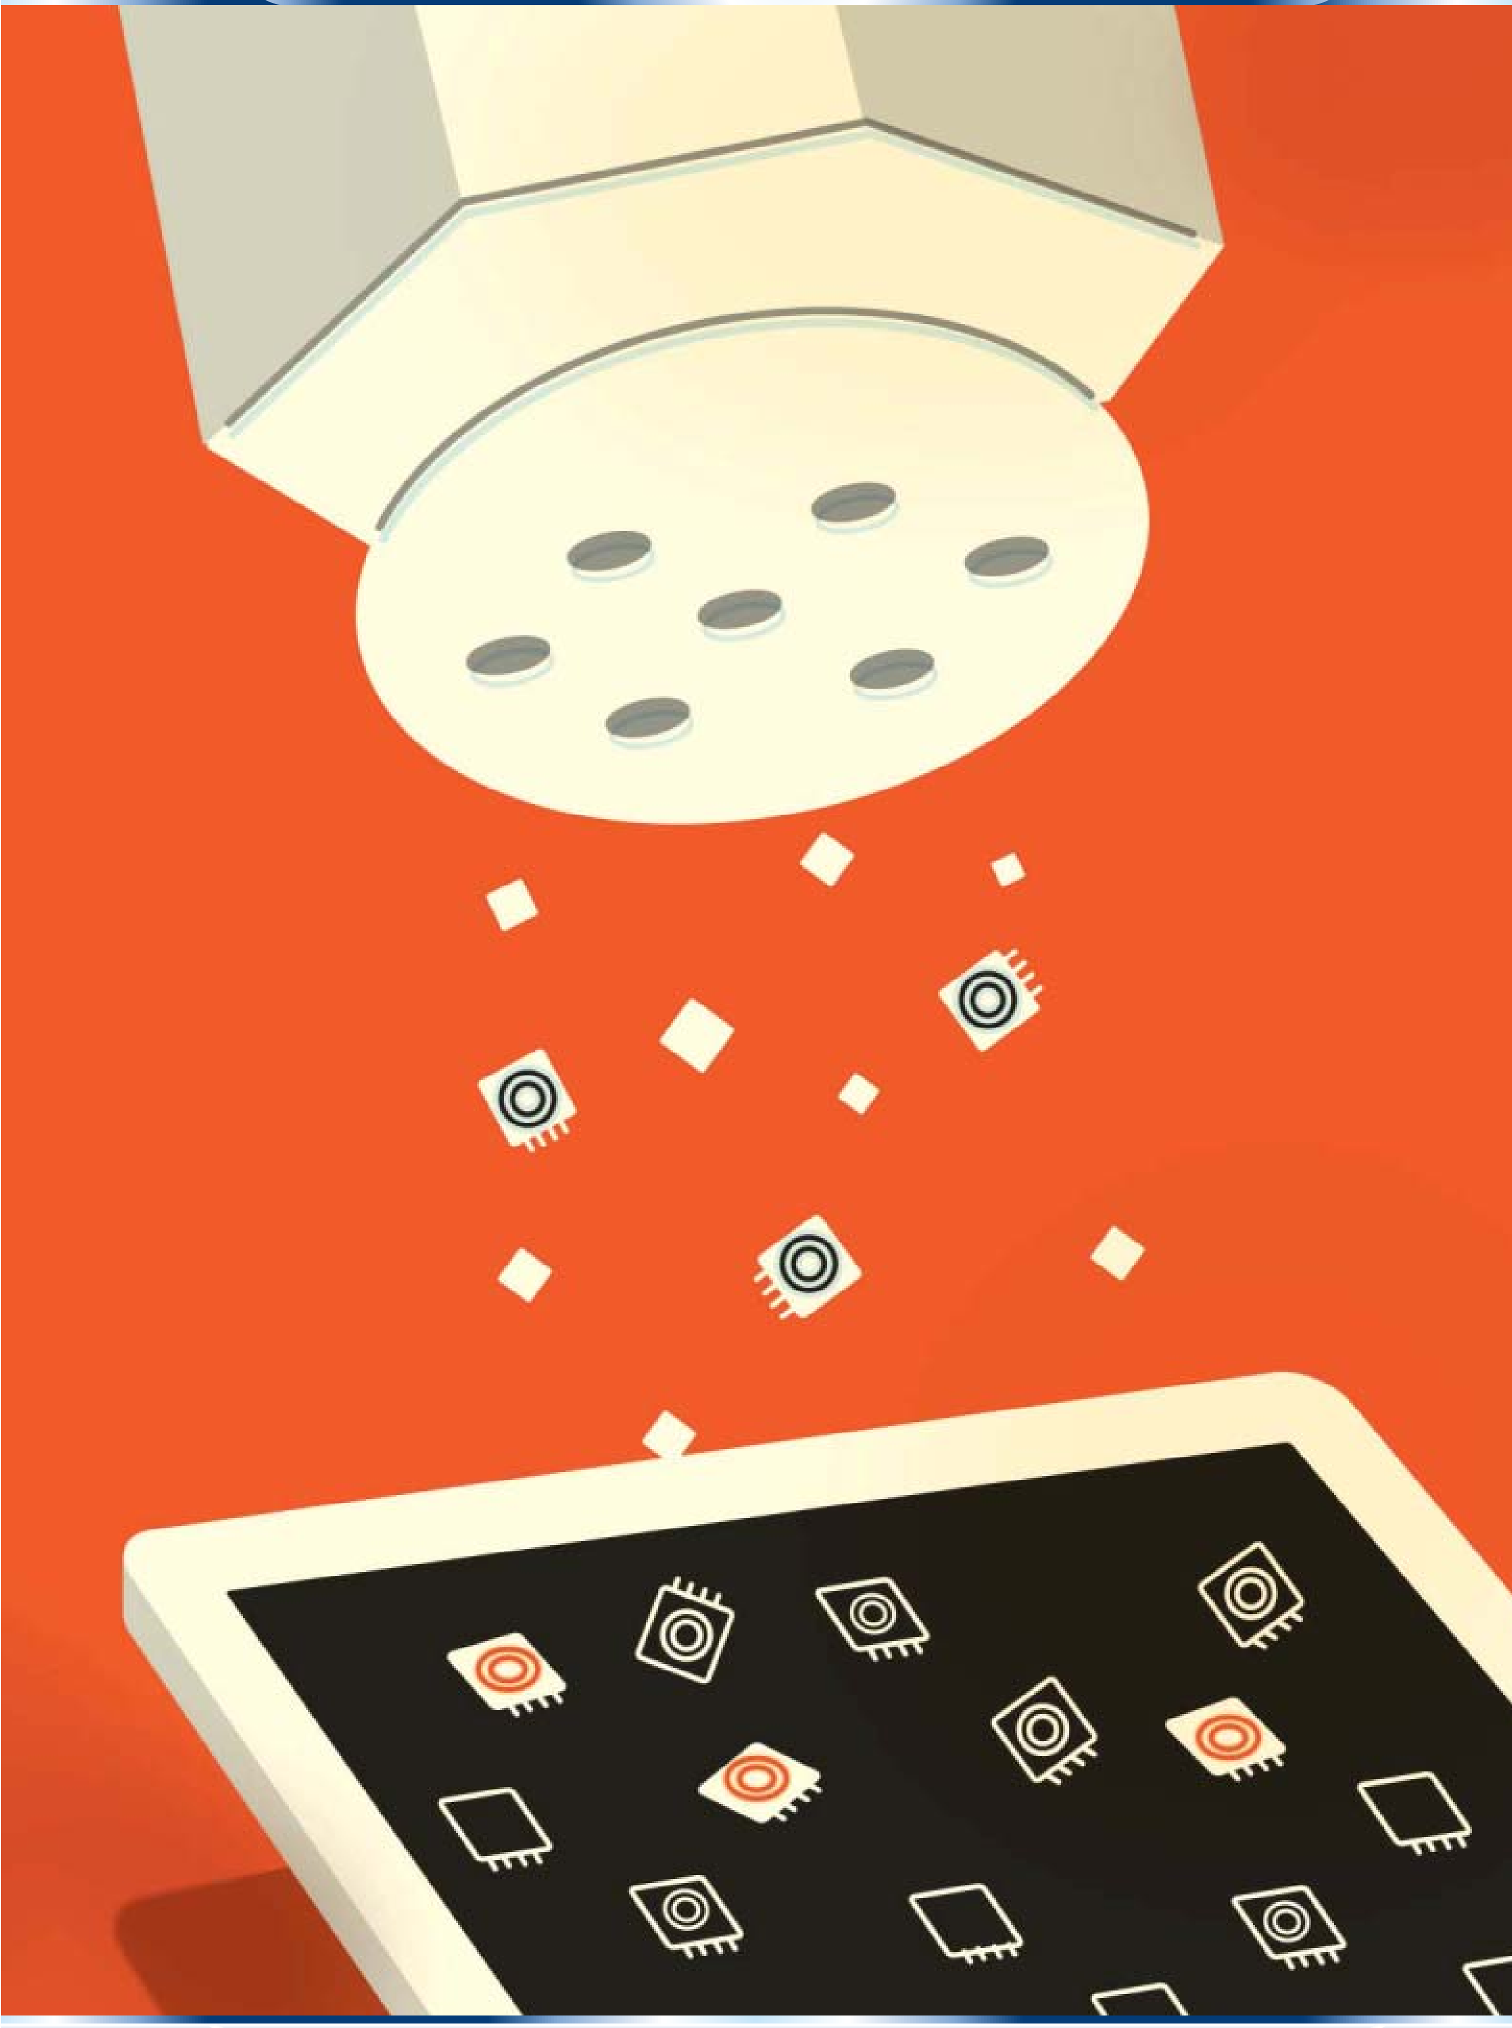

In 1975, future Hall of Famer Roger Staubach had the football but little else in a playoff game against the Minnesota Vikings. Behind by four points at midfield with 24 seconds to go, the Dallas Cowboys quarterback closed his eyes, threw the ball as hard as he could, and said a Hail Mary. (For you soccer fans, this would be like David Beckham taking a shot on goal from midfield late in injury time.)

His prayer was answered. Staubach's receiver collided with a Viking defender just as the ball arrived but nevertheless managed to pin the football against his leg, scoring the touchdown that took the Cowboys to the Super Bowl. (Imagine Beckham's long ball beating the goalie.) Ever since that game, a desperate pass with little chance of success has been labeled a Hail Mary.

Thirty years later, the semiconductor industry threw the equivalent of a Hail Mary pass when it switched from making microprocessors run faster to putting more of them on a chip—doing so without any clear notion of how such devices would in general be programmed. The hope is that someone will be able to figure out how to do that, but at the moment, the ball is still in the air.

Why take such a gamble? In short, because there wasn't much of an alternative.

For decades, microprocessor designers used the burgeoning number of transistors that could be squeezed onto each chip to boost computational horsepower. They did this by creating microprocessors that could carry out several operations at once—for example, fetching the next instruction from memory while the current one was being executed. And chipmakers continually upped microprocessor clock rates, something the diminishing size of transistors readily allowed.

But around 2003, chipmakers found they could no longer reduce the operating voltage as sharply as they had in the past as they strived to make transistors smaller and faster. That in turn caused the amount of waste heat that had to be dissipated from each square millimeter of silicon to go up. Eventually designers hit what they call the power wall, the limit on the amount of power a microprocessor chip could reasonably dissipate. After all, a laptop that burned your lap would be a tough sell.

Designers now accept that although transistors will still get smaller and more numerous on each chip, they aren't going to operate faster than they do today. (Indeed, peak clock speeds are lower now than they were five years ago.) And if you tried to incorporate all those transistors into one giant microprocessor, you might well end up with a device that couldn't compute any faster than the chip it was replacing, which explains the shift to assembling them into multiple microprocessor cores instead. Although each core may have modest computational abilities, you'll have many of them at your disposal.

