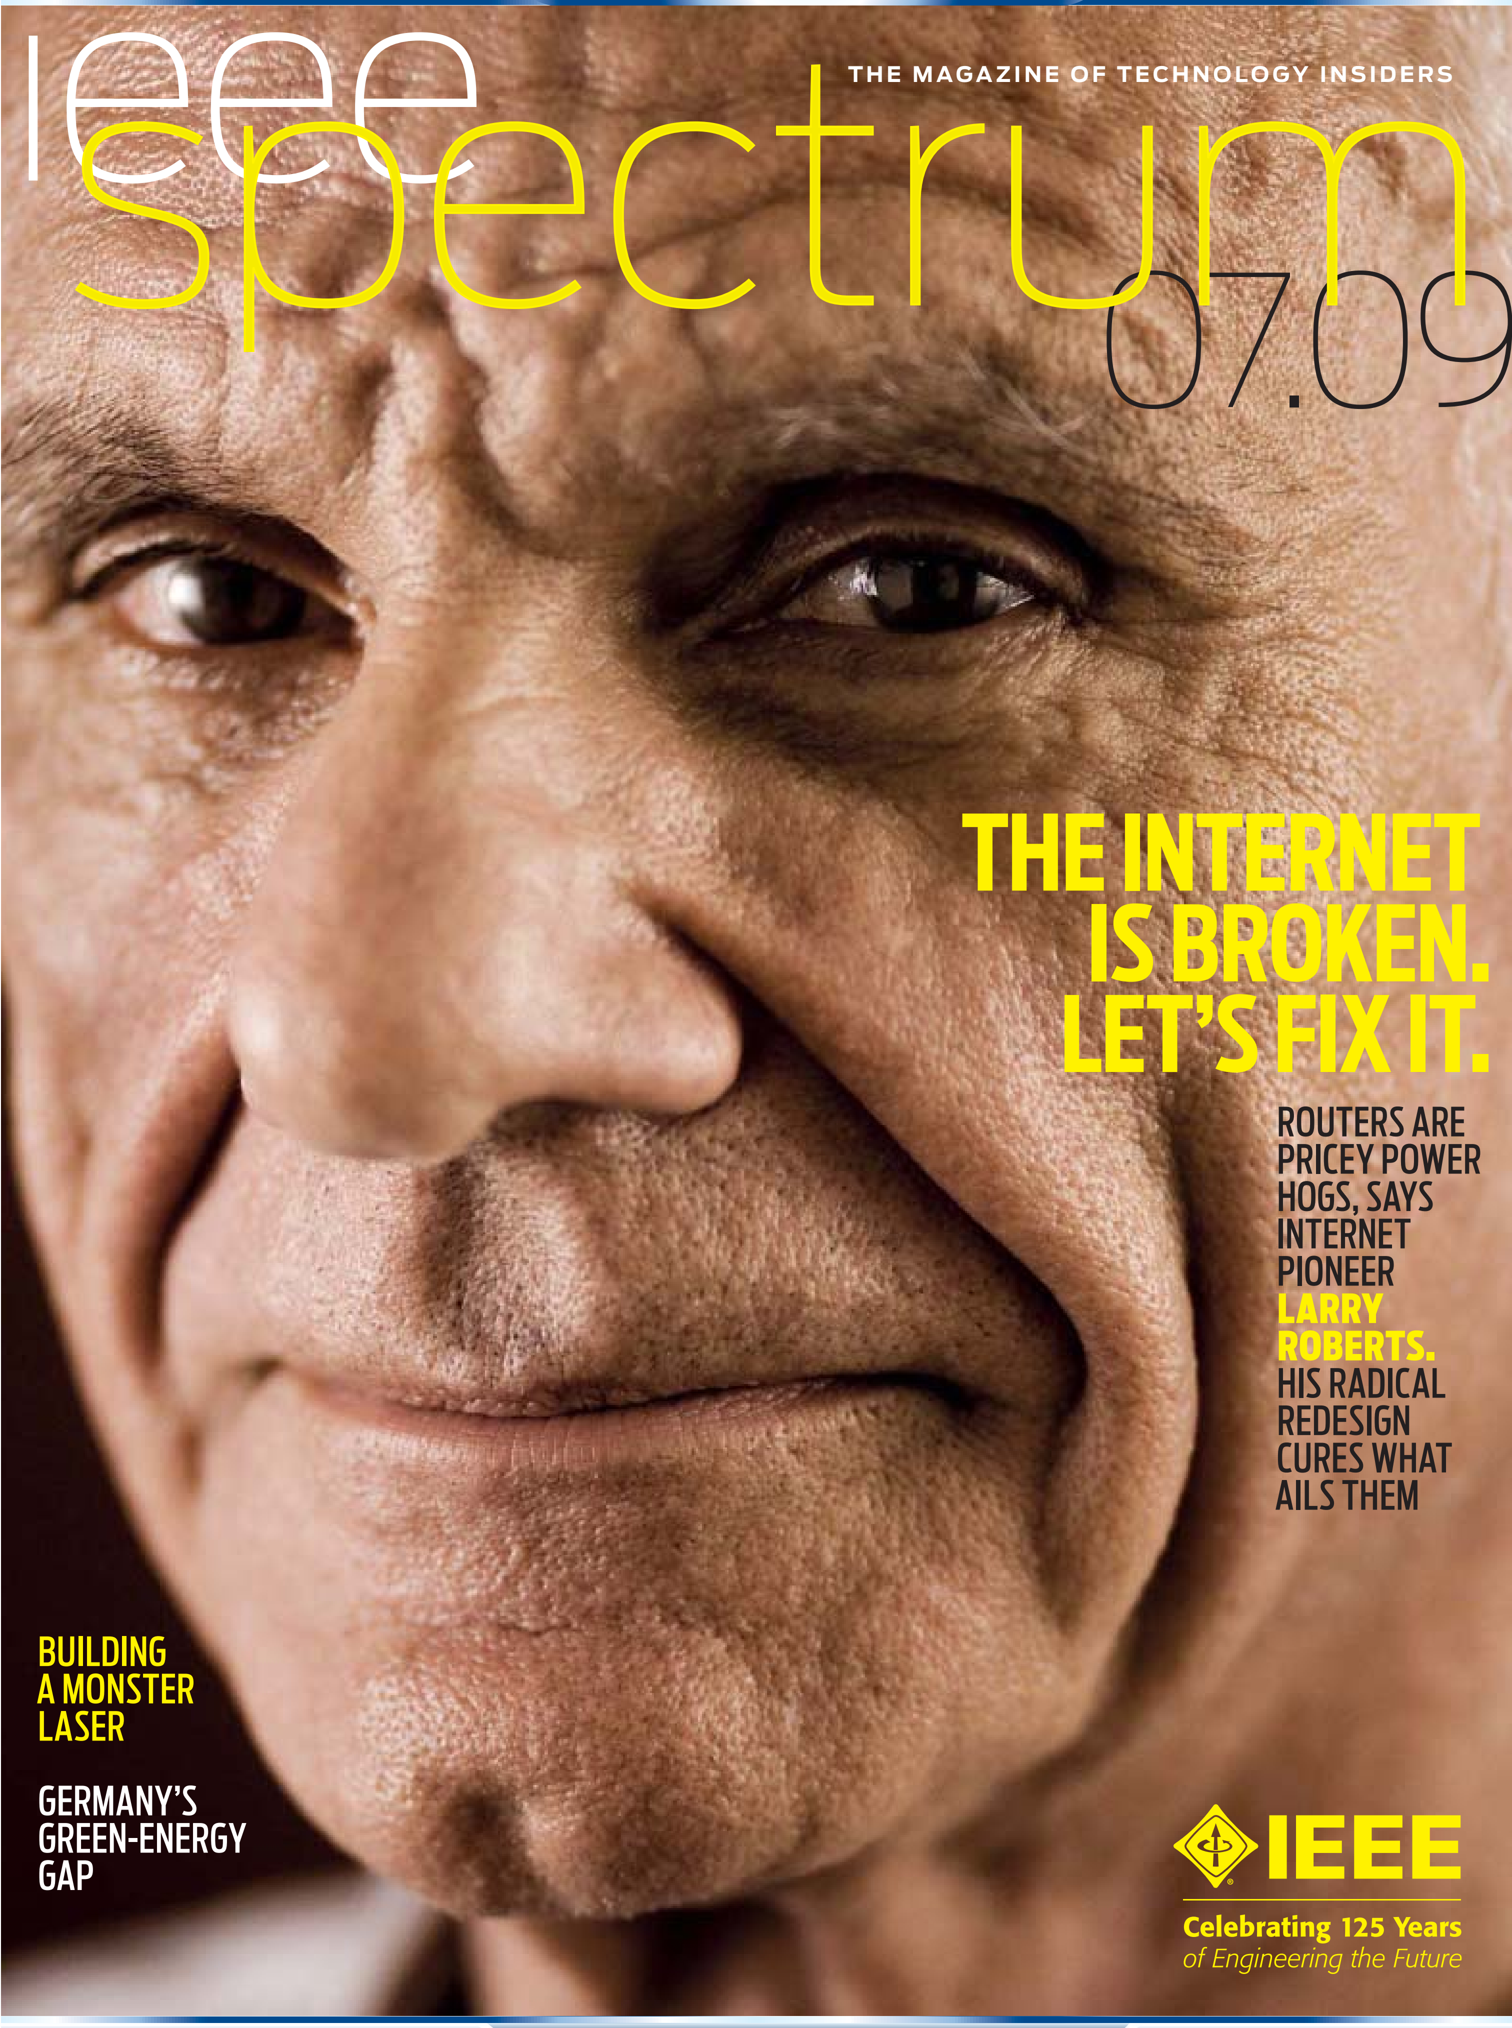

IEEE Spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
07.09

THE INTERNET IS BROKEN. LET'S FIX IT.

ROUTERS ARE PRICEY POWER HOGS, SAYS INTERNET PIONEER **LARRY ROBERTS.** HIS RADICAL REDESIGN CURES WHAT AILS THEM

BUILDING A MONSTER LASER

GERMANY'S GREEN-ENERGY GAP

IEEE
Celebrating 125 Years
of Engineering the Future

COMSOL MULTIPHYSICS®

Improving Prediction of Semiconductor Lifetimes

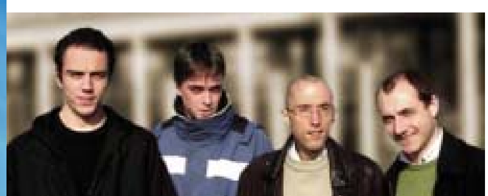

The modeling team at STMicroelectronics. From left: Sébastien Gallois, Romuald Roucou, Vincent Fiori and Florian Cacho.

The "electron wind" is the prevalent failure mode for integrated circuits interconnects. STMicroelectronics is using COMSOL Multiphysics® to predict semiconductor lifetimes and save time in the qualification process when bringing new advanced CMOS technologies onto the market.

A MULTIPHYSICS CHALLENGE

With the scaling down of semiconductor devices, current density in the metal interconnects joining individual transistors increases. Until now, lifetime models have been based on empirical methods. Thus, an evaluation of potential failure modes is very important as STMicroelectronics brings new advanced CMOS technologies to the market. In devices fabricated with these process technologies, interconnect copper lines can be just 100nm thick and roughly the same height. The prevalent failure mode for interconnects is electromigration, which is the net transport of material caused by conducting electrons colliding with metal ions. Over time, a number of metal atoms are knocked from their original positions due to this phenomenon commonly known as an "electron wind." The subsequent vacancies or holes in the crystal structure are due to the migration of metal atoms, and over time, these can accumulate to form minute voids that lead to open circuits and device failure.

