ieee spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
08.09

THE LED's DARK SIDE

A mysterious malady called "droop" is dimming the prospects for light-emitting diodes as the ultimate illuminators

How DEAN KAMEN took his island off the grid

Good vs. bad photo fakery: It's all in the mathematics

Self-healing chips

U.S.A $3.99
CANADA $4.99
DISPLAY UNTIL
3 SEPTEMBER 2009
0 74470 64849 9 08>
www.spectrum.ieee.org

IEEE
Celebrating 125 Years
of Engineering the Future

volume 46 number 8 international

ieee spectrum 08.09

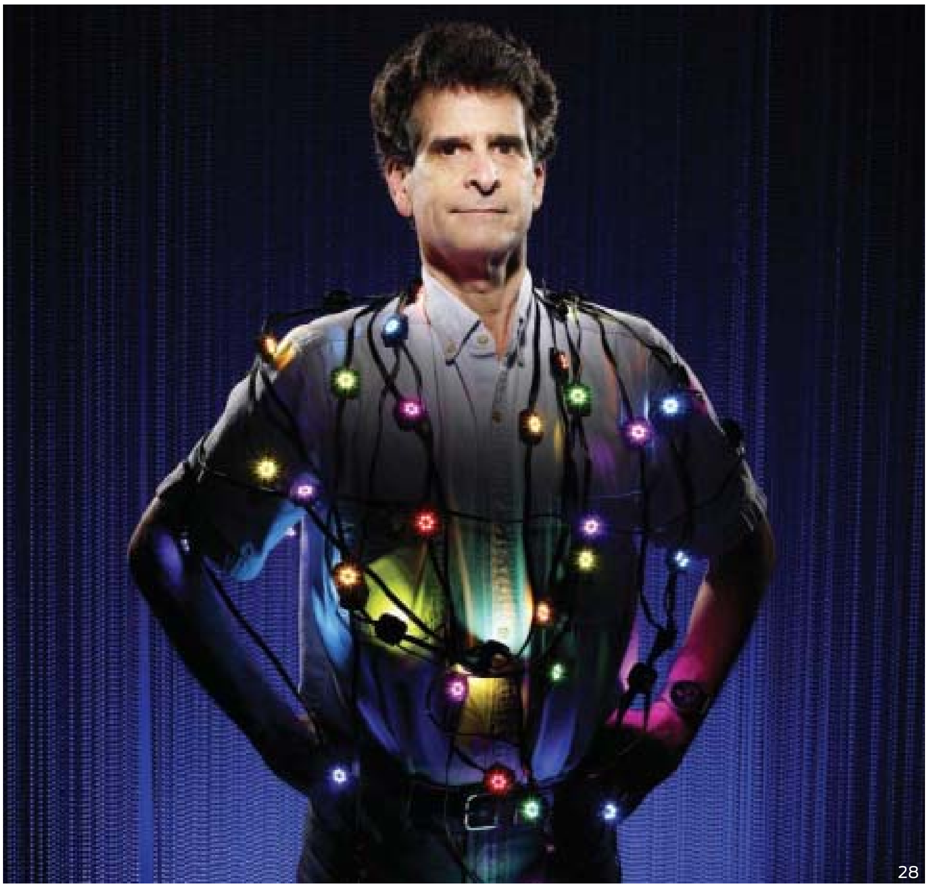

28

BYE-BYE, BULB: Inventor Dean Kamen is ablaze with solid-state lighting. His mission is to end the reign of the incandescent bulb.

COVER: BRYAN CHRISTIE DESIGN
PHOTO: DAVID YELLEN

volume 46 number 8 international

ieee spectrum 08.09

PHOTO: NOKIA

WWW.SPECTRUM.IEEE.ORG

AVAILABLE 1 AUGUST ON IEEE SPECTRUM ONLINE

TV TO GO

Digital television in the United States was designed long before the boom in portable devices; it was meant to be watched on a stationary receiver in the home, not on the go. Start moving that receiver and the picture quickly falls apart, freezes, or goes to a blue screen. Some cellphone providers offer limited TV to go, but for a price. Outside the United States, however, latecomers to digital TV were able to build mobility into their systems from the start, thanks to advances in receiver technology. Can the United States catch up? Following a technology development and standardization effort of unprecedented speed, mobile digital TV in the United States may now be ready for prime time.

ONLINE FEATURES:

CARBON FOOTPRINT CALCULATOR:
Want to know the amount of carbon emissions you personally control? Our calculator will give you an accurate answer, whether you live in Xi'an or São Paulo, Oxford or Silicon Valley.

LUNAR VACUUM CLEANER:
In an exclusive video, we visit a next-generation rover that uses compressed gas to collect soil from the moon.

ALSO ONLINE:

- Webcasts
- Podcasts
- News
- Videos
- Blogs
- Jobs
- Career Accelerator
- IEEE Xplore® digital library
- Interviews
- Opinions
- More!

TECH INSIDER WEBINARS

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

Calibre Solutions for Advanced Design Rule Checking

Creating Powerful Data-Capture Solutions

Episode II: Sense, Think, Act for Unmanned Robotic Systems

Multiphysics Simulation

WWW.IEEE.ORG/THEINSTITUTE

AVAILABLE 6 AUGUST ON THE INSTITUTE ONLINE

IEEE HONORS TECH LEADERS

Read about the recipients of this year's IEEE medals and awards at the Honors Ceremony. Their work has contributed to deep-space exploration, improvements in digital-storage technologies, and advances in radar detection.

GAIN A COMPETITIVE EDGE IN A CROWDED JOB MARKET

Learn effective job search techniques, strategies for building your confidence, and tips on how to make the best possible impression in interviews from the "How to Win in a Competitive Job Market" webinar, held by IEEE-USA.

HAMMERING OUT SMART-GRID STANDARDS

Read about the work of the IEEE 2030 group, which aims to create a how-to guide that defines a broad set of interoperability standards for smart grids, to be used in the computer, communications, and power industries.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2009 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 the Society of National Association Publications.

back story

The Spectrum Dumpling Gang

DEAN KAMEN is a throwback to the era of Edison, Tesla, and Westinghouse, when a charismatic engineer-entrepreneur could be famous. And when Kamen [right], the multimillionaire inventor of the Segway, wants to escape the limelight, he flies his helicopter to his private island off the coast of Connecticut.

Kamen has allowed few journalists to see that retreat, North Dumpling Island. But he did invite *IEEE Spectrum*'s Sally Adee and Francesco Ferorelli, along with two photographers [above], for a weekend in May. They got an exclusive preview of Kamen's amazing digs—documented in "Empire Off the Grid," in this issue—which were then in the middle of a high-tech renovation.

The backstory behind this Back Story: Kamen's lair includes a U.S. Coast Guard lighthouse. Last year, when the Coast Guard cut the cable that supplied power to his island, Kamen decided to revamp the compound as a superefficient, LED-lit, zero-net-energy nation.

But when Adee, Ferorelli, and company got their preview, the road to renovation was definitely going through a bumpy stretch. The boathouse in which the four stayed was so cold that Ferorelli and Adee slept, teeth chattering, in every layer of clothing they'd brought with them. In the main house the next morning, they found electrical engineer and LED specialist Fritz Morgan, an IEEE member, flinging towels onto the stove, mopping up a flash flood that had poured through the kitchen ceiling.

Kamen hadn't done a thing to his island's mansion since he bought the place in 1986. The walls of his bedroom suite were covered with '70s-era paneling, and the bathroom shower featured a floor-to-ceiling one-way mirror like the ones used to shield the identities of police witnesses. The idea is that you can shower unabashedly while taking in the stunning views of Fishers Island Sound.

But the island's delights more than made up for its privations, Adee asserts. At sunset, Kamen gave her a personal tour in his helicopter. "He might have done a couple of extra steep loops when she was in there," Ferorelli says. ❑

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. 45, no. 8 (INT), August 2009, p. 9, or in *IEEE Spectrum*, Vol. 45, no. 8 (NA), August 2009, p. 13.

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; William Sweet, w.sweet@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

INTERN Rosaleen Ortiz, rortizde@ieee.org

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

ART & PRODUCTION

SENIOR ART DIRECTOR Mark Montgomery

ASSOCIATE ART DIRECTOR Michael Solita

ASSISTANT ART DIRECTOR Brandon Palacio

PHOTO EDITOR Randi Silberman

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, Stephen L. Diamond, Hiromichi Fujisawa, Kenneth Y. Goldberg, Susan Hackwood, Bin He, Erik Heijne, Charles H. House, David H. Jacobson, Christopher J. James, Ronald G. Jensen, Mary Y. Lanzerotti, Ruby B. Lee, Tak Ming Mak, David A. Mindell, C. Mohan, Fritz Morgan, Andrew M. Odlyzko, Leslie D. Owens, Barry L. Shoop, Larry L. Smarr, Harry L. "Nick" Tredennick III, Sergio Verdu, 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.

BRIDGET BATCH

contributors

DANIEL BEJAR, who illustrated "CPU, Heal Thyself" [p. 36], is a "huge fan of technology," but "making an image of microprocessors or CPUs was more challenging than I thought," he says. He used traditional printmaking, gouache paint, and "a dash of the computer." Among the Brooklyn-based artist's clients are *Newsweek, Esquire,* and *The New York Times.* Originally from Chicago, he migrated to Los Angeles to do artwork for the movie industry. His portfolio includes covers for *Wired, Time,* and the 1000th issue of *Rolling Stone.*

DAVID BLAAUW is a professor at the University of Michigan, in Ann Arbor, where **SHIDHARTHA DAS** was his student. Blaauw joined the faculty in 2001, after eight years at Motorola, spurred by his desire to "explore wild and crazy things." One result is described in "CPU, Heal Thyself" [p. 36]. Coauthor Das is now an R&D engineer at ARM in Cambridge, England, where he continues to work on ways to maintain chip health—and has gotten healthier himself by walking and biking around town. "In Ann Arbor," he says, "everybody drives."

HANY FARID is a computer science professor at Dartmouth College, in Hanover, N.H. About 10 years ago, he randomly picked up a book in the library on the federal rules of evidence and was astonished to learn that courts made no distinction between photos from negatives and digital photos, even though the latter can be easily altered. "Digital photography was just getting started," Farid says. "I realized this was really going to be a huge problem." Since then digital photo forensics has become the focus of his research, which he describes in "Seeing Is Not Believing" [p. 42]. He now frequently testifies in U.S. courts and advises law enforcement agencies. "All that from picking up a book," he says.

KEVIN COOLEY, photographer for "Empire Off the Grid" [p. 28], took a helicopter ride over Dean Kamen's private North Dumpling Island with the *Star Wars* theme as a sound track. "It was amazing to listen to it as we flew around Long Island Sound at dusk," says Cooley.

MICHAEL ELINS, who created the image of a down-and-out Bill Gates for "Seeing Is Not Believing" [p. 42], specializes in digital illustrations that incorporate his own photography.

RICHARD STEVENSON, author of "The LED's Dark Secret" [p. 22], got a Ph.D. at the University of Cambridge, where he studied compound semiconductors. Then he went into industry and made the things. Now, as a freelance journalist based in Wales, he writes about them. Between assignments, he builds traditional class A hi-fi amplifiers, as opposed to the class D type favored by *IEEE Spectrum*'s Glenn Zorpette. "If we were to share an office," Stevenson says, "many hours would be lost to discussions of the path to hi-fi nirvana."

CLOCKWISE FROM TOP LEFT: DANIEL BEJAR; JOSEPH MEHLING; KATHARINE STEVENSON; MICHAEL ELINS; KEVIN COOLEY

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

spectral lines

The Kindlers Are Coming to Get Me

IT'S SUMMERTIME where I live, the first day of a nice long weekend. I swing by the library to return some books and scoop up a well-reviewed new novel, en route to an important and much anticipated rendezvous with a shaded lounge chair and a tall glass of ice tea.

The library, however, is closed.

Another woman and I arrive at the locked doors at about the same time, and turn away. "Darn," she says. "I need a book for my book club assignment and I was hoping to read it this weekend."

"No kidding," I reply. "I'm flat out of books."

"Oh," she says, whipping a new 6-inch Kindle out of her purse, "I never have that problem," and she proceeds to give me a 10-minute pitch about the many marvels that are Kindle. She was making this now-rare trip to the library, she tells me, because the one book she really needs, something printed a while back, isn't yet available on Kindle, much to her dismay.

Now, I spend vast amounts of time in front of a computer. I also read lots of books, mostly fiction, and sometimes the nonfiction du jour. I read them in hard- and softcover, borrowed from the library or friends, or purchased from Amazon or a garage sale or a secondhand bookshop. But I don't Kindle.

For me, getting away from the electronic screen is part of the appeal of books. Books flip a little switch in my brain that tells it to kick back and relax, because what I'm doing when I turn those paper pages is, more often than not, purely recreational.

You think that'd be okay with people. My gym began tracking workout stats electronically years ago, but no one seems to mind that I still use the paper tracking cards, neatly filed under my last name in a drawer near the entrance. I'm no Luddite, but here, too, when I go to the gym it's for a minivacation, not more screen time.

More and more, however, I'm running into Kindlistas of all ages, sizes, and readerly persuasions. The access-denied woman in front of the library is just the latest. Other true believers whose e-book–enlightened paths have crossed my paper-strewn one have been determined to convert the page-flipping, book-toting tree killers among us. Including me. And they just won't give up.

I try to be polite. "How convenient for you," I murmur. But meanwhile, I'm thinking that even if I wanted to read books in an electronic format, which I don't, having to purchase each and every book for US $9.99 or more would be a deal breaker, given that most of the books I read are borrowed, not bought. Conservatively, I'd say I go through three books a week. That's about $1560 a year if I ever cross over to the e-book side. Ouch.

I finally ask this latest Kindle proselytizer about the cost. Does she find herself spending more on reading material? Well, yes, she confesses, she's made a deal with her husband to drop all her newspaper subscriptions to help subsidize her book Kindling. And no, she hasn't simply moved her newspaper subscriptions to the Kindle; she's just pretty much stopped reading print and online newspapers altogether.

I'm sure my library lady won't be the last Kindler who'll try to convince me to throw off my book shackles and walk with dignity into the digitally grayscaled light. And maybe, someday, when Kindles are even more portable, and colorful, and much cheaper, and a lot less breakable, and powered by batteries with months-long lives, and maybe, just maybe, when every document in the world can be within my reach for a few dollars in a few seconds, I will finally succumb to their electronic charms. But until then, if I see someone approaching with a Kindle in hand and a bedazzled look in her eyes, I'll quickly open my book and whisk myself away to another time and place.

—TEKLA S. PERRY

A version of this column appeared in IEEE Spectrum Online's Tech Talk blog on 9 July.

AMAZON.COM

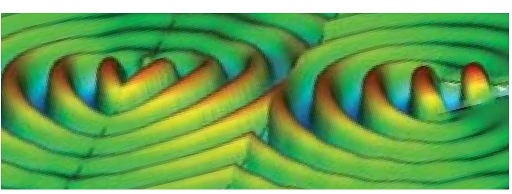

forum

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, with the article, go to http://www.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.

TO BOLDLY GO...OR NOT

REGARDING YOUR "Special Report: Why Mars? Why Now?" [June]: On 20 July, NASA will mark the 40th anniversary of the *Apollo 11* lunar landing. For one shining moment, Earth stood still and humanity took one giant leap into the unknown. Whatever happened to the American dream of exploring the universe? I look up at the moon and wonder when we will go back. I look up at Mars and wonder whether we'll ever get there. Human exploration of the moon, Mars, and beyond would restore national pride and bring purpose to our space program. Space exploration is an investment in our future and a benefit to humanity. I hope Americans press on to make the unknown elements of space known to our world. President Obama and NASA administrator nominee Charles Bolden can put NASA back on course. Let's keep the dream alive and boldly step back on the moon by 2020!

RICK SCHREINER
San Marino, Calif.

DID I miss the article on why we should colonize Mars in the first place and the moral implications of doing so? If history is any indicator, Mars will be the dumping ground for what we don't want here: nuclear waste, pollution, prisoners, the persecuted, and the general offscouring of Earth. Eventually, as we do with every other natural resource in our hands, we will exploit it until we completely ruin it. Perhaps we should squander our billions fixing the problems we've created here before hauling off to go ruin another perfectly good planet.

JAMES METZGER
IEEE Member
Vienna, Va.

I REALLY ENJOYED this latest issue of *IEEE Spectrum* dedicated to a single topic. I am, however, disappointed that you did not present a more balanced view on why we have not gone back to manned space programs. Most of the articles seem to suggest that this is a failure of NASA and the international community. That really undermines the great progress that NASA and the space programs of many other countries have made in unmanned space exploration, new propulsion systems, and autonomous robotics. One could argue that all these are prerequisites to a return to manned space exploration.

One of the most critical remaining obstacles to manned space missions is the lack of an efficient form of energy that isn't chemical or fossil based. Launching a manned space program with the intent to make Mars habitable while we rely mainly on fossil fuels is putting the cart in front of the horse. Maybe Russia, China, France, Germany, India, and Japan are doing the smart thing to invest equally in fusion power research. The United States has cut its funding for this research and for thermonuclear power generation. So I'm not at all surprised that we have not flown to Mars yet, since we still have the unfinished business here on Earth of securing a sustainable power source for such aspirations.

ARMIN KITTEL
IEEE Member
Centreville, Va.

AND THE WINNER IS...

AS I read "The Million Dollar Programming Prize" [May] I groaned with empathetic embarrassment when I realized that Netflix has committed the cardinal sin of nonlinear modeling. By permitting teams to perform multiple runs on the test set, it has allowed its training data to become contaminated. Although the results will be great examples of curve fitting, they will have questionable utility as generalized models that can perform any better than the Cinematch algorithm. Granted, Netflix has tried to prevent "gaming" by limiting the iterations, but allowing access to the test set even once causes contamination.

BARRY FLOWER
IEEE Senior Member
Sydney

Coauthor Jim Bennett responds: The contest has two hidden data sets that must be predicted. When a submission is made, the results on one of the sets (the "quiz" set) are reported, but the results on the other (the "test" set) are not. It is the score on this hidden test set that is used to determine the winner of the contest.

CORRECTION

In "Moonstruck" [June], research was attributed to David L. Shuster and Ben Weiss that should have been credited to Donald Bogard, Pratt Johnson, and Robert Pepin, who showed in the 1980s and 1990s that meteorites on Earth were from Mars, based on Viking 1's sampling of the Martian atmosphere in 1976. Later work by Shuster and Weiss concerned the temperature of a Martian meteorite at the time of its ejection.

more online at www.spectrum.ieee.org

update

Jordan's Radioactive Water Problem

A critical $1 billion engineering project in Jordan could be complicated by radium

JORDAN IS in a tight spot. The virtually landlocked country is 80 percent desert, and the remaining 20 percent loses most of its rainfall to evaporation. The Dead Sea and the Jordan River, which feeds it, are drier than ever. With its population swelling with Iraqi migrants, water is Jordan's foremost concern.

Most of the country receives water service once a week at best, and unexpected disruptions force the Ministry of Water and Irrigation to deliver water by truck. "When a country has its back against the wall, you take the least damaging solution," says Munther Haddadin, the country's former water minister.

The water ministry has decided that the best way to get water to the capital of Amman is to mine it, tapping into what the department says is some of the cleanest, purest water in the world. The water sits in the pores and holes of the Disi aquifer, an expanse of sandstone some 500 meters beneath the desert in southern Jordan and northwestern Saudi Arabia.

Having just secured the final US $200 million in loans needed from European development banks in May, the government will soon begin building a 325-kilometer pipeline across the country, from the heart of the desert to Amman. The plan is to pump 100 million cubic meters of water from 55 wells in Disi each year. The water will travel about 1300 meters uphill, requiring about 4 kilowatt-hours of energy to deliver each cubic meter, according to Othman Kurdi, the

WATER BELOW: Jordan plans to mine water from the Disi aquifer beneath the desert.
PHOTO: CHRISTOPH ROSENBERGER/GETTY IMAGES

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

update

engineer in charge of the Disi Water Conveyance Project. At that rate, the power required to pump a year's worth of water is equivalent to the output of a 45-megawatt power plant, or about 4 percent of the country's electricity production.

"It's not rocket science, but it's a megaproject and a challenge in every sense," Kurdi says. Indeed, the Disi Water Conveyance Project is riddled with complications. Recent research has revealed that the water may not be as pure as project planners had said, and that could make the scheme more complex and costly—and even take a toll on public health. Further, the pipeline project is just a stopgap measure that will leave Jordan permanently poorer in natural freshwater resources while the country pursues an even larger, costlier, and more energy-intensive solution that remains decades away.

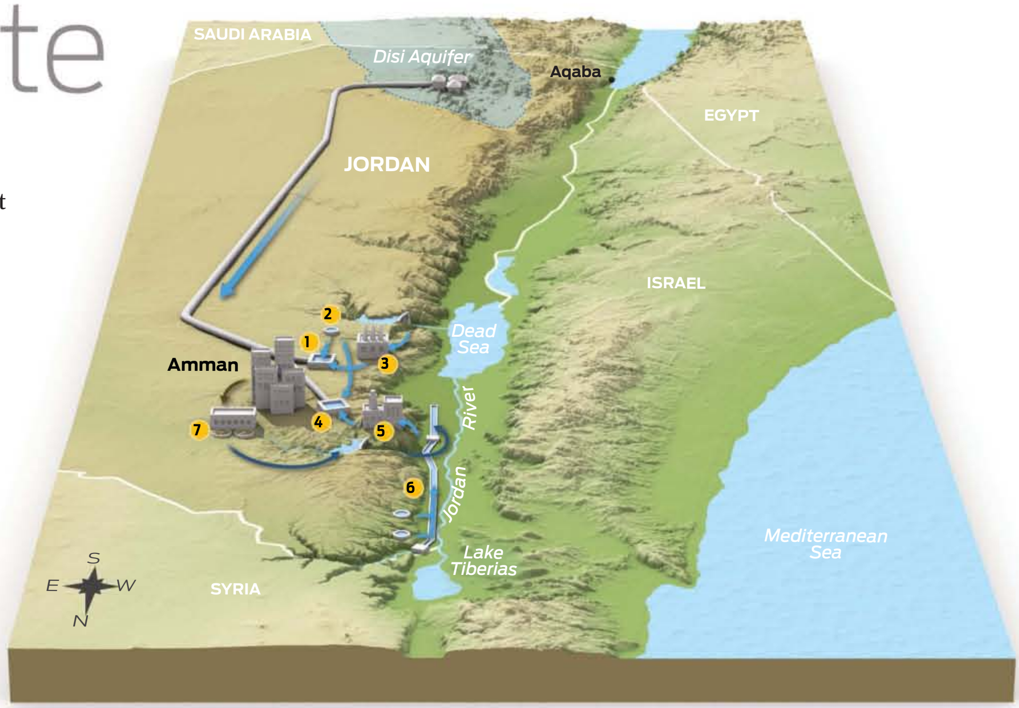

JORDAN'S WATER CYCLE: After the Disi pipeline is built, 100 million cubic meters of water will travel from the desert to two reservoirs each year. About 40 million cubic meters will go to the Abu Alanda reservoir **(1)** and mix with some surface water from Wala **(2)** and treated brackish water from the Zara Ma'en desalination plant **(3)**. The other 60 million cubic meters will go to the Dabouq reservoir **(4)** and blend with Wala surface water and the output of the Zai Treatment Plant **(5)**, which treats water from the King Abdullah Canal **(6)**. Amman's used water is sent to the As Samra Wastewater Treatment Plant **(7)** and later used for irrigation.