Such novel chips are called multicore microprocessors—or sometimes many-core microprocessors when a large number of cores are involved—to distinguish them from traditional single-core designs. In a sense, the core has become the new transistor, to borrow a phrase from Chris Rowen, president and chief technology officer of Tensilica, in Santa Clara, Calif. That is, from here on out, chip designers will concentrate on how to gang together lots of cores, just as the previous generation of microprocessor engineers thought about the circuitry they were creating at the level of individual transistors.

The trick will be to invent ways for programmers to write applications that exploit the increasing number of processors found on each chip without stretching the time needed to develop software or lowering its quality. Say your Hail Mary now, because this is not going to be easy.

WHEN THE PRESIDENT and CEO of Intel, Paul S. Otellini, announced in 2004 that his company would dedicate "all of our future product designs to multicore environments," why did he label this "a key inflection point for the industry"? The answer is clear to anyone familiar with the many now-defunct companies that bet their futures on the transition from single-core computers to systems with multiple processors working in parallel. Ardent, Convex, Encore, Floating Point Systems, Inmos, Kendall Square Research, MasPar, nCUBE, Sequent, Tandem, and Thinking Machines are just the most prominent names from a very long list of long-gone parallel hopefuls. Otellini was announcing that despite this sobering record, software applications in the future will run faster only if programmers can write parallel programs for the kinds of multicore microprocessors that Intel and other semiconductor companies have started shipping.

HARRY CAMPBELL

But why is parallel processing so challenging? An analogy helps here. Programming is in many ways like writing a news story. Potentially, 10 reporters could complete a story 10 times as fast as a single reporter could ever manage it. But they'd need to divide a task into 10 equally sized pieces; otherwise they couldn't achieve a full tenfold speedup.

Complications would arise, however, if one part of the story couldn't be written until the rest of it was done. The 10 reporters would also need to ensure that each bit of text was consistent with

EASY AS PI

Standard algorithms for determining the value of pi rely on sequential calculations [top], but you could also use a parallel approach [bottom].

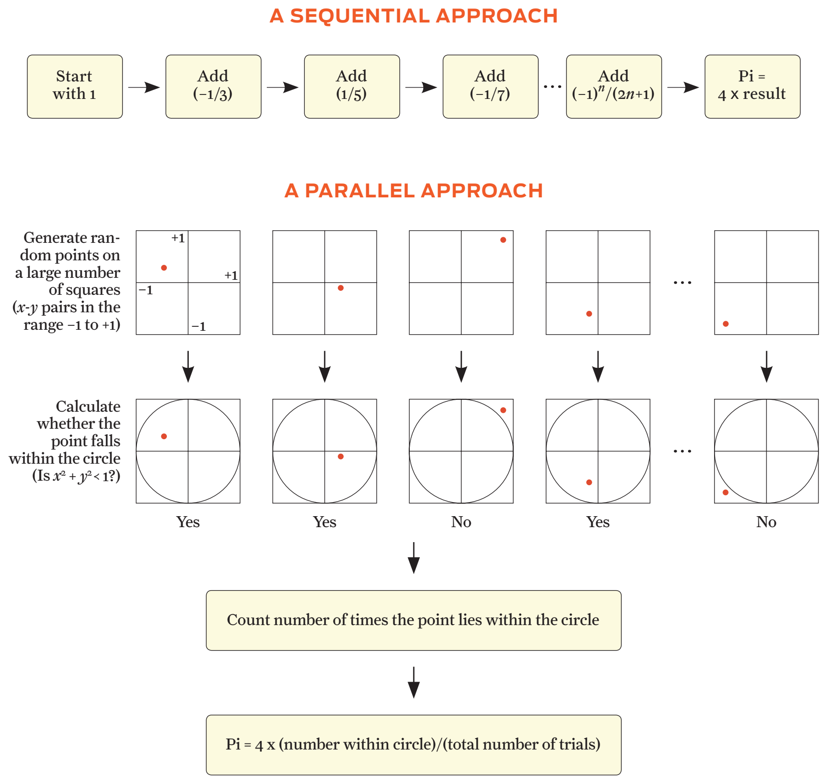

what came before and that the next section flowed logically from it, without repeating any material. Also, they would have to time their efforts so that they finished simultaneously. After all, you can't publish a story while you're still waiting for a piece in the middle to be completed. These same issues—load balancing, sequential dependencies, and synchronization—challenge parallel programmers.

Researchers have been trying to tackle these problems since the 1960s. Many ideas have been tried, and just about as many have failed. One early vision was that the right computer language would make parallel programming straightforward. There have been hundreds—if not thousands—of attempts at developing such languages, including such long-gone examples as APL, Id, Linda, Occam, and SISAL. Some made parallel programming easier, but none has made it as fast, efficient, and flexible as traditional sequential programming. Nor has any become as popular as the languages invented primarily for sequential programming.

Another hope was that if you just designed the hardware properly, parallel programming would become easy. Many private investors have been seduced by this idea. And many people have tried to build the El Dorado of computer architecture, but no one has yet succeeded.

A third idea, also dating back to the 1960s, is to write software that will automatically parallelize existing sequential programs. History teaches that success here is inversely proportional to the number of cores. Depending on the program, there will likely be some benefit from trying to automatically parallelize it for two, four, or even eight cores. But most experts remain skeptical that the automatic parallelization of arbitrary sequential code is going to be beneficial for 32, 64, or 128 cores, despite some recently published advances in this area.

All in all, things look pretty bleak. Nevertheless, there has been progress in some communities. In general, parallelism can work when you can afford to assemble a crack team of Ph.D.-level programmers to tackle a problem with many different tasks that depend very little on one another. One example is the database systems that banks use for managing ATM transactions and airlines use for tracking reservations. Another example is Internet searching. It's much easier to parallelize programs that deal with lots of users doing pretty much the same thing rather than a single user doing something very complicated. That's because you can readily take advantage of the inherent task-level parallelism of the problem at hand.

Another success story is computer graphics. Animated movies or ones with lots of computer-generated special effects exhibit task-level parallelism in that individual scenes can be computed in parallel. Clever programmers have even found parallelism in computing each image. Indeed, the high-end graphics processing units (GPUs) used to accelerate games on a PC can contain hundreds of processors, each tackling just a small piece of the job of rendering an image. Computer scientists apply the term "data-level parallelism" to such applications. They're hard enough to program, but in general they're easier than applications that don't offer this inherent parallelism.

Scientific computing provides a third success story—weather prediction and car-crash simulations being two well-known examples. These are long-running programs that have lots of data-level parallelism. The elite teams that create these programs are often combinations of Ph.D. computer scientists and people with doctorates in the sciences relevant to the application. Desktop applications rarely have that much intellectual horsepower behind them.

GIVEN THIS STARK landscape, you might not expect this latest foray into parallel computing to be greeted by success. But there are reasons for optimism. First off, the whole computer industry is now working on the problem. Also, the shift to parallelism is starting small and growing slowly. Programmers can cut their teeth on dual- and quad-core processors right now, rather than jumping to 128 cores in one fell swoop.

One of the biggest factors, though, is the degree of motivation. In the past, programmers could just wait for transistors to get smaller and faster, allowing microprocessors to become more powerful. So programs would run faster without any new programming effort, which was a big disincentive to anyone tempted to pioneer ways to write parallel code. The La-Z-Boy era of program performance is now officially over, so programmers who care about performance must get up off their recliners and start making their programs parallel.

Another potential reason for success is the synergy between many-core processing and software as a service, or cloud computing as it is often called. Google Search, Hotmail, and Salesforce are some here-and-now examples of such services, in which the application you need runs in a remote data center rather than on your own computer. These services are popular because they reduce the hassle to both users and providers. The user needs only a browser and needn't fool with software installation, upgrades, and patches. Software providers are happy, too, because their applications run only inside a data center where they control the environment. This has allowed their developers to improve their software much faster than can programmers writing traditional "shrink-wrap" applications, which must run on a host of different computers with many combinations of hardware and software installed.

Expert programmers can take advantage of the task-level parallelism inherent in cloud computing. To service millions of users, these programmers divvy up the work to run on thousands of computers. Because their software already uses many processors, it's easy for the people mounting such operations to embrace many-core chips. Indeed, these cloud-computing providers see many-core as a welcome way to reduce costs rather than as a disruptive technology. So expect the coming proliferation of many-core processors to boost today's rapidly growing zeal for cloud computing.

DESPITE THESE REASONS for hope, the odds are still against the microprocessor industry squarely completing its risky Hail Mary pass and finding some all-encompassing way to convert every piece of software to run on many parallel processors. I and other researchers at the main centers of parallel-computing research—including Georgia Tech, the University of Illinois, Rice University, Stanford, and the University of California, Berkeley—certainly don't expect that to happen. So rather than working on general programming languages or computer designs, we are instead trying to create a few important applications that can take advantage of many-core microprocessors. Although none of these groups is likely to develop the ultimate killer app, that's not the intention. Rather, we hope that the hardware and software we invent will contain some of the key innovations needed to make parallel programming straightforward. If we're successful, this work should help to usher in whatever application ultimately wins the "killer" distinction.

HARRY CAMPBELL

For example, my colleagues and I at Berkeley's parallel computing laboratory—the Par Lab—have decided to pursue just a few target applications. One is speech recognition, or perhaps I should say speech understanding. Our hope is to improve speech-recognition software dramatically so that a computer can recognize words spoken in crowded, noisy, and reverberant environments. That would surpass today's crude speech-recognition software and allow such things as real-time meeting transcription. Such software exists now, but those programs generate a frustratingly large number of mistakes.

One problem we're facing in this effort is that microprocessors with large numbers of cores are not yet being manufactured. So we have nothing to run our experimental software on. And a prototype many-core microprocessor would take years to design and millions of dollars to fabricate. We could, in principle, emulate such chips in software. But software running at the level of detail needed to evaluate a 128-core design could take days to simulate a few seconds, which means that iterations between hardware and software improvements would be excruciatingly slow.

We can, however, skirt this roadblock by using field-programmable gate arrays (FPGAs) to simulate future computers. FPGAs are integrated circuits that contain large collections of circuit components that can be wired together on the fly using a special language to describe the desired hardware configuration. And they can be rewired as many times as needed. So they offer the best of both worlds, having the flexibility of software but also the ability to run 250 times as fast as software simulators. This prospect inspired the Research Accelerator for Multiple Processors (RAMP) project, a collaboration of nearly a dozen universities and companies, to create and share a common infrastructure to accelerate research in many-core designs.

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

DO YOU SPEAK MULTICORE?

There's no lack of languages designed to support parallel processing—this is just a selection of them. Still, they don't make parallel programming easy or straightforward.

ABCPL	Cilk	Go	Millipede	SALSA
ActorScript	Clojure	Id	MultiLisp	Scala
Ada	Curry	Janus	Modula-3	SISAL
Afnix	DAPPLE	JoCaml	Nimrod	SR
Alef	E	Join Java	Occam	Stackless
Alice	Eiffel	Joule	Orc	SuperPascal
APL	Emerald	Joyce	Oz	VHDL
Axum	Erlang	LabView	Pict	XC
C*	Fork	Limbo	Polaris	Zounds
Chapel	Glenda	Linda	Reia	ZPL

How many RAMP configurations and specialized programming environments will be needed is still anyone's guess. In 2004, Phillip Colella of Lawrence Berkeley National Laboratory claimed that seven numerical methods would dominate scientific computing for the next decade. Starting with that claim, he, I, and a small group of other Berkeley computer scientists spent two years evaluating how well these seven techniques would work in other fields of endeavor. We ended up expanding the list to 12 general methods. Think of these as fundamental computational patterns, ones that are as different as the various motifs found in, say, Persian carpets—trees, spirals, paisley, and so forth.