A scanning electron microscopy image of a copper line interconnect where a void has developed due to vacancy accumulation.

A MULTIPHYSICS SOLUTION

STMicroelectronics chose COMSOL Multiphysics because it was able to efficiently handle all the physical factors that influence electron migration in metallic interconnects. The model couples several physics: standard diffusion due to concentration gradients; the "electron wind" driven by a chemical potential difference; hydrostatic stress and heat-induced atomic diffusion.

The model results show that the location of the void nucleation can be determined by the occurrence of a critical vacancy concentration. The model helps predict maximum concentration as a function of applied current, initial stress, temperature and, above all, the line's geometry. As a result a preliminary predictive model for the lifetime of metal interconnects was developed. These modeling results are important because accelerated tests must be very predictive and accurate. STMicroelectronics is currently using COMSOL Multiphysics to develop better predictive failure models that save time in the qualification process.

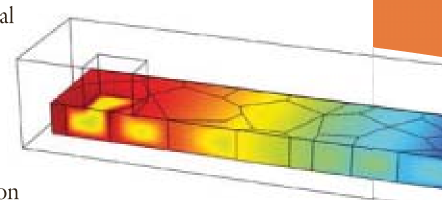

A vacancy concentration plot generated by STMicroelectronics modeling team using COMSOL Multiphysics.

Learn more at www.comsol.com/stm

COMSOL

ieee spectrum

volume 46 number 7 international

07.09

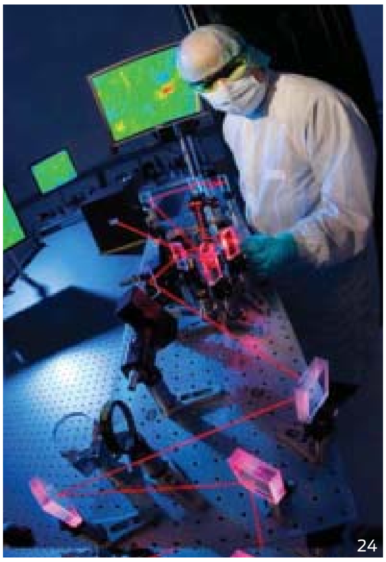

24

36

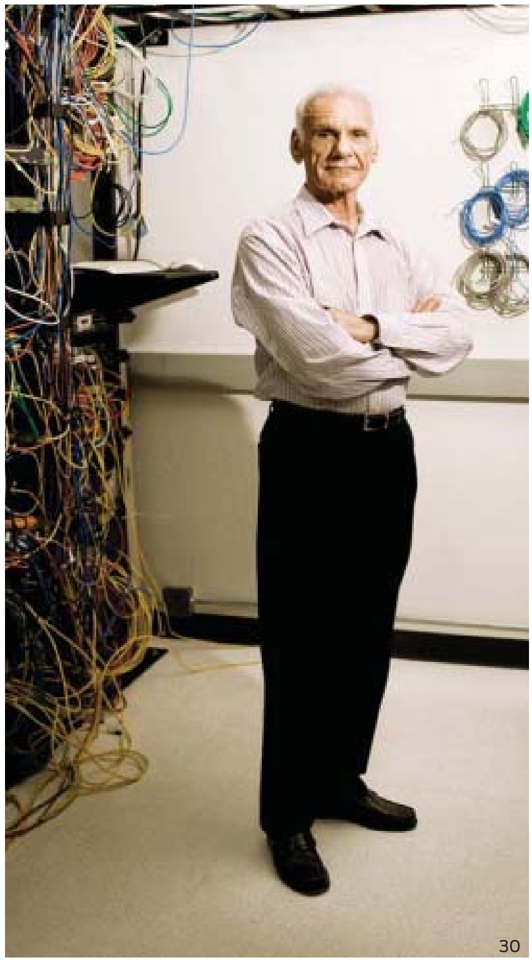

30

LASER FOCUSED: Northrop Grumman research creates a weapons-grade laser; Lawrence Roberts is wired up; it's all too quiet on Germany's green-energy front.