ILLUSTRATION: EMILY COOPER

By pumping the Disi aquifer, Jordan will be depleting its only strategic reserve of water, a move also being considered by other developing nations that are poor in both energy and water resources. Unlike rivers and lakes that refill with rainfall or melting snow, once this so-called fossil water is pumped, it leaves Jordan forever. Much of the water in the Disi aquifer essentially hasn't moved since it began dripping into the ground during the Pleistocene era, some 30 000 years ago. "For developing countries in that region, they have no other choice—using this water is the only way to survive the water crisis," says Avner Vengosh, a geochemistry professor at Duke University, in Durham, N.C.

Policymakers and water experts had been debating the merits of draining Disi through much of the project's planning. But in February the debate suddenly shifted, when Vengosh published a report in the journal *Environmental Science & Technology* describing the Disi water as highly radioactive. He and his coauthors collected samples from 37 wells in the Disi area used mostly for agriculture and mining activities. They found that in all but one well, the concentrations of radium-226 and radium-228 isotopes exceeded the levels considered safe by the World Health Organization and even the more relaxed European Union and U.S. water standards. In some spots, the radiation levels were observed to be 30 times the WHO's thresholds. Long-term exposure to radium is believed to increase the risk of developing bone cancer.

Vengosh theorizes that the isotopes entered the water from the surrounding sandstone through a physical process known as recoil. Thorium-232 and thorium-230, the parents of radium-228 and radium-226, respectively, exist naturally in the porous sandstone that holds the Disi water. When one of those thorium atoms radioactively decays to emit an alpha particle (two protons and two neutrons bound together), some energy is also released that causes the new atom to move in the opposite direction of the ejected particle. In some cases, the recoil can cause this new atom to get pushed out of its host material into a surrounding medium—in this case, from sandstone into water.

The Ministry of Water and Irrigation contends that the radiation is not a problem. The Disi pipeline will send the water to two large reservoirs outside Amman, where the fossil water will be diluted with 105 million cubic meters of treated surface water, Amman's current supply. According to Susan Kilani, a ministry official in charge of water quality,

825.5 TERAFLOPS Measured processing power of the new No. 3 supercomputer in the world, JUGENE. The computer and its sister, 10th-ranked JUROPA, both housed at Forschungszentrum Jülich, in Germany, were the only two non-U.S.–based machines to make it into the new top 10 list.

the quantity is "in excess of what we need for blending."

Judging by Vengosh's data, this doesn't appear to be the case. Dilution would double the total volume of water, which means that the Disi water's radiation can be no more than double the desired threshold in order to comply with international benchmarks. Very few of the wells tested by Vengosh and his colleagues met that criterion. Relying on blending would limit the amount of usable water in the aquifer and curtail the life span of the pipeline project—or expose the population of Amman to heightened levels of radium.

Aqaba, Jordan's small port city, has had an instructive experience with Disi water. The city has been pumping 15 million cubic meters of water from Disi each year to a collection reservoir outside Aqaba. Samples taken at the wells come out radioactive. But without any diluting or further treatment of the isotopes, the water is somehow pristine by the time it reaches the city, according to Aqaba's utility.

Imad Zureikat, the Aqaba Water Co.'s general manager, maintains that Aqaba's water is tested at several international laboratories and adheres to international standards. "My kids are drinking water from the tap," Zureikat says. "We don't play with people's health. Everywhere you go in Jordan, you can drink from the tap."

The locations and depth of the wells likely play a role in Aqaba's good fortune [see "Jordan's Red Sea Desalination Plan," *IEEE Spectrum*, July 2009]. Near Aqaba, at the southern tip of Jordan, the aquifer lies closer to the surface. Vengosh's samples from these shallower wells showed much more variability in their radium content. So cautiously choosing well sites may indeed bring the water ministry within sight of international standards for the Amman project.

Elias Salameh, a professor at Jordan University, in Amman, has been monitoring the elevated alpha-particle activity in Disi water. His data is unpublished, because the presence of radium in groundwater is just not news, he says. In the United States, for example, New Jersey has relied on water from a radium-tainted aquifer for many years. He points out that should blending fail to neutralize the water, the ministry can treat it using reverse osmosis, by forcing

DESERT BLOOM: A well in the Disi aquifer near Aqaba, in the south, provides water for agriculture.
PHOTO: AVNER VENGOSH

the water through a membrane that prevents the passage of radium. Ion-exchange purification, in which the water is fed through columns of porous materials whose pores work as capture sites for the radium, is another option. "We can't choose another land, another country, so we have to do our best," Salameh says.

Of course, any purification treatment would come at a price. Haddadin, the former water minister, studied the cost of the Disi project, along with the collection and treatment of its wastewater, and concluded the cost of water service would reach 10 percent of the income of the average Amman resident. Adding a treatment facility, and the needed disposal of radioactive waste, would drive the cost up even more.

Whatever the price, Jordan will almost certainly find a way to use the Disi water. As Nizar Abu-Jaber, a geology professor at Yarmouk University, in Irbid, Jordan, sees it, the "availability of water takes precedence over radioactivity." Water-stressed countries are frequently forced into difficult, or at least expensive, choices. In Libya, a massive project known as the Great Man-Made River delivers fossil water from the Sahara Desert and distributes it along entirely new waterways, while Saudi Arabia has plowed money into becoming the world's largest producer of desalinated seawater.

Amman, for its part, is undergoing a population boom, spurred by an influx of an estimated million immigrants from Iraq. "The Disi project will not solve the problem of water in Jordan," says Kurdi, the project's leader. "It will just maintain the status quo." With any luck, the project will at least enable the government of Jordan to increase each person's share of the kingdom's water, to meet what it considers to be the daily demand of 120 liters per person per day. (Abu Dhabi residents, by contrast, go through an average of 550 L per day.)

By early 2013, Kurdi estimates the water should be flowing to Amman. The aquifer will be pumped for about 25 years, he says, until a subsequent water project is in place. That project, known as the Red-Dead Canal, would send water north from the Red Sea, with half of the water replenishing the shrinking Dead Sea and the other half to be desalinated for consumption.

To get within sight of the ballyhooed canal, Jordan must first find a way to keep its population—and economy—humming along for the decades the country will need to build it. "We are going to implement this project, insha'Allah, God willing, as we say," Kurdi says. "Because this is what we need." —Sandra Upson

news briefs

Hydrogen Helper
Scientists at the University of Delaware say that carbonized chicken feathers could be a cheap way to store hydrogen for fuel cells. According to chemical engineering professor Richard Wool, the heat-treated feathers could hold more hydrogen than costlier competing technologies, such as metal hydrides and carbon nanotubes.
PHOTO: ERIC ISSELÉE/ISTOCKPHOTO

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

update

news briefs

Eye on the Prize
In late June a coalition of research teams announced that it had qualified for the US $1 million prize that Netflix has offered to whomever improves its movie recommender system by at least 10 percent. "We haven't won the million yet, just crossed the necessary threshold, moving the competition into a 30-day 'last call' period," says Yehuda Koren of Yahoo Research, a member of the coalition. Koren and two colleagues at his former employer, AT&T Labs Research, won an earlier stage in the competition (see "The Million Dollar Programming Prize," *IEEE Spectrum*, May 2009). Koren says his team would reveal its latest algorithmic tricks "if and when we win."

Future of Low-Power Chips in Doubt

Big flaw found in transistor noise theory

Engineers at the U.S. National Institute of Standards and Technology (NIST) say that the basic theory explaining the origin of a certain type of noise produced by very small transistors is totally wrong.

Known as random telegraph noise, this aberrant signal is becoming a problem for static RAM and flash memory, and it will also become a threat to future low-power logic circuits as their dimensions continue to shrink and the voltage at which they operate decreases. Without a theory, engineers will find it difficult to reduce the noise, predicts Jason Campbell, a National Research Council postdoctoral fellow at NIST. "If you don't know where it's coming from, you don't know how to fix it," he says.

"It's critical to raise this issue," says Genadi Bersuker, a fellow at Sematech in Austin, Texas. "The whole spectrum of processes involved in noise is much richer than the simplified description everyone tends to use."

The prevailing theory says that as an electron passes through the transistor's channel it can temporarily get stuck in the insulation above. That happens because of a quantum mechanical phenomenon called tunneling, which makes it possible for the electrons to jump through the insulation as if it weren't there. According to the theory, electrons randomly tunnel into defects in the insulation and then tunnel out. The random moving-stuck-moving pattern was thought to cause the noise.

The duration of the tunneling process is related to the thickness of the insulation, and decades ago, when the theories were formulated, insulation was micrometers thick. But now it is measured in nanometers. Campbell measured production-quality nanometer-scale test transistors, made by TSMC, that were much smaller than those in high-end microprocessors and whose insulation was only 1.4 nm thick. He expected to see tunneling that took nanoseconds to picoseconds. Instead, the process took the milliseconds to seconds you'd see in a much bigger transistor.

"What people thought is very likely not correct," says Campbell, who presented the results at two IEEE-sponsored conferences in April and May. —Samuel K. Moore

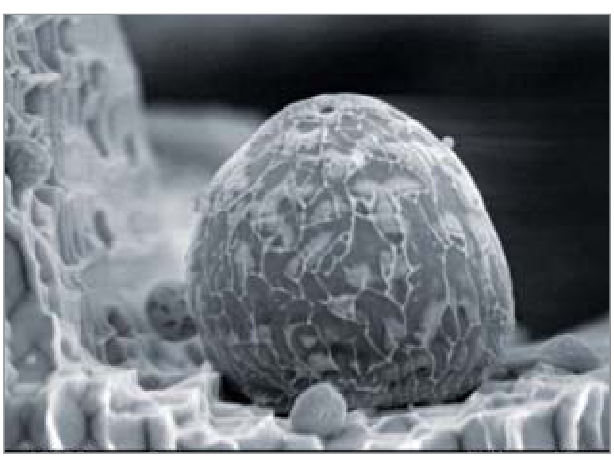

DJEMEL LELLOUCHI

Dinosaur Egg

Djemel Lellouchi, of microelectromechanical systems reliability firm Nova MEMS, in Ramonville, France, found this oddity when investigating the failure of an optical switch due to an electrostatic discharge. The image won the annual Art of Failure Analysis photo contest at the IEEE International Symposium on the Physical and Failure Analysis of Integrated Circuits, in Suzhou, China, from 6 to 10 July. The egg was just one of many microscopic curiosities collected by failure-analysis experts this year, including a sulfide flower, a tungsten snake's tail, and an infant's face on gold. For the top 10 entries, see http://www.spectrum.ieee.org/slideshow/semiconductors/devices/slideshow-the-art-of-failure-2009.

US $370 MILLION Estimated cost of converting drawings and software from the U.S. shuttle program from English units to the International System of Units, or SI. Because NASA will use legacy designs in its next generation of shuttles and it would cost too much to convert them, the agency says it plans to design the new craft using English units as well.

SCOTTISH SEA POWER: Scotland has the potential to contribute 25 percent of Europe's tidal power.
PHOTO: EARTH-VISION.BIZ/WAVE DRAGON

£10 Million Sea Power Challenge

Scotland sets terms for Saltire Prize

Scotland is finalizing the terms of a contest in wave and tidal energy that takes inspiration from the prize that prompted Lindbergh's transatlantic flight in 1927 and successors like the X Prizes and the Virgin Earth Challenge. The aim of the contest is to make ocean energy more than just a technical curiosity and, not so incidentally, give the country's inventors and entrepreneurs a boost in an area where they have some obvious advantages—suitable geography, friendly government policies, and a head start in engineering.

Dubbed the Saltire Prize, after the cross that is the central element in Scotland's flag, the prize of £10 million (about US $16 million) will be awarded in five years. Contestants will need the time to devise and demonstrate their technology because, by all accounts, Saltire is a very challenging challenge, so much so that only a Scottish company may be able to win it.

"To win the prize," the rules say, "entrants must demonstrate, in Scottish waters, a commercially viable wave or tidal energy technology with a minimum output of 100 gigawatt-hours over a two-year period—using only the power of the sea. The winner will be judged to be the best overall technology after consideration of cost, environmental sustainability, and safety."

Although the final guidelines won't be released until the third quarter of this year, by early July the contest had already attracted the interest of more than 100 potential competitors, says Claire M. Smith, Saltire's spokesperson. About 30 percent are from the United States, she reports, and "U.S. companies have just as much chance as anyone else."

But Roger Bedard, the lead ocean-energy researcher at the Electric Power Research Institute, in Palo Alto, Calif., is skeptical. This is a prize "that won't be won by the United States," he says. Bedard points out that complex U.S. regulations do not favor getting new technology into the water fast, and that technical challenges are often underestimated. A tidal project in New York City's East River is a case in point. In the waterway that connects Long Island Sound with the city's harbor, the first turbine blades installed promptly broke in the powerful current.

What is more, Bedard continues, a plant producing 100 GWh for two years is not small by ocean energy standards. Taking into account the fact that an average plant would be generating electricity at capacity only about a third of the time, the prize calls for a tidal facility of roughly 18 to 20 megawatts. But the largest plant operating at present, in Aguçadoura, Portugal, is just 2.25 MW. On reflection, says Bedard, "I bet nobody is going to do it."

Some of those that come closest will almost certainly be Scottish. The Portuguese plant was built by Pelamis Wave Power, in Edinburgh, a world leader. Meanwhile, ScottishPower Renewables is getting set this month to ask the Scottish and Irish assemblies for approval to develop three 1- to 20-MW sites, with turbines developed in Norway, according to Scotland's *Herald*. The European Marine Energy Center is located amid the country's Orkney Islands.

The Scots are shooting to get 31 percent of their energy from renewables by 2011 and 50 percent by 2050, which are among the world's most ambitious targets. Their country has the potential to generate 25 percent of Europe's tidal energy and 10 percent of its wave energy, says Saltire's Smith. "Scotland is a small country that likes to punch above its weight," she told a United Nations conference in June.

—William Sweet

update

How Software Found the Air France Wreckage

A new drift simulator pinpoints ocean accidents

When Air France Flight 447 from Rio de Janeiro disappeared over the Atlantic in the early hours of 1 June, search and rescue teams had to look for survivors in an area almost as big as Great Britain. Not knowing exactly where the Airbus 330-200 went down, French and Brazilian authorities turned to new software developed for the U.S. Coast Guard that uses the location of debris to look back in time to estimate the most likely site of an accident.

The Coast Guard's "reverse drift" modeling program takes into consideration the types and locations of objects found floating in the water, the time of an accident, and environmental conditions—such as the speed and direction of ocean currents and winds—to calculate the location where an accident may have occurred.

The Air France recovery operation was the first real-world test of this reverse-drift capability, which was added to SAROPS, the Coast Guard's Search and Rescue Optimal Planning System, only in early May. It was developed by Metron, Applied Science Associates, and Northrop Grumman.

In the past, drift models could only look forward in time. If a person falls off a cruise ship, for example, a forward-drift computer simulation can determine the approximate area where the current and winds will carry him. But that works only if you know where the person fell off the ship.

The exact coordinates of an accident at sea are not always known, as was the case for Flight 447. That's when reverse-drift modeling comes in. Reverse drift helps provide the missing variable—the location of the accident—that rescue planners need in order to look for possible survivors and wreckage. A seat cushion and other small debris were the first clues that operators fed into the simulation program in the early days of Flight 447's search operation, officials say.

Once the reverse-drift simulation gave officials an idea of where the plane went down, they used forward-drift simulation to pinpoint where to look for debris, bodies, and possible survivors, all of which would be carried away from the crash site at different rates by wind and currents.

"The computer model was very accurate, and that's when [the rescue teams] started to find the wing debris and the bodies," says Geoff Pagels, search and rescue specialist at the Coast Guard rescue coordination center in Portsmouth, Va., who was in daily contact throughout June with officials at the rescue coordination center in Gris Nez, France. By 15 June, search teams had found the bodies of 50 of the 244 victims and more than 400 pieces of debris from the plane, which were scattered over a 200 000-square-kilometer area. "Each time they find something, they keep coming back to us and ask us to run more simulations," Pagels says.

At sea, every object is affected by winds and currents. The effect of the current is simple. "The object goes where the water goes," says Tom Kratzke, programmer at the scientific consulting firm Metron, who wrote the code for the reverse- and forward-drift models. "The effect of the wind is much more complicated," he says.

The program uses surface area and nine other parameters of typical objects to determine how the wind will drive them.

Even though the SAROPS system is equipped with a sophisticated environmental data server, which gathers ocean wind and current information from multiple sources, environmental conditions are never fully known. So the program makes 10 000 educated guesses about the different paths that objects may take and combines them to find the most likely location.

The reverse drift gives more accurate results with each new object found but produces a larger search area as time goes by, says Jack Frost, manager of development for SAROPS. "Time is the enemy," he says.

—Rosaleen Ortiz

WRECKAGE FOUND: Software simulated how wreckage from Flight 447 would drift. *PHOTO: ERALDO PERES/AP PHOTO*

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

Foundry Giant Sees the Light

What does TSMC know about solar cells and LEDs? More than you'd expect

What do you do when you dominate your market? You find other markets to dominate. That's the plan for Taiwan Semiconductor Manufacturing Co. (TSMC), the world's biggest semiconductor foundry, according to its founding chairman, who retook the reins of the company in mid-June. Morris Chang, the recently reappointed chairman and CEO, says that a better future lies ahead for the firm if it pushes into green energy—specifically solar power and LEDs.

The immediate question from watchers of both industries is: What does a silicon foundry know about solar cells and LEDs? Light-emitting diodes aren't even a silicon technology, as they rely on gallium arsenide and other compound semiconductors. With a 2008 net income of NT $99.93 billion (about US $3 billion), TSMC can surely afford to buy expertise. "We will pursue smart brains from outside TSMC and adopt a merger-and-acquisition strategy," Chang says.

But a close look at TSMC shows that the company already has some of the technological foundations for work in both areas. According to a 2008 report by tech consulting firm Semiconductor Insights analyzing U.S. patents in solar technology, TSMC is the second largest patent holder in solar cells. Granted, most of those patents are primarily concerned with TSMC's existing business, with solar cells mentioned as a secondary use often in boilerplate language in the patent claims. But some, such as a 2006 patent application for a method of patterning indium-tin-oxide films—used as the transparent electrode in solar cells, among other things—are more directly relevant.

TSAI TIME: TSMC must break into LEDs and solar to grow faster. PHOTO: SAM YEH/AFP/GETTY IMAGES

In the LED arena, TSMC has indirectly invested in a number of start-ups, including LED-lighting maker BridgeLux, LED-driver firms Auramicro and Exclara, and LED thermal-management company Liquidleds. The investments were made through VentureTech Alliance, a venture capital company owned mostly by TSMC and run by several former TSMC executives.

Closer to home, local media reports reveal that the company is bringing metal-organic chemical vapor deposition (MOCVD) equipment from old silicon-chip manufacturing plants back online for LED trial production. TSMC has also acquired a 1.2-hectare plot of land in Hsinchu Science Park for its expansion.

The solar and LED businesses will come under a new unit headed by Rick Tsai. Until Chang replaced him in a move that surprised many analysts, Tsai was TSMC's CEO.

The main reason for the green-tech push is to counter the slowing growth of the semiconductor foundry business. According to Chang, without new business operations, TSMC's revenues, which stood at US $10.6 billion for 2008, might grow by 3 percent annually to between $14 billion and $15 billion by 2018. However, Chang hopes solar and LEDs will push growth to 4.5 percent per year and add $2 billion in revenue per year by 2018.

TSMC's move is in step with government initiatives to boost the production value of the local solar photovoltaic industry from NT $101.1 billion in 2008 to NT $450 billion (US $13.6 billion) by 2015 and that of the LED industry from NT $46 billion in 2008 to NT $540 billion (US $16.3 billion) by 2015. The government expects China to be a big consumer of its green-tech wares. In early June, a cross-Strait conference on the LED industry was held in Taipei, and Chinese delegates asked Taiwanese LED makers to participate in a project in China to light up 10 cities with energy-saving LED streetlights.

Some in Taiwan are pleased with TSMC's new direction. Cheng-Chung Lee, a professor of optics and photonics at the National Central University and chairman of Taiwan's 200-member Optical Engineering Society, says there is a gap between the industry and research circles in Taiwan that the TSMC move might help close. He believes that when challenged by a rival such as TSMC, local companies would have no choice but to acquire new technologies developed in universities and research institutes.

—Yu-Tzu Chiu

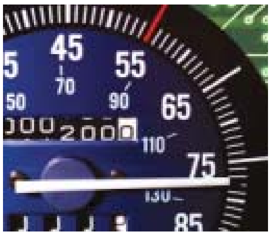

news briefs

Silicon Odometer

You can get an idea of how well a car will perform by looking at the mileage on its odometer. University of Minnesota electrical engineering professor Chris H. Kim wants to do essentially the same for silicon chips. His lab has developed a silicon circuit that tracks the three signs of aging that can slow down a chip's performance. The hope is that integrating such a circuit into a microprocessor could help the microprocessor compensate for its own aging. See http://www.spectrum.ieee.org/semiconductors/devices/an-odometer-for-silicon-chips.

PHOTO: GETTY IMAGES

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

careers

THE RISE AND FALL OF THE QUANTS

Can Wall Street do without the techno-geeks who designed its complex derivatives markets?

THE BEARS are outrunning the bulls these days, but Wall Street still relies on quants like Ph.D. electrical engineer Juhua Zhu, now a vice president at Morgan Stanley, to crunch the numbers. PHOTO: DAVID YELLEN

When the U.S. financial system melted down, fingers were quickly pointed at the "quants"—the physicists, mathematicians, and engineers who had devised the computer programs, statistical tools, and financial instruments that were supposed to help investors manage risks. Critics said that it was the flawed assumptions of those financial models that brought banking to the brink of Armageddon.

You might guess that Wall Street is now shunning physicists, mathematicians, and engineers, but you'd be wrong. Talented people with quantitative backgrounds are more welcome than ever, says Petter Kolm, deputy director of New York University's master's program for mathematics in finance and formerly a quant for Goldman Sachs. Before the financial mess started in the fall of 2007, he says, investment banks typically hired fresh MBA grads as entry-level analysts. "Many investment banks I talk to today say they want to replace that portion of people with MBAs with people who have quantitative analysis skills, such as engineering and math."