Some of the 12 computational motifs are embarrassingly parallel. Take, for example, what are called Monte Carlo methods, which examine many independent random trials of some physical process to determine a more general result. You can do a lot with this approach. You could, for instance, determine the value of pi. Just compute what happens when you throw darts at a square board. If the darts hit random points on the square, what fraction of them fall within the largest circle you can draw on the board? Calculate that number for enough darts and you'll know the area of the circle. Dividing by the radius squared then gives you a value for pi.

Other motifs can be a lot harder to carry out in parallel, such as the common problem of sequencing through a series of well-defined states, where the rules for transitioning from one state to another depend on the values of various external inputs. A sequential computer calculates which state to assume next based on which state it is in and the inputs presented at that moment. Having multiple cores available doesn't do much to speed up that process, so the only opportunity to run through the sequence of states faster is to figure out ahead of time the state transitions that might be coming up. But that requires the computer to guess which state it might soon find itself in and how the inputs might change in the meantime. And when you guess wrong, it takes so much time to recover that you'll go even slower than you would have without any guessing. The hope is that you'll guess correctly most of the time, so that on average you come out ahead. Figuring out how to program such speculation about state transitions is tricky, to say the least.

IN 1995, I made some public predictions of what microprocessors would be like in the year 2020. I naively expected that the information technology community would discover how to do parallel programming before chipmakers started shipping what I then called "micromultiprocessors." From the perspective of 2010, I now see three possibilities for 2020.

The first is that we drop the ball. That is, the practical number of cores per chip hits a ceiling, and the performance of microprocessors stops increasing. Such an outcome will have a broad impact on the information technology industry. Microprocessors will likely still get cheaper every year, and so will the products that contain them. But they won't pack in more computational oomph. Consider netbooks as the first step down this cost-reduction path. Such an evolution will only accelerate the shift to cloud computing, because the servers that are doing the real work will be able to take advantage of the parallelism of many-core microprocessors, even if desktops and handheld computers cannot.

Another possibility is that a select few of us will be able to catch today's risky Hail Mary pass. Perhaps only multimedia apps such as video games can exploit data-level parallelism and take advantage of the increasing number of cores. In that case, the microprocessors of 2020 may look more like the GPUs from Nvidia, Advanced Micro Devices, and Intel than the traditional microprocessors of today. That is, the GPU will be promoted from a sideshow to the main event. It's unclear whether such applications by themselves will be able to sustain the growth of the information technology industry as a whole.

The most optimistic outcome, of course, is that someone figures out how to make dependable parallel software that works efficiently as the number of cores increases. That will provide the much-needed foundation for building the microprocessor hardware of the next 30 years. Even if the routine doubling every year or two of the number of transistors per chip were to stop—the dreaded end of Moore's Law—innovative packaging might allow economical systems to be created from multiple chips, sustaining the performance gains that consumers have long enjoyed.

Although I'm rooting for this outcome—and many colleagues and I are working hard to realize it—I have to admit that this third scenario is probably not the most likely one. Just as global climate change will disadvantage some nations more than others, what happens to the microprocessor industry will probably be uneven in its effect. Some companies will succumb to the inability of microprocessors to advance in the way they have in the past. Others will benefit from the change in the new playing field of computing.

No matter how the ball bounces, it's going to be fun to watch, at least for the fans. The next decade is going to be interesting. ❑

CELLPHONE CRIME SOLVERS

Could the murder victim's BlackBerry lead to her killer? Increasingly, the answer is yes BY RICHARD P. MISLAN

It was a dark and stormy night. In a narrow alley on the west side of Chicago, a dead woman is slumped inside her Mercedes, clutching her mobile phone to her head, as if she were still in conversation. Rain pours through the broken car window, blown out by the bullet that went through the woman's temple. A small crowd is standing around the car as Chicago PD detective Nick Fasano arrives on the scene, responding to a 911 call made by a witness waiting in a restaurant line. Fasano knows he's looking at a homicide. For one thing, there's that bullet in her head. He immediately realizes that another sort of witness to this crime might be on the other end of that phone connection.

He reaches through the open car window to grab the phone and thumb through its recent call history. Then he stops himself. He knows better than to disturb a crime scene. And he's never seen that particular model of phone—he could potentially push the wrong buttons and destroy evidence. He needs to get that device to a forensic lab, where the information can be extracted properly, in a way that preserves not only the contacts, call histories, text messages, e-mail, images, and videos but also their admissibility in court.

DAN SAELINGER


33
MOTOROLA
CLR

If the fictitious Fasano were a television detective, forensic examiners would arrive in moments with high-tech tools and search the phone on the spot. Fasano would have the information he needed in 15 minutes and he'd solve the entire case within the hour. Here's the reality: Detectives don't carry forensic tool kits that let them search mobile devices. Instead, they photograph the scene and then remove the phone. But this procedure is riddled with pitfalls.

A detective who finds a mobile phone at a crime scene immediately has a decision to make—whether to turn the phone off or leave it on. If he turns it off, the investigator in the lab may have to deal with a password prompt when the phone restarts (60 percent of people password-protect their phones, according to a 2009 study by Credant Technologies). If the detective leaves it on, the phone could receive calls and text messages during the drive to the lab, which could force the device to overwrite information inadvertently. It's even possible that someone connected to the crime may hit the phone with a text or e-mail "bomb" that floods the phone's memory with messages that crowd out all other previous calls from the log.

To prevent that, the detective could leave the phone on and place it in a metal-mesh shielding bag to block all signals. Indeed, such products are readily available. But even that tactic would cause problems. Placing the phone in such a bag would erase vital location information stored by certain kinds of phones—after searching for a signal for a while, the phone would give up and zero out the register that holds location data. Bagging the phone would also drain the battery faster, because the phone, trying in vain to lock on to a tower, would boost its transmitting power to the maximum. And these preservation bags aren't completely impervious to wireless signals. Drive within a few dozen meters of a cellphone tower and all bets are off.

Back to Fasano (who, along with the others named in this article, is a fictional character based on actual cases I've come across in my work consulting for federal, state, and local law enforcement agencies). Fasano photographs the scene and takes the cellphone. Because the phone is already on and at least 75 percent charged, he leaves it on, without a shielding bag, taking the risk that incoming calls may overwrite evidence. He races to the forensic lab.

At the lab, Fasano tells the forensic examiner, Marla McKenna, about the phone and its relationship to the crime scene. McKenna tells Fasano the exam will take several hours. Then she photographs the front and back of the phone and takes a close-up of the screen. She notes the name of the manufacturer—HTC—and the carrier branded on the phone's case. She looks through her selection of nearly 100 different data cables, hunting for one that will fit. She'll need a charging cable soon as well.

She's not sure of the model number; that's usually stamped under the battery, and she's not ready to remove the battery and kill the power. She goes online to look at phones by HTC, trying to home in on the model by matching its styling with the pictures on the screen. She decides it's an HTC Magic—an Android-powered touch phone. Only now can she determine what hardware and software tools to use to examine this phone.

Believe it or not, there are actually quite a few forensic tool kits available that can suck out the binary contents of cellphones. The group includes the Universal Forensic Extraction Device from CelleBrite Mobile, XRY and XACT from Micro Systemation, Secure View from Susteen, Aceso from Radio Tactics, Device Seizure from Paraben Corp., MobileEdit Forensic from Compelson Labs, and the Forensic Suite from Oxygen Software. In choosing a forensic

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

FORENSICS TO GO: CelleBrite Mobile's forensics kit includes cables and a handheld computer that can pull pictures, videos, text messages, call logs, and other data from cellphones and PDAs at a crime scene.
PHOTOS: CELLEBRITE

tool, McKenna must consider quite a few options, because there is no one package that can disgorge the data inside each and every mobile phone on the market. Some kits were specifically developed to meet the needs of the examiner, but most are offshoots of consumer products that allow people to transfer photos, address books, ringtones, and other stored information to a new phone from an old one, or to synchronize calendars and contacts with files stored on a personal computer.

Such consumer products, known as synchronization tools, generally don't meet all the needs of a forensic specialist, because they don't protect the phone data from tampering, which means that the data might not be admissible in court. Tools developed specifically for the examination of evidence don't make it impossible to tamper with the data, but they make it easy to prove that tampering did or didn't happen. They do this by means of a mathematical technique called hash functions. When the forensic software pulls the data from the phone into the computer, it automatically runs a set of mathematical operations on the data, using those operations to generate a series of numbers. Later, if attorneys or judges question the quality of the evidence, the software again runs the operations on the data file and generates a new series of numbers. If the two sets of numbers don't match exactly, it's likely the data changed along the way—which would mean its authenticity could be challenged in court.

The data we've been talking about so far comes from the phone's active memory, stored in static RAM chips built into the phone. This active memory contains the user's contacts, call history, text messages, images, videos, e-mail, and cached Web pages, as well as basic information about the phone needed to connect it to the network. Phones may also have removable memory cards, usually in the MicroSD format. Most forensic tools pull this data as part of the active memory; the card may also be removed from the phone later and read as if it were a flash-based hard drive. For some phones, that's where the story ends. But phones built to the Global System for Mobile Communications (GSM) standard have an additional storage area—a removable smart card known as the subscriber identity module (SIM). GSM is the most common phone standard outside the United States and is also used by AT&T and T-Mobile USA inside the country. The SIM contains the phone number, along with other authentication and security information that allows the phone to connect to the network. It also acts as a secondary storage bin for contacts, text messages, call history, and other information that the user might want to take with him if he switches to another phone; it's up to the user whether the phone sends that information to the built-in memory or to the SIM.

The SIM can be a great place to look for evidence, because deleting a text message or contact from the SIM doesn't necessarily mean that the corresponding data is lost. Instead, it's simply flagged as deleted, making it no longer accessible to the phone; it doesn't really disappear until the number of stored messages exceeds the SIM's capacity, which is typically 20 or 40 messages. Then only the oldest message is pushed out when a new message comes in.

CRIME SOLVERS' TOOL KITS

A sampling of products on the market

Free!

BitPim, built by an open-source community, pulls basic data from many CDMA phones manufactured by LG, Samsung, and Sanyo. *http://bitpim.org*

US $800 to $1500

Oxygen Forensic Suite analyzes over 1500 basic cellphones, smartphones, and PDAs, in both logical and physical formats. *http://oxygensoftware.com*

Paraben Device Seizure covers over 2000 cellphones, smartphones, PDAs, GPS devices, and SIM cards. *http://paraben.com*

Extracting these hidden messages from the SIM sometimes requires even more software tools. Ideally, a forensic lab would have enough different tools on hand to cover all the cellphone carriers and models sold in its region. But the typical forensics lab can afford only a small proportion of these tools; they're just too expensive, with prices often in the tens of thousands of dollars. To make the situation even worse, these tools can handle only certain specific sets of data. For any one type of phone, the lab must purchase one piece of software to pull the contacts, call history, and text messages and a second software tool to pull the images, videos, and ringtones. And as a final blow to the lab budget, the tools must be updated frequently to handle new phone models, new versions of phone operating systems, and other technologies.

Let's say McKenna's lab currently has 12 different software and hardware tools, each using a variety of communication protocols specific to various device manufacturers. These protocols, such as BREW, OBEX, FBUS, and SyncML, to name a few, allow the forensic hardware and software tools to communicate with the various phones and access their data on both the SIM card and the main memory.

The next step is to move the phone data to a computer that's set up to preserve the evidence in a legally acceptable way. Cellphone data is obviously different from blood spatters in the back of a car, but for either to be used as evidence in a court of law, it must be gathered according to standard methods acceptable to that court.