COVER: JONATHAN SPRAGUE/ REDUX

PHOTOS, CLOCKWISE FROM TOP LEFT: NORTHROP GRUMMAN; JONATHAN SPRAGUE/ REDUX; PETER FAIRLEY

volume 46 number 7 international

ieee spectrum 07.09

SPECTRUM.IEEE.ORG

THE WEB SITE: NEW! IMPROVED!

***IEEE Spectrum* has a new look online. But the changes are more than skin deep. So now that you're looking at the new home page [above], what should you do next?**

■ Browse the latest, most popular stories: Use the arrow buttons to flip through the top stories with big photographs. Or scroll down and click the tabs to view the most popular stories of the week.

■ Check out a topic of interest: In the black navigation bar, click to browse one of 11 categories or more than 60 subcategories, each filled with a mix of blogs, articles, and multimedia content.

■ Find exactly what you're looking for: Use the improved search function, which lets you search the entire site or narrow your search down to a particular content type or blog.

■ Interact: Leave a comment, anywhere—on articles, videos, or blog entries. Then subscribe to the comment RSS feed to keep track of what other visitors are saying.

■ Gorge on feeds: Click the RSS button in your address bar and get exactly what you want. You can subscribe to a single blog, a specific category of content, or just get one feed with everything. It's up to you!

■ Watch and listen: Click the "Multimedia" link at the top of the page to check out videos, podcasts, and slide shows.

■ Get related content: When you're reading an article, look to the right—the site automatically provides links to related *Spectrum* articles and relevant papers in the IEEE Xplore® digital library.

TECH INSIDER WEBINARS

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

Windows Embedded: The Platform for Connected, Efficient Industrial Devices

Creating Powerful Data-Capture Solutions

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

Multiphysics Simulation

WWW.IEEE.ORG/THEINSTITUTE

AVAILABLE 7 JULY ON THE INSTITUTE ONLINE

TECHNOLOGY FOR RURAL CLASSROOMS

Bringing educational materials to students in impoverished and remote areas is the goal of IEEE Fellow Krishna Palem. Learn about the LED solar-powered tablet PC, called the I-slate, that he and his team developed to replace a blackboard.

GETTING YOUR NAME OUT IN CYBERSPACE

The more your name appears on the Internet in connection with your field, the better your chances of landing a job, getting a plum assignment, or just making a new and useful contact, career experts say. IEEE members have an advantage over others because they can create an online presence in a number of ways. Find out how.

TECH IN THE MOVIES: JAMES BOND

James Bond movies are famous for their futuristic technologies and gadgets. Frederik Nebeker of the IEEE History Center discusses the technologies behind some of the top Bond gadgets over the years.

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.

TOP: LOIS HECHT; BOTTOM: JAMES D. GATHANY/CDC

back story

Lasers, Great and Small

As children, we fantasize about lasers that could blast a spaceship into smithereens. As adults, drinking gin and tonic outdoors on a summer evening, we'd settle for a much smaller laser to vaporize mosquitoes.

Jeff Hecht [above] is just the man to tell you how either one would work. He was putting the finishing touches on this month's "Ray Guns Get Real," about the U.S. military's program to create a massive solid-state laser, when some interesting news came out of Washington state: Researchers there had successfully tested a "mosquito flashlight," designed to kill the insects before they can snack on you. A team of astrophysicists formerly with Lawrence Livermore National Laboratory, including Lowell Wood, a protégé of Edward Teller's, had managed to build a laser accurate enough to fry mosquitoes in midflight.

Hecht was delighted. The veteran technology correspondent has been writing about lasers, optoelectronics, and solid-state physics since 1974 for such publications as *Laser Focus World*, *New Scientist*, and the *Bulletin of the Atomic Scientists*. In fact, he had come up with this same mosquito-killing scheme himself and submitted the idea in the form of a short science-fiction story to the journal *Nature*. The story, called "Directed Energy: More Than a Flash in the Pan," was published in 2006—a full year before Wood suggested his mosquito zapper.

The idea is destined to become more than a sci-fi figment. Bill Gates recently announced he would put the considerable weight of his charitable foundation behind the project, with the goal of eradicating malaria.

"Is it really possible that I came up with such a wild and crazy idea before Lowell Wood—the master of wild and crazy ideas?" Hecht asks. The Pentagon's research effort into laser "death rays" is celebrating its 50th birthday this year, but without delivering anything capable of torching a tank or melting a missile. So perhaps it's fitting that the first casualties of the laser death ray should be mosquitoes. "The way the funding for the military lasers is going," Hecht says, "I think the mosquito death ray just might make it into production first." ❑

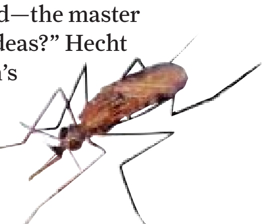

CITING ARTICLES IN IEEE SPECTRUM

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. 7 (INT), July 2009, p. 9, or in *IEEE Spectrum*, Vol. 45, no. 7 (NA), July 2009, p. 13.

contributors

ELISE ACKERMAN, technology writer for the *San Jose Mercury News*, began her tech-reporting career after covering the Clinton sex scandal in the 1990s. "Space missions and artificial intelligence are just more interesting than Monica Lewinsky's blue dress," she says. But Ackerman was never a stranger to technology. The daughter of a computer scientist, she grew up with *IEEE Spectrum* on the coffee table. In "Interplanetary Internet Tested" [p. 9], she looks at how a team from the University of Colorado at Boulder is working to extend Earth's Internet structure to space.

PETER FAIRLEY, a contributing editor, writes about energy for *Spectrum*. When he started research for "Germany's Green-Energy Gap" [p. 36], he anticipated that the nation's efforts to replace coal- and nuclear-fueled electricity with power from offshore wind turbines might provide a road map for other countries. He discovered, though, that Germany's green-energy push has stalled. His article reveals why. Fairley also writes about a strategy to reduce carbon-dioxide emissions from power plants by mixing biomass with coal, in "King Coal Eats Its Vegetables" [p. 12].

CLIVE FEATHER sifts through *Professor Stewart's Cabinet of Mathematical Curiosities* in "Puzzles by the Drawerful" [p. 21]. He says that if he had to pick a favorite trick from the book, the one he probably had the most fun with was "learning how to not slice off my fingers with a piece of string." Feather is an Internet regulatory expert and an amateur historian of the London Underground.