The culture of Wall Street is already relying less on traditions and personal connections and more on complex technologies. Juhua Zhu, who has a Ph.D. in electrical engineering from Princeton and is now a vice president at Morgan Stanley, saw the change coming shortly after joining the firm four years ago. She says it was soon clear that "traditional trader jobs would vanish as many transactions would be transferred to computers."

Those changes are carving out a need for more math whizzes and data crunchers. "Quantitative jobs demand research talent—people who can read any text in a technical field and reach a high level of expertise in a short amount of time," says Alp Atici, a Columbia University math Ph.D. who works as a quantitative researcher at hedge fund Citadel Investment Group. "People with Ph.D.'s in science and pure math are usually accustomed to much harder and deeper research texts."

One of the first quants was Robert Merton, who started out as an applied mathematician at Caltech before switching to economics at MIT. In 1969 he tried to work out the pricing of stock options with the help of stochastic calculus—a branch of mathematics used to model random systems such as Brownian motion, the movement of particles in liquid. A high point for quants came when Merton and Myron Scholes won the 1997 Nobel Prize in Economics for their option-pricing model. (When it comes to low points, Merton, the quant trendsetter in so many ways, was ahead of his time again. He and Scholes were members of hedge fund Long-Term Capital Management's board of directors when that company lost more than US $4 billion in 1998; the Federal Reserve, fearing a liquidity crisis, put together a restructuring plan that presaged its aggressive liquidity interventions in 2008.)

The trickle of physicists and mathematicians eschewing low-paying academic jobs in favor of Wall Street bonuses turned into a flood in the Reagan years. The 1980s fads for junk bonds

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

and leveraged buyouts came and went, but the quants stayed. Their direct role in today's recession started in the booming 1990s housing market as they helped banks package mortgages, credit cards, and other credit assets, slice up the packages, and sell them as instruments known as asset-backed securities. That took risky assets off companies' balance sheets, freed up capital, and let the companies borrow more money. Quants are credited with creating models that helped investors understand, manage, and price the risk associated with these securities.

Asset-backed securities are a great innovation if used sparingly. But this did not happen. Emanuel Derman, industrial engineering and operations research professor at Columbia and author of the memoir *My Life as a Quant: Reflections on Physics and Finance*, says banks "got too big for their boots and borrowed too much money." The quants' mathematical machinations didn't so much dilute risk as hide it. And then came the breaking point, one that quantitative models did not take into account: record numbers of subprime borrowers defaulting on their mortgages.

So was this a lack of foresight on the quants' part? The simple truth is, there is no one right model, Derman says. Inputs to financial models are related to how people will behave in the future. But patterns of behavior change over time.

"If you're designing a global positioning system, the distance from here to Waterloo doesn't change," he says. "Financial engineering isn't based on financial science. It's scientific methods applied to human variables."

But why didn't quants raise the alarm about risks associated with their innovations? NYU's Kolm says that business decision makers, not quants, were calling the shots and did not always care about the risks. At a major investment bank, he says, risk managers who tried to warn about risks were unpopular: "Either you played along or you left because you were the bad guy at the party." It's also becoming clear that mortgage fraud played a big role in distorting the data that the models used to predict foreclosures—a case of garbage in, garbage out.

Despite the role that physicists and engineers played in the economic crisis, the relationship doesn't seem to have soured. Quants aren't being recruited less or fired more—though there are fewer jobs overall as some companies, like Lehman Brothers, disappear, and others, like Bank of America and Merrill Lynch, merge. Beverly Principal, assistant director of employment services at Stanford, says that finance companies are still coming to campus, while a large number of students are still interested in financial careers. Career service representatives at other schools echo this sentiment.

Companies have started to be selective, and people with advanced degrees have the upper hand. Credit Suisse Securities has spent the past year trying to attract people with doctorates in physics, math, and engineering from top schools, says Ilias Tagkopoulos, who got his Ph.D. in EE from Princeton last summer. A trifecta of skills—programming, mathematical proficiency, and an ability to communicate—led to his job at Credit Suisse as a relationship manager. Most members of his team are also Ph.D.'s from top schools, he says.

The lack of jobs, though, has started to make itself felt. Ming Zhong, a portfolio manager at Lazard Asset Management, an investment bank in New York City, says that when he graduated from Columbia with a master's in EE in 2004, 70 percent of the jobs available through the campus-recruiting program were finance related. Now, he says, graduates are having a hard time finding internships, even unpaid ones. "One year ago, I'd say keep trying, always have hope. Six months ago, I suggested students look around. In the spring, I told them, 'You should think about changing careers.' "

But things won't always be so bleak. When the economic pendulum eventually swings back, predicts Morgan Stanley's Zhu, "the whole pie will have shrunk, but the portion of jobs for people with a technical background will have continued to grow."

—Prachi Patel

WHY BE A QUANT? What draws analytical and technical minds to finance? The high pay certainly helps. However, many say the attraction isn't simply monetary. Quants see their hard work pay off as profits or losses, and perhaps surprisingly, many feel they've done more good in finance than they would have in science.

ALP ATICI *Ph.D. in math, 2006, Columbia University* CURRENT POSITION: *Quantitative researcher, Citadel Investment Group, Chicago*

"In the academic world, you are valued by how others perceive your work, whereas in quantitative finance the measures are far more objective. For any ambitious person, the meritocracy is very real. There is no upper limit on rewards for good effort. A finance career was the best choice of my life. I love the work I do, the people I work with, and what I contribute to the world."

SILVIO BORELLI *M.S. in financial engineering, 2004, Columbia University* CURRENT POSITION: *Director of credit analytics at Moody's Corp., New York City*

"One of the things that attracted me most about finance was the constant need to be on top of news worldwide and analyze its potential effects on countries, companies, currencies, and economies. It's a job that invites and demands a broader view of the world. It has definitely been a rich experience."

DAVE FERGUSON *Ph.D. in robotics, 2006, Carnegie Mellon University* CURRENT POSITION: *Quantitative researcher at Two Sigma Investments, New York City*

"I was excited about the challenges in computational finance—it is an exciting field in that the results of your research can be immediately applied and objectively evaluated. It has been fascinating, especially given the current situation in the world of finance."

ILIAS TAGKOPOULOS *Ph.D. in electrical engineering, 2008, Princeton University* CURRENT POSITION: *Relationship manager, Credit Suisse Securities, New York City*

"For my Ph.D., I was looking at dynamic systems and behavior. One of the biggest, most complex dynamic systems that's very close to chaos is the stock market. For a guy like me, it's very stimulating, intriguing, and rewarding."

JUHUA ZHU *Ph.D. in electrical engineering, 2006, Princeton University* CURRENT POSITION: *Vice president, Morgan Stanley, New York City*

"I'm a pragmatic person and sometimes worried if my research will be useful in real life. In the financial world, results are instantly tested by profit and loss. Second, I feel that a finance job implies a more exciting life. Following the news and seeing its impact on the market is exciting for me."

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

tools & toys

media

Quantum Rock
Back in the 1960s, Neal Doughty [below] divided his time between electrical engineering classes at the University of Illinois and a rock band he had just started called REO Speedwagon. The keyboardist still feeds his inner geek; his latest toy is a German-made virtual-analog synthesizer, and he's reading popular books on quantum physics during the band's current U.S. tour with Styx and 38 Special.

"The guys were tired of hearing about the Large Hadron Collider on long bus rides," Doughty says. "Of course, they don't mind having a mad scientist on board when it comes to free support for their Macs."

For tour dates, go to http://www.speedwagon.com.
—Susan Karlin

FireSword-IV
LED flashlight, Elektro Lumens
US $400

THE FUTURE'S SO BRIGHT

A flashlight made from high-output LEDs is not for the fainthearted

Anyone who thinks that LED illumination is for wimps should look at—not into—one of Wayne Johnson's new flashlights. With 4 four-chip LED modules, it puts out roughly the same light as a 200-watt incandescent bulb. Parabolic reflectors point most of that light into a beam that can illuminate the landscape dozens of meters away.

Johnson, a former computer programmer, reports that he first got into custom flashlights because he liked to take nighttime walks. But confrontations with furiously barking dogs he couldn't even see or skunks he could make out only when well within spraying range made his evening ambles more chancy than he liked. Incandescent flashlights had too little battery life, and LED units cast a negligible beam. When the first 1-watt LED modules appeared, he was hooked.

For more than five years he's supported himself with Elektro Lumens, a company that makes both high-output LED conversion kits for conventional flashlights and whole custom-built models for the deep pocketed.

Johnson's newest flashlight, the FireSword-IV, looks a little like an art deco Olympic torch: The body is turned from a solid piece of aluminum, with fine cutting-tool marks that provide a nonslip grip. The head is turned from a larger-diameter chunk, with deeply incised fins that dissipate the heat thrown off by both the LEDs and the batteries, which are special lithium-ion rechargeables with relatively low internal resistance. The flashlight weighs just under a kilogram with batteries, and it casts a beam that's bright even in daylight.

The roughly 9.5-watt modules that give the flashlight its power come from Cree, in Durham, N.C., one of several companies competing in the high-intensity LED market. Though residential fixtures are also available for renovations and new construction, for the most part today's superbright LED lights go to commercial and institutional customers, for which the high cost of a single module—and the need to revamp existing lighting fixtures—is less important than LEDs' energy savings and fantastically long life.

Going out with a FireSword-IV at night is an experience that's not easy to describe. Streetlights and car headlights cast relatively wan beams in comparison. The only drawback to that much light right in your hand is that nearby objects may be so brightly illuminated that anything past 20 meters or so starts looking pretty dark.

There's another, subtler drawback that doesn't strike you until you've been using the light for 15 or 20 minutes: Some of those watts coursing through the LEDs are bathing your surroundings in visible photons, but plenty of them are still turning into heat. Even with those cooling fins, the head of the flashlight and the forward sections of the handle become increasingly warm to the touch. If you stroll long enough, a FireSword-IV could well become too hot to handle.

Nevertheless, the device is a triumph of lighting machismo and amply makes the point that LED illumination can, under certain conditions, vanquish incandescent or fluorescent lights [see "The LED's Dark Secret," elsewhere in this issue]. And for someone who regularly needs enormous light output in a tiny, cordless package, slightly dimmer Maglite replacement heads are available that might end up costing about the same as less outré rechargeable work lights. —Paul Wallich

LEFT: MARK WEISS; RIGHT: PAUL WALLICH

books

World Power

How Yankee ingenuity harnessed the future

The United States rose to global power on the strength of its technology, and the lifeblood of that technology has long been electricity. By providing long-distance communication and energy, electricity created the modern world, and it did so in a single lifetime. Yet properly understood, the age of electricity is merely the second phase in the age of steam, which began a century earlier.

"It is curious that no one has put together a history of both the steam and electric revolutions," writes Maury Klein in his preface to *The Power Makers: Steam, Electricity, and the Men Who Invented Modern America*. Klein, a noted historian of technology, spins a narrative so lively that at times it reads like a novel. *Power Makers* is chiefly about the men of the golden age of invention. Their names are a roll call from IEEE's hall of fame: Watt, Morse, Bell, Edison, Westinghouse, Tesla, and others. Equally important, though, are the industrialists—such as J.P. Morgan, Charles Coffin, and Samuel Insull—who turned groundbreaking discoveries into vast commercial enterprises.

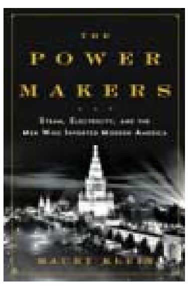

The Power Makers: Steam, Electricity, and the Men Who Invented Modern America
By Maury Klein; Bloomsbury Press, 2008; 560 pp.; US $29.99; ISBN: 978-1-59691-412-4

The story begins in Scotland, where Watt perfected "the machine that changed the world." Klein writes, "America did not invent the steam engine, but once they grasped its possibilities they embraced it eagerly and put it to more uses than anyone else."

Meanwhile, over the course of the 19th century, electricity went from quaint curiosity to indispensable commodity. Morse devised a code for sending messages over an electromagnetic circuit. Bell then gave the telegraph a voice. Edison perfected an incandescent bulb that brought electric light into the American home.

Critically, Edison realized that success depended on mass electrification, which he demonstrated in New York City. With help from the Serbian émigré Tesla, Westinghouse's firm developed a rival system using alternating current, which soon became the dominant form of delivery.

What moved the United States to industrial preeminence, however, were the businessmen who quickly provided the acumen and capital to turn private power facilities into public utilities and to extend local power-delivery networks into massive distribution grids. Morgan's backing helped create the General Electric Co., which Coffin forged into a conglomerate to rival Westinghouse's own. Insull, who had long been Edison's chief lieutenant, left for the greener pastures of the Midwest, founding the Commonwealth Edison Co. in Chicago and expanding it into rural areas, creating a model for the electrification of all America.

To frame his epic story, Klein devises the character of Ned, a fictional witness to the progress brought about by the steam and electric revolutions in America during one man's lifetime. It's a pleasing contrivance that helps turn a long narrative into an engrossing one.

—Kieron Murphy

tools & toys

World's Coolest Personal Computer

By dunking the guts of its computers in liquid, Hardcore Computer, in Rochester, Minn., says it can dampen (yes, that's the word it uses) temperature spikes, raising both reliability and performance through "the fastest sustainable overclocking available." Overclocked video cards and other components run at roughly half their normal temperatures.

Cofounder Chad Attlesey says that the speed of its flagship computer, the Reactor, is matched by hobbyists using liquid nitrogen but that their systems can't run with overclocking for long periods of time. "The original Hardcore prototype has been running error free in my office for over four years."

Hardcore's proprietary polymer-based Core Coolant doesn't conduct electricity. Therefore it can safely bathe the entire Reactor core without risking short circuits or corrosion.

Powerful circulating-style pumps from Laing Thermotech (better known for its spa pumps) and thin-fin skived copper heat sinks dissipate heat via a radiator with four 120-millimeter fans on its plenum. For rack systems, the pumps convey the warm fluid so waste heat can be dissipated or put to work.

The company is initially targeting thick-wallet gamers, but Attlesey says that its overclocking can still make the Reactor a bargain—despite prices that start at US $3000 and top out at $11 000—for medical imaging or military simulation. "Analyses that take 24 hours, we've been able to run in 2," he reports.

—Daniel P. Dern

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

invention

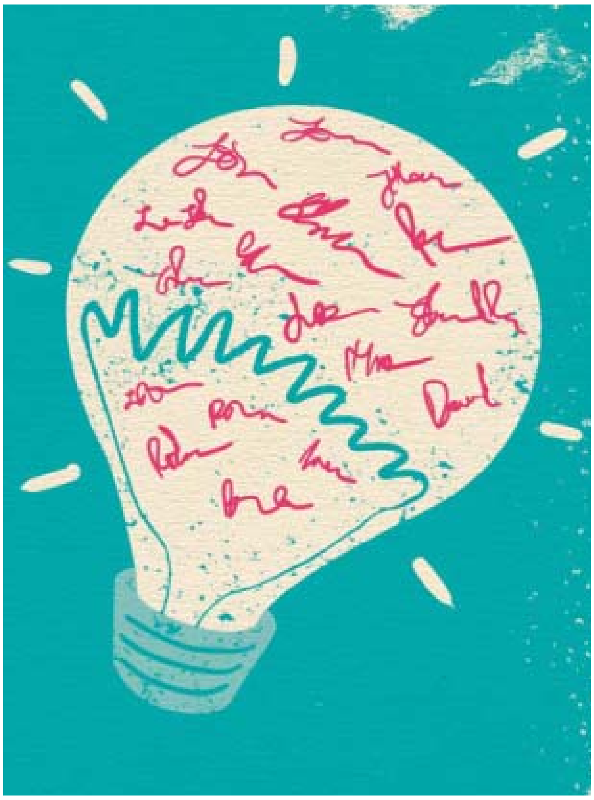

EDWARD McGOWAN

WHAT MAKES YOU AN INVENTOR?

A new patent case raises the question of whether a contribution is "significant"

PATENT APPLICANTS often name someone as an inventor who isn't—or they may fail to name someone who is. Some companies, as a matter of policy, list the CEO as an inventor in all their patents "out of respect." Sometimes colleagues are named because they will be rewarded monetarily for patents where they are listed as inventors. And conversely, actual inventors may be omitted, either accidentally or on purpose. When that happens, it can be hard to enforce patent rights.

Consider this scenario: Company A has a patent and sues Company B for infringing it. Company B then discovers that X should have been listed as an inventor, so B pays X for his rights in the patent. Now B owns the patent jointly with A and cannot be sued for infringement.

So who counts as a coinventor? That's the question in *Nartron Corp. v. Borg Indak*, a case recently decided by the U.S. Court of Appeals for the Federal Circuit, which handles all U.S. patent appeals. Nartron holds a patent for a car seat with lumbar support, which it had designed for another company, Schukra U.S.A. Nartron sued a third firm, Borg Indak, for patent infringement.

In its defense, Borg argued that a Schukra employee had suggested an "extender" component for the lumbar support adjuster, that the extender had appeared in one of Nartron's patent claims, and that the Schukra employee had not been listed as an inventor. Borg noted that the law required all owners of a patent to join in any infringement suit, that Schukra's employee was such a co-owner, and that the employee had in fact not joined in the suit. Therefore, Borg argued, it could not be sued for patent infringement.

The Nartron case turned on whether the Schukra employee was an inventor, defined as anyone who contributes to any claim in a patent application. What "contributes" means, though, is not always clear.

Last year, the original trial court agreed with Borg that the Schukra employee was a coinventor and dismissed Nartron's lawsuit. In March of this year, the Federal Circuit reversed the decision after finding the Schukra employee's contribution was "insignificant." Extenders were well known, the higher court found, and including the extender as part of the claimed invention was "merely the exercise of ordinary skill in the art." Over the course of three different cases, the court has stated:

> One who simply provides the inventor with well-known principles or explains the state of the art without ever having a firm and definite idea of the claimed invention as a whole does not qualify as a joint inventor. A person will not be a coinventor if he or she does no more than explain to the real inventor concepts that are well-known in the current state of the art. Moreover, a joint inventor must contribute in some significant manner to the conception or reduction to practice of the invention and make a contribution to the claimed invention that is not insignificant in quality when that contribution is measured against the dimension of the full invention.

According to the Federal Circuit, then, coming up with an idea claimed in a patent application doesn't necessarily make you an inventor. Unfortunately, such pronouncements are mere guidelines, and there is still no litmus test for deciding who is and who isn't an inventor.

There are a couple of saving graces. First, the list of inventors given in a patent is presumed by law to be correct. So anyone who alleges that a patent should have named someone as an inventor must provide clear and convincing evidence. In many cases, such evidence is hard to come by. One interesting feature of the Nartron case is that everyone agreed that the Schukra employee alone did in fact contribute the idea of the extender. The only point of contention was whether the contribution was significant. Expect to see more patent battles over inventorship, then, turning on what it means to be "significant." Fortunately, if mistakes are made, credit for inventorship can be corrected even after the patent is issued. —KIRK TESKA

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

technically speaking BY PAUL McFEDRIES

Brave Neuro World

We live in an information society. What's the next form of human society? The neuro-society.
—Zack Lynch & Byron Laursen, *The Neuro Revolution: How Brain Science Is Changing Our World*

JULIEN PACAUD/COLAGENE

IN APRIL 2009, the journal *Nature* published the results of a poll that asked people whether they were using beta blockers and drugs like Ritalin—not for their original medical purposes but to boost their brain power. Of the 1400 people from 60 countries who responded, one in five—an eyebrow-raising proportion—reported they had done so. About 80 percent said that healthy adults should be able to take such drugs for nonmedical purposes.

These surprising results touched off a flurry of off- and online harrumphing and tut-tutting, which one scientist dismissed as mere **neurogossip**. The mental, physical, legal, and ethical pros and cons have, of course, been well debated over the past year or so, but all this talk about **brain boosting** has also generated a tidy collection of new words and phrases (such as *brain boosting*). The use of pharmaceuticals to enhance memory, focus, and mental stamina in healthy brains is known generally as **cognitive enhancement**; the pharmaceuticals themselves are often called **cognitive enhancers**.

As you can imagine, the *neuro-* prefix gets quite a workout in these circles, with the equivalent terms being **neuroenhancement** and **neuroenhancer**, a word that rhymes with *Neuromancer*, the title of a seminal 1984 book by cyberpunk novelist William Gibson. The term **neuro** made the leap from prefix to adjective recently with the publication last month of *The Neuro Revolution*, coauthored by Zack Lynch, who also coined the phrase **neurosociety**.

In a letter to *Nature* published in December 2007 (under the terrific title "Professor's Little Helper," a clever shout-out to "Mother's Little Helper," the Rolling Stones' paean to housewives on prescription drugs), the scientists Barbara Sahakian and Sharon Morein-Zamir wrote, "The drive for self-enhancement of cognition is likely to be as strong if not stronger than in the realms of 'enhancement' of beauty and sexual function." In other words, forget cosmetic surgery; the next fad is likely to be **cosmetic neurology**.

The "off label" (that is, outside its original scope) use of a drug such as methylphenidate (aka Ritalin), which is normally prescribed to treat attention-deficit hyperactivity disorder (aka ADHD), is sometimes called **brain customization**, but these days it is more likely to be referred to as **mind hacking**. In this kind of neuroenhancing, the drugs are usually administered in the form of **smart pills**, and the effect they create has been called **smart-in-a-pill**. Mind hackers use these drugs, marketed euphemistically as "study aids," to increase their **mental horsepower**.

Some people want their brains neuroenhanced because they find themselves forgetting things or aren't quite as sharp as they used to be. They don't necessarily want **supercognition**; they just want their brains to be normal again. However, the **neurodiversity** movement is based on the belief that there is no such thing as "normal" when it comes to the human mental landscape—the **neurotypical** person simply does not exist. People display a wide variety of neurological behaviors and abilities, and most of us exhibit some form of mental "disorder" from time to time, albeit in nondebilitating—or subclinical—form: mild depression, temporary anxiety, and so on. We accept that the world is populated with people who are tall and small; big boned and bird boned; ecto-, meso-, and endomorphic; and, the theory goes, also diverse when it comes to neurological traits like forgetfulness. Are we so concerned with eternal youthfulness (**perma-youth** in the vernacular) that we need to turn to **smart drugs** to achieve at least the mental side of that goal?