Now that McKenna has selected software tools and connected the cable, she waits—and hopes—for the tool to connect to the phone. The tool may try several different communication protocols to do this. Finally, after several attempts, the software starts transferring data from the phone to the computer. Depending on the memory capacity and the processor speed of the phone, this could take anywhere from 5 minutes to 3 hours.

Unfortunately for the beleaguered McKenna, this is a smartphone with a memory packed with more than 750 contacts, at least 1000 text messages, and more than 2000 e-mails, along with scores of images and videos. She's in for a long night.

Getting data from a mobile phone is not like backing up a hard drive—it's more like surfing the Internet. The software pulls the data from the phone in context, starting with the contents of the various phone books, its own number, and contacts, as well as missed, dialed, and received calls. Then it moves on to the message stores, including the draft, in-box, and out-box files, before wrapping up with the music, images and videos, and so on.

And that's the reason matching the software to the phone is so critical. For example, one common set of software commands for the Motorola, Nokia, Samsung, and Sony Ericsson phones all start with the characters "AT," which call the device to "attention." The convention is a holdover from commands that controlled Hayes Communication brand modems as far back as the 1970s. Other phone manufacturers rely on proprietary communication protocols. And smartphone manufacturers put their own twists on phone communications. Apple, for example, uses the iTunes software, and RIM uses its own BlackBerry desktop manager.

Ideally, some forensic tools create a duplicate image of the data on the hard drive of the examiner's computer, copying the data bit by bit. Getting this kind of data from phones requires a special cable that connects to the phone's circuitry through tiny contacts that are typically hidden under the battery. In the worst cases, though, no software is available for the specific phone in front of the examiner.

For Fasano, there is hope. An Android phone stores information in much the same way as a hard drive—as a set of files within nested folders. Several of the tools available can pull these files onto the computer, either bit by bit, creating an exact copy of the data on the computer's hard disk, or as a series of files and folders, preserving the original nested structure. Were this an ordinary cellphone, the software would instead sort the data into simple sets of related information, like contacts, text messages, images, and so on.

After about 2 hours, the forensic software finishes transferring the files to the computer and analyzing them, categorizing the different types of information. The software then generates a 192-page report detailing the recovered evidence. Finally, McKenna goes through the screens by hand, spot-checking the information to make sure that the software did a complete and accurate job; she may later have to testify to this fact in court.

After saving the report and respective e-mail, image, and video files to a DVD, McKenna calls Fasano. She has good news for him. She explains that in those 192 pages there are close to 2000 text messages, which may shed light on at least 100 different conversations the victim had over the last year and a half with people among the 750 names found in her phone's contacts list. Besides all that, Fasano will also get access to other important pieces of evidence from the phone, such as the call history, received and sent e-mails, Internet browsing history, and the browser cache.

Meanwhile, across town, several local, state, and federal agencies are conducting a drug bust. Jay Webb, the lead narcotics agent, is using a one-of-a-kind tool kit developed out of the Cyber Forensics Lab at Purdue University. (Full disclosure: Besides consulting for law enforcement agencies, I am a Purdue professor associated with that lab.) During the bust, Webb and his men capture 13 different phones from a group of men who were processing crack cocaine until…well, very shortly before the raid began. Right at the scene, moments after the bust, Webb connects each of those phones to its corresponding cable and to his laptop computer to quickly get the most important evidence: contacts, call history, text messages, specific images, and videos. The software requires no forensic training or even knowledge of the model of the phone he's dealing with. It can detect enough information about the phone to decide for itself what protocols to use. Within minutes Webb is tracking down other buyers and sellers based on information retrieved from the first of the 13 drug dealers' phones.

It's a basic scene that's playing out with increasing frequency all over the world, but not always with the happy result that Webb achieves. Forensic labs everywhere are being slammed with an influx of mobile phones found at crime scenes—hundreds of thousands annually in the United States alone. So there's a real need for investigators in the field to be able to triage the key evidence before having to take it to a lab for a full forensic exam. The problem is so acute that some ambitious detectives are taking training courses and purchasing forensic tools in an attempt to get a jump on their own labs. Sadly, the time and energy expended in this effort are often wasted. The initial software costs of up to US $25 000 and ongoing costs of up to $5000 annually (for licenses and updates) are high, and many times the tools sold to these detectives are too advanced for their on-the-spot intelligence needs.

Often a first responder just needs quick information from these devices, such as the call history, the contacts, or the text messages. This information is immensely useful in helping the on-scene investigator determine the next steps in his investigation. The contact lists represent a useful group of suspects, allies, and family members; the call history tells who was most recently in communication with the victim or suspect; and the text messages reveal those quick and easy private conversations that often give specific details of what was going on perhaps moments before a murder or a raid. Images and videos further fill out the portrait that the information on the phone has started to sketch about the owner (or former owner).

These new on-the-scene forensic tools are slowly making their way into the hands of law enforcement officials. However, with each new phone come new services, new features and, sometimes, new operating systems. The manufacturers struggle to keep their tools up to date. So most detectives, like Fasano, still have to wait for lab results.

Fasano's story ends badly. When he returns to the lab several hours later, he picks up the report. As he sorts through its 192 pages, he realizes that many of the text messages came from one contact distraught over the end of a relationship; one message reads, "You'll be sorry." Fasano grabs for his cellphone to call the county prosecutor to get an arrest warrant, but by the time the detective arrives at the suspect's apartment, the place is empty, showing signs of a hurried departure. After several fruitless months, he drops the investigation.

Fast-forward five years. Fasano's department, having recognized the importance of cellphone evidence, now has the latest tools, and a similar story of crime will end this time in punishment.

In a narrow alley on the west side of Chicago, a dead woman is slumped inside her Mercedes, clutching her mobile phone to her head. Rain pours through the broken car window, blown out by the bullet that went through the woman's temple. Chicago PD detective Nick Fasano arrives on the scene, snakes a cable through the window, and connects the phone to a handheld computer. Fasano scrolls through the torrent of information to zero in on a series of agitated text messages that grow increasingly threatening. He notes the number of the sender and forwards it to police headquarters, where detectives finger the owner of the number and send a squad car to his apartment to pick him up, just as he arrives, spattered in blood. The suspect confesses, and Fasano marks the case closed. ❏

TO PROBE FURTHER The Small Scale Digital Device Forensics Journal *(http://www.ssddfj.org/current.asp) publishes a wealth of articles about the theory, research, and practice of mobile-device forensics. The Electronic Evidence Information Center (http://www.e-evidence.info/blogs.html) collects links to discussions about digital-device forensics.*

$2500 to $3500

SecureView2 covers over 2200 devices with analysis tools that create a graphical interface to the data. *http://mobileforensics.com*

$7000

CelleBrite UFED (see photos, "Forensics to Go") covers more than 2500 devices with a stand-alone system that doesn't require a separate computer. *http://www.cellebrite.com*

$7500 to $10 000

Micro-Systemation's XRY analysis tool may recover deleted data, with tamper-proof reports. *http://msab.org*

$25 000

Logicube CellDek is an all-in-one mobile forensics kit that points users to the right cables. *http://www.logicube.com*

The World's Best GALLIUM NITRIDE

A little Polish company you've never heard of is beating the tech titans in a key technology of the 21st century

BY RICHARD STEVENSON

GALLIUM DUST
Ammono's first gallium nitride crystals were tiny, and metallic impurities gave them a brownish tint.

ROBERT LASKA (2)

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

WANT TO REVOLUTIONIZE THE ELECTRONICS INDUSTRY, become a multimillionaire, and earn your place as an immortal in the tech pantheon? Your job is simple: Figure out a cost-effective way to make really good, reasonably large crystals of pure gallium nitride.

With such crystals as the foundation for the growth of devices made of the same material, manufacturers would have a far richer yield of the violet lasers on which the optoelectronics industry increasingly depends. For example, the short wavelengths of these lasers are needed to read the hyperfine, data-rich line that rings the discs in Blu-ray players and in the latest game machines. Better gallium nitride would also let automakers make the power-handling circuitry in their hybrid electric vehicles more efficient, improving mileage and possibly even affordability. And with a fabulously good crystal foundation, LEDs could perform better, speeding the demise of the century-old incandescent bulb.

So far, though, gallium nitride crystals of good size and archangelic purity have been beyond the grasp of all but one of the companies that have worked for years to create them. That company's based not in Japan, Korea, or even the United States, but in Poland. Meet Ammono, the greatest success story in materials science you've never heard of.

GALLIUM JEWEL
After nearly two decades of refinement, Ammono's growth technique now yields wondrously fine hexagonal crystals up to 2 inches across.

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

THE COMPANY got where it is today by bucking the common wisdom in the industry. Instead of growing crystals with vapor deposition, the approach that all the leading gallium nitride substrate manufacturers take, it grows them the way the Earth does: under high heat and pressure.

Ammono, in Warsaw, is building up its stock of superhigh-quality gallium nitride crystals measuring 2 inches (51 millimeters) at their longest dimension. In a year or so it expects to have enough to start slicing some of them, salami-style, to produce wafers that can be turned into the round substrates on which semiconductors are grown. Admittedly, 51 mm is puny—a sixth the diameter of standard silicon wafers. But this size dominates today's market for gallium nitride substrates, which are used as the foundation for making violet lasers. Analysts' estimates of the market vary wildly, but everybody agrees it's upward of US $100 million and that its double-digit rate of growth won't end anytime soon.

It gets better: If Ammono increases the crystal size to 100 mm or more, major players in the silicon industry should start knocking on the company's door, hoping to exploit other advantages of gallium nitride besides the color of the light it emits. For instance, it conducts heat far better than silicon does. By making large substrates from gallium nitride rather than from silicon, you can provide a better foundation for the diodes and transistors that convert battery power into a form that a hybrid electric car can use. That's because the high currents heat up the chip, and if it's made of silicon, it'll need a water-cooling system of its own. Chips made of gallium nitride can simply share the cooling system of the internal-combustion engine, cutting costs and increasing the energy efficiency of the car. This potential market will grow along with sales of hybrid electric vehicles; according to analyst Philippe Roussel from the French firm Yole Développement, chip production will gobble up 800 000 of these 100-mm substrates in 2015.

THAT KIND of success would fulfill a long-held ambition of Ammono's president, Robert Dwiliński, who pioneered the development of the firm's novel gallium nitride crystal growth process nearly two decades ago. The inspiration for its efforts can be traced back to a weekly seminar that Dwiliński attended as a physics undergraduate at the University of Warsaw. He vividly remembers a talk given by Izabella Grzegory, a researcher at the High Pressure Research Center of the Polish Academy of Sciences, also in Warsaw. She lauded the strengths of a family of nitrides that promised to enable production of LEDs with outputs ranging from the ultraviolet to the infrared. And she claimed that their performance would be stunning.

Unlocking the true potential of these devices required technological progress, including the development of a good substrate. Gallium nitride devices today are often built on substrates of sapphire, silicon carbide, or even plain silicon. But in each of those materials, the atoms in the crystalline lattice are spaced differently from those of gallium nitride, introducing a strain of the sort you'd get if you tried to stack goose eggs on top of a layer of chicken eggs. What you really want is a substrate sliced out of a large, pristine crystal of gallium nitride itself. In fact, such a foundation always gives the best results, and it is a prerequisite for laser manufacture.

Dwiliński was intrigued by the challenge of making the first gallium nitride substrates, a quest that immediately became the central theme of his life. He resolved to pursue a Ph.D. in the subject, and to do that he needed a thesis advisor. Finding one wasn't easy, because Polish scientists tend to eschew commercially oriented research. Fortunately, one professor at the University of Warsaw did not: Maria Kamińska was willing to mentor him, and together they pursued several methods for growing gallium nitride crystals.