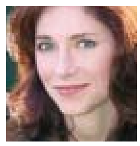

SUSAN KARLIN lists among her achievements acting, drawing, traveling to every continent on Earth, and writing for publications such as *The New York Times*, *Entertainment Weekly*, and *Spectrum*. For this issue she follows the trail of a coffee-making cellphone in "Phone-y Brew" [p. 18] and reports on an electrical engineer who helps museums spot fake van Goghs in "Art Fraud Forensics" [p. 19].

LAWRENCE G. ROBERTS led the team of scientists who developed ARPANET, the predecessor of the Internet, in the late 1960s. In "A Radical New Router" [p. 30], he proposes dumping his old router design for one that will better handle the complexities of today's audio and video Web traffic. Not that Roberts watches movies or listens to music online. "I don't get a lot off the Internet in that way," he says. "I don't even watch TV."

JAMES TURNER, Java expert and freelance writer, says he once wrote code for a portable Linux-embedded bar-code scanner with Wi-Fi to catalog his thousands of science-fiction books. He then printed the information on library cards and put them in a "genuine two-drawer card-catalog box I got off eBay." In "Time Is on Your Side" [p. 20], Turner describes his latest DIY project, building a straight-forward network-based clock.

CLOCKWISE FROM TOP LEFT: ELISE ACKERMAN; CARIN BAER; ANAGRAN; JAMES TURNER; LINDA FEATHER; RANDI SILBERMAN

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

spectral lines

IEEE Spectrum Online: Content Is Liberated!

LUCKY YOU: A quarter of a billion Web sites now compete for your attention. Meanwhile, desk jockeys are uploading 20 hours of video to YouTube every minute—that's 33 minutes of content a second. Twitter and other social networks, with their constant stream of updates, suck up whatever spare time you might have left.

It was against this tumult that we set out, 20 months ago, to revamp IEEE Spectrum Online (spectrum.ieee.org). The motivation was straightforward: To thrive, we've got to make our Web site a part of your daily media diet.

The new site was unveiled in early June, with everything you need to troll through our vast, nutrient-rich information or quickly zero in on a specific bit of knowledge. Instead of updating the site daily, we're blasting out new stories in multiple mediums several times a day, which you can now follow by using our RSS feeds. We've liberated our blog and multimedia content from the silos where they've languished far too long. Every hour, the latest entries from our blogs—Automaton, EnergyWise, Risk Factor, Sandbox, and Tech Talk—pop up all over the site, including in our topical, newly refurbished channels and subchannels.

In addition to our stalwarts Semiconductors and Computing, we've added new channels we think you'll find compelling: Aerospace, Green Tech, and Robotics, as well as the career-oriented At Work and the lifestyle-based Geek Life. These channels let you drill down into deep reserves of *Spectrum* content, much of which you may never have seen before.

We've also expanded our multimedia offerings by adding audio slide shows, which showcase our staff's expertise in gathering compelling audio—used in our broadcast-quality podcasts—along with stunning photography, which you can now enjoy on practically every page of our site.

And everywhere you go, starting with our easy-to-navigate home page, you will be guided by the choices you and your fellow visitors make via our traveling table of contents: Most Viewed, Most E-mailed, and Most Commented-on boxes appear everywhere to help guide you to the hottest content. Building a community around our site was so important to us—and to you, our loyal visitors—that we've improved the ease with which you can leave comments and join the conversation around topics of the day. Even the sponsors of our webinars are stoking the community vibe by offering series of live events on topics like robotics, which are drawing return visitors and building connections among engineers with similar interests.

A great Web site gives the illusion that it was effortless to create. Nothing could be further from the truth. The long process of improving our Web site started with usability studies of the existing site. As we watched through a one-way mirror, people attempted to perform tasks like signing up for a webinar, posting comments to a blog, or finding an article. It will come as no surprise that most Web sites have major "user experience" problems. Our site proved to be no exception.

Sufficiently chastened, we did some soul searching and then moved on to making it better. We had to rethink the way people find what they want, talk about it with each other, and keep coming back for more. For us, that meant building the site so that the content we produce on the back end automatically winds up in the right spots on the front end, freeing our writers and editors to do what they do best—keep you abreast of the latest technological innovations and provide you with a context in which to understand their implications.

Which brings us to our second usability study takeaway: We need to make sure Spectrum Online allows you to explore the content-rich world of both *Spectrum* and IEEE. In doing so we hope to promote IEEE to new audiences of potential members while connecting the dots for current members and other users. We're already automatically linking to related content from the IEEE Xplore digital library on pages throughout the site. And we're planning to bring you related content from IEEE's societies and conferences in the very near future.

We'd love to hear what you think about our latest effort and your suggestions for making things better. And for a list of all the folks who worked so hard to make the new site possible, go to spectrum.ieee.org/static/thank-you.

—HARRY GOLDSTEIN & JOSHUA J. ROMERO

forum

LETTERS do not represent opinions of IEEE. Short, concise letters are preferred. They may be edited for space and clarity. To post your comments online, go to 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.

NOW PROCESSING: "CHIPS"

I DON'T KNOW who to thank for "25 Microchips That Shook the World" [May]—especially the sidebar "My Favorite Chip." Wonderful article. Who would have anticipated that such great stories and such interesting information could have come out of one neat little idea. Kudos.

LARRY TOMENGA
IEEE Affiliate Member
Carrollton, Texas

HOW IN the world could you not include the Intel 4004? Not only was it the first microprocessor, but its invention and development is also one of the great stories of the semiconductor industry. Oh well, at least it's well documented on the Internet.

ROBERT GILCHRIST HUENEMANN
IEEE Life Member
Hollister, Calif.

I CAN SHED some light on the timeline for the Intersil ICL8038 waveform generator, as I was vice president of R&D at Intersil prior to cofounding Maxim Integrated Products in 1983. In the early 1970s, Jack Gifford and I contacted Hans Camenzind about doing a contract design for Intersil. Initially, we were thinking of an improved timer based on the success of Hans's 555 design. However, after some discussions the idea of a waveform generator evolved, which Hans designed for us. The part was introduced around 1973 and became an immediate success with hobbyists and low-cost-instrument makers. It was manufactured for the next three decades, a testimony to the soundness of both the product definition and its design.

DAVID FULLAGAR
Los Gatos, Calif.

BRIAN R. SANTO might consider correcting his statement that the ICL8038 chip was used in Moog synthesizers. I've been repairing Moog synthesizers for more than 33 years and have yet to locate any Moog that used this IC.

KEVIN LIGHTNER
Wrightwood, Calif.

The author responds:
About the ICL8038 being in the Moog, all I can say is what Bogart said about coming to Casablanca for the waters: "I was misinformed." I apologize for the error.

THE MOST important ICs ever have to be the Texas Instruments 7400 series. This set of chips formed the major stepping-stone to the digital age. Their compactness, functionality, speed, and logical completeness gave designers flexibility for creative and affordable designs. Future ICs followed the lead—and the rest is history.

JACK STERETT
Portland, Ore.