Other naysayers reject these pharmacological tricks as cheating. They deride them as **academic steroids** and the practice as **brain doping**. Rather than welcome an ever more efficient and productive neurosociety, they wring their hands at the prospect of a stressed-out, always-on, 24/7 society. And **neuroethicists** worry that the goal of **better brains**, while worthy in itself, may create social disparities if only the well-off can afford nootropic (literally "mind affecting") drugs. Who's right? I guess I'll go have another espresso and think about it. ❑

THE LED'S DARK SECRET

Solid-state lighting won't supplant the lightbulb until it can overcome the mysterious malady known as "droop"

BY RICHARD STEVENSON
ILLUSTRATIONS BY BRYAN CHRISTIE DESIGN

THE BLUE LIGHT-EMITTING DIODE, arguably the greatest optoelectronic advance of the past 25 years, harbors a dark secret: Crank up the current and its efficiencies will plummet. The problem is known as droop, and it's not only puzzling the brightest minds in the field, it's also threatening the future of the electric lighting industry.

Tech visionaries have promised us a bright new world where cool and efficient white LEDs, based on blue ones, will replace the wasteful little heaters known as incandescent lightbulbs. More than a dozen countries have already enacted legislation that bans, or will soon ban, incandescent bulbs. But it's hard to imagine LEDs dislodging incandescents and coming to dominate the world electric lighting industry, unless we can defeat droop.

In flashlights, in backlights for screens in cellphones and now televisions, and in a bunch of other applications, white LEDs already constitute a multibillion-dollar market. But that's just a US $5 billion niche compared to the overall lighting industry, whose sales next year should reach $100 billion, according to the market research firm Global Industry Analysts. The trick will be to make LEDs turn electricity into light efficiently enough to offset their relatively high cost—roughly 16 cents per lumen, at lightbulb-type brightness, as opposed to about 0.1 cents or less for incandescents.

Look at the competition and you'd think the job was easy. Today's garden-variety incandescent bulbs aren't much different from the ones Thomas Edison sold more than a century ago. They still waste 90 percent of their power, delivering roughly 16 lumens per watt. Fluorescent tubes do a lot better, at more than 100 lm/W, but even they pale next to the best LEDs. The current state-of-the-art white LED pumps out around 250 lm/W, and there's no reason why that figure won't reach 300 lm/W.

Unfortunately, these LEDs perform at their best only at low power—the few milliamps it takes to backlight the little screen on your mobile phone, for instance. At the current levels needed for general lighting, droop kicks in, and down you go, below 100 lm/W.

THE FIRST-EVER REPORT OF light emission from a semiconductor was by the British radio engineer Henry Joseph Round, who noted a yellowish glow emanating from silicon carbide in 1907. However, the first devices at all similar to today's LEDs arrived only in the 1950s, at Signal Corps Engineering Laboratories, at Fort Monmouth, in New Jersey. Researchers there fabricated orange-emitting devices;

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

green, red, and yellow equivalents followed in the '60s and '70s, all of them quite inefficient.

The great leap toward general lighting came in the mid-1990s, when Shuji Nakamura, then at Nichia Corp., in Tokushima, Japan, developed the first practical bright-blue LED using nitride-based compound semiconductors. (Nakamura's achievement won him the 2006 Millennium Technology Prize, the approximate equivalent in engineering of a Nobel Prize.) Once you've got blue light, you can get white by passing the blue rays through a yellow phosphor. The phosphor absorbs some of the blue and reradiates it as yellow; the combination of blue and yellow makes white.

All LEDs are fabricated as aggregated sections, or regions, of different semiconductor materials. Each of these regions plays a specific role. One region serves as a source of electrons; it consists of a crystal of a compound semiconductor into which tiny amounts of an impurity, such as silicon, have been introduced. Each such atom of impurity, or dopant, has four electrons in its outer shell, compared with the three in an atom of gallium, aluminum, or indium. When a dopant takes a place that one of these other atoms would normally occupy, it adds an electron to the crystalline lattice. The extra electron moves easily though the crystal, acting as a carrier of negative charge. With this surfeit of negative charges, such a material is called *n*-type.

At the opposite end of the LED is a region of *p*-type material, so called because it has excess positive-charge carriers, created by doping with an element such as zinc or magnesium. These metals are made up of atoms with only two electrons in their outer shell. When such an atom sits in place of an atom of aluminum, gallium, or a chemically similar element (from group III in the periodic table), the lattice ends up an electron short. That vacancy behaves as a positive charge, moving throughout the crystal like the missing tile in a sort-the-number puzzle. That mobile vacancy is called a hole.

In the middle of the sandwich are several extraordinarily thin layers. These constitute the active region, where light is produced. Some layers made of one semiconducting material surround a central layer made of another, creating a "well" just a few atoms thick—a trench so confined that the laws of quantum mechanics rule supreme. When you inject electrons and holes into the well by applying a voltage to the *n*- and *p*-type regions, the two kinds of charge carriers will be trapped, maximizing the likelihood that they will recombine. When they do, a photon pops out.

To make an LED, you must grow a series of highly defined semiconductor layers on a thin wafer of a crystalline material, called a substrate. The substrate for red, orange, and yellow LEDs is gallium arsenide, which works wonderfully because its atoms are spaced out identically to those of the layers built on top of it. Hardly any mechanical strain develops in the semiconductor's crystalline lattice during fabrication, so there are very few defects, which would quench light generation.

Unfortunately, blue and green LEDs lack such a good platform. They're called nitride LEDs because their fundamental semiconductor is gallium nitride. The *n*-type gallium nitride is doped with silicon, the *p*-type with magnesium. The quantum wells in between are gallium indium nitride. To alter the light color emitted from green to violet, researchers vary the gallium-to-indium ratio in the quantum wells. A little indium produces a violet LED; a little more of it produces green.

Such LEDs would ideally be manufactured on gallium nitride substrates. But it has proved impossible to grow the large, perfect crystals of gallium nitride that would be necessary to make such wafers. Unipress, of Warsaw, the world leader in this field, cannot make crystals bigger than a few centimeters, and then only by keeping the growth chamber at a temperature of 2200 °C and a pressure of almost 20 000 atmospheres.

So the makers of blue LEDs instead typically build their devices on wafers of sapphire, whose crystalline structure does not quite match that of the nitrides. And that discrepancy gives rise to many defects—billions of them per square centimeter.

It is amazing that such LEDs work at all. Any arsenide-based red, orange, or yellow LED that contained as many defects would emit absolutely no light. To this day, researchers, including Nakamura himself—who moved to the University of California, Santa Barbara (UCSB) in 1999—can't agree on the cause of the phenomenon. Perhaps the solution to this problem may also explain droop.

THE EXPLANATION WON'T COME EASILY. When researchers set out to find the cause of droop in nitride LEDs, one of their first suspects was heat, which they knew could cause droop in arsenide LEDs. There, heat imparts so much energy to the electrons and holes that the quantum well can no longer trap them. Instead of recombining, some of them escape, only to be swept away by the electric fields in the device. But researchers dismissed this possibility after noting that nitride LEDs suffered from droop even when driven by short, pulsed voltages spaced far enough apart to let the devices cool down.

Another theory was proposed as far back as 1996 by Nakamura. He argued that everything could be explained by the structure of the quantum well. Nakamura and his colleagues looked at LEDs with a transmission electron microscope and were surprised to find light and dark areas within the quantum well, suggesting that the material there was not uniform. They then investigated the crystalline structure more closely, using X-ray diffraction, and found that the quantum well had indium-rich clusters (bright) next to indium-poor areas (dark).

Nakamura conjectured that because the indium clusters were free from defects, the electrons and holes would be trapped in them, making bright emission possible, at least at low currents. Continuing with this line of reasoning, Nakamura's team argued that LEDs' high efficiency at low currents stemmed from a very high proportion of electron-hole recombination in defect-free clusters. At higher currents, however, these clusters would become saturated, and any additional charge carriers would spill over into regions having defects dense enough to kill light emission. The saturation at high current, they suggested, accounted for the observed droop.

This theory has fallen out of favor in recent years. "To start with, we saw indium-rich clusters in InGaN quantum wells, just as the rest of the world did," explains Colin Humphreys, the head of the Cambridge Centre for Gallium Nitride at the University of Cambridge, in England. But then he and his team began to suspect that their electron microscope was causing the very thing it was detecting. So the group carried out low-dose electron microscopy. "We looked at the first few frames—a very low exposure—and saw no indium clustering at all. But as we exposed the material to the beam, these clusters developed," he says. They concluded that the clustering was merely an artifact of measurement.

LED Architecture

AT THE HEART of every white LED is a semiconductor chip made from nitride-based materials. The chip is traditionally positioned on top of the cathode lead. Applying several volts across this device makes the chip emit blue light. Passing the light through a yellow phosphor yields white light. Modern, high-power LEDs are variants of this architecture, featuring more complex packages for superior thermal management.

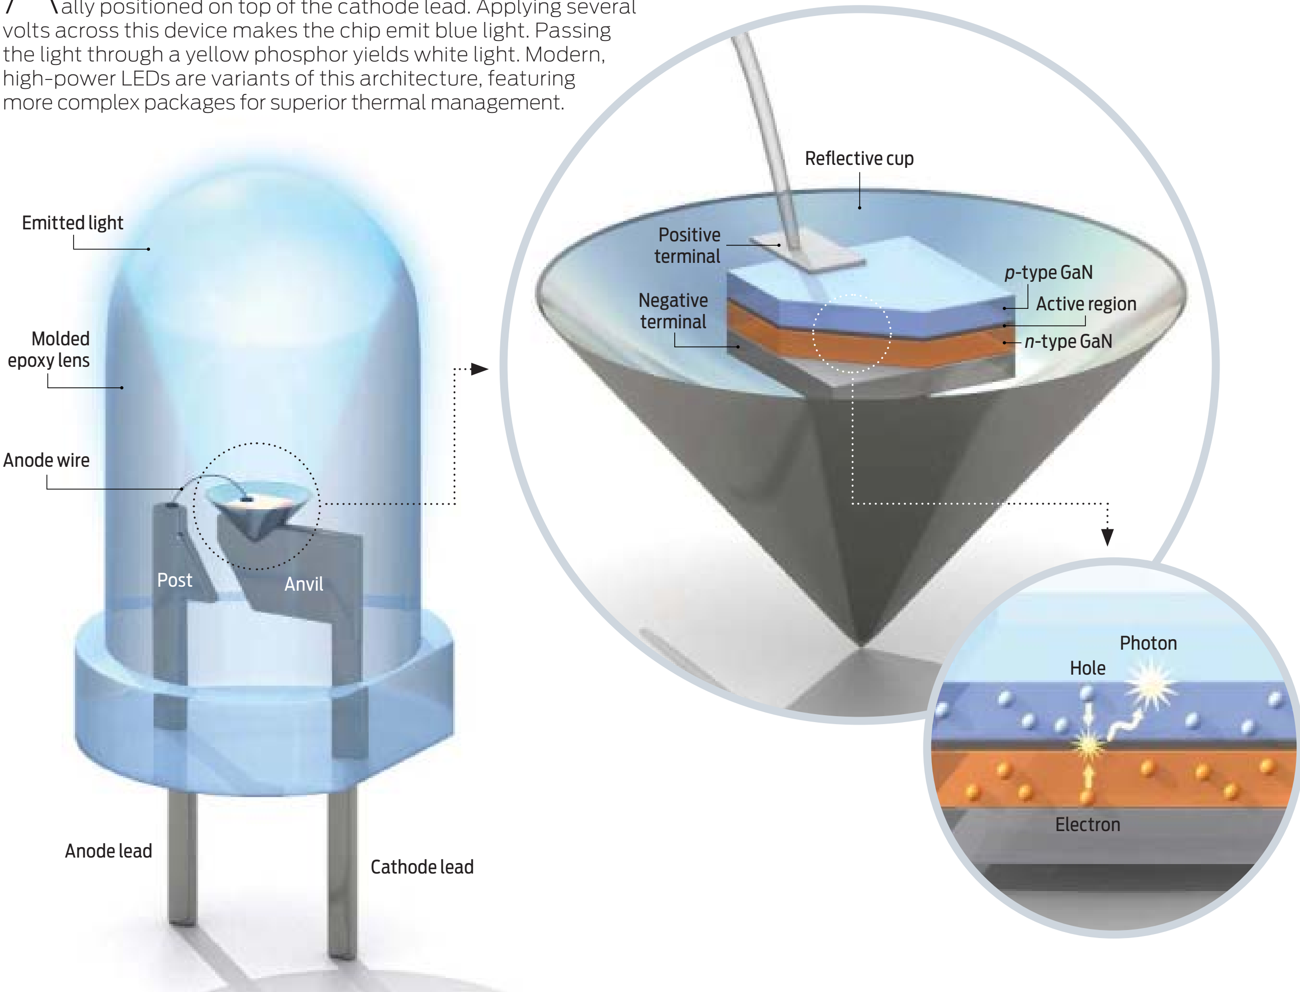

Combatting Droop

DROOP—the loss of efficiency at high power—afflicts conventional nitride LED structures. These feature an active region with gallium indium nitride quantum wells and GaN barriers, and an electron-blocking layer to keep electrons in this region. Researchers at Rensselaer Polytechnic Institute have reduced droop with new active regions, made first by combining GaInN wells and aluminum gallium indium nitride barriers and, more recently, by pairing GaInN wells with GaInN barriers. Meanwhile, Philips Lumileds has also developed a structure that is less prone to droop, thanks to a far thicker quantum well.

Conventional structure

Rensselaer Polytechnic Institute's polarization-matched structure

Lumileds' double heterostructure

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

IN 2003, HUMPHREYS PRESENTED that jaw-dropping finding at the Fifth International Conference on Nitride Semiconductors, in Nara, Japan. It wasn't well received. Many delegates contended that something must have gone wrong with the Cambridge samples. So Humphreys's group went back and studied a wider variety of specimens, including LEDs supplied by Nichia. Their work only reinforced their view that the clusters were formed by electron-beam damage.

In 2007, Humphreys's Cambridge team, together with researchers at the University of Oxford, described how they had attacked the problem with what's known as a three-dimensional atom probe. This device applies a high voltage that evaporates atoms on a surface, then sends them individually through a mass spectroscope, which identifies each one by its charge-to-mass ratio. By evaporating one layer after the other and putting all the data together, you can render a 3-D image of the surface with atomic precision.

The resulting images confirmed, again, what the electron microscope had shown: There is no clustering. Discrediting the cluster theory was an important step, even though it left the research community without an alternative explanation for droop.

Then, on 13 February 2007, the California-based LED manufacturing giant Philips Lumileds Lighting Co. made the stunning claim that it had "fundamentally solved" the problem of droop. It even said that it would soon include its droop-abating technology in samples of its flagship Luxeon LEDs.

Lumileds kept the cause of droop under wraps for several months. Then, at the meeting of the International Conference of Nitride Semiconductors, held September 2007 in Las Vegas, it presented a paper putting the blame on Auger recombination—a process, named after the 20th-century French physicist Pierre-Victor Auger, that involves the interaction of an electron and a hole with another carrier, all *without* the emission of light.

The idea was pretty radical, and it has had a mixed reception. *Applied Physics Letters* published Lumileds' paper only after repeated rejections and revisions. "In my experience, it was one of the most difficult papers to get out there," says Mike Krames, director of the company's Advanced Laboratories.

KRAMES'S TEAM USED A LASER to probe a layer of gallium indium nitride, the semiconductor used for quantum wells in a nitride LED. They tuned the laser to a wavelength that only the gallium indium nitride layer would absorb, so that each zap created pairs of electrons and holes that then recombined to produce photons. When the researchers graphed the resulting photoluminescence against different intensities impinging on the sample, they produced curves that closely fit an equation that described the effects of Auger recombination.

The bad news is that you can't eliminate this kind of recombination, which is proportional to the cube of the density of carriers. So in a nutshell, if you've got carriers—which of course you need to generate light—you've also got Auger recombination. The good news, though, is that Lumileds has shown that you can push the peak of your efficiency to far higher currents by cutting carrier density—that is, by spreading the carriers over more material. The company does so with what's known as a double heterostructure (DH), essentially a quantum well that's 13 nanometers wide, rather than the usual 3 or 4 nm. It still shows quantum effects, although they are not so pronounced, and the design is less efficient than the standard one at low currents. Still, it excels at higher currents. The Lumileds team has created a test version that delivers a peak efficiency slightly higher than that of a conventional LED.

Promising though this new crystalline structure may be, it is difficult to grow. Perhaps this is why Lumileds has yet to incorporate the design into its Luxeon LEDs. "There are multiple paths to dealing with droop, and we've investigated most of these paths," says Krames. "We have new structures in the pipeline, DH as well as non-DH, and we will move forward with the best structure."

NOT EVERYONE IS CONVINCED that Auger recombination is the cause of droop. One such skeptic is Jörg Hader, a University of Arizona theorist, who works with former colleagues in Germany at Philipps-Universität Marburg and at one of the world's biggest LED manufacturers, Osram Opto Semiconductors, in Regensburg.

"All [Lumileds] showed was that they can fit the results with a dependence that is like Auger," claims Hader. "It's a fairly weak argument to see a fit that fits, and see what might correspond to that fitting." In his view, there's a good chance that the Lumileds data could also be fitted with other density dependencies, as well as the cubed dependence that is classically associated with Auger recombination.

Less Leakage

POLARIZATION FIELDS may cause LED droop. Such fields are claimed to drive electrons out of the active region and into the *p*-type layer, where some recombine without emitting light [top]. A "polarization matched" structure [bottom] has a far weaker internal field and therefore suffers less electron leakage, leaving more electrons to recombine with holes.

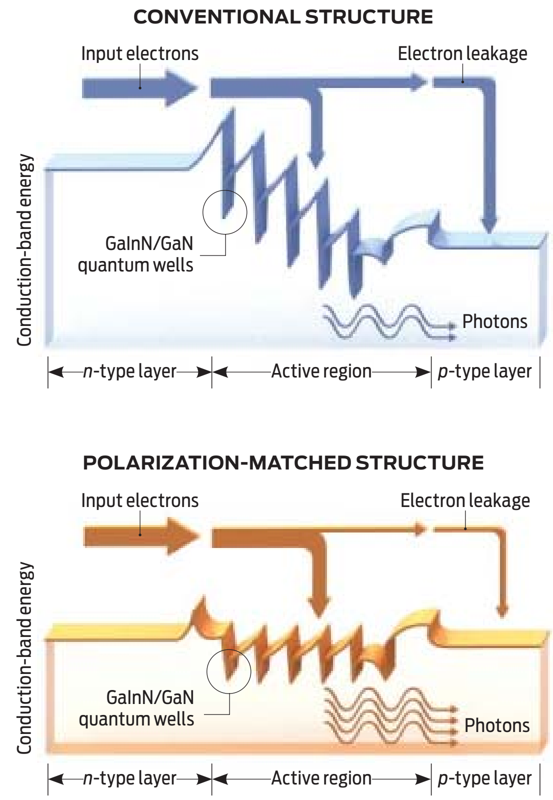

Hader has calculated the magnitude of direct Auger recombination for a typical blue LED. The equations that describe this interaction of an electron and a hole with a third carrier date back to the 1950s, but that doesn't mean that they are easy to solve. Hader says he took no shortcuts. Instead, he accounted for all physical interactions in a program tens of thousands of lines long, a program that in its initial form would have taken several years to run. However, Hader says he's learned what he can omit safely in order to get the running time down to just 1 minute. He says the model shows that Auger losses are too small to account for LED droop, although he does allow that droop might be caused by other processes related to Auger recombination. These processors are more complicated because they also involve defects in the material or thermal vibrations (phonons, in quantum terms) of the semiconductor crystal.

Krames criticizes Hader's calculations for leaving out the possibility that electrons might occupy higher energy levels, known as higher conduction bands. But Hader believes that including these bands would hardly affect his conclusions.

This May, computer scientists at UCSB brought new evidence to bear on this debate. Chris Van de Walle's team included a second conduction band in their calculations of Auger recombination in nitrides and concluded that Auger contributes strongly to droop. However, they modeled only the bulk materials, not realistic quantum wells, for which Van de Walle admits his methods cannot handle the calculations, at least not on today's computers.

Hader does not doubt the general shape of the UCSB results. However, he points out that the value Van de Walle's team has taken for the second conduction band substantially differs from that given in certain academic papers. Using these published values would have profound effects on any estimate of the magnitude of Auger recombination. The conclusions of Hader and Van de Walle highlight the lack of consensus among theorists over the cause of droop.

Meanwhile, a group headed by E. Fred Schubert at the Rensselaer Polytechnic Institute, in Troy, N.Y., has proposed yet another theory. His team, in collaboration with Samsung, blames droop on the leakage of too many electrons from the quantum well.

Interestingly, Schubert's team, like the researchers at Lumileds, drew its conclusions by pumping light into the nitride structures and observing the light that those structures emitted in response. But Schubert and company investigated full LED structures, and they compared the results they'd obtained from optical pumping with light output generated when a voltage was applied, as it is in normal operation. As expected, droop kicked in when the device was pumped electrically. But the researchers saw no sign of droop in the photoluminescence data.

They then brought in Joachim Piprek, a theorist from the NUSOD Institute, a device simulation consultancy in Newark, Del. He used a computer model to simulate the behavior of a blue LED and found that the strong internal fields characteristic of nitrides must be causing electrons to leak out of the wells.

Now Schubert and his colleagues have produced direct evidence to back up their argument for leakage. They took an LED unconnected to any circuit and hit it with light at a wavelength of 405 nm, which is absorbed only in the quantum wells. The researchers detected a voltage across the diode, implying that carriers must leave the wells, contradicting Lumileds' theory.

Schubert's team has tried to control electron leakage by redesigning the LED. By carefully selecting the materials for the active region—switching from the conventional gallium nitride barrier to an aluminum gallium indium nitride version—they have been able to eliminate the charges that tend to form wherever distinct crystalline layers meet. They say such "polarization matching" consistently cuts droop, raising power output by 25 percent at high currents.

Schubert believes that the electrons that leak out of the wells recombine with holes in the *p*-type region. If he could detect this recombination, it would certainly add weight to his explanation. "We've looked for that luminescence," says Schubert, "but we have not seen it." He's not surprised, though, because *p*-type gallium nitride is a very inefficient light emitter, and the LED's surface is nearby, so surface recombination at the top contact is also likely.

However, it is possible to detect electrons in the *p*-type region by modifying the standard LED structure, and researchers at UCSB have done just this. This team, led by Steven DenBaars and Nakamura, did the job of fitting the *p*-type region with an additional quantum well, one that emits light of a color different from that of the main LED. At a workshop in Montreux, Switzerland, in the fall of 2008, the group reported that they had found just this sort of emission.

Although this experiment proved that electrons do flow into the *p*-type region, it can't tell us where they came from. And while Schubert's theory of electron leakage could explain the results, there may well be other things that can also account for them. We can't even rule out Auger recombination as the dominant mechanism, because the proportion of electrons flowing into the *p*-type region is still to be quantified.