GEMLIKE PERFECTION: Robert Dwiliński, president of Ammono [with crystal], and cofounders Leszek Sierzputowski [left], Roman Doradziński [right], and Jerzy Garczński [far right]. PHOTOS: ROBERT LASKA (2)

Most semiconductor substrates today are manufactured by a process invented in 1916 by the Polish scientist Jan Czochralski. You begin with a tiny, high-quality crystal to seed the growth process, then you rotate that seed crystal inside a melt of the same material. As you slowly pull out the seed, the molten material cools and solidifies around it, creating a cylinder that is tapered at the ends. That cylinder is called a boule, and it is what technicians slice to make wafers. The trouble is that gallium nitride won't succumb to the Czochralski process below a temperature of 2225 °C and a pressure of 64 000 atmospheres (6.49 gigapascals), comparable to conditions very deep within the Earth's mantle. "It is almost impossible to build such a system," says Dwiliński.

So he looked at other methods, such as combining gallium-based solutions and high nitrogen pressures in small vessels, an approach that can form gallium nitride at a more manageable 1500 °C. Such efforts were already under way in the group where Grzegory worked, and while working on his Ph.D., Dwiliński teamed up with her and her coworkers. But he soon realized that although the quality of the gallium nitride crystals was outstanding, their dimensions were never going to be big enough for the mass production of commercial devices. Even today, this superhigh-pressure technique can produce crystals no larger than 20 mm, and the logic of high-pressure manufacturing makes it demonically hard to scale up the process. The thickness of the walls of the pressure chamber must increase by the cube of the increase in the size of the inner chamber; on top of that, the assembly must stay strong as temperatures rise to 1500 °C or more.

So Dwiliński began investigating alternative growth processes. He sought help from two friends he'd known in school and in the Boy Scouts who were working on their own Ph.D.s. Leszek Sierzputowski was an expert in chemical processes, and Roman Doradziński was a theorist skilled in thermodynamic calculations. Sierzputowski worked at Warsaw University of Technology, and he roped in a colleague, Jerzy Garczński, who had expertise in the design of high-pressure growth systems.

All four of them were attracted by the efficiency, low temperatures, and reasonable pressures associated with the commercial process used to make quartz crystals, which in a single run can turn out 1200 kilograms of product, most of which is made into crystal oscillators of the sort used in inexpensive wristwatches. Could this tried-and-true process be adapted for gallium nitride? The group resolved to find out.

AMMONO'S AUTOCLAVES: In a single run, they now produce over 70 2-inch crystals of gallium nitride.

MORE THAN a century ago, industrialists began making quartz crystals by simply mimicking the way nature does it. Today's quartz manufacture is based on the process they developed. It begins with an autoclave—a vessel capable of withstanding high pressures and temperatures—filled to the top with hundreds of quartz seed crystals. Underneath those seed crystals is the feedstock, silicon dioxide gravel extracted from Brazilian mines. To begin making a single large crystal, you first put some water into the chamber. Then you shut the autoclave and turn on the heat, boiling the water and driving up the pressure within the container. At about 400 °C, the pressure hits several hundred atmospheres, and the water is now neither a gas nor a liquid but a supercritical fluid. Think of it as a dense gas that moves freely and very quickly.

This supercritical water dissolves the silicon dioxide and, driven by convection, transports the dissolved material to the upper, cooler part of the autoclave, where the quartz seeds sit. Here the supercritical solution cools, reducing the solubility of the dissolved silicon dioxide and forcing it out of solution. This precipitate adheres to each of the quartz crystals, acquiring their crystalline form and thus increasing their size. To speed up what would be a glacial pace of growth, chemists add either alkaline metal hydroxides, particularly sodium hydroxide, or salts, such as sodium carbonate. These materials introduce negative ions that lead to a series of reactions that ultimately increase the solubility of the feedstock, which in turn makes it easier for dissolved material to move out of solution and onto the crystals, speeding their growth.

To adapt the process for gallium nitride, Dwiliński and his colleagues made three alterations: First, they replaced the silicon dioxide feedstock with gallium nitride, relying on alkali metal amide mineralizers; second, they swapped water for ammonia; third, they upped the temperature to 550 °C and the pressure to 5000 atmospheres. The one thing that they couldn't do, though, was prime the process with a handful of gallium nitride seeds, which do not occur in nature. They hoped that temperature variations within the autoclave would spur the formation of tiny crystals, which could serve as seeds later on.

It didn't work, partly because oxygen gas leaked into the autoclave and contaminated the growth. Then a tip from Herbert Jacobs, a professor at the Technical University of Dortmund, Germany, set them on the right track. Jacobs, too, was growing rare-earth nitrides using ammonia and heat, and he helped the Polish workers devise an oxygen-proof autoclave.

With their new, leak-proof autoclave in hand, the researchers performed their first trial run in 1993. It didn't work at first, but in a few months they were regularly opening their autoclave to find a grayish liquid inside. Filtering this revealed a pale yellow powder containing tiny gallium nitride crystals. Initially, these crystals were only a few micrometers in size—far too small to see individually with the naked eye—but that didn't stop Dwiliński from rejoicing. He knew he was onto something, because when these crystals were zapped with a laser, they lit up, a phenomenon known as photoluminescence, which couldn't have occurred if they'd had many impurities. Probing with an X-ray source showed that they also had the correct hexagonal crystal structure. In other words, the crystals were first-rate.

The team had no money, and there were no domestic venture capitalists to appeal to. So they decided to plead their case with the leading gallium nitride developers in the United States, Europe, and Asia. The bold move paid off: In the fall of 1999, Shuji Nakamura, then employed by Nichia Corp., the Japanese firm he was boldly pushing into nitride optoelectronics, flew in to see them and was impressed with what he saw. On the advice of Nakamura—the inventor of the gallium nitride laser—Nichia would fund a joint research project to develop ammonothermal gallium nitride growth, and in return it would take a stake in Ammono's intellectual property, as well as access to the crystals that were made.

So Ammono was under way, with the four researchers cofounding the company and renting a room in Warsaw from the Industrial Chemistry Research Institute. However, just a few months later, they heard the news that shocked the entire nitride community: Nakamura was leaving Nichia for the University of California, Santa Barbara.

The Nichia-Ammono relationship survived Nakamura's departure. Before he left, Ammono had already found a great ally in Nichia's Yasuo Kanbara, an open-minded and enthusiastic researcher who had already been mulling over the idea of ammonothermal growth of gallium nitride. He started visiting Ammono for a week every few months, a schedule he kept for several years. He also helped by sending Ammono's materials back to Japan, where researchers characterized and tested it.

Ammono's first challenge was to improve the design of their autoclaves, which tended to produce a lot of tiny gallium nitride crystals but also allowed them to dissolve slowly back again into the ammonia solution from whence they had come. To avoid a "two steps forward, one step back" process, in 2000 Ammono developed a new autoclave that more precisely controlled the temperature profile in the internal chamber. By optimizing the difference in temperature between the various parts of the autoclave, the researchers could make the most out of a peculiarity of the interaction of gallium, ammonia, and the alkali metals. Generally, a warm solvent can dissolve more material than a cool one, which is why hot tea dissolves sugar more easily than iced tea. But if you put gallium in ammonia, with the addition of small amounts of alkali metals, you get what's called retrograde solubility instead: The solubility decreases with increasing temperature.

How the Autoclave Works

At 500 °C or more, ammonia boils, forming a supercritical solution, attacking the gallium nitride feedstock, dissolving it, and transporting it to the gallium nitride seeds. There the solubility is lower, so gallium nitride leaves the solution and grows on the seeds, enlarging them.

ILLUSTRATIONS: JOHN RICE

THIS PECULIARITY has a major impact on autoclave design. Here's how: The high-temperature, seed-containing zone must be placed physically below rather than above the low-temperature zone containing the feedstock. That way, the ammonia solvent attacks the gallium nitride feedstock at the top of the autoclave, first dissolving it and then transporting it to the higher-temperature seed crystals below. There, where the temperature is higher, the solubility is lower, and gallium nitride precipitates from solution and grows on the seed.

What this means is that the seed can be hotter than the feedstock. At this high temperature the atoms are more mobile and therefore more likely to take up the correct position in the lattice, rather than getting stuck in the wrong place and marring the crystal. Another advantage of this arrangement is that unwanted impurities in the crystal are kept to a minimum, because they are actually *less* soluble at the lower temperatures prevailing in the upper part of the reactor. There they are forced out of solution and onto the surface of the autoclave.

Dwiliński's next step was to scale up the crystals. He started by taking crystals produced in one growth run and putting them back as seeds in the next. But when the crystals reached a certain size, the autoclave constrained them from getting any bigger. So for the past 10 years Ammono has built a succession of ever-larger autoclaves.

At some point between 2000 and 2003—the company refuses to say exactly when—it was able to make its first 1-inch-diameter gallium nitride substrate. Many other companies would have crowed about it in public, but Ammono kept quiet, sharing the philosophy of its Japanese backer. "People know that if Nichia announces something, it's not just an idea," says Dwiliński. "They achieve something, develop it to production, have ready product, and then advertise this achievement together with something to sell."

Ammono kept striving to make bigger crystals. At some point, the autoclaves became so large that the company needed larger rooms to accommodate them. (Autoclaves are long and thin; the largest ones are about 3 meters high, and you may need another 3 meters of free space above that to take them apart for reloading.) In 2003, Nichia agreed to finance a move to more spacious quarters in return for shipments of substrates. Ammono found a plot 25 km north of Warsaw and built a two-story complex of offices, manufacturing facilities, tools to polish substrates and characterize them, and even a couple of guest rooms.

Today, with 50 people on its payroll, Ammono is still a tiny player in a minuscule niche of the semiconductor industry. But it makes what are widely regarded as the world's best gallium nitride substrates, at sizes of 25 mm and 38 mm. They're not big enough to use in commercial production, but companies still want them to investigate what's possible with the best gallium nitride available today.

Ammono is now trying to accumulate enough 2-inch seeds to begin manufacturing 2-inch substrates, the minimum size for laser manufacturing lines. Only substrates of that size and of surpassing quality—known as laser grade—will do. The company must also convince its customers that its material is significantly better and cheaper than what they are using today. The target is clear: A single 2-inch laser-grade crystal now sells for $5000.

Dwiliński says that Ammono's 2-inch material will be priced competitively, but he insists that the cost will fall as production ramps up, economies of scale kick in, and the manufacturing technology matures. He figures that the industry wants to pay $1000 for a 2-inch piece of gallium nitride, and here's the good part: He reckons that the target is within the reach of his ammonothermal growth method but impossible with the industry's standard method. Dwiliński says there's no reason why his costs shouldn't fall, in the long run, to those common for gallium arsenide: $200 for a 4-inch substrate, according to the global consulting firm Strategy Analytics.

THE INCUMBENT technology for manufacturing gallium nitride substrates is called hydride vapor phase epitaxy (HVPE). First, you heat a substrate, typically gallium arsenide or sapphire, to around 1100 °C, after which you waft a mixture of gaseous compounds containing nitrogen and gallium onto its surface. Here they decompose to release gallium and nitrogen atoms, which form a gallium nitride film that can be peeled off and sliced into substrates.

One of the strengths of HVPE is its fast growth rate, but with speed comes a higher rate of defects. "There is talk of crystalline growth at 800 micrometers per hour," Dwiliński says, "but people involved say that you need to work in the range of 50 to 100 μm per hour to avoid highly textured surfaces." (That is, you must avoid any unwanted deviation from the ideal, which is a surface that's flat on the *atomic* scale.) This rate is still about 10 times as fast as what you get in ammonothermal growth, but the total amount of product you can get through the factory isn't necessarily higher, because the autoclaves can be filled with hundreds of seeds at a time. In comparison, the formation of a gallium nitride crystal by HVPE tends to involve growth on one substrate at a time. What's more, ammonothermal growth can outperform HVPE on other fronts: Autoclave maintenance is simple because it doesn't need cleaning, and all the gallium is converted into product, versus less than 15 percent for HVPE.