TIME TO GIVE CREDIT

I ENJOYED SHERRY Sontag's review of the video game *Time Engineers* ["Time Waits for No Engineer," Tools & Toys, May]. The game's development was funded by a grant to Valparaiso University through the Eli Lilly Foundation, and I served as project director. While Ray Shingler was subcontracted to create the game environment and all graphics, I want to be sure that my colleagues in the College of Engineering at Valparaiso, in Indiana, are recognized for contributing to the game. Five faculty members (including two IEEE members) developed the eight different engineering activities. They designed and tested the software for the analysis of each activity and wrote all the background information describing the engineering principles. The project was truly a joint venture between Shingler and Valparaiso University.

ERIC W. JOHNSON
IEEE Member
Valparaiso, Ind.

MEMORY REFRESH

IN THE article about Robert Dennard ["Thanks for the Memories," May], the definition of dynamic RAM, or DRAM, is wrong. The article says it's called "dynamic" because the act of reading the bit erases it, but the correct name for that behavior is "destructive read." The term "dynamic" is used because the memory must be refreshed to avoid losing data, as the article does note. Dynamic RAM is an alternative to static RAM, which does not need to be refreshed.

STEVEN ROTHMAN
IEEE Member
Bolton, Mass.

The editor responds:
We agree, regret the error, and thank reader Rothman.

CORRECTION

Veteran space reporter James Oberg is tall, but not quite as tall as we claimed he was in our June issue. The Back Story gave his height as 213 centimeters. In fact, he is 205 cm tall.

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

more online at www.spectrum.ieee.org

update

Interplanetary Internet Tested

Delay-tolerant network will weave a different kind of web in space

THE MANY paths a message can take through the Internet make that network robust and efficient—and the envy of those whose job it is to design communications schemes for the far-flung spacecraft leaving Earth each year. After more than a decade of development, NASA is in a rush to have a communications network ready by 2011 that can efficiently carry data between Earth and the multiple probes, rovers, orbiters, and spacecraft exploring the solar system—effectively binding them together to form an interplanetary Internet. Tests performed on the International Space Station last May were the second of three tryouts of the network's key technologies, called Delay Tolerant Networking, or DTN, protocols.

The DTN protocols will extend the terrestrial Internet into space by overcoming a number of obstacles, including the extraordinary length of time it takes packets to move between separate hops in a deep-space network, the intermittent nature of network connections, and bit-scrambling solar radiation.

"The communication delays are huge, and they are variable, because the planets are in orbit around the sun," says Vint Cerf, co-inventor of the Internet's TCP/IP protocol and a key member of a group of computer scientists who began working on DTN in 1998. On Earth, packets move from source to destination in milliseconds. By contrast, a one-way trip from Earth to Mars takes a minimum of 8 minutes. The constant motion of celestial bodies means that packets have to pause and wait for antennas to align as they hop from planet to probe to spacecraft.

So sending communications in space is very different from

SKY NET: Using computers on the International Space Station, engineers tested the protocols for a future interplanetary Internet in May 2009.
PHOTO: NASA

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

update

doing so on Earth, where the stable topology of the Internet is taken for granted.

"What we have to do instead is to tell all the nodes that these are the changes that are going to occur," says Scott Burleigh, a software engineer at NASA's Jet Propulsion Laboratory, in Pasadena, Calif., and one of the original developers of DTN. "You are going to be able to communicate from A to B at this data rate starting at 12:30 and ending at 3:30, and then you are not going to be able to communicate on that link anymore...until next Tuesday."

An initial test of DTN in space last October was successful. The code was loaded on a comet-studying spacecraft called Deep Impact as that probe headed out for a flyby of Comet Hartley 2. "We turned on the software on the spacecraft and on about a dozen nodes on Earth and just left it running, completely automatic for about a month," Burleigh says. During the test about 300 images were transmitted over distances that stretched up to 24 million kilometers. Although a couple of bugs were found, no packets were dropped, and no bits got corrupted. The software even survived the unintentional reboot of one of the Earth-based antennas. "The protocols underneath it were able to recover the data and actually get stuff through," says Keith Scott, a principal engineer at Mitre Corp., in Reston, Va., who has been working on DTN with Burleigh and others.

A key to DTN is a technique called "store and forward." Basically, every node hangs onto the data it receives until it can safely pass it on. On Earth, the data would simply get dumped if there was a problem and be retransmitted by the source.

The second test, conducted by Kevin Gifford at the Payload Operations Control Center at the University of Colorado, Boulder, used computers on board the ISS to send images to Earth.

For a third test, scheduled for early October and involving the Deep Impact spacecraft, engineers will introduce a security protocol as well as a new file-transfer protocol. After that, DTN will be "pretty much ready for deep-space research," says Jay Wyatt, the NASA program manager who has been coordinating the project. At that point, the researchers are hoping other space agencies will try it also.

Mitre's Scott chairs a working group at the Consultative Committee for Space Data Systems, an international organization that recommends standards for spacecraft communications. They are considering adopting DTN. Then, mission by mission, a network would grow, weaving an interconnected Web between the planets, the space station, and spacecraft. —Elise Ackerman

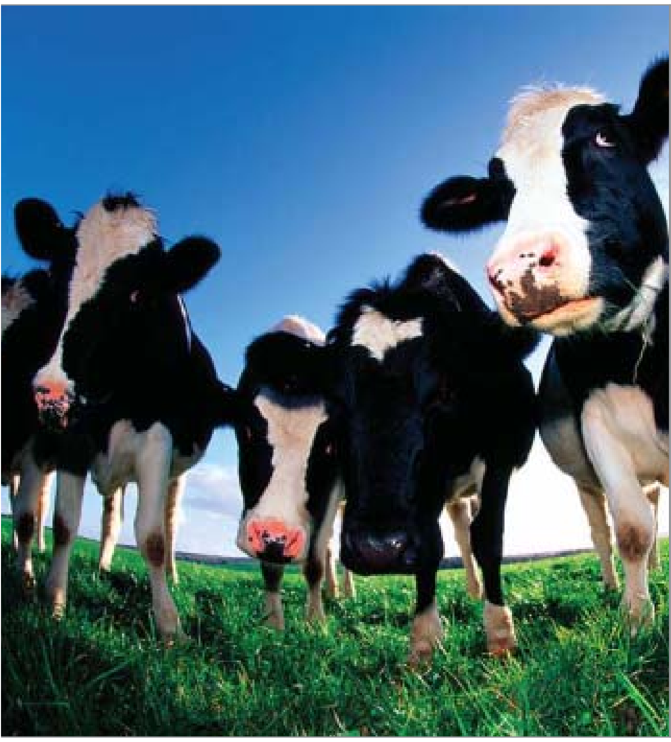

Of Cows and Power Lines

A team of researchers from Germany and the Czech Republic has discovered that, all factors being equal, cattle and two species of deer tend to align themselves along a north-south axis using some innate magnetic sense and that this preferred alignment is disturbed when they graze under high-voltage power lines.

For their research, Professor Hynek Burda and his team at Germany's University of Duisburg-Essen pored over images from Google Earth. When they spotted cattle that were at least 150 meters away from power lines, the researchers detected a significant alignment among the animals along the north-south axis. By contrast, images of cattle grazing within 50 meters of high-voltage overhead power lines showed no preferred body alignment. The exception was when the power lines ran east to west. In that case, the cattle tended to align along the direction of the power lines.