EACH THEORY HAS ITS CHAMPIONS. Theoreticians at Philipps-Universität Marburg support Auger recombination, mainly the phonon-assisted form, as the main cause of droop. So does Semiconductor Technology Research, a device-modeling company based in Richmond, Va. Meanwhile, Hadis Morkoç's group at Virginia Commonwealth University seconds Schubert's support of electron leakage, which they attribute to the poor efficiency with which holes are injected into the quantum well.

Confused? Join the club—and realize that this controversy is precisely what you'd expect to find in a field that has suddenly begun to make great progress. Even if we don't have a universally agreed-upon theory to account for droop, we do have a growing arsenal of proven weapons to fight it—Schubert's polarization-matched devices, Lumileds' wide quantum well structures, as well as designs that improve hole injection, among others. Too bad that we still can't agree on how they work.

The industry will move forward. LEDs are just starting to supplant fluorescent as well as incandescent lighting. Someday, in our lifetimes, incandescent filaments will finally stop turning tens of gigawatts into unwanted heat. Smokestacks will spew less carbon into the global greenhouse. And we won't have to get up on stepladders to change burned-out bulbs nearly so often as we do today.

And around that time, when you're reading this magazine by the light of an LED, perhaps the theorists will have watertight explanations for the experimentalists, and we'll know the answer to the burning question that remains: What causes droop? ❑

TO PROBE FURTHER For a list of the papers published by Applied Physics Letters *on LED droop, go to http://www.spectrum.ieee.org/semiconductors/optoelectronics/the-leds-dark-secret.*

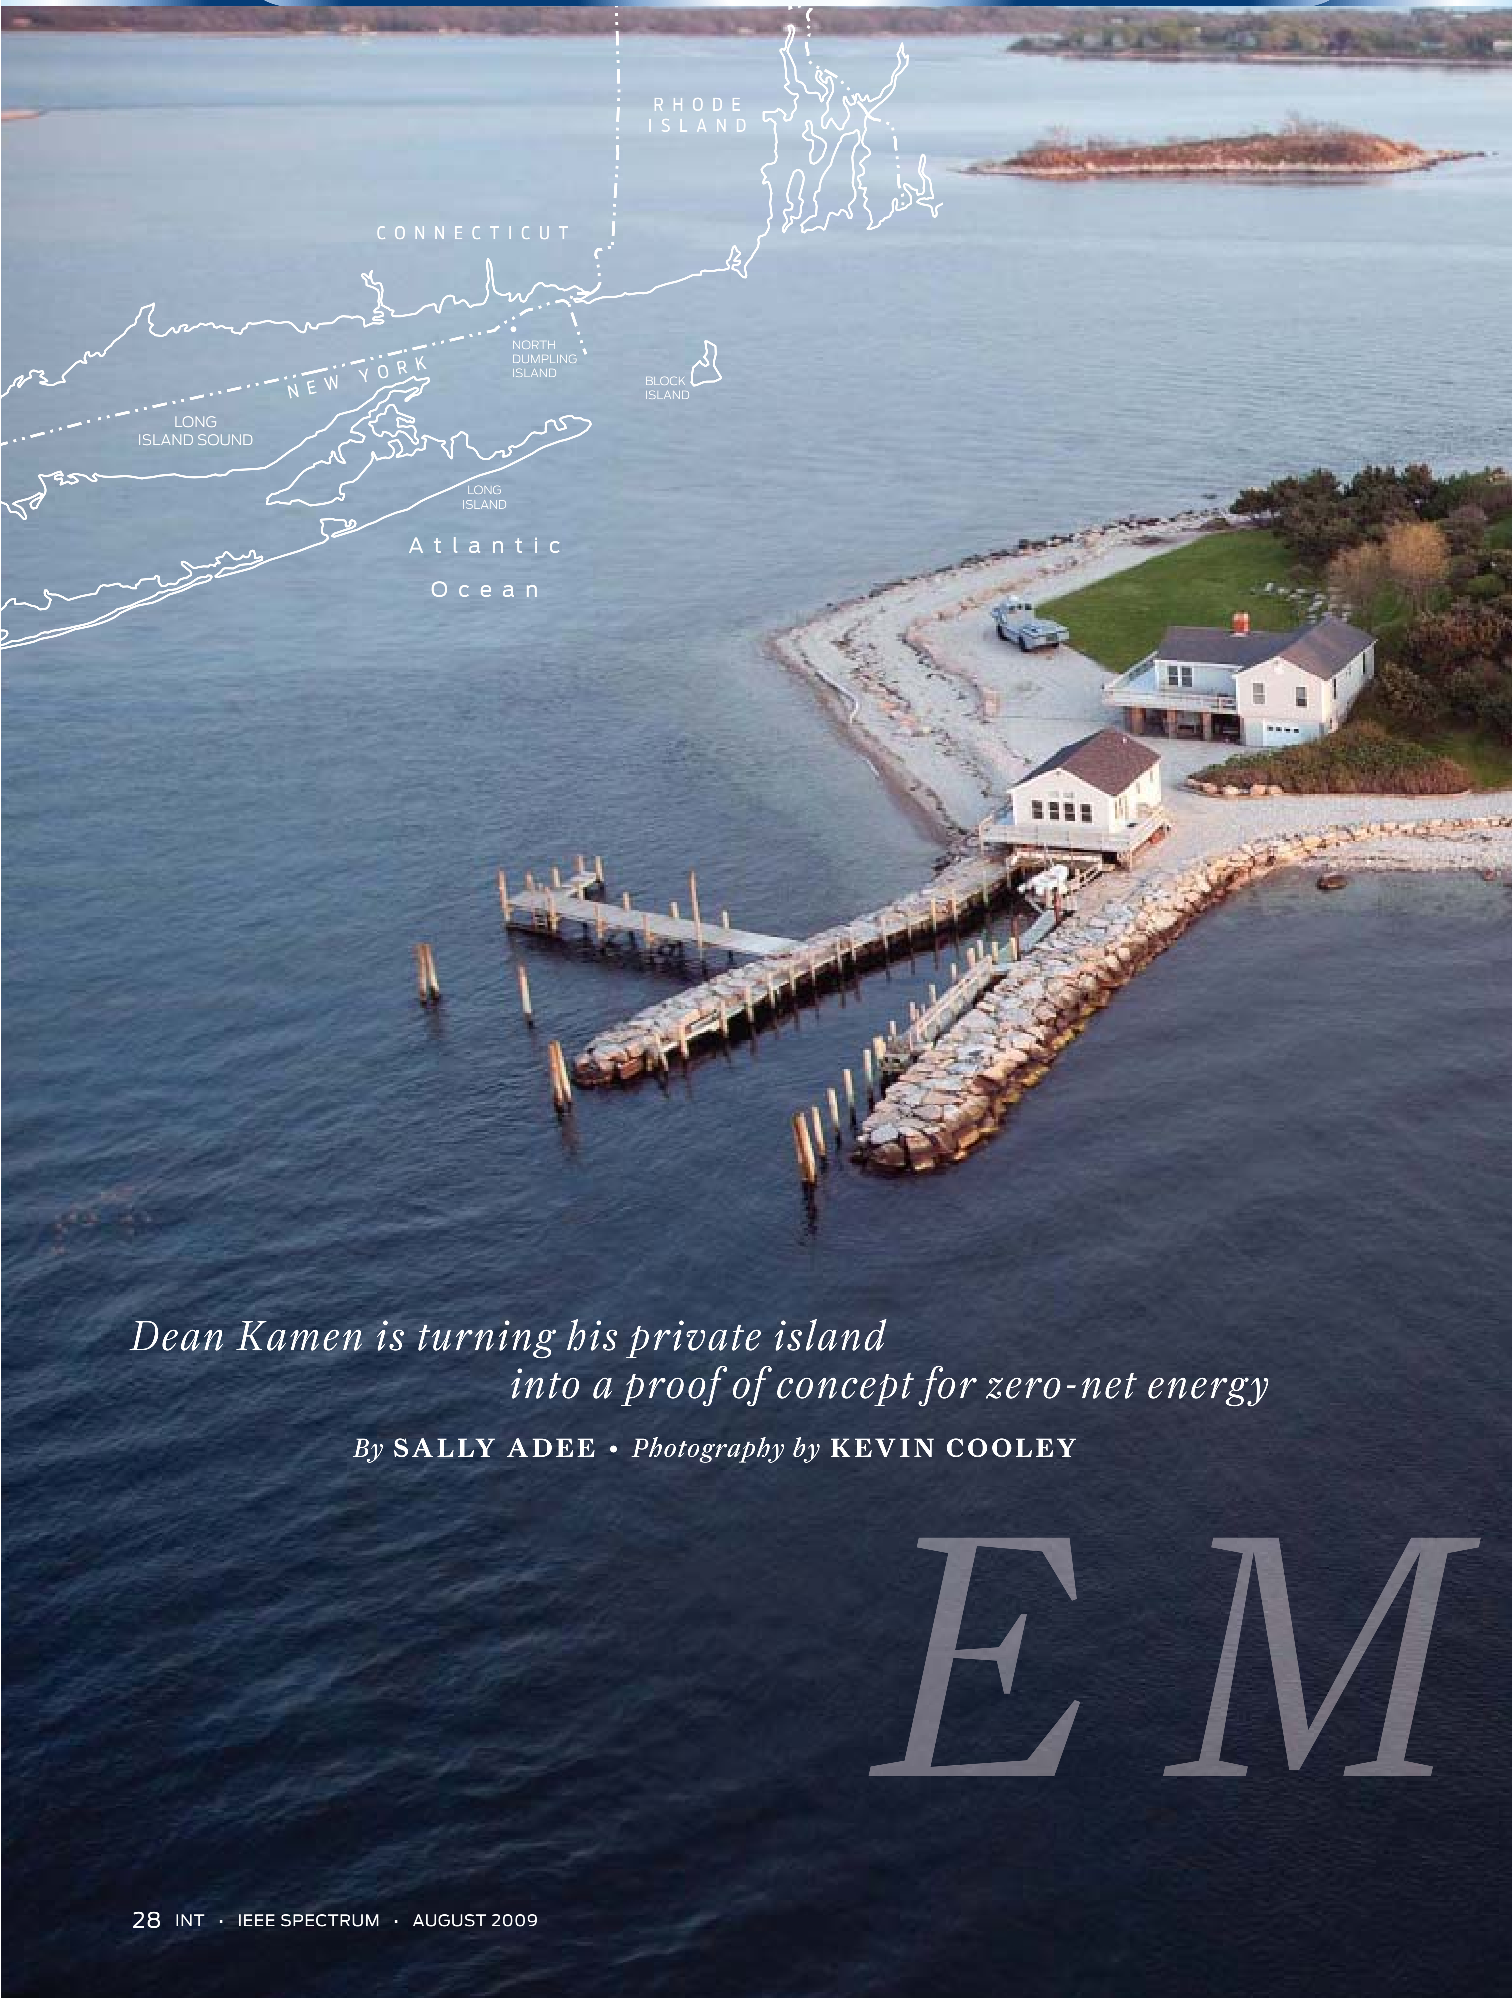

Dean Kamen is turning his private island into a proof of concept for zero-net energy

By **SALLY ADEE** • *Photography by* **KEVIN COOLEY**

EM

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

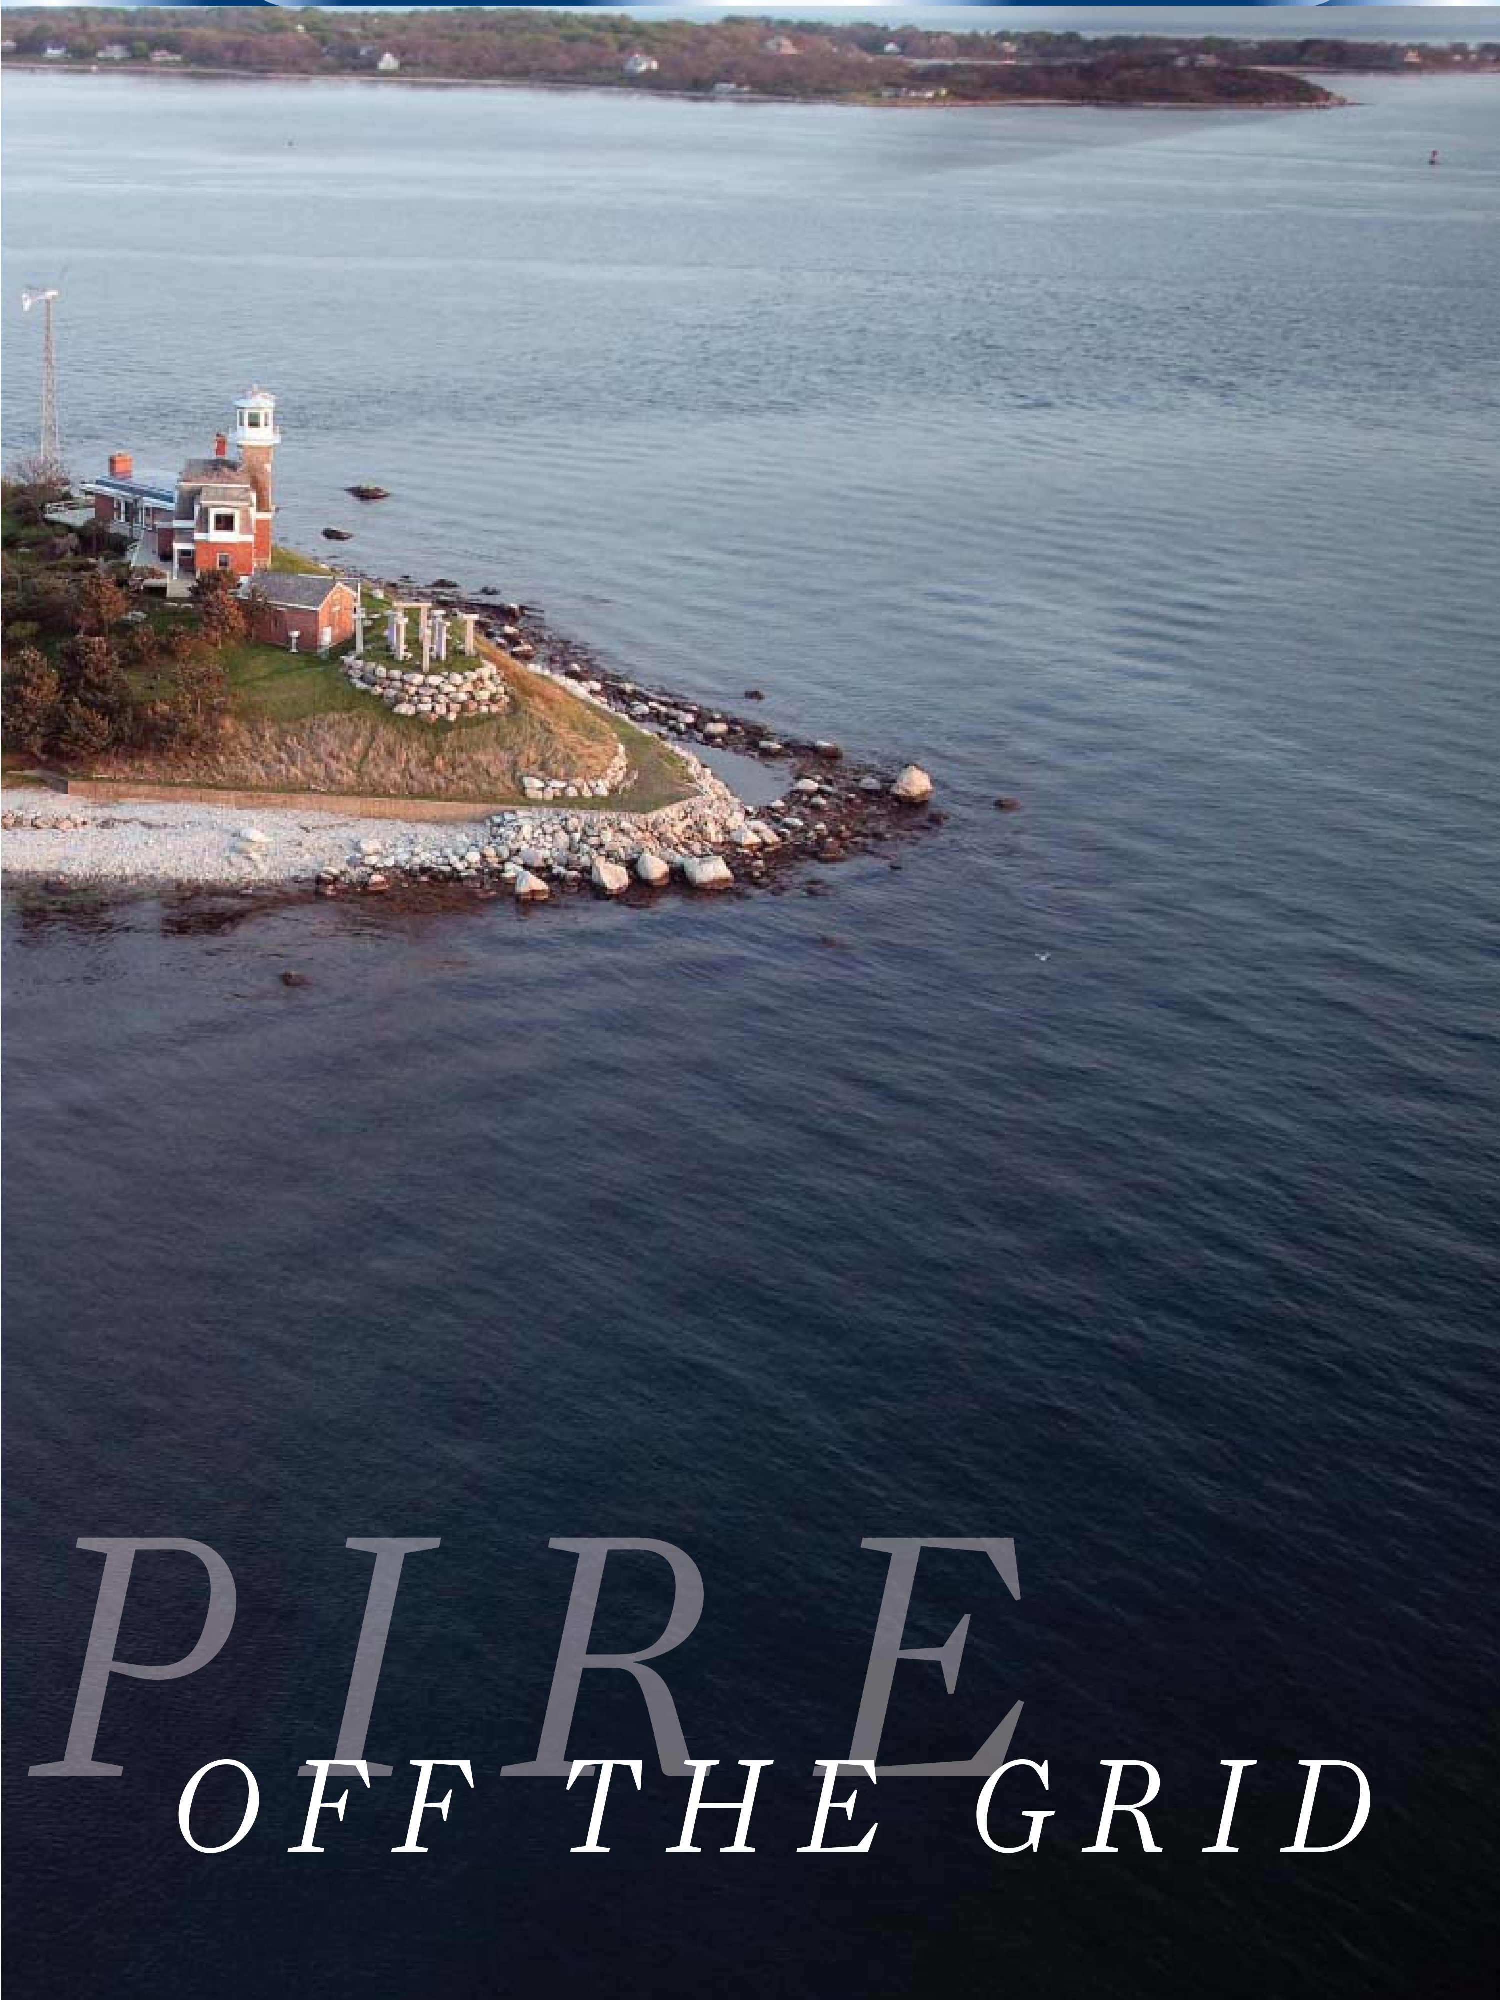

PIRE OFF THE GRID

DEAN KAMEN, arguably the world's most famous living inventor, is in his favorite place. It's the basement of the main house on his private island, a pristine concrete bunker outfitted with a command station that can monitor and control all aspects of the production and consumption of electric power on the island. The island, off the coast of Connecticut, is called North Dumpling Island, but Kamen, who invented the Segway and is something of an eccentric, calls it the Dumplonian Empire. He, of course, is Lord Dumpling.

Hanging in the foyer of the house is a charred, yellowing copy of North Dumpling's ancient constitution, a nearly word-for-word match with the U.S. Constitution. The island has its own currency, called—what else?—the Dumpling, dispensed somewhat inconveniently in units of pi. Though the money and the constitution might be tongue in cheek, Kamen's latest endeavor is anything but: North Dumpling is becoming the world's first energy-independent nation.

Kamen seceded from the electrical grid last October, when the U.S. Coast Guard, which operates a lighthouse on Kamen's island, informed him that it was cutting the power to the submarine cable feeding his island. The Guard would instead power the lighthouse with solar panels, with little electricity to spare for the rest of the island. So Kamen did what any self-respecting multimillionaire engineer would do. He went off the grid.

"Off the grid" is a term more likely to come up in connection with grizzled hippies or disturbed militiamen than with famous inventors. But Kamen is determined to use his North Dumpling empire to show that zero-net energy is not only possible now but that it will be very appealing in the future. He has resolved to live off the renewables on the island: a wind turbine and three monster arrays of solar panels, plus Kamen's beefed-up version of the Stirling engine. It all adds up to a peak generation capacity of just under 25 kilowatts. To coordinate the different sources, Kamen designed an intelligent system that knows, down to an individual solar panel or light source, how much energy is being produced and consumed on the island and, based on that information, negotiates the relationship between the two in real time. It is this system that he monitors and controls from his basement command center.

LED ISLAND: No filaments allowed.

Kamen has many strategies to minimize his energy use, but the big breakthrough is the island's lighting. With the help of his friend Fritz Morgan, chief technology officer at Philips Color Kinetics, in Burlington, Mass., Kamen has replaced every bulb and fixture on his island with light-emitting diodes. That switch cut the power he needed to light his lair by 50 to 70 percent.

LEDs last about 30 times as long as incandescent bulbs and three to four times as long as compact fluorescents. And they're gorgeous. They wash the island in colors that cycle through the rainbow (and they can be set to a disco beat, but that's another story). Just as stunning is the system's engineering marvel that lets Kamen control every circuit on his island empire from anywhere in the world.