Ammonothermal also trumps HVPE in material quality. Differences between the atomic spacing of the gallium nitride crystal and the substrate on which it's grown cause HVPE-grown gallium nitride to bend. Ammono's process, however, can achieve a substrate that's 100 times as flat as HVPE-grown material and has two or three orders of magnitude fewer defects, a mere 5000 per square centimeter. These defects are a major pain for laser manufacturers because they quench luminescence and shorten device lifetime.

Recent advances with gallium nitride lasers also strengthen Ammono's position. Over the last few years, researchers all over the world have been trying to extend nitride laser emission to the green part of the spectrum [see "Lasers Get the Green Light," *IEEE Spectrum*, March 2010]. One of the most promising ways to do so is by growing lasers on different planes of the gallium nitride crystal, known as the semipolar and nonpolar planes, which Ammono's superflat crystals can yield with particular ease. That's because the crystals' greater flatness allows them to grow, defect-free, to greater thickness—and that thickness is needed to exploit the nonpolar and semipolar planes.

Ammono's long-term commercial success hinges not only on offering the market a better product for today but also on fending off threats from nascent technologies. One of these, known as the sodium-flux method, has many admirers, including Dwiliński himself, but he still doesn't think it can match ammonothermal's growth. This seed-based rival process involves molten gallium, which is mixed with sodium to increase the amount of nitrogen that can be dissolved in this solution at typical pressures of 70 to 80 atmospheres. Two-inch crystals have been produced by this technique, and they do have a lower defect density than HVPE-grown seeds. However, variations in the ratio of gallium to nitrogen in the solution make it hard to grow many crystals in a given run.

Ammono has finally stepped out of stealth mode and started to court publicity. Technical conference presentations began in 2007, when the company caused a stir by appearing to come out of nowhere and claiming to have produced large, very-high-quality crystals. This came as quite a shock to delegates, who were used to hearing about a new idea at one conference, followed by reports of incremental progress over the next few years. But most of the nitride community now accepts that the company's material is very good.

ONE QUESTION hanging over Ammono is whether its superior substrates lead to superior devices. Nichia clearly knows the answer, but the notoriously secretive company isn't giving anything away. However, Dwiliński maintains that the answer is a resounding yes and that the substantial device benefits will soon become evident. Ammono is collaborating with some research institutions that fabricate devices, and the results generated by this effort should be published in the academic press later this year.

The company plans to begin shipping 2-inch substrates in the second half of next year. It also expects to produce its first 1-inch semipolar and nonpolar material, which will provide a great foundation for making blue and green lasers. Ammono expects to introduce 3-inch gallium nitride substrates in 2013 and 4-inch substrates in 2015.

Let's hope that this fascinating little Polish company finds a backer or two to speed the introduction of big substrates. Of course, it would be good news for the makers of game players, laser TVs, and hybrid electric cars. More important, it would be a triumph for the little guy who beats the world by thinking in new ways. ❑

FROM BRICKS TO BITS

BUILDING THE LEGO UNIVERSE ONLINE

By DAVID KUSHNER

PHOTO: HOLLY LINDEM; LEGO SCULPTURE: NATHAN SAWAYA

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

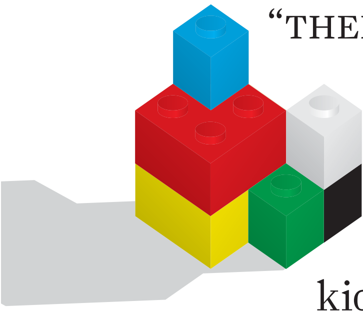

"THERE'S A DARKLING, DUDE! Come on!" It's after school in Louisville, Colo., and two boys are playing just like generations of kids before them—with Lego bricks. As their little yellow Lego guys battle some evil monsters called the Darklings, the boy in the blue shirt yelps to his friend, "You gotta start building a bridge now! Hurry!" A bridge takes shape as interlocking Lego bricks snap together.

But this is no ordinary playdate. The boys are sitting in a windowless room behind a two-way mirror. Cameras monitor their every move, and a scruffy technician taps notes into a laptop. The biggest difference, though, is the Lego blocks themselves. They aren't made of plastic. They're made of pixels.

The boys are play-testing *Lego Universe*, an online computer game due out in October. Ten years in the making, the game is being created here at a company called NetDevil, with supervision from the Lego Group. It's a children's product, but it's also serious business. *Lego Universe* marks the legendary company's first foray into massively multiplayer gaming, and for the iconic building-block maker, it's a major gamble. With the game, the creators are aiming to pull off three incredibly difficult feats: translating the creative, distinctly tactile Lego experience into a virtual arena; creating an online environment that's both kid-friendly and kid-safe; and opening up a new market for the US $2 billion toy company.

And because this is an international project and brand, addressing these issues is a global challenge. On launch, subscriptions will be available to over 20 countries from Austria to the United States. Success for this project—estimated to cost over $10 million to create, though the Danish company does not make the budget publicly available—is by no means assured. The company has had occasional troubles before, such as trying to create its Lego theme parks.

Most massively multiplayer games cater to an older crowd, where norms of behavior need not be strictly enforced. But *Lego Universe* seeks to capture the grade school set and the estimated $20 billion children's game market. It's one thing to unleash bawdy mayhem in a fantasy game like *EverQuest* and quite another to nurture a virtual play space where kids can be kids—and not fall prey to predators or bullies or other unwelcome Internet lurkers. So with *Lego Universe*, the company's reputation as one of the most venerated and parent-friendly toy brands is at stake.

In addition, Lego must stay faithful to its meticulous engineering. It can't afford to disappoint users looking for the familiar interplay of studded plastic pieces that families have been clicking together for generations. Everything built in the game must be buildable in real life.

ALMOST REAL: The game designers wanted the objects in *Lego Universe* to have the same texture and sheen as the Lego bricks that children have played with for generations. *IMAGES: LEGO*

But unlike traditional Lego play, the online version will offer unprecedented opportunities for players to share and interact. The sprawling Lego fantasyland will be able to support more than half a million "brick heads" from around the world. Each player will start by assembling a personal Lego miniature figure to serve as his or her avatar. Players can then venture into the live *Lego Universe*, where forces of chaos and destruction—monsters such as the Darklings—threaten to destroy the Land of Imagination.

As players explore brightly colored lands such as the Avant Gardens and the Gnarled Forest, they'll get to build rocket ships and skyscrapers and animals and whatever else they can imagine. With over 2000 types of pieces in 26 colors, the variations are seemingly endless. Players will also receive their own property, which they can populate with their constructions. Using the simple icon-based programming that Lego developed for its Mindstorms robotics program, players can even animate their creations—a rocket ship that flies or a Ming dynasty vase that rotates when visitors get close. And of course, online virtual creations can be made physical—for a price. Using a service called Design by Me, a player can have a kit of physical blocks shipped directly from the Lego factory and then rebuild the virtual Lego construction for real.

"It's definitely going to be a challenge," says Michael Gartenberg, a partner with Altimeter, a technology research firm based in New York City. "They want and need to protect their brand; they need to make sure this experience will be kid-friendly." But given the popularity of Lego, there's little room for error. "They're taking a fairly big risk and they have to get it right immediately," Gartenberg says. "It's not something they can check out and hopefully have it work and work out bugs. " In the end, though, it all comes down to winning over kids like the two boys behind the two-way mirror.

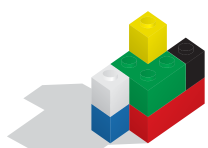

IN THE BEGINNING, there was the brick. And that universally recognized building block remains at the core of *Lego Universe*. Of course, an online brick has no substance; you can't hold it in your hand or feel the satisfying snap of connecting two blocks together. Beyond replicating that sound, Lego doesn't attempt the impossible. Instead, the game puts all its emphasis on creating an authentic simulation of the building blocks themselves. "We couldn't begin work on this game until we figured out how to make a virtual brick," says Ronny Scherer, studio director of *Lego Universe*.

The brick has been through transformations before. When Danish carpenter Ole Kirk Christiansen began making toys in 1932, they were all made of wood. He called his company Lego after the Danish phrase for "play well," *leg godt*. With the adoption of injection molding machines in the mid-1940s, Christiansen and his son Godfred realized kids could play better with cellulose acetate, and they created plastic bricks that could be connected. The company began churning out brightly colored "automatic binding bricks," later made from sturdier acrylonitrile butadiene styrene plastic, which is used to this day.

Lego has often enhanced the play experience with new technologies. Lego Mindstorms brought the bricks to life electromechanically and has hooked millions of kids on robotics. Lego video games, including *Lego Indiana Jones* and *Lego Rock Band*, are now mainstays on home gaming consoles. In 2004, Lego introduced its Digital Designer software, a free program that lets a player design Lego models on-screen.

Lego began exploring the idea of a massively multiplayer online game in 1999, but prospective developers repeatedly told the company that there was no way to digitally replicate the bricks—and all their studs and crannies—with the necessary

fidelity. In 2006, Lego finally awarded the project to NetDevil, an online-game company founded in 1997.

NetDevil was known among game aficionados for the nimble and realistic physics of games like *Jumpgate*, *Auto Assault*, and *Warmonger*, but it had never done a children's game before. NetDevil programmers were initially stumped when they heard that players would likely be using graphics cards that were several generations behind the times. "Kids tend to play on older computers," explains Ryan Seabury, NetDevil's creative director. "They get the hand-me-downs." How on earth would they be able to render an entire world of bricks for players using graphics cards built in 2003?

Digital objects aren't solid; they consist of numerous flat polygons connected into a 3-D shape. A simple-looking object isn't necessarily simple to render on screen. While an avatar in an online game such as *World of Warcraft* might consist of 2500 triangles, a detailed rendering of a one-centimeter-square Lego brick can have over 250 000, thanks to the perfectly round studs and other fine details.

"A digital brick has to look like a real brick," says NetDevil technical director, Eric Urdang, holding up a bright red rectangular block. "It has to have that certain sheen and imperfections."

NetDevil developers figured that they could have a maximum of 500 000 polygons on screen before the frame rate slowed and animations began to look jerky. When programmers initially built a pirate ship, they found that it alone contained 430 000 polygons. "The more polygons, the less you can do," says Urdang, "so we had to find smart ways to make this efficient."

NetDevil spent two years experimenting with ways to decrease the polygon counts. Techniques such as bump mapping—in which 3-D textures are simulated with 2-D images—proved too unwieldy for the underpowered graphics processors they were targeting. Instead, the solution came from examining the real-life bricks. When Lego blocks are snapped together, the interlocking studs disappear. So rather than draw every unseen portion of the brick, the computer could be instructed to draw only the visible parts. "There's lots of polygons you can throw away," says Urdang.

Though it sounds simple enough, the implementation was anything but. Urdang and his team of 40 engineers applied a technique called hidden surface removal, which involved going into the 3-D model of each Lego object and manually erasing the hidden portions. The team then spent another two years automating the process, until what had taken hours by hand could be done in seconds. That 430 000-polygon pirate ship got squeezed down to just 13 000 in a fraction of a second.