More at http://www.spectrum.ieee.org/energy/environment/of-cows-and-power-lines

PETER CADE/GETTY IMAGES

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

10 902 METERS The deepest dive ever, performed 31 May 2009 at the Mariana Trench in the western Pacific Ocean. The victorious vehicle was Nereus, an unmanned submarine tethered to the surface by an optical fiber. The sub, built by engineers at the Woods Hole Oceanographic Institution, can operate either remotely or autonomously.

SIZE MATTERS: Is there room in an iPhone world for the Mobile Internet Device?
PHOTOS: LEFT, CLARION; RIGHT, APPLE

Little Mass Appeal for Intel's Mobile Internet Device

A device to fill the gap between smartphones and netbooks struggles to find a market

THE MOBILE-computing-device market is a little like the car market. On the high end, you've got your big, fully loaded luxury models—that would be your large-screen laptops. On the low end, you've got your park-anywhere, use-minimal-power gizmos—your smartphones. In between, you've got a range of sedans and low-cost compacts—your tablet PCs, subnotebooks, and netbooks.

Seems like plenty of variety in that lineup, but Intel—for more than a year now and hot on the heels of its promotion of netbooks—has been evangelizing a new entrant, a device it calls the MID computer, for Mobile Internet Device. Others refer to it as the ultracompact computer or the tweener. As described at multiple industry events by Intel, such a computer would fill a niche, being more portable and less power hungry than a netbook, yet having a bigger screen and providing more functions than a smartphone. Intel's vision, originally a device about the size of a standard videocassette, seems now to be shifting to embrace multiple shapes and sizes.

It's no coincidence that Intel has the perfect processor to power such a computer—the Atom, launched last year. It's powerful enough to run full versions of standard software, like Microsoft Office. It's efficient enough to do so for hours without recharging. (You'll rarely see a netbook user today without a power cord.) So indeed, you could build a small, powerful, efficient MID computer.

But if you build it, will they come?

The answer, says Dave Blakely of Ideo, the Palo Alto, Calif., firm that did the product design for the original Apple mouse and the Palm V—is maybe.

"It is an entirely viable product for a small slice of the mobile computing population," says Blakely, Ideo's senior director in charge of technology strategy. "But that is a fractured slice. I can't name a single significant market segment or single killer app."

As examples of those slivers of the computer market, Blakely points to women who carry purses big enough to make the size difference between a smartphone and a MID computer irrelevant. They might pick the MID for its better movie viewing. He also thinks that road warriors who do a lot of videoconferencing might also prefer a device with a bigger screen than what's on a smartphone.

"I can name lots of potential users, but none add up to massive success," Blakely says.

Pankaj Kedia, director of global ecosystem programs at Intel, agrees that there are clusters of types of users rather than a single big market segment. He describes the MID computer as a flexible concept rather than a specific size and shape. "The devices are very personal, subject to personal taste," he says.

Says Kedia, "It may be pocketable, a smartphone replacement. For that I want the screen size to be 4 inches, to fit in my jeans pocket. For navigation, in your car, research has shown that a 5-inch display is more optimal. In the enterprise space, interviewing clients, looking at data, you need a 6- to 7-inch display."

Indeed, points out Ken Dulaney, a vice president and distinguished analyst with Gartner Research, Intel is right to back down on its original vision of the MID as a new class of all-purpose lifestyle computers, because that just isn't going to happen. "Intel badly miscalculated this space," he says. "The iPhone showed that people want a slightly smaller product that can be used as a convenient phone; you can't put a MID in your shirt pocket or hold it up to your head."

Still, it might not be smart to completely write off this device. "It has a shot," says Blakely. "We've seen a history of products with brand-new form factors appear in the world: There were people who thought the Sony Walkman would never make it." —TEKLA S. PERRY

news briefs

BENDABLE MEMRISTORS

Researchers at the National Institute of Standards and Technology report that they've developed a flexible memory that acts much like the new fundamental circuit element known as the memristor. The memristor's resistance changes according to how much current passes through it, and it maintains its resistance when the power is turned off. The NIST device still functions even after being flexed 4000 times.
PHOTO: NIST

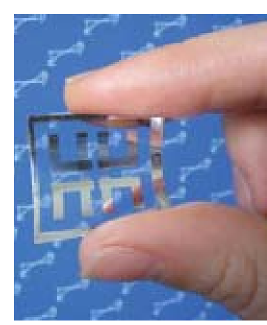

update

King Coal Eats Its Vegetables

Blending in biomass makes coal-fired generators seem a little cleaner

NIGEL RODDIS/REUTERS

CAPTURING CARBON-dioxide emissions and sequestering them deep underground is often presented as the sole means of significantly cutting the carbon emissions from a coal-fired power plant. In fact, this as-yet-unproven scheme is being beaten to the punch by a comparatively simple alternative: blending biomass into the plant's coal feed.

Adding agricultural leftovers, wood chips, or even dried sewage replaces some of the power plant's fossil carbon with renewable carbon. Such "cofiring" of coal and biomass is beginning to boom, thanks to preferential pricing and tax incentives supporting renewable energy, standards mandating utilities to obtain more of their power from renewable resources, and carbon caps that make utilities pay for CO_2 emissions. Europe, which already employs all three of these policies, is leading the way.

In the United Kingdom, cofiring has become one of the fastest-growing sources of power, following the introduction of the country's renewable portfolio standard in 2002. The UK's first cofiring operation started up the same year, and by 2006 cofiring was generating more than 2.5 terawatt-hours of electricity annually and displacing 2.6 percent of the power sector's coal use. More cofiring is coming, as Drax Power, operator of the UK's largest generator (4000 megawatts), in North Yorkshire, boosts the proportion of power the utility generates through cofiring from 2.5 percent to 12.5 percent. Drax boasts that from next year its renewable generation will rival that of 600 wind turbines, slashing annual CO_2 emissions by over 2.5 million metric tons.

Cofiring is attractive due to its low installation price relative to most alternative power technologies. Coal plants accommodate up to 20 percent biomass with a modest addition of equipment to store and handle the renewable fuels. The upgrade costs for cofiring range from US $50 to $300 per kilowatt. That's less than a third of the price tag for an onshore wind farm and one order of magnitude cheaper than building a standalone biomass generator.

Biomass burned with coal also delivers more energy per ton of fuel than biomass burned alone. The typically smaller biomass-only generators convert just 25 percent of their fuel energy into electricity. Conventional coal plants average 36 percent thermodynamic efficiency, and the efficiency lost when blending in modest proportions of biomass is small.

If carbon capture and storage does ultimately prove economic and safe, what then for cofiring? Cofired plants that include CO_2 capture provide a net extraction of carbon from the atmosphere. "If you combine the two, you have a sort of CO_2 vacuum," says Marc Londo, a biofuels expert at the Energy Research Centre of the Netherlands, in Amsterdam.

On the flip side, says Londo, cofiring may look less beneficial if carbon capture fails to pan out. That's because cofiring extends the economic viability of coal-fired power while diluting support for alternatives such as solar and wave power. For this reason, the UK has capped cofiring's contribution to its renewables mandate. Londo predicts that Europe may similarly limit cofiring's role in its mandates. As Londo puts it, "It's a bad idea to keep an old and relatively inefficient coal plant alive just via biomass." —PETER FAIRLEY

THREE The highest number a new genetically engineered strain of bacteria can count to. In their quest to create a computational tool kit for use in biotechnology, bioengineers at Boston University and MIT developed a genetic network that counts a type of molecular event in a microbe.