KAMEN FLITS by helicopter between North Dumpling and his main residence, a majestic house in Manchester, N.H., near his Deka Research and Development Corp. in a row of refurbished mill buildings. If you can't hitch a ride to the island with Kamen in his helicopter, then you board a boat for the 20-minute ride from the Noank Shipyard, near Mystic, Conn. The boat docks on the northeastern edge of the island, made up of 1.2 hectares of landscaping and smooth pebbles, at an ancient pier of gray-washed New England wood that leads to a boathouse. A few steps from the boathouse, on a rocky beach facing east, is the caretaker's cottage, near a rack of blue-black solar panels. Up a grassy hill and to the west is the stately main house, a shingled, weathered, asymmetrical Colonial that's wrapped around a lighthouse towering above.

The island's four arrays of solar panels have a combined peak output of 12.2 kW. A 10-kW wind turbine that looks like a cartoon missile, made by Bergey Windpower Co., perches on a lattice tower 25 meters above the island. The Stirling engine in the basement can contribute another 2 kW if necessary.

Kamen deploys this patchwork of sources and generators using a system he designed with a team of engineers at Teletrol Systems, another of his companies, also based in New Hampshire. The software acts like a central nervous system, monitoring loads and choreographing the various energy sources to make sure energy consumption on the island never outpaces generation.

Against one spotless corner in the basement is a big con-

KAMEN UNPLUGGED: Dean Kamen sits in the cupola of the 160-year-old lighthouse that started it all.

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

sole made of sleek, pale bamboo—"Renewable!" Kamen crows—smoothly inlaid with a flat touch-screen display. Here he can see, in real time, exactly how much power is being produced by his renewables and how much is being siphoned by every single process on the island, down to the individual LED fixtures.

Along the opposing wall, the custom control electronics for the Teletrol system hum quietly in glinting server racks. Basically, the system establishes priorities among the power drains: At the very top are the necessary functions—lighting LEDs currently in use, heating or purifying water, or cooling the house. Heating is not really an issue, because the island isn't accessible on the coldest days, when the swell in the surrounding seas can reach 9 meters. The house is cooled centrally with ultraefficient central heat exchangers.

After addressing immediate needs, the system diverts any extra power to a bank of batteries. "That way, we don't have to waste any voltage generating energy to charge them," says Kamen.

The formal dining room gleams with the soft yellow light of two flickering candelabras. But we're standing in the next room, around the kitchen counters, eating cold-cut sandwiches on hamburger rolls and washing them down with warm beer. Kamen is telling us about a speech he gave as the keynote speaker at a high-profile event organized by Minneapolis-based Medtronic, the medical-device colossus. Waiting backstage as the master of ceremonies and Medtronic senior VP Stephen Oesterle extolled the many inventions of Benjamin Franklin, "I thought, wait a second, I thought I was the one who was giving this talk," Kamen croaks, his tone in a perpetual tug-of-war between acerbic and affable. And then, as Kamen was ushered onstage to give his speech, he heard, "...and that's why we think Dean Kamen is the reincarnation of Benjamin Franklin!"

"Of course it's all crap," Kamen says. But then, eyebrows rising along with his index finger, he amends himself. "Benjamin Franklin reinvented glasses as bifocals. *I* also took something that had been around forever and reinvented it." That would be the wheelchair. Kamen's version, the iBot, can hop curbs and raise its occupant to eye level with a standing person. Kamen also invented a wearable insulin pump and a prosthetic arm.

After all this it should be no surprise that Kamen wields some influence. Campaigning in New Hampshire in late 2007, Barack Obama made sure to send an emissary to Deka. But Kamen says he talks to politicians about only two things: bringing electricity and clean water to the developing world, and FIRST (For Inspiration and Recognition of Science and Technology), a program he developed that injects science and technology into elementary, middle, and high school classrooms in the United States. The vehicle? Robots. Contenders are divided up into robot-building teams. They face off in elimination rounds throughout the year, culminating in a championship in which a winning team is crowned. FIRST gives technical fields the cachet that was once reserved for athletics.

The FIRSTies worship Kamen. They proclaim their devotion on message boards, where they identify themselves by their team numbers ("I'm on 339!" "Team 862 Represent!" "Third year on team 88!"). The most devoted extend that love even to Kamen's mother, who has attended these events religiously since Kamen founded the organization 20 years ago.

Kamen, who hobnobs with George H.W. Bush and ice-cream activists Ben and Jerry and flies around in a souped-up Enstrom 480 helicopter (the souping-up is of his own design, natch), has been called a rock star inventor. But it's hard to reconcile the moniker with his unabashed geekery and top-to-toe denim uniform. No Savile Row for Kamen, though he could probably afford to buy Savile Row—the street, not a suit.

The same lack of pretense is apparent on North Dumpling, which is going through an awkward phase on this cool May evening. Here in the kitchen, high-end stainless appliances, including a massive side-by-side refrigerator and a vast range hood, bump elbows with frumpy, vaguely olive cabinetry. The cabinets are not long for this world.

In fact, the whole North Dumpling complex is in the midst of a complete face-lift. The point is to showcase lighting, the most visible aspect of Kamen's rethinking of the way we use energy.

"Do I have to fly anyone to the mainland?" Kamen asks around the small kitchen. If somebody needed some toothpaste or something, Kamen would chauffeur them

MOOD LIGHTING: At night, outdoor LED fixtures wash the island in a moody blue.

STONE WASHED: The miniature Stonehenge replica is perpetually awash in disco colors from the marine-rated LED fixtures.

to the New London, Conn., airport in the helicopter. No big deal. He's had over 10 000 hours of flight time. For him, it's not much more special than driving a car.

Nobody needs anything.

"Then give me a beer!" he says to Morgan, the only man on the island with a pressed shirt. Morgan, a neat blond in starched maroon plaid and khakis, is on a sabbatical from Philips Color Kinetics while he helps Kamen with the project. He knows everything there is to know about the LED systems, of course, because they're all made by his company. He reels off models and ratings without pausing for breath.

Morgan estimates that the cost of all this high-tech lighting would have been about US $100 000. But that figure doesn't mean much, he adds, because some of the devices installed here aren't even on the market yet. Morgan insists that the price will fall steeply as LED lights become prevalent.

Consider, for example, the ColorBlast Powercore lights that shine on the island's miniature re-creation of Stonehenge. (Yes, when Kamen bought the island it came with a little Stonehenge. As with the original, no one knows who built it.) The lights are rectangular devices with alternating rows of eight red, blue, and green LEDs. "When different fractions of the LED lights are lit, any of 16.7 million possible colors can be created," Morgan says. They create additive color the way a computer screen uses different mixtures of red, green, and blue pixels. The fixtures cost about $1000, Morgan says, including the marine-rated housings that protect them from the brutal New England winter storms. Average power consumption on North Dumpling is about 2500 watts. With every single light on the island shining—including the main house, Stonehenge, the guesthouse, and the caretaker's cottage—total consumption peaks at 5000 watts.

An average household dedicates approximately 20 percent of its energy budget to lighting. Here's how that 20 percent can be manipulated by replacing incandescents with LED lighting. A traditional 60-watt incandescent bulb lasts about 1500 hours and has only a 2.5 percent efficiency because it wastes most of its energy as heat. Its brightness rating is about 16 lumens per watt, which is a measure of how much luminosity your bulb produces for every watt it eats up. A compact fluorescent lightbulb, or CFL, does much better: A 13- to 15-W bulb has about an 11 percent efficiency, emitting around 75 lm/W. The less-wasteful CFL lasts about 10 times as long as an incandescent (a little under two years).

On the downside, CFL bulbs contain mercury, which means they must be treated and disposed of as hazardous waste. Now consider LEDs. The best of these can pump out a whopping 150 lm/W (from the "raw" chips, not the integrated lamp or system) and function at about 22 percent efficiency. With typical use, that's at least 11 years without changing a bulb.

Perhaps a dozen countries, including the United States, Australia, and Cuba, have passed laws banning incandescent bulbs, most kicking in by the mid-2010s. With the environmental concerns raised by the toxic content of CFLs, the switch to LEDs is inevitable. Within 10 years, Kamen believes, almost everyone will have some means of locally generating electric power. It's the only logical solution, he thinks, to an overtaxed, undermaintained national grid that's vulnerable to any disruption, be it an ice storm or a terrorist attack.

But in Kamen's native United States (or rather his "neighboring nation," as he reminds me), LEDs contend with a mixed reputation, a legacy of their evolution. Though they consume much less power than their incandescent counterparts, their light can be lurid and flat. They irradiate more than they illuminate.

At Philips, Morgan has been working to improve the quality of LED light from the hard bluish-white of early LED flashlights. "How do we replace them in a way that has all the benefits of an incandescent but none of the problems?" he recounts. "How do we design the optics so we get one soft shadow, as opposed to a bunch of harsh shadows?" The result was a unit with a cluster of LEDs inside, flanked by reflectors. The light even has a dimmer switch actuated by a microprocessor at its base.

The extent of the challenge is on view in Kamen's study—lovingly lined with deep brown bookshelves, complete with a ladder to ascend to the topmost shelf in the high-ceilinged room. It appears to be bathed in the sweet, warm glow of recessed incandescents, but in fact these fixtures use LEDs. In lighting lingo, their bulbs are a new kind of PAR 38 lamp (parabolic aluminized reflector). Philips expects to launch the product next year. In fact, the sole surviving fluorescent bulb, in the hallway

OFF GRID: Clockwise from top left: The Stirling engine is the backup generator on the island, a 193-year-old design tweaked by Dean Kamen to meet his needs; the Teletrol system's display tracks all the energy coming and going on the island; the beating heart of the empire is in the server racks, which also have a Web front end, letting Kamen—and anyone else—monitor the island's energy generation and consumption from anywhere in the world; North Dumpling Island runs partly on three top-of-the-line solar-panel arrays, one of which is pictured here; the vast bank of specialty deep-cycle batteries ensures that neither snow nor rain nor gloom of night will deprive the island of power.

next to Kamen's study, proves a telling contrast. The hall light is dingy and sulfurous. In the library, the effect is indistinguishable from that of incandescents—clean, bright, and warm.

The range illustrated by the lights on the island—the disco luminescence of the patterns on Stonehenge; the warm glow of the study; the cool, official mood of the living room; the twinkling coziness of the bar downstairs; the hypnotic blue of the outside lighting that guides the visitor up the byzantine path snaking between houses; and the efficient illumination in the basement—are why this island is the perfect showcase for everything an LED can do, and with so little power draw.

To be sure, going off the grid is anything but cheap. Kamen installed the wind turbine in 1992, but doing the same thing today would likely set your garden-variety island-owning multimillionaire back between $41 000 and $57 000 (including the tilt-up tower that's necessary for construction because it's pretty hard to get a crane onto a 1.2-hectare private island). The Evergreen Solar ES series that Kamen had installed in four clusters on the island combine 180- and 205-W solar panels.

With the switchover, Kamen was able to pare the energy consumption of the lighting from 7500 watts down to about 3000—a factor of almost two-thirds. "There are more circuits related to lighting than all the other appliances in the house combined," Kamen says. "If each incandescent consumes 5 or 10 times as much power as an LED, think about the total energy savings."

Of course, to take full advantage of the potential savings of LEDs, a homeowner would, like Kamen, have a Teletrol building controller to control the LEDs. These systems start at several thousand dollars and go way up from there. The controller is incredibly granular. Tapping through a succession of screens on the display in the basement, Kamen eventually gets to a plot of the day's energy generation. Colorful lines represent the different incoming sources: a black line tracks shifts in solar intensity and a paired red line shadows it precisely, representing the current feed to the island's rechargeable batteries through the charge converters. The wind turbine shows up as a green line below the juddering red and black lines, and it has racked up only 10 kWh, as attested to by the mild blue sea outside. By contrast, although it's only a quarter past four, the solar panels have already generated a total of 35.97 kWh.

The red and black lines hover at zero during the nighttime hours, creep up a bit around dawn, then rise steeply at around 7 a.m. They hover high on the graph for several hours and then, at the 1 p.m. point, begin to shudder through a series of dips and rises. "Those are clouds passing overhead," Kamen points out.

But what happens when the island gets neither sun nor wind? That's when Kamen's Stirling engine kicks in. Kamen's reimagining of the 193-year old design is about half the size of a refrigerator, and its job is to convert heat energy into work.

A Stirling engine contains a sealed cylinder, one end hot, the other cool, with a sliding piston that moves the gas back and forth. On the hot side, the gas expands and exerts pressure on a piston; on the cool side, it contracts. Unlike a gasoline or diesel engine, Kamen's Stirling is an *external* combustion device. The fuel that heats the gas inside the engine never touches the engine itself. Unlike spark-ignition or compression-ignition vehicle engines, a Stirling doesn't require that the fuel be mixed with air in a particular ratio. It can run on pretty much any liquid or gas fuel, including diesel, methane, gasoline, or olive oil. Got cow manure? That'll work, too [see online sidebar, "Electricity From Cow Flops," http://www.spectrum.ieee.org/energy/renewables/electricity-from-cow-flops]. The design is ideal in theory, converting fuel to heat and heat energy into mechanical energy, but one problem remains: The machine generates electricity at an efficiency of only 20 percent. The rest is wasted as heat.

But Kamen says the concept of "waste heat" changes with context. In winter that heat can be diverted to the house, the clothes drier, and the island's domestic water supply. "Think of it as a furnace that's also making electricity," Kamen explains. Harness the heat from the Stirling while it generates electricity and you can get an efficiency darn near 100 percent, he says.

His own machine typically produces 1 kilowatt of electricity, enough to power 128 LED lights, each equivalent to a 60-W incandescent. The engine sits in a far corner of the basement. At the moment, it's inactive because the solar panels are generating a steady, incoming stream of 2 kW.

The Teletrol system runs an iron-fisted dictatorship over all of these processes. If there's any excess power, the system uses it to feed the vast bank of custom batteries lining an entire wall of the basement, 24 big red plastic containers that resemble gasoline canisters. These batteries weigh 144 kilograms apiece, cost between $1285 and $2000 a pop, and must have their chemistry checked periodically with a hydrometer.

They're also truly and deeply rechargeable. They are lead-acid batteries optimized for what's called a deep discharge cycle, meaning that their charge can be repeatedly and substantially depleted by 80 or 85 percent. For contrast, most battery chemistries typically tolerate at most 5 to 10 percent depletion on a regular basis. The batteries were made by a company called Surrette Battery Co., headquartered in Springhill, N.S., Canada, and Salem, Mass., which specializes in cells optimized for solar and wind-power charging. Each battery comes with a 10-year warranty.

"Every once in a while, we know we're going to have to depend on these batteries to run the entire island," Kamen says. They can power the island for three days continuously in the unlikely case that all sources of energy simultaneously go off-line.

The Teletrol display shows the state of the batteries' charge overnight waning ever so slightly, as minimal, basic processes dip into their charge. The blue line representing the batteries' state of charge has been rising gradually over the course of the day, as the excess power from the solar panels fills them up.

After the batteries are fully charged, the water in the hot water heater is hot, the lights are off, and no activity is drawing the incoming energy, the system is intelligent enough to survey the house and ask itself whether there is enough purified water. Beneath the island, vast holding tanks are filled with seawater that has been purified to medical grade by Kamen's vapor-compression distillers, two of which sit in the basement. Inside their four-foot, glass-paneled cabinets, snaking coils of what looks like garden hose wrap around complicated machinery.

Kamen calls his invention the Slingshot and refuses to discuss its design. He asserts that it can derive absolutely pure water from any wet material, even raw sewage or toxic waste. "It's 40 times as efficient as a standard still would be," Kamen says. "The goal is to make about 1000 liters of water a day using less total energy than a handheld hair dryer." He allows himself a little smile.

To be sure, the system has some wrinkles that need to be ironed out before it can be flaunted for the world's press. Earlier in the day, when we first meet Kamen, he wants to show me how his system works. He excitedly pulls out his laptop, and we sit down in the grassy shade of his helicopter's tail boom as he tries to show me the laptop-based remote version of the display controller on the basement bamboo console. The only problem is that Dumpling's wireless network is down. Kamen pecks at his keyboard for about half an hour while I try to find a spot on the screen that isn't obscured by fingerprints.

Soon, however, North Dumpling's control system is online. The system will do two things. It will let Kamen log into it when he's leaving work at Deka, in New Hampshire, to rouse the sleeping house, lighting and heating the house and the water to his specifications in time for his arrival. But even more important, the system will allow people all over the world to peek into the inner workings of the Dumplonian Empire. Virtual visitors worldwide will be able to monitor in real time how much energy the island is generating and consuming. "You'll be able to see whether we are consuming or producing power," he says, "and where it's going at any moment."

Today only the LEDs are ready for prime time. By September, the entire island will be, too. "It's going to make a powerful statement," Kamen says. "Whether you're interested in the environment or your own budget, the answer is one thing—stop using incandescent lights." ❑

WIND POWER: Dean Kamen's Bergey wind turbine provides 10-kilowatt peak generation on gusty days.

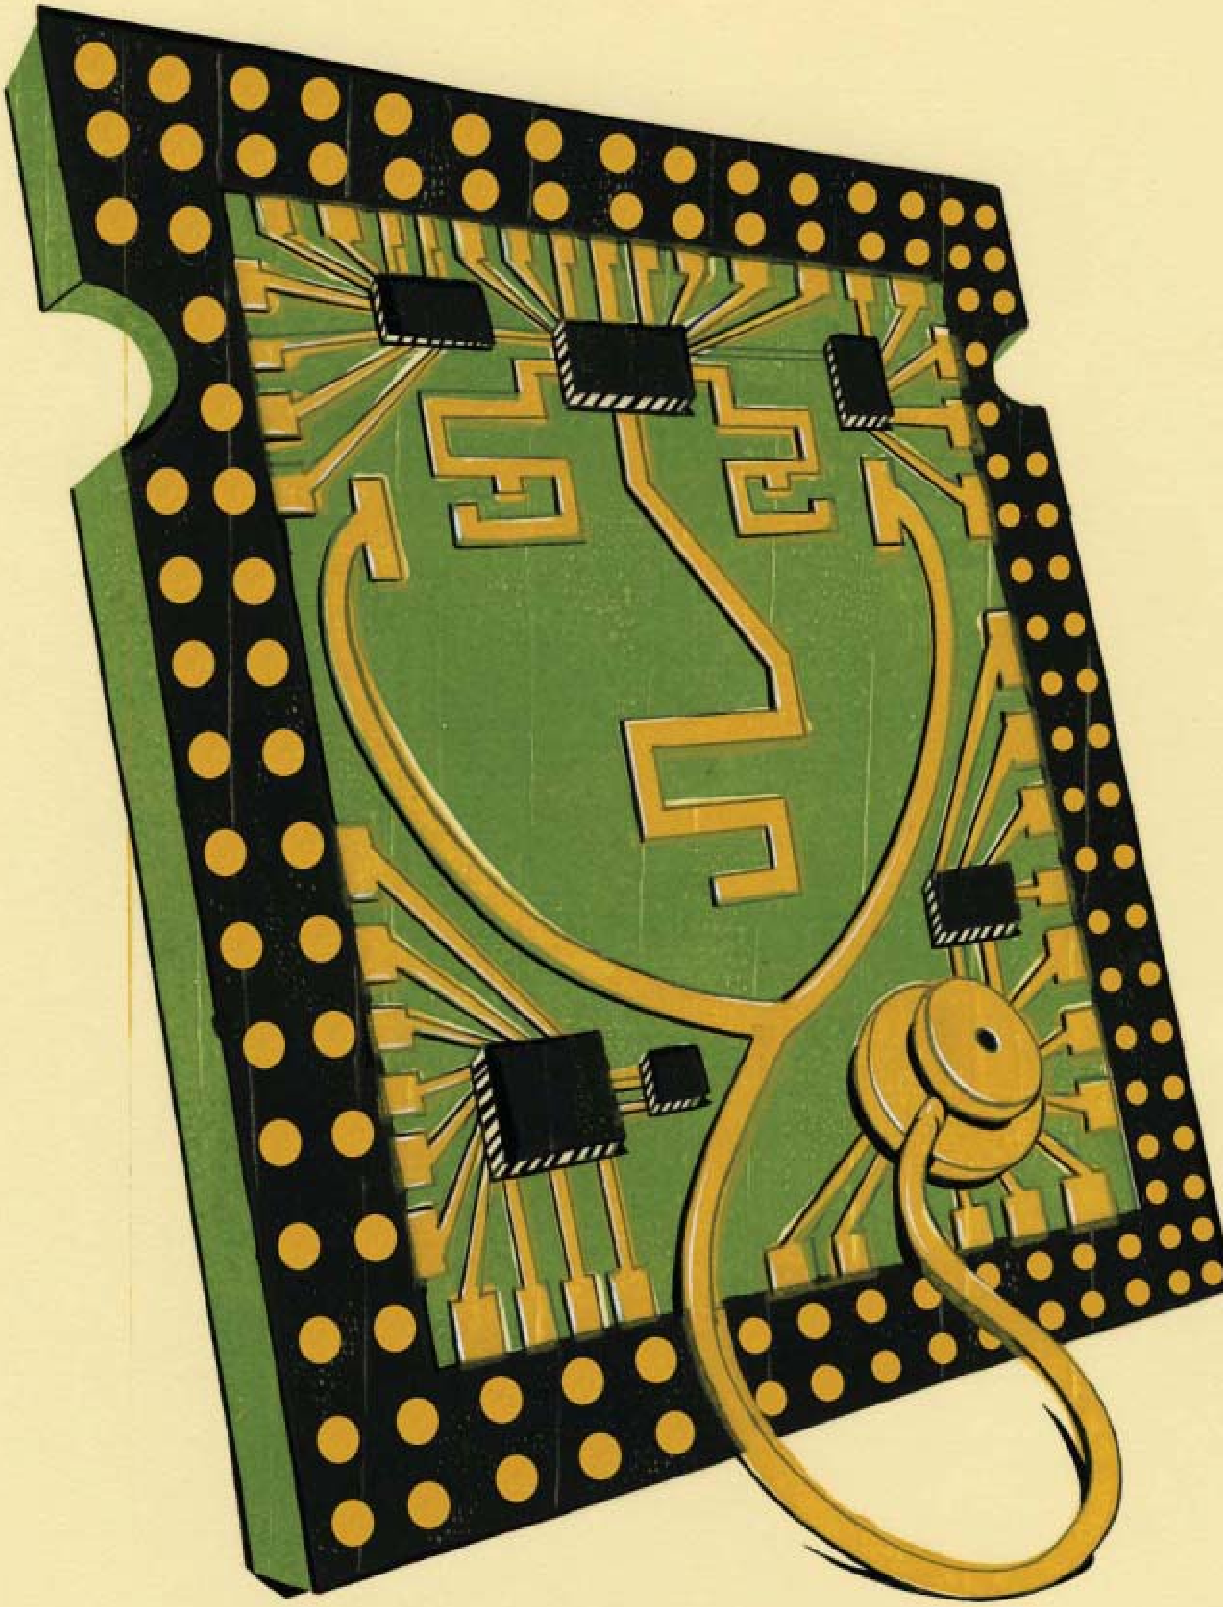

A FAULT-MONITORING MICROPROCESSOR DESIGN CAN SAVE POWER OR ALLOW OVERCLOCKING

BY DAVID BLAAUW & SHIDHARTHA DAS

CPU, Heal Thyself

IN THE OLD DAYS, COMPUTER VENDORS WOULD OFTEN pull a fast one. They would tell you their system had the latest microprocessor when it actually had a cheaper, slower version running faster than the chip's rating permitted. So the shiny, new 500-megahertz system you thought you were buying might contain only an overclocked 300-MHz CPU. But the computer worked fine; indeed, it might have operated perfectly for years, with you none the wiser. And you perhaps replaced it only because a good buy on a 1-gigahertz machine eventually came along.