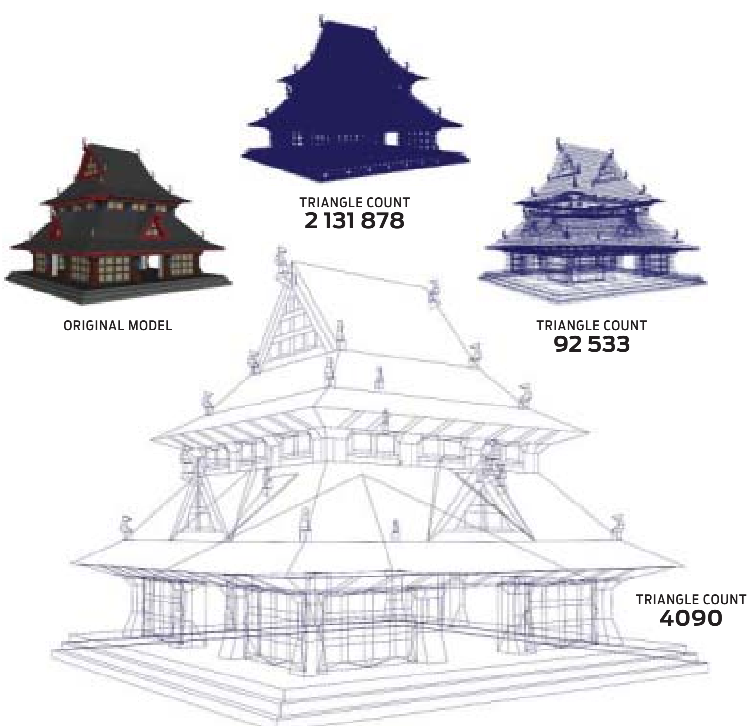

SHAPE-SHIFTING: When NetDevil artists built models from virtual Lego bricks, they discovered that it required too many polygons. They initially removed hidden triangles by hand but later automated the process. Ultimately, their goal was to generate a model that could be adequately displayed on computers with limited graphics hardware, yet maintain the visual detail and resolution of the original. IMAGES: NETDEVIL

ONE DAY LAST SUMMER at NetDevil, the *Lego Universe* team gathered around the table in the company's sprawling conference room for a project code-named Naughty Duck. Lego ducks are an inside joke at the company. Upon arrival at the Danish headquarters, new employees are given the fowl test, which involves building a duck out of just six Lego pieces. The NetDevil challenge had a twist: Using the same six duck pieces, team members were asked to create the most obscene object imaginable. "We got a lot of phalluses," says Seabury.

The experiment was no laughing matter. *Lego Universe*, like *World of Warcraft* or *EverQuest*, isn't just a product; it's a social space. While Lego's core demographic is 8- to 12-year-olds, some of the most ardent fans are adults. But as cautious parents everywhere know, allowing unknown adults to mix with young kids in an unsupervised, unfettered environment is simply not wise. That's as true online as it is in the real world. Most online worlds for kids such as Webkinz or Club Penguin address this issue by granting users very limited freedom to express themselves. But as I'm frequently told throughout my visit, Lego is all about creativity, and the company wants to foster that same spirit

NATURE LOVERS: The concept art for *Lego Universe* visualized an environment where Lego characters interact in natural-looking landscapes.
IMAGE: NETDEVIL

online. At the same time, says Mark Hansen, senior director of *Lego Universe*, "Children's safety is high priority. We don't want to put kids at risk."

So how do you balance creativity with safety? Lego is taking several approaches. First, it's charging for subscriptions to *Lego Universe*, which creates inherent accountability, because each account can be traced back to an individual. That alone isn't adequate to rein in the riffraff; if a player subscribes using a gift card or fills out the personal information incorrectly, the account tells Lego very little.

The second safety backstop is having online moderators. Though Lego hopes eventually to automate the moderation process, for now it will be conducted by a designated team of employees around the world. At any given time, roughly 100 *Lego Universe* moderators—roughly one for every 500 players—situated around the globe will be overseeing the game.

It turns out that safety is as much an engineering challenge as one of enforcement. "Positive social behavior can be engineered into the game," Seabury says. NetDevil and Lego thought long and hard about how to implement free-form building in the online game. They certainly didn't want to repeat the experience of Electronic Arts, which struggled to keep up with the provocatively shaped creatures that users created in its online game *Spore* (which earned the nickname Sporn).

Hence the Naughty Duck test. At various points in *Lego Universe*, players will encounter a similarly random pile of bricks and be challenged to arrange them into something useful—say, an elevator or a bridge. The test is called a showcase. By limiting each showcase to a small number of bricks, Lego hopes to limit abuses to the vaguely suggestive. Also, any new showcase creation is viewable at first only by its creator. If a player is proud of her handiwork and wants to show it off to others, she can submit it to Lego's moderators for approval. Things that Lego disapproves of extend well beyond phalluses; they also include religious symbols such as crosses and six-pointed stars.

The company is now looking at how to handle users who create letters that can be linked together to spell profane words in a variety of languages. It's possible that *Lego Universe* players will not be allowed to make shapes representing their ABCs, at least not in public. But on their own property, which requires special access, players will be freer to create (though even that content will still be subject to moderation). And when two players "friend" each other, the content they share becomes trusted and doesn't necessarily get the same level of moderator scrutiny.

It's not just Lego constructions that must be moderated. *Lego Universe* also has chat rooms for its players. Other virtual worlds for kids limit conversation to preselected chat phrases like "This is fun!" or "What's up?" Lego didn't want to be quite that restrictive. Instead, a user can compose messages from any of the roughly 2000 acceptable words on a "white list." Each language has its own list. Here, obscenity is not the chief concern; bullying and predatory behavior are. At the moment, for example, Lego will not allow numeric characters in messages, to try to prevent users from giving out addresses and phone numbers. Moderators will be checking in as well, which means that Lego needs employees who are proficient in numerous languages.

The company will also employ a behavior rating system for the game, using algorithms that will assign each player a "goodness score" (based on such achievements as building and sharing well-made objects) and a "badness score" (triggered by obscenity and bullying). "If a player does inappropriate activity, they'll be moderated more often," says Scherer.

If it all sounds impossible to control, Lego admits that on a certain level, it is. The easy answer would be simply to crack down—limit the chat, limit the blocks, limit players' ability to share. But that would go against the company's philosophy. "The platform is to inspire creativity," says Seabury, "and creativity doesn't come if you block everyone's ability to communicate."

BACK AT NETDEVIL, the two play-testing boys head out, and a girl and her mom arrive for their turn. In the lobby, the girl gazes in awe at the life-size Lego figurines dressed up in holiday outfits, standing on clouds of puffy white cotton amid Lego reindeer.

The mother and daughter stop to peek in on a guy assembling a giant *Star Wars* Millennium Falcon out of Lego bricks. He's one of several full-time employees assigned to test-build objects in the game to make sure they can be constructed in real life. Behind him are rows and rows of multicolored tubs of Lego pieces. With 10 million pieces in house, NetDevil has one of the top five Lego collections in the world. "Wow," says the girl, eyes widening.

It's exactly that sense of wonder that Lego hopes to foster and protect online—despite skeptics who question just how effective virtual Lego can be. "The computer game misses a very important component of play for a child: human interaction," says Mona Delahooke, a child psychology expert with the American Psychological Association. "Building Lego creations with a parent or friend involves social communication—commenting on each other's work, sharing nonverbally by looking and smiling, giving and taking feedback, and so on. My concern with computer Lego is that it takes away the human interface—which is play, children's natural language of learning."

But even as its creators iron out the final technical and security hurdles, *Lego Universe* is already set to grow. It will launch with six thematic worlds, with two to four new expansions each year. Ultimately, the goal is to break down the fourth wall between plastic and pixels, so that someone can buy a box of Lego bricks in a store and play with them online or off. "The design of *Lego Universe* is that all things Lego past, present, and future can fit into this world," says Scherer. ❑

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

Faculty Positions at the School of Engineering

Bar-Ilan University (Israel)

The School of Engineering at Bar-Ilan University, which will celebrate its 10th anniversary next year, is in the process of expanding its programs and its faculty body. It therefore has immediate openings for several tenure-track faculty positions at all levels. Applications will be accepted from highly qualified individuals with outstanding competence in areas related to electrical, computer, bio-, and industrial engineering. In particular, applicants with interest and demonstrated innovative research in the following core areas of our School of Engineering are encouraged to apply:

Electrical Engineering: Signal and Image Processing, Communications, Micro and Nano Electronics, Electro Optics and Photonics

Computer Engineering: Architecture, Embedded Systems, Computer Networking, Real-Time Systems, VLSI Design and Electronic Design Automation

Bio-Engineering: Cell and Tissue Engineering, Functional Biological Imaging, Bioelectrical and Biomedical Measurements, Interfaces between Cells and Microelectronic structures, MEMS and Nano-Engineering

Industrial Engineering: Information Systems, Production and Service Systems, Operations and Project Management, System Analysis, Electronic Commerce, Game theory, Quality Assurance and Reliability, Production Planning & Control, Supply Chain Management

Responsibilities include undergraduate and graduate education, developing and sustaining externally funded research programs, graduate student supervision and mentoring, and professional/institutional service. Besides having earned a PhD in a field commensurate with their field of expertise, candidates are expected to have been a post-doc at a reputable university, to have demonstrated excellence in research and the ability to develop an independent research program, and to have a strong interest in teaching at both the undergraduate and graduate levels. Adjunct senior faculty positions are also available.

Applications should include a letter of application with a professional resume, statement of research and teaching goals, and the names and contact information of at least five references. Applicants should send all material to Mrs. Dina Yeminy, Administrator, School of Engineering. **dina.yeminy@mail.biu.ac.il**

Tradition of Excellence

Bar-Ilan University

the data

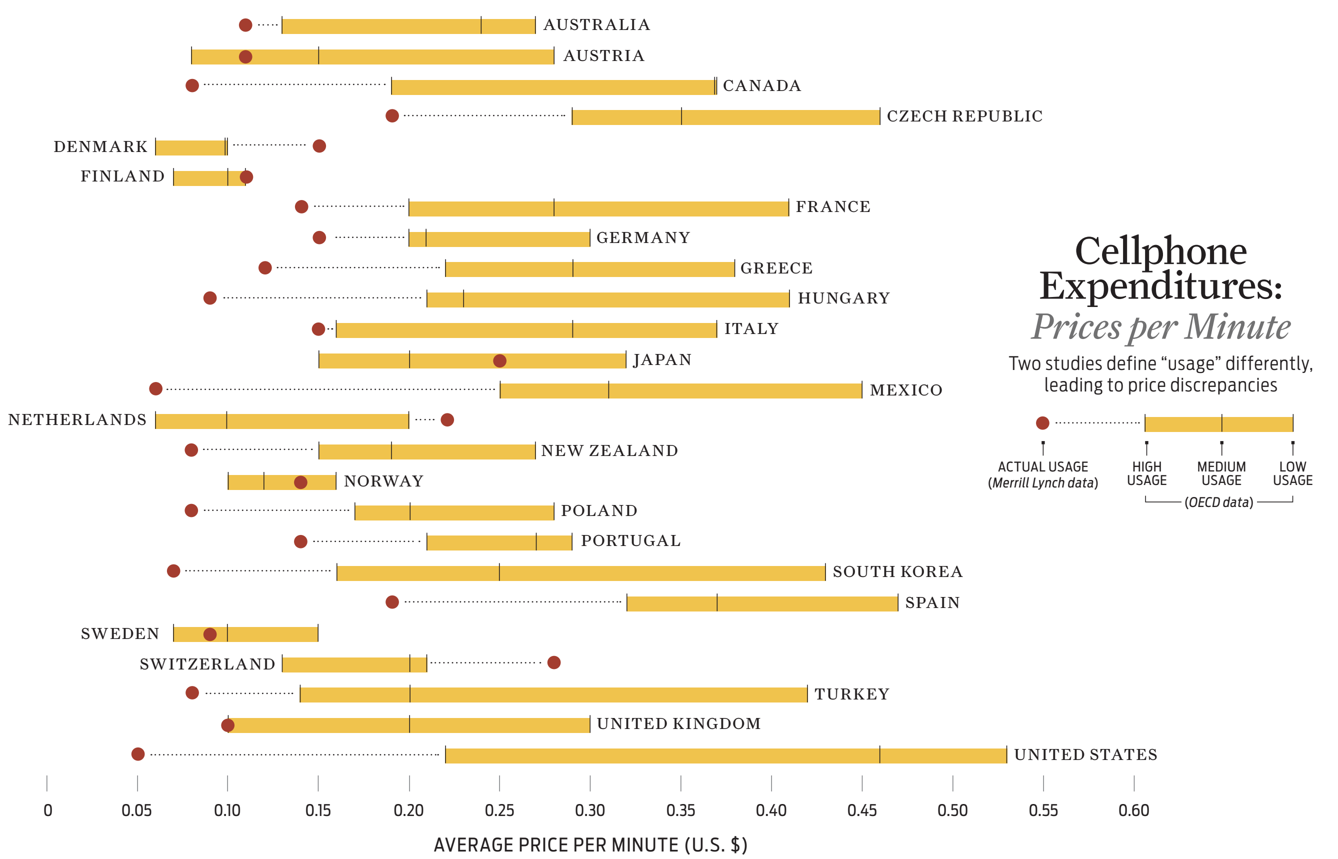

Who Pays More for Wireless? It Depends on Whom You Ask

According to the Organisation for Economic Co-operation and Development (OECD), the price of wireless phone service is lower in the Netherlands than in any other developed country, and it's highest of all in the United States. However, Merrill Lynch & Co. recently issued a ranking order that's nearly the reverse, placing the U.S. price second from the bottom (only Hong Kong's, not shown, was lower) and the Netherlands price near the top.