Is This the Moment for Broadband Over Power Lines?

Smart grids and the push for rural connectivity propel power-line data communications

A LONG-TIME DARK horse in the race to extend broadband access to the far corners of the United States—broadband over power lines, or BPL—may have finally found its golden moment.

BPL uses radio frequencies to impose the high-speed data signals on top of the AC that power lines carry. The broadband data signals are therefore prone to disturbance, particularly from the many voltage transients that power lines experience.

One reason for the renewed interest in BPL is the Obama administration's pledge to provide greater Internet access to underserved Americans, even those living in rural areas, where other means of providing broadband typically aren't economical.

Last February's American Recovery and Reinvestment Act will provide US $2.5 billion in loans and grants through one agency and $4.7 billion through another to expand broadband connections for residents of rural and underserved areas (as well as for public-safety agencies).

BPL provider International Broadband Electric Communications (IBEC), in Huntsville, Ala., is one firm likely to tap into this new government money. IBEC works with customer-owned electric cooperatives to give their members BPL, which costs less than other options because so much of the necessary infrastructure is already in place.

Brent R. Zitting, IBEC's chief technical officer, says that his company has been able to provide BPL in places with as few as three houses per kilometer at

RANDI SILBERMAN

data rates as high as 3 megabits per second for the company's premium service, which costs home users $89.95 per month. Most customers choose the $29.95 option, however, which gives only 256 kilobits per second—an order of magnitude less than the download speeds cable offers.

Still, when you're used to a 56-Kb/s dial-up connection, the availability of 256 Kb/s is well appreciated, says Zitting. He expects that with upcoming government loans and grants, his company will be able to expand BPL coverage to even more areas.

"We're following the same pattern as rural electric delivery back in the 1930s," he says, referring to the government programs during the Great Depression.

Electric utilities' increasing interest in obtaining real-time data on how end users are consuming electricity has also driven the resurgence of interest in BPL. Getting such data, of course, requires both suitably intelligent meters and a communications network to relay the information back to the utility.

In Boulder, Colo., for example, Minneapolis-based Xcel Energy is using BPL in combination with short-range radio links for its SmartGridCity pilot project. The links send data from power meters, hot-water heaters, thermostats, and renewable-energy systems.

According to Daniel Sangines, a communications engineer who until recently worked on SmartGridCity, the data flows along the power lines for about a kilometer before it's siphoned off the line and into an optical fiber or cellular-based backhaul system. Attempting greater BPL distances would require multiple repeaters to deal with signal attenuation, reducing the bandwidth unacceptably, he explains. In Boulder, the smart-grid network can provide a typical home with at least a 5-Mb/s link. Rather than sell the bandwidth, Xcel is reserving it for future smart-grid applications.

While BPL seems a natural match for the problem that electric utilities are trying to solve, not all smart grids are using it. Some, such as the vast grid just proposed for Miami and environs, will likely use wireless ties rather than BPL. Still, given the new funding opportunities for rural broadband and the burgeoning web of smart grids, BPL is sure to keep a foot in the ongoing race to network people and their power-hungry appliances. —DAVID SCHNEIDER

news briefs

ZAP!

Engineers at the Fraunhofer Institute for Laser Technology, in Aachen, Germany, have invented a laser system that can drill 3000 holes per second in a solar cell during the manufacturing process. The holes improve the system of electrical contacts and thereby the cell's efficiency.

PHOTO: FRAUNHOFER INSTITUTE FOR LASER TECHNOLOGY

1 BILLION YEARS Theoretical length of time a memory cell made from a carbon nanotube and a nanocrystal of iron can hold data, according to researchers in Berkeley, Calif.

update

SEA SOURCE: The city of Aqaba will get its water supply from the ocean.
PHOTO: SANDRA UPSON

Jordan's Red Sea Desalination Plan

At the Aqaba Water Co., artificial intelligence saves water

For a small city in a water-poor country, Aqaba got lucky. The city of 110 000 sits on the Red Sea, on Jordan's only coastline. Residents get their water from 20 wells in the nearby Disi aquifer, a store of ancient groundwater that straddles the border with Saudi Arabia. While the rest of Jordan makes do with a few hours of water service a week, Aqaba's supply is uninterrupted.

Some of that happy circumstance is self-generated, with computer-controlled water infrastructure and a new artificial intelligence system that will soon manage most of it. Now Aqaba is becoming a linchpin in Jordan's grand water strategy. By the end of next year, the city plans to start building Jordan's first seawater desalination plant, which will provide 10 million to 15 million cubic meters of water per year, matching Aqaba's current usage. The reverse-osmosis treatment plant will be the first step in an ambitious plan to build a canal to send water from the Red Sea to the shrinking Dead Sea, generate hydropower on the canal, and install another desalination facility along the way. "Our project will be a pilot for Jordanian engineers to gain experience and prepare for larger desalination projects in the future," says Imad Zureikat, general manager of Aqaba Water Co., the city's utility.

With investors across the Gulf pouring money into Aqaba, the utility expects demand to double in five years. Business-friendly regulations introduced in the last eight years have begun to turn the port city into a miniature Dubai, bringing new spikes in water use, Zureikat says.

To meet that burgeoning demand, the Aqaba Water Co. is implementing new management techniques. The utility is already a model of efficiency for the country, if not the region. In stark contrast to the surrounding desert, Aqaba's verdant palm trees and grassy patches thrive on treated wastewater, or "gray water," distributed through a separate network—a classic conservation scheme that many advanced industrialized countries lack. The rate at which water is lost through leaks or metering failures is less than half that for the rest of the country, which loses 43 percent of the water entering its networks.

Naem Saleh, the water company's technical and engineering director, attributes the difference to the company's move, in 2004, to implement a supervisory control and data acquisition (SCADA) system that reads the network's flow characteristics every few milliseconds. "The first thing in our strategy was to automate and computerize everything," Saleh says. Working with the software company Oracle to integrate the individual controls, the utility is now adding artificial intelligence to react swiftly to shifts in demand. Its decision-making algorithms will be trained on the last two years of operational data. When complete, the system will control the pumps that fill the reservoirs and adjust pressure levels in the pipes to match consumption—for example, to keep the pressure around 500 kilopascals during the day and 200 kPa at night, when pressure can build up and cause pipes to burst. It would also maintain pH and chlorine levels and, eventually, automate the work of about 20 field technicians.

With Aqaba using mostly desalinated water, the water in the Disi aquifer can be preserved for Amman. A separate project will transport the Disi water 325 kilometers uphill to Jordan's capital. That supply will buy Jordan the time it needs to build larger desalination plants and truly harness the Red Sea. Nisreen Haddadin, an engineer managing the Ministry of Water and Irrigation's master plan, sums it up concisely: "This is our dream." —Sandra Upson

This is the first in a two-part series on water technology in the Middle East.

the big picture

DOPPELGÄNGER DANCE

Imagine how you'd feel meeting a version of yourself that never ages and will never die. That's what the female lead character in the new play *Robots* must deal with. The machine is one of three (including a butler and a dog) created by the male lead, who's on the brink of swearing off human contact in favor of his preprogrammed companions.

The play's robots are the result of a 10-year collaboration in Lausanne, Switzerland. BlueBotics, a firm specializing in autonomous robots, came up with a laser-based guidance system and a scheme that uses something like a Musical Instrument Digital Interface protocol for controlling the humanoids' every move. François Junod, a renowned maker of lifelike mechanical moving figures, supplied the faux-female machine's fine motor coordination, allowing "her" to make subtle gestures. (*Robots* ran from 1 to 16 May at a theater in Servion, just outside of Lausanne.)