How did that poor 300-MHz processor cope with such abuse? The short answer is that the manufacturer had set the clock speed low to ensure that its products would function without fault despite the inevitable variations among chips and among their different operating environments. Shady overclockers took advantage of that conservatism, inviting unpredictable failures when they eliminated the chipmaker's prudent safety margins.

Lately, overclocking has gone mainstream. You can, for example, find competitions on the Web in which hardware hackers vie for top honors in this domain. Even chip manufacturers themselves are doing it in public trials to show off how blazingly fast their processors can run under the right conditions—like when they are being cooled with liquid helium to within a few kelvins of absolute zero.

Engineers at Advanced Micro Devices, of Sunnyvale, Calif., did just that this past April to prove that the company's Phenom II CPU could break the 7-GHz barrier. In theory, they could have used the same approach to reduce the voltage at which this chip runs at its normal clock speed. That in turn would have significantly diminished the power it consumed.

While saving a few watts is not so important in a desktop system, it's critical for smartphones, mobile Internet devices, and other such gadgetry, which now have to handle glitzy graphics, video, Web access, and gaming without burning too quickly through their tiny batteries. And reducing the amount of power that a CPU uses translates to an enormous amount of money saved for the companies that deploy vast numbers of microprocessors in large-scale server farms.

The problem is that if you use anything less than their normal voltage, some of these chips, on rare occasions, will fail to produce the correct results. That might happen when a laptop is turned on after being left to bake in a hot car, for example.

DANIEL BEJAR

The resulting miscalculations could be catastrophic—or maybe not.

What if a microprocessor could check its output and correct any error on the fly? Suppose further that the chip could slow itself down or turn up its voltage slightly when it noticed it was flubbing up too often. Experiments we and our colleagues at the University of Michigan in Ann Arbor have carried out show that adding those capabilities to a microprocessor can slash energy use by more than a third.

THIS POWER-SAVING TRICK can be very effective, but only if the chip is intentionally designed to fail at times. That's because every chip is different—both in how it comes off the production line and how it is ultimately used—so you need to push a given chip to its limits to truly know what those limits are. But the thought of operating something without a comfortable safety buffer makes most engineers shudder. Perhaps that's one reason why for the past decade or so circuit designers have attempted to conserve battery power in less radical ways.

One easy approach, called clock gating, disables the clock signal in the circuitry that isn't working on a given operation. That way, the bypassed transistors won't use energy while switching on and off. A variation of this theme not only disables the clock but also cuts off the power being fed to the unused components. Doing so can substantially improve the energy efficiency of cellphones and similar mobile devices, which typically idle for long periods, interrupted by short bursts of activity.

In 2003, we began exploring more ambitious ways to reduce the power required to run a microprocessor—work done together with Todd Austin and Trevor Mudge, our colleagues at the University of Michigan, and with the help of many students. We were aware that chip designers routinely compensate for manufacturing variations, as well as for high-temperature or low-voltage conditions (which can vary even within a single chip), by specifying an operating voltage that's higher than it really needs to be. We knew, too, that manufacturers have had to become more and more conservative in this regard, because it has been increasingly difficult to control the operating characteristics of transistors as they get smaller—a factor of two safety margin is not uncommon. Yet only a few of these transistors will ever experience problematic temperature or voltage conditions for very long. So most of the time, the built-in safety margin just squanders energy.

How, we wondered, could we minimize this waste? One day while we were chatting about possible ways to do that, one member of our research group noted that our basic goal was to shave the safety margin to a minimum. So, fittingly enough, we named the hardware modifications we were working on Razor.

Our idea—inspired by the adage "If you're not failing some of the time, you're not trying hard enough"—was to reduce the operating voltage until the chip would sometimes stumble; then we'd give it a way to recover. Although this tactic had been proposed before for a few very specialized applications (for within-chip communications and in certain digital-signal processors), our Razor project was the first to apply it to a general-purpose microprocessor.

POWER SAVINGS IN THE PIPELINE

The authors' Razor circuitry saves power by reducing the microprocessor's operating voltage. This slows the processor's many transistors, increasing the chance of a timing error, but Razor includes a safety net. Consider one logic stage of a pipelined microprocessor running normally [below, left]. In this example, logical 1s are transformed to 0s, although the signal lines do not change all at the same time. If the transistors involved in the operation switch too slowly [right], incorrect results are copied, but the subsequent change in the output indicates a timing error.

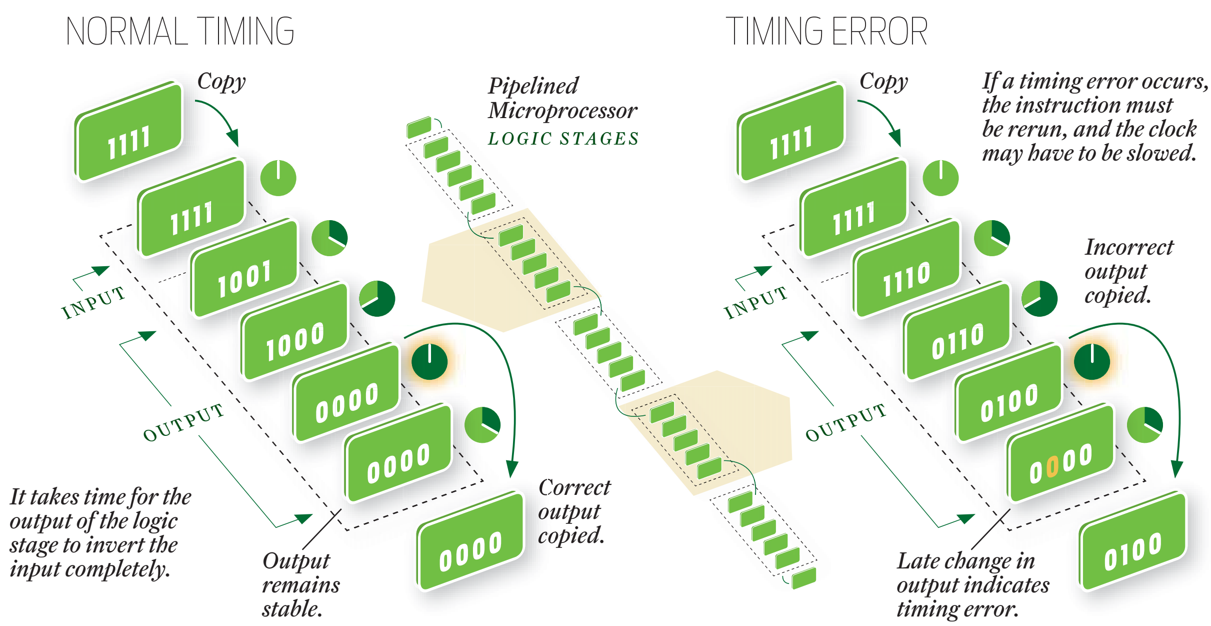

We knew up front that it would take some energy to monitor for failures and correct them—an overhead that would need to be kept small. But we didn't know how common those failures would be as we reduced the operating voltage. Would the first failures be sporadic or would nearly all the instructions go haywire? Armed only with some very approximate theoretical analysis and crude experiments, we took it mostly on faith that the failure rate would be reasonably low—say, one instruction in 10 000, making error correction feasible.

What we found after we built and tested our first Razor chip was that as the voltage dropped to where the chip just started to err, it in fact erred far less often than we had guessed—typically only once in every 100 million operations. The rarity of failures meant that we could more than make up for the cost of error correction on the millions of instructions that executed properly at low voltage.

The Razor system we devised manages the trade-off between system voltage and error rate. It monitors for failure and automatically tunes the supply voltage to achieve the error rate that saves the most energy. In a sense, the chip maintains its own health.

WHAT'S THE MOST ENERGY-EFFICIENT way for a microprocessor to determine that it has messed up? And how can it reliably correct its mistakes? To understand the system we've engineered to do those things, you need to know a little about how modern CPUs work.

To speed processing, most of these chips use a strategy called instruction pipelining. Although the name conjures up a water pipe, the better analogy is to a bucket brigade, where one person fills a pail with water

DANIEL BEJAR

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

and passes it to a second person, who then passes it to a third, and so forth. All the while, the first person is filling and handing off more buckets.

A pipelined microprocessor owes its high speed to the same strategy of breaking down each operation into a series of discrete steps. For a simple processor, there are often five: Fetch the instruction to be carried out from memory, decode it, execute it, determine the address in memory where the result is to be written, and write it there. High-end microprocessors might extend this strategy to a couple of dozen separate pipeline stages.

Pipelining works only because these different functions can all be carried out at the same time. For example, while one of the programmed instructions is being executed, the following one can be decoded, and the one after that can be fetched from memory. Each step is carried out by a specialized circuit that takes the input provided to it, reacts to it in some fashion, and then presents the results to the next stage in the logic pipeline.

As with an actual bucket brigade, these operations need to take place with a regular rhythm. Here, the microprocessor's clock provides the necessary timing. At some designated instant—say, when the clock signal switches from low voltage to high voltage—each processing stage makes a copy of the data on its input lines. Each stage then works with its copy to produce a result.

The time it takes for the input of any stage to be translated into the corresponding output depends on how long it takes the different transistors involved to switch states. The processor's clock is normally set to run slowly enough to ensure that the output will be correct by the time the clock next switches from low to high—which is to say, when the output from one stage becomes the input for the next one. As long as the transistors are finished switching states by the time the next low-to-high clock signal comes around, everything works well.

Now suppose you turn down the supply voltage so that the microprocessor's many transistors can't switch logic states quite so fast. One or more slowpoke transistors within some critical calculation pathway may cause an output to switch states *after* the clock has commanded the following stage of circuitry to copy the data presented to it. Working with the wrong input data, that next stage would, of course, produce an erroneous output, which would wreck whatever operation is flowing through the chip's instruction pipeline. This could easily cause the application—or even the whole computer—to crash. Razor provides a way to avoid such a fiasco.

With our latest version of Razor, each copying circuit is modified so that it includes a transition detector, which is sensitive to changes in the output for a short period of time after each tick of the clock. If the output is not yet valid at the clock tick, the next logic stage will be working with the wrong data. But catastrophe can still be averted, because the correct data will arrive slightly later, triggering the transition detector, which flags the event as a timing error.

When this occurs, a special error controller executes the problematic instruction again. Although it rarely happens in practice, it's possible that this particular instruction will produce an error on the next attempt, too—maybe even on many repeated attempts. To avoid such a deadlock, the controller we've designed tries only a handful of times. If the error persists, the controller circuitry cuts the processor's clock frequency in half during the next attempt to ensure adequate time for the error-free operation of the problematic instruction. The correction process might seem cumbersome, but as the first iteration of our Razor system has shown, this chain of events occurs so infrequently that it slows the average computation speed by only a fraction of a percent.

HITTING THE SWEET SPOT

The Razor circuitry is used to maintain an error rate of 0.04 percent, which keeps energy use at a minimum while barely affecting computation speed [right].

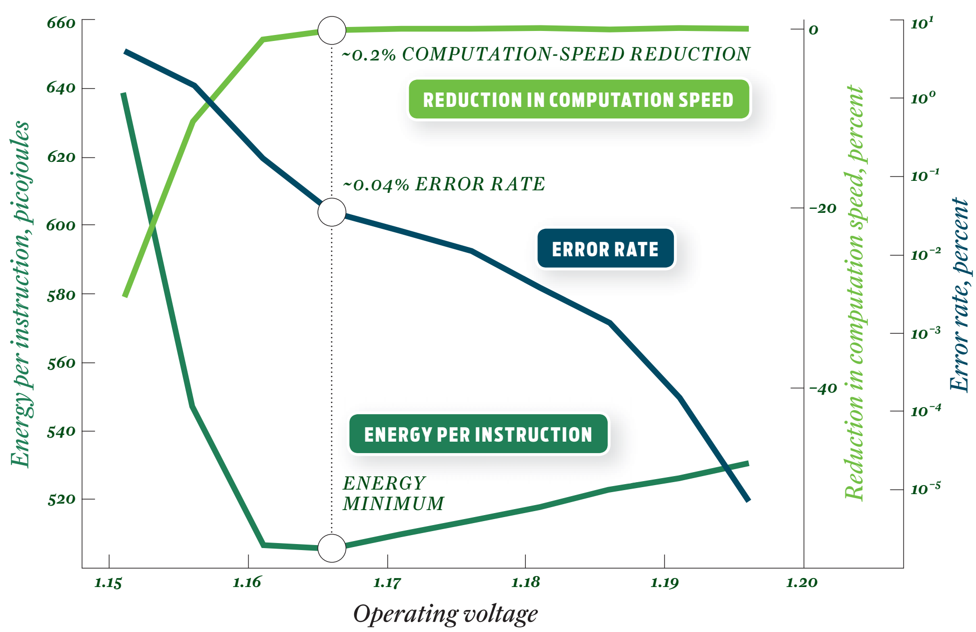

Ironically enough, the biggest challenge in designing the Razor system has been to prevent the microprocessor's circuitry from working too quickly. The reason this can be a problem is that the transition-detection circuitry is dumb: When it sees a signal line change state shortly after a clock tick, it doesn't know whether this is old data from the previous clock cycle arriving late or new data from the current clock cycle arriving early. So the transition detector could mistakenly flag the early arrival of valid data as an error. And such an event might well occur again and again during attempts at recovery, even with a slower clock.

To prevent this from happening, we had to introduce some extra delay into the microprocessor's speedier circuitry. This ensures that the output of a given pipeline stage doesn't switch states while the transition detector connected to it is still sensitive to changes. Adding delay to the faster circuits consumes some power, but it doesn't diminish the processor's overall speed, which remains limited by how quickly the slowest circuits can operate.

You might guess that we also needed to add some very complex circuitry to the microprocessor to enable it to repeat an operation after a timing error occurs. In fact, we did very little, because most of the replay circuitry was already there. It's required to deal with one of the subtle drawbacks of instruction pipelining—the dependence that one instruction often has on the outcome of the previous instruction. That's a problem for a pipelined microprocessor, which must begin processing the second instruction before the result of the first one is known.

In such instances, the microprocessor often guesses what the result will be. The answer could determine, for example, whether to jump to some other part of the program. If the processor guesses correctly, all is well. If not, the microprocessor executes the instruction once more using the correct result as input. This was just the mechanism we needed to force the microprocessor to replay operations when a timing error occurs.

THE INITIAL VERSION OF RAZOR included both error-detection and correction circuitry within each copying circuit, more than doubling the number of transistors needed to copy a stage's input data. The most recent version of Razor, which we described publicly in 2008, adds only error detection to the various pipeline stages. It relies entirely on the microprocessor's existing replay circuitry for error correction. This approach significantly reduces the area on the chip and the amount of energy the new circuitry consumes.

We've also limited the additional area and power needed by adding Razor elements only to those portions of the circuitry that are prone to fail when the operating voltage is reduced. This requires a careful analysis to be done up front for each microprocessor design. But that effort pays off handsomely, because only 10 to 30 percent of a processor's copying circuits typically require Razor's protection against failure.

Although we've tried to minimize the additional power used as much as possible, the Razor circuitry does use a small amount of energy at all times. So if you were to run the modified microprocessor at its standard operating voltage, it would use slightly more power than normal. But the Razor circuitry doesn't do that; it lowers the operating voltage until the error rate is about 0.04 percent. This saves far more energy than the additional transistors consume. Running the chip at still lower voltages would, however, make the error rate surge, requiring many replays, which would slow computation appreciably and use more energy overall. So the Razor circuitry senses the error rate and adjusts the supply voltage to keep the chip at its optimum operating point—which depends not only on the environment and manufacturing variations but also on the calculations being performed at that time. As they say, your mileage may vary, but on average you can expect a 35 percent savings in energy use.

Since our team first described Razor, in 2005, several semiconductor companies have started to investigate this approach. For example, at the IEEE International Solid-State Circuits Conference in February, researchers from Intel presented a Razor-inspired mechanism for improving high-performance chips. Instead of reducing the supply voltage to save energy, the Intel team keeps that voltage constant and uses a Razor-like technique to boost computation speed by 25 to 32 percent. And one of us (Das) now works full time with other researchers at ARM, a British company that designs and licenses reduced-instruction-set computer processors, in an effort to harness Razor's ability to make processors faster or less power hungry.

Razor is one of the more far-out attempts to push microprocessors to their limits—but we're sure it's not going to be the last of its kind. As transistors continue to shrink, we can expect chip designers to invest a growing portion of their transistor budget in circuitry to monitor, analyze, predict, correct, and adapt. Computers might never obtain anything resembling consciousness, but the microprocessors they contain will surely become increasingly self-aware. ❑

TO PROBE FURTHER

An overview of adaptive techniques for microprocessor design (including Razor) is presented in the chapter "Architectural Techniques for Adaptive Computing," by Shidhartha Das, David Roberts, David Blaauw, David Bull, and Trevor Mudge, in *Adaptive Techniques for Dynamic Processor Optimization: Theory and Practice*, Alice Wang and Sam Naffziger, editors, Springer, 2008.

A more detailed description of the early version of Razor is available in "A Self-Tuning DVS Processor Using Delay-Error Detection and Correction," by S. Das, D. Roberts, S. Lee, S. Pant, D. Blaauw, T. Austin, K. Flautner, and T. Mudge, *IEEE Journal of Solid-State Circuits*, Vol. 41, no. 4, April 2006.

The newer strategy for Razor is described more fully in "RazorII: In Situ Error Detection and Correction for PVT and SER Tolerance," by S. Das, C. Tokunaga, S. Pant, W.-H. Ma, S. Kalaiselvan, K. Lai, D. M. Bull, and D.T. Blaauw, *IEEE Journal of Solid-State Circuits*, Vol. 44, no. 1, January 2009.

by HANY FARID

SEEING IS NOT BELIEVING

Doctoring digital photos is easy. Detecting it can be hard

JUST DAYS AFTER SARAH PALIN'S SELECTION last August as the Republican vice presidential candidate, a photo of a bikini-clad, gun-toting Palin blitzed across the Internet. Almost as quickly, it was revealed as a hoax—a crude bit of Photoshop manipulation created by splicing an image of the Alaska governor's head onto someone else's body. From start to finish, the doctoring probably took no more than 15 minutes.

Altering digital imagery is now ubiquitous. People have come to expect it in the fashion and entertainment world, where airbrushing blemishes and wrinkles away is routine. And anyone surfing the Web is routinely subjected to crude photographic mashups like the Palin hoax, whose creators clearly aren't interested in realism but in whatever titillation or outrage they can generate.

But other photo manipulations demonstrate just how difficult it has become to tell altered images from the real thing. For example, in 2005 Hwang Woo-Suk, a South Korean professor, published a paper in one

ILLUSTRATION: MICHAEL ELINS;
BILL GATES: PETER MACDIARMID/GETTY IMAGES

ROCKET DUSTUP: A July 2008 photo [left] shows four Iranian missiles streaking skyward. But only three of those rockets actually left the ground; a fourth was digitally added. The altered image was first posted on the Web site of Sepah News, the media arm of Iran's Revolutionary Guard, and then published by media outlets around the world. Careful observers pointed out that portions of the faked rocket's exhaust plume and dust cloud had obviously been duplicated from its neighbors'. Sepah News soon replaced the faux photo with the original [right] without explanation.

COURTESY HANY FARID

of the most prestigious scientific journals, *Science*, claiming groundbreaking advances in stem-cell research. But at least 9 of the 11 uniquely tailored lines of stem cells that Hwang claimed to have made were fakes. Much of the evidence for those 9 lines of stem cells involved doctored photographs.

Apparently, Hwang's fabrication was not an isolated occurrence. Mike Rossner, then the managing editor of *The Journal of Cell Biology*, estimated that 20 percent of the manuscripts his journal accepted contained at least one image that had been inappropriately manipulated. Since then, a number of scholarly journals have implemented new fraud-detection procedures, such as software that makes it easier to compare images within or between documents. The incidence of image fraud in scholarly publishing has not declined, though; indeed, it seems to be on the rise.

A more recent example of photo tampering came to light in July 2008. Sepah News, the media arm of Iran's Revolutionary Guard, celebrated the country's military prowess by releasing a photo showing the simultaneous launch of four missiles. But one of those missiles had, in fact, failed to launch. The truth emerged after Sepah circulated the original photo showing three missiles in flight—but not before the faked image appeared on the front pages of the *Chicago Tribune*, the *Financial Times*, and the *Los Angeles Times*. If the world could be fooled by such a photo, then what's to prevent any country or militant group from using doctored images to intimidate?

To be sure, photographic alterations have existed about as long as photography itself. But before the digital age, such deceptions required mastery of complex and time-consuming darkroom techniques. Today anyone with a modicum of computer skills can call on powerful and inexpensive software to alter digital images. And as sophisticated forgeries appear with alarming frequency, people's belief in what they see has been eroded.

Over the past few years, the field of digital-image forensics has emerged to combat this growing problem and return some level of trust in photographs. By using computer methods to look at the underlying patterns of pixels that make up a digital image, specialists can detect the often-subtle signatures of manipulated images that are invisible to the naked eye.

NEARLY EVERY DIGITAL FORGERY starts out as a photo taken by a digital camera. The camera's image sensor acts as the film. It consists of a two-dimensional array of photoelectric elements that become electrically charged when exposed to light, which is why this type of light sensor is called a charge-coupled device, or CCD. The amount of charge is proportional to the light's intensity, so the electrical pattern captured by the CCD faithfully represents the light pattern striking the sensor.

Although exquisitely sensitive to intensity, the CCD elements can't detect the light's wavelength—that is, its color. So a device called a color-filter array is overlaid on the CCD, enabling each element to record a limited range of wavelengths corresponding to red, green, or blue.