How can both of these studies be true? It is a classic example of comparing apples to oranges—in this case, multidimensional apples and oranges. By choosing the dimensions for comparison carefully, you can reach almost any conclusion you want.

The OECD computes the price for three fixed baskets of service. The "high use" basket includes 246 minutes of calls per month and 55 text messages. But by American standards, this is low use. The average American wireless user talks for 830 minutes a month and sends about 400 text messages. Why? For one thing, most cellular pricing in the United States is like much other U.S. pricing—you can get a lot more minutes for just a few dollars more, just as you can get a much bigger drink at McDonald's for a few cents more.

George Ford, chief economist at the Phoenix Center for Advanced Legal & Economic Public Policy Studies, in Washington, D.C., says that telecommunications companies are simply doing what economic theory says they should: responding to market conditions in their own countries. "American consumers are better off with the price structure in the United States, and Dutch consumers are better off with the pricing structure in the Netherlands," he says. Any study that condenses the complex wireless market into a single "average price" paid by an "average user" is bound to misrepresent the actual diversity of user behaviors and plans that are tailored to those behaviors.

—Dana Mackenzie

Sources: Organisation for Economic Co-operation and Development, "Communications Outlook 2009"; Bank of America/Merrill Lynch, "Global Wireless Matrix 2Q 09"; OECD press release, "Mobile phone calls lowest in Finland, Netherlands, and Sweden, says OECD report," 11 August 2009.

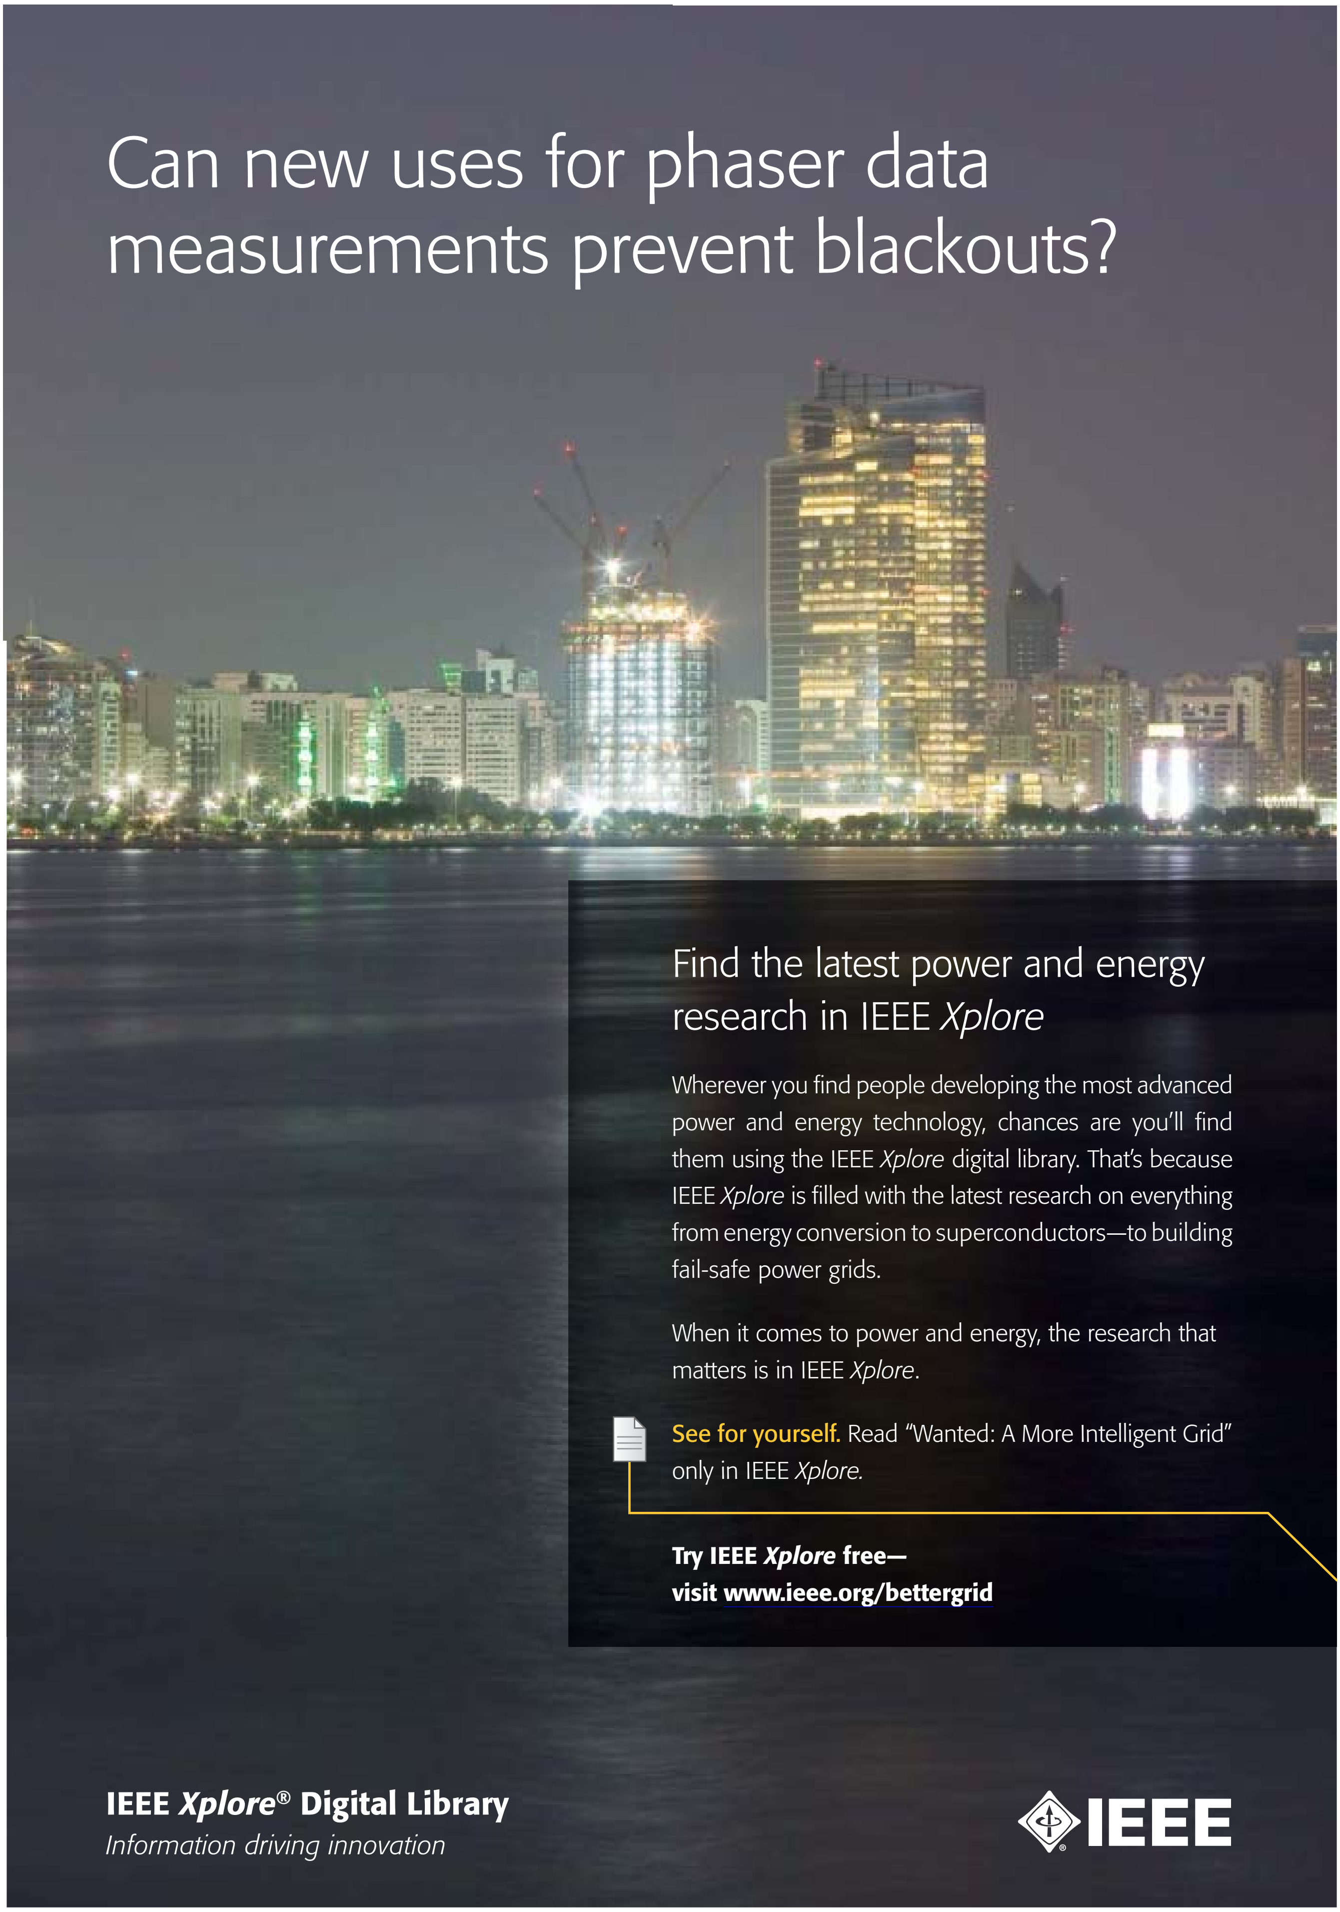
Can new uses for phaser data measurements prevent blackouts?
Find the latest power and energy research in IEEE *Xplore*
Wherever you find people developing the most advanced power and energy technology, chances are you'll find them using the IEEE *Xplore* digital library. That's because IEEE *Xplore* is filled with the latest research on everything from energy conversion to superconductors—to building fail-safe power grids.
When it comes to power and energy, the research that matters is in IEEE *Xplore*.
See for yourself. Read "Wanted: A More Intelligent Grid" only in IEEE *Xplore*.
Try IEEE *Xplore* free—
visit www.ieee.org/bettergrid
IEEE *Xplore*® Digital Library
Information driving innovation
IEEE


AA-1