After taking a picture, the camera transfers the pattern of electrical charges to the camera's memory, where it is represented as an array of pixels. A 6-megapixel camera, for example, has a CCD sensor with 6 million elements and takes digital images of up to 6 million pixels each. The charge or light intensity is translated into a number, 0 being the minimum and 255 the maximum. In a full-resolution color image, each pixel is assigned three such numbers, one for the intensity of red, one for green, and one for blue. But as noted above, the color-filter array initially assigns each CCD element just one color. So the camera fills in the missing color values by interpolating across neighboring pixels. These three values can yield more than 16 million colors.

Digital images can be stored in a number of formats. The most basic is raw format, in which the pixel values are stored exactly as they're recorded by the CCD, with no interpolation. This format is efficient,

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

as only one number is stored per pixel, but it requires any subsequent photo-editing software to perform the interpolation. The remaining image formats fall into one of two categories: nonlossy and lossy. Nonlossy formats, such as TIFF, PNG, and BMP, compress an image file by representing redundant or repetitive data using a kind of digital shorthand; when the file is subsequently expanded, the redundant data can be retrieved, so there is no loss of information. The lossy formats all compress their files by permanently removing data. The GIF format, for instance, limits the number of colors in a compressed image from millions to typically a few hundred. The JPEG format, perhaps the most popular lossy format, compresses by removing some color and image details.

THE FIRST RULE in any forensic analysis must surely be "preserve the evidence." So you might think that lossy image compression, which deletes information, would be a forensic analyst's worst enemy. In fact, it's a great aid: The unique properties of lossy compression can be exploited to track manipulations.

Take the ubiquitous JPEG format. It uses a compression algorithm that transforms the underlying pixel values into a map of low, middle, and high frequencies, where the low frequencies correspond to areas where the color changes very little (a blue sky and white clouds, for instance) and the high frequencies correspond to rapidly changing colors (as in a flamboyant Hawaiian shirt). Human eyes are less sensitive to the minute details of the high frequencies, so in JPEG files these areas are compressed more than the lower frequencies.

The JPEG image format also specifies how compression and memory consumption are balanced. This balance is represented as a matrix, called a quantization table, with each value specifying how much each of 64 distinct frequencies in the image, in each of three channels specifying brightness and color, has been compressed.

If you've shopped for a digital camera lately, you know that different models of cameras often produce very different results, even if they have the same overall pixel count. That's because camera manufacturers balance image compression and quality in subtly different ways, which creates differences in the quantization table. That means the photos taken by a given model of camera will have a signature of sorts embedded within each JPEG file it produces. The quantization tables used by Photoshop and other photo-manipulation software are also distinct, so you can tell whether one of those programs has been used to alter the image. This allows for a crude form of digital-image "ballistics." In many cases, you can figure out what type of camera the photographer used to take the shot.

Another related technique that my group at Dartmouth College, in Hanover, N.H., is now studying makes use of the thumbnail image that every digital camera automatically creates along with each full-resolution photo. The thumbnail, which has a resolution of about 160 by 120 pixels, is the tiny image you see when you preview a photo you've just taken. To create the image, the camera takes the full-resolution image, filters it, selectively removes pixels, filters it again, and then adjusts the brightness and contrast. Our research shows that this image processing relies on algorithms that appear to vary between different camera models. In our experiments, we've been able to estimate the parameters used to create a given thumbnail. The next step will be to build a database of thumbnail parameters from a large array of camera makes and models, which can then be used to authenticate the source of an image. In addition, the thumbnail is itself saved as a JPEG, using a different quantization table, and this information can be used to further refine the camera's signature.

FONDA SPEAKS TO VIETNAM VETERANS at Anti-War Rally" reads the headline, and the accompanying photograph, purportedly from 1970, shows a young Jane Fonda sharing a stage with a fresh-faced John

HOW HE DID IT

Michael Elins specializes in illustrations that blend digital imagery with his own photography. His image of a down-at-the-heels Bill Gates, an homage to a *Saturday Evening Post* cover from 1924, started with a stock photograph of the Microsoft billionaire. Elins then took photos of a male model dressed as a hobo (note the knife held by an assistant) and of Zippy the dog. He blended the three photos together using professional editing software, digitally adding the flames and smoke to create the makeshift campfire. The scenario may be improbable, but the image is deceptively realistic.

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

IDOL CURIOSITY: The cast of the TV show "American Idol" posed for this shot at different times. You can see it in their eyes—literally. The white specks, or specular highlights, indicate the direction of the lighting. Randy Jackson and Paula Abdul [seated] were likely photographed together, while Simon Cowell and Ryan Seacrest [standing] each posed separately. The yellow arrows at right indicate the lighting source for Cowell [left], Seacrest [middle], and Jackson and Abdul [right].

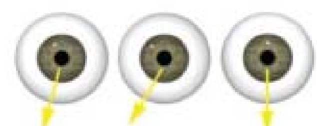

Kerry. The image, used to discredit Kerry during his unsuccessful U.S. presidential campaign in 2004, was a fake, composited from two unrelated photos taken at different places in different years.

To create such a composite, it is often necessary to resize, rotate, or stretch portions of an image. Let's say you're creating a composite image by grafting one person's head onto another person's body. It's unlikely that the relative sizes of the two images match exactly, so you'll have to enlarge or shrink one of them. In the process, you'll alter the underlying pattern of pixels in a distinct and detectable way.

Consider a small 3- by 3-pixel patch. Each of those pixels has a number corresponding to its brightness. Now let's say you want to enlarge, or up-sample, that patch by a factor of two. Enlarging an image requires adding extra pixels; in this case, an extra row of pixels would be added after each original row, and you'd end up with a 3- by 6-pixel patch. The computer software automatically assigns each new pixel's brightness by averaging the values of its immediate neighbors. As a result, the new pixels are perfectly correlated to their neighbors. Such correlations are unlikely to occur naturally, so a forensics expert detecting their presence knows the image has been manipulated.

But what if the resized image has been enlarged to a lesser degree or reduced in size? In those cases, the periodic correlations among the pixels are trickier to spot, but they do exist. My group has developed a computer

AMERICAN IDOL: ASSOCIATED PRESS; BELOW: COURTESY HANY FARID

LEADING MAN: To create a more heroic portrait of himself, Italian dictator Benito Mussolini ordered the horse handler removed from the original 1942 photograph.

program that can detect such patterns by iteratively looking for pixels that are correlated to their neighbors. If detected, the correlations are then used to determine which portion of the image has been resized.

IN APRIL 2005, months before the romance between actors Brad Pitt and Angelina Jolie had been confirmed, *Star* magazine featured a cover photograph of the two über-celebrities strolling down the beach. The photo was a fake. The telltale sign was that the lighting on Pitt's and Jolie's faces was inconsistent with a single light source—in this case, the sun. To judge by the photo, you might conclude there had been two suns shining that day.

The amount of light that strikes a surface depends on the 3-D orientation of the surface relative to the 3-D position of the light source. But if you use photo-editing software to alter the image, you're dealing with a 2-D image, so it can be difficult to match the lighting conditions exactly. Studies show that our eyes are often insensitive to such lighting inconsistencies. But where human eyes fail, computers excel.

My group has developed a technique that can estimate the direction of the light source in an image by looking at a given object's 2-D surface contour, such as a person's jawline and chin. There, the surface orientation to the light source is always perpendicular to the contour.

By measuring the brightness and orientation at several points along the contour, we can estimate which direction the light is coming from. Then we can compare the lighting direction for that object to those of other objects in the photo. Any inconsistency in the lighting direction is evidence of tampering.

THE ASSOCIATED PRESS PLANNED to run a photo of the cast of the television series "American Idol." A photo editor at the wire service had doubts about the photo's authenticity, however, and contacted my lab for a second opinion.

When my colleagues and I examined the image closely, we immediately noticed that the small white specks of reflected light in each person's eyes, known as specular highlights, were inconsistent. To us, that was an obvious sign that the cast members had posed at different times and that the individual photos had been melded together.

The eyes are a beautiful tool for digital forensics, because they act as tiny windows into the world in which the photo was taken. By looking at the shape, color, and location of the specular highlights, we can learn quite a bit about the lighting that was used to take the photograph.

The location of the bright spot on the eye, for example, can indicate where the light source was positioned; multiple spots indicate more than one light source. The precise position of the specular highlights depends on both the curve of the eyeball and the relative orientations of the eye, the camera, and the light. The curve of the eye, it turns out, is remarkably similar from person to person, and there are very accurate 3-D models of eyeball shape. To calculate the relative orientation between a person's eyes and the camera, we can compare the shapes of the circular boundaries between the iris and the white part of the eye, known as the limbus. For example, if the person is directly facing the camera, the limbus in each eye will appear to be perfectly circular. As the orientation of the eyes changes relative to the camera, the limbus becomes more elliptical.

We can then use the shape and orientation of the limbus to estimate the direction to the light. Any inconsistencies in the lighting are evidence of tampering.

EVEN AS EXPERTS CONTINUE to develop techniques for exposing photographic frauds, new techniques for creating better and harder-to-detect fakes are also evolving. As in the battle against spam and computer viruses, it seems inevitable that the arms race between the forger and the forensic analyst will continue to escalate, with no clear victor. Improved image forensics will never be able to eradicate or prevent digital tampering, but these techniques can make it more time-consuming and difficult for forgers to ply their trade. Tomorrow's technology will almost certainly enable digital manipulations that today seem unimaginable, and the science of digital forensics will have to work hard to keep pace. It is my hope that these new techniques, along with a greater awareness of the technological possibilities and sensible updates in policy and law, will help the media, the courts, and the public contend with the exciting but often baffling events of our digital age. ❑

TO PROBE FURTHER The author's home page at Dartmouth College (http://www.cs.dartmouth.edu/farid) has more information about digital photo tampering as well as more examples.

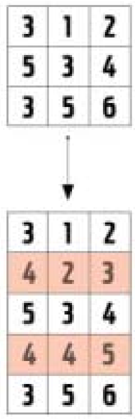

BIGGER: To enlarge an image, editing software adds extra pixels [pink], assigning each new pixel a value that is perfectly correlated to its neighbors'—a telltale sign of image tampering.

RUBBED OUT: Cuban dictator Fidel Castro and Carlos Franqui fell out over the Soviet intervention in Czechoslovakia. Franqui went into exile, and Cuban authorities had his image expunged from photographs.

COVER UP: O.J. Simpson's 1994 mug shot following his arrest for murder was digitally darkened on *Time*'s cover, but not on *Newsweek*'s. *Time* apologized and issued a second cover but said the intent wasn't racist.

Faculty Positions in
Intelligent Systems and Control
Zhejiang University, China

The Institute of Cyber Systems and Control (CSC), i.e., the former "Institute of Advanced Process Control (APC)", at the Department of Control Science and Engineering, Zhejiang University invites applications for 5 positions as Full Professor and Associate Professor in the fields of all areas of intelligent systems and control, with particular interests in robotics, biosensors, network controls and optimization. The applicants should have a PhD Degree in the related areas with outstanding academic credentials that demonstrate their ability to conduct independent and successful research in their areas of expertise. Good responsibility and team work spirit are required. Competitive salary and benefits will be offered according to qualifications and experience.

For full consideration, applicants should send a cover letter, complete curriculum vitae with a list of main publications and a proven record of academic accomplishments, statement of research interests and future plan, and the names and contact information (phone and email address) of three references to the secretary of search committee, Dr. Xu, Weihuo (email address: ***whxu@csc.zju.edu.cn***). Copies of key publications may also be submitted. Applications will be open until the positions are filled.

Zhejiang University located near beautiful West Lake in Hangzhou is recognized as one of the top universities in China. The Institute of CSC being the core body of State Laboratory and National Engineering Center is ranked one of the top research organizations in China.

1918
TALLINNA TEHNIKAÜLIKOOL
TALLINN UNIVERSITY OF TECHNOLOGY

Tallinn University of Technology, the Faculty of Information Technology invites applications from outstanding scientist for two fulltime faculty positions:

Full Professorship in Dependable Embedded Systems

Areas of particular interest include, but are not limited to: embedded digital systems, dependability, reliability, fault tolerance, architectures and platforms, embedded software, system-level design issues, modelling, analysis and synthesis. Ideally, the candidate's research would connect with and complement the TUT's research in related areas.
More information: *www.ati.ttu.ee, gert.jervan@pld.ttu.ee*

Full Professorship in Biomechanics

Areas of interest: biomechanics, biolocomotion, soft body modelling and deformation registration, artificial muscle mechanical modelling and control, mechanical design of bio-inspired devices, bio-inspired control, learning and adaptation, locomotion in fluids, experimental and computational fluid mechanics in the context of biolocomotion and any other areas related or complementary to the research in TUT Center for Biorobotics.
More information: *www.biorobotics.ttu.ee, maarja.kruusmaa@biorobotics.ttu.ee*

Duties: General responsibility for research and research education, teaching and supervision at graduate and postgraduate level and teaching at advanced and basic level. Participation in leadership and administrative roles of laboratory projects involving responsibility for budgets and personnel may be relevant. Demonstrated ability to attract external grants will be expected. Previous experience gained from research in industry or the public sector will be particularly relevant. The applicant should be able to demonstrate good leadership skills and high capability of cooperation.

Applications should be sent by September 30, 2009 to: Tallinn University of Technology, Department of Computer Engineering, Raja 15, 12618 Tallinn, Estonia. **Attn.:** *Gert Jervan.*

THE HONG KONG
POLYTECHNIC UNIVERSITY
香港理工大學

DEPARTMENT OF COMPUTING

The Department invites applications for Professors/Associate Professors/Assistant Professors in Database and Information Systems/Biometrics, Computer Graphics and Multimedia/Software Engineering and Systems/Networking, Parallel and Distributed Systems. Applicants should have a PhD degree in Computing or closely related fields, a strong commitment to excellence in teaching and research as well as a good research publication record. Applicants with extensive experience and a high level of achievement may be considered for the post of Professor/Associate Professor.

Please visit the website at ***http://www.comp.polyu.edu.hk*** for more information about the Department. Salary offered will be commensurate with qualifications and experience. Initial appointments will be made on a fixed-term gratuity-bearing contract. Re-engagement thereafter is subject to mutual agreement. Remuneration package will be highly competitive. Applicants should state their current and expected salary in the application. Please submit your application via email to ***hrstaff@polyu.edu.hk***. Application forms can be downloaded from ***http://www.polyu.edu.hk/hro/job.htm***. **Recruitment will continue until the positions are filled.**

Details of the University's Personal Information Collection Statement for recruitment can be found at ***http://www.polyu.edu.hk/hro/jobpics.htm***.

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

the data

Europe Comes to Bury CO_2, Not to Praise It

CARBON CAPTURE and storage (CCS) made its world debut in Northern Europe in 1996, when Norway's state-owned oil and gas company started burying carbon dioxide that had been stripped out of natural gas at its Sleipner West offshore field. Rather than venting the CO_2—and paying Norway's punishing US $50 per metric ton carbon tax—the firm pumped it into a saline aquifer 1 kilometer below the seabed. Sleipner has sequestered over 11 million metric tons of CO_2. Extensive testing says that the gas is staying put.

Success at Sleipner and at several smaller projects has sparked the interest of Norway's meganeighbor to the south, the 27-member European Union, which has a goal of cutting greenhouse-gas emissions 80 percent from 1990 levels by 2050. Utilities and a few industrial firms propose to incorporate CCS into dozens of projects—primarily coal-fired power stations, which generate nearly a quarter of the EU's CO_2 emissions. This spring EU leaders short-listed 13 projects to share €1.05 billion ($1.46 billion) in economic stimulus funds. Each proposes to capture at least 85 percent of its CO_2 and would test one of several technologies.

CCS is no panacea. According to European Commission projections, by 2020 the process will still cost operators of coal-fired power plants more than carbon allowances that would let them release CO_2 into the atmosphere (even at a projected permit price as high as €41 per metric ton). Thus the EC expects power firms to sequester only 7 million metric tons per year in 2020—less than 1 percent of the CO_2 they currently generate.

The EU must also convince a wary public that buried CO_2 will stay buried for good, protecting the densely populated communities above it. —*Peter Fairley*

SOURCES: European Commission; Scottish Centre for Carbon Storage, University of Edinburgh, http://www.geos.ed.ac.uk/ccsmap; StatoilHydro; Intergovernmental Panel on Climate Change

ILLUSTRATION: BRYAN CHRISTIE DESIGN

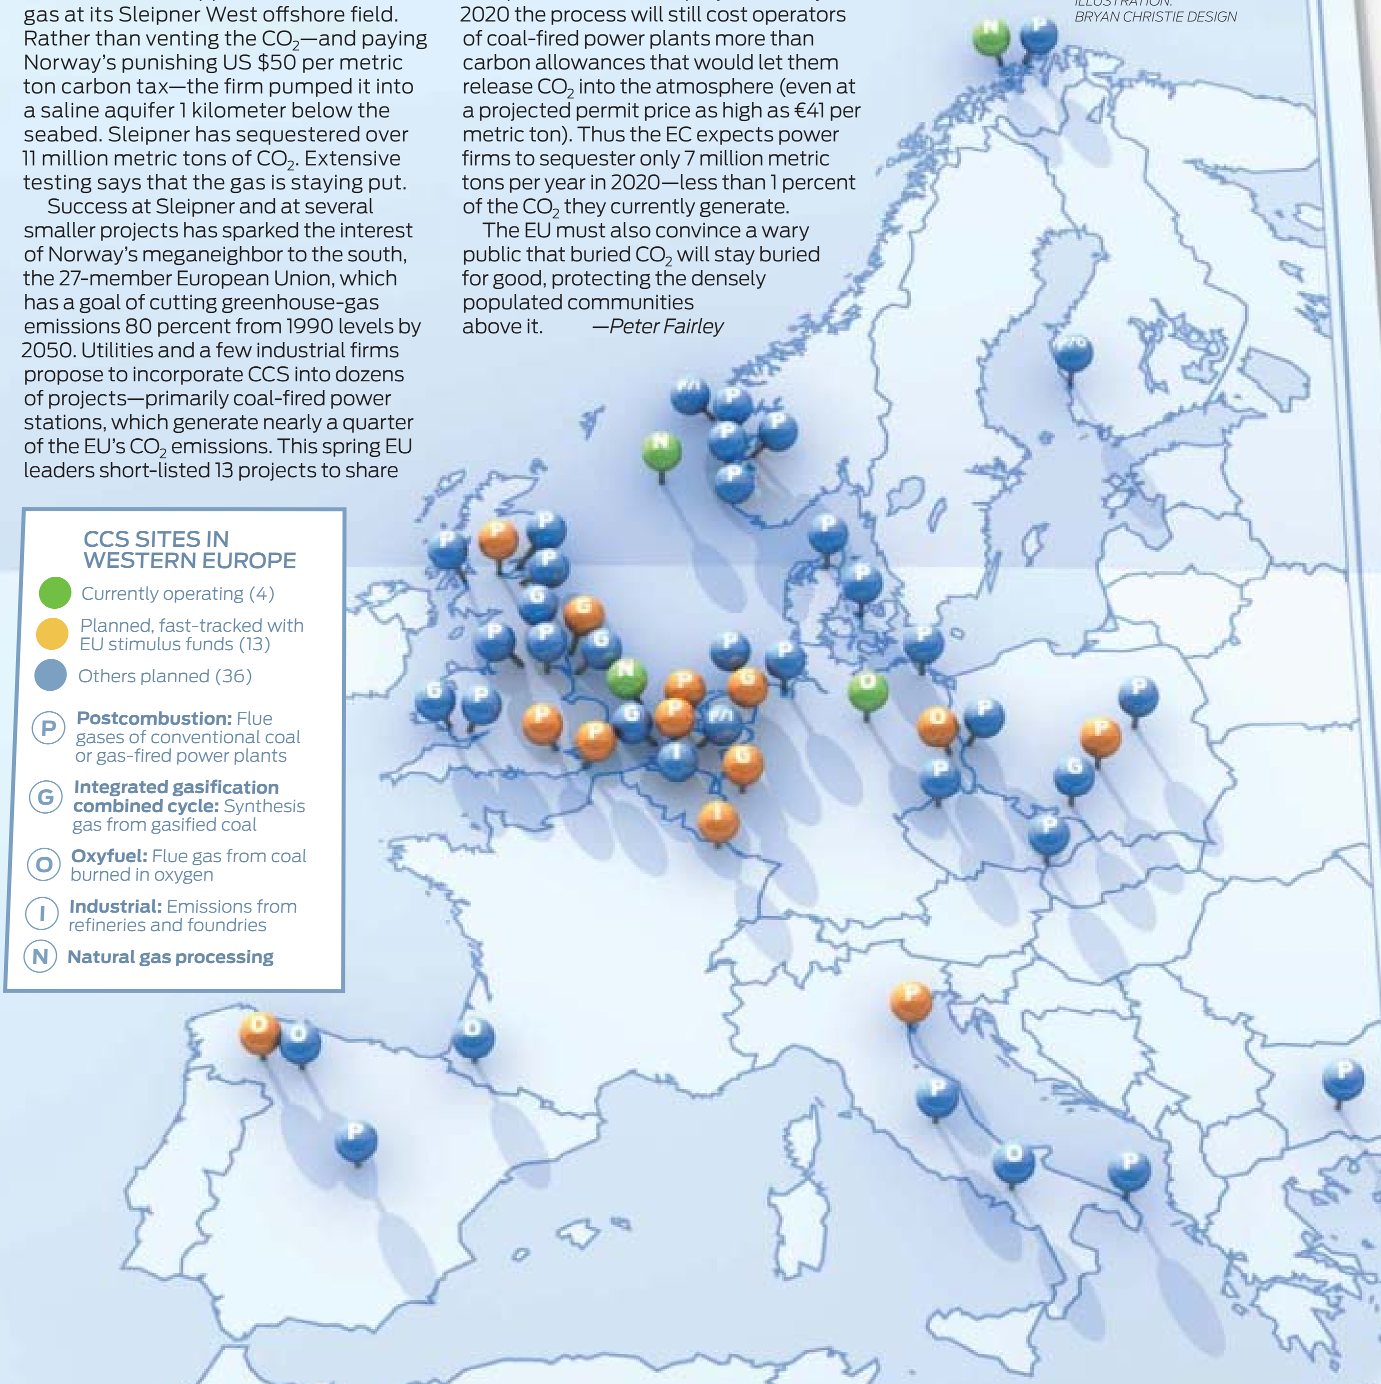


The MathWorks™
Accelerating the pace of engineering and science
あなたは
MATLAB
を話しますか?
Over one million people around the world speak MATLAB. Engineers and scientists in every field from aerospace and semiconductors to biotech, financial services, and earth and ocean sciences use it to express their ideas. Do you speak MATLAB?
Cells in mitosis: high-throughput microscopy for image-based screens. Provided by Roy Wollman, Univ. California, Davis.
Article available at mathworks.com/ltc
MATLAB®
The language of technical computing
©2009 The MathWorks, Inc.