PHOTO: VALENTIN FLAURAUD/REUTERS/LANDOV

geek life

Phone-y Brew

How a fake coffee-making phone is rebranding a Canadian province

At first, the Pomegranate NS08 seems plausible as a sleek next-generation cellphone. But the list of features goes on just a tad too long to be true: a built-in translator (cool), a coffeemaker (please, please), a harmonica (huh?), and finally... an electric shaver (okay, you got me!).

But the Pomegranate exists only in a YouTube video and a Web site designed to draw you to another Web site, that of the "Come to life" initiative in Nova Scotia, Canada, a partnership between the government and 300 private associations, schools, and businesses located in the coastal province, which is due east of Maine. The idea is to rebrand a rustic vacationland as a bastion of innovative minds, creative businesses, and a balanced quality of life.

In the first six months after its September launch, the Web site registered more than 1 023 000 visits from 198 countries, led by the United States, Canada, Germany, Italy, and Britain. An astonishing 64 percent of visitors clicked through from the fake phone to the Nova Scotia pages.

At least one featured business, Ross Screenprint, in Antigonish, received calls for new business, and a Greek film festival wants to screen the site's opening phone sequence. But success will be measured not by buzz but by results.

"Our goal was to reach key influencers—media, senior politicians, CEOs, people addicted to their BlackBerrys," says Stacey Jones-Oxner, communications advisor of the Nova Scotia campaign. "A new gadget that did everything was the best way to grab those folks, so we created one."

The site leads users through increasingly fantastic features until you click "release date" or "I've seen enough." (You'd think I'd have reached that point when water was poured through

WILDLY FUTURISTIC: The Pomegranate phone can brew coffee, translate foreign languages, and be played like a harmonica. *IMAGES: PROVINCE OF NOVA SCOTIA*

the phone to brew coffee, but I kept watching!) Then the punch line pops up: "Someday you'll be able to get everything you want in one device. Today you can get everything you want in one place." The site then pitches the Nova Scotia lifestyle, businesses, education, culture, and local entrepreneurs.

A creative team consisting of Web programmers, video producers, and an ad agency spent nine months putting together the campaign. They stamped real pomegranates with the campaign's URL and passed them out in Ottawa, Toronto, and Boston. Then word of mouth took over.

While no one has seriously tried to buy the Pomegranate, Jones-Oxner adds, "We did get a call from Germany telling us that if we decided to actually make this phone, they would distribute it. So we may be onto something here."

—Susan Karlin

careers

ART FRAUD FORENSICS

An engineer helps curators foil forgers

How many engineering jobs let you take a van Gogh off the wall and hold it in your hands? The kind C. Richard Johnson Jr. landed. He's both an electrical engineering professor at Cornell University, in Ithaca, N.Y., and an adjunct research fellow at the Van Gogh Museum, in Amsterdam. As such, Johnson says, he can "speak the language of people on both sides."

And when the two sides talk, they mainly talk about fraud and how to detect it. Two years ago, Johnson organized a conference at the museum that brought together researchers from Pennsylvania State University and Princeton, in the United States and Maastricht University in the Netherlands. Together, they processed high-resolution images with specially designed signal-processing algorithms to help sort fake van Goghs from real ones at the brushstroke level. It was the first time that image-processing teams at different universities could compare authentication approaches on the same paintings. Another workshop will follow next year at the Museum of Modern Art, in New York City.

"Fraud detection is a 'sexy' topic, which is why it was an early focus of my activities," says Johnson. "But we're 10 to 15 years away from the computer having any authority in it. So now my colleagues and I are pursuing a wide variety of issues of interest to conservators and art historians, where signal processing can provide assistance that reaches well beyond just the detection of frauds."

Johnson's current focus is on canvas thread counts—the number of horizontal threads crossing a vertical line 1 centimeter long—to identify paintings from the same roll of canvas. "Placing a questioned painting on the same canvas roll as a painting known to be from a particular artist supports authentication to an artist who bought canvas in rolls, as van Gogh often did," he says. "When canvas is prepared with a lead white ground, the grooves between the threads are filled with radio-opaque material," says Johnson. "This registers in an X-ray as an intensity pattern that reveals the individual threads, permitting a calculation of the weave density." The pattern is then analyzed with a Fourier transform, the same technique that radio engineers use to break down a signal into a series of simple sine waves.

PAINTING BY NUMBERS: C. Richard Johnson [left, center] uses signal-processing algorithms to authenticate canvases believed to be painted by van Gogh.
PHOTOS: VAN GOGH MUSEUM

The team is distributing the software free to museums. The Van Gogh Museum already uses the data generated to identify paintings from the same canvas roll by determining how the sections were arranged on the roll before being cut for use.

Johnson stumbled into art as he wandered through Berlin museums during a college year abroad while earning a bachelor's in electrical engineering from Georgia Tech. Later, while working on his Ph.D. in EE at Stanford, he took a class in the Dutch masters, which rekindled his passion. In 1977, he became the first Ph.D. student to graduate from the university with a minor in art history.

He went straight into academia, teaching at Virginia Tech until 1981, when he moved to Cornell. He was named an IEEE Fellow in 1989 for his work in digital control and signal processing.

"This kind of research is not something to recommend to Ph.D. students. There are no jobs, no one's eager to fund this, and it's career killing for any pretenure academic," he says, laughing. "But for me, it's like having a backstage pass. I go to a conservation studio and can take a van Gogh out of its frame and examine it." —Susan Karlin

hands on

TIME IS ON YOUR SIDE

Or at least it will be, once you build this network-based clock

There's no longer any excuse for having the wrong time. Computers, set-top boxes, and even some wristwatches can get their time from the U.S. Naval Observatory or some other authoritative source. But what about a device you build yourself? Many will come with a timing chip that you can use as a counter, but they won't tell you what time it is in the real world or self-adjust for daylight saving time. What you want is a microprocessor on a computer board that will query a network for the correct time and pass it on to the rest of your device. This month's project is just that, a clock on a single-board computer that uses the Network Time Protocol (NTP) to give you millisecond accuracy for about US $115.

Single-board computers (SBCs) are cool because they pack a lot of functionality into a very tight package. They're scary, too, because they involve you in technologies that are usually taken care of by off-the-shelf computing hardware. For one thing, SBCs don't usually come with such niceties as displays and USB ports. And when you get down to the smallest of the small, you are pretty much dealing with a processor, some flash memory, a few input/output ports, and a way to burn an image to an electrically erasable programmable read-only memory (EEPROM).

It can get pretty complicated, so a kit makes for a nice introduction. Tuxgraphics, of St-Laurent, Que., Canada, has one for €50.50, or about US $70. You'll get an SBC based on the ATmega168 processor, an 8-bit reduced-instruction-set computer chip running at 20 megahertz. The SBC has 16 kilobytes of flash memory, 5 kilobits of EEPROM, and a whopping 1 KB of RAM, so you're not going to be booting up Linux or running Photoshop here. But what the board lacks in memory, it more than makes up for in hardware—it has an Ethernet port, digital I/O ports, eight analog inputs, and an LED socket.

The second major component in the box is a 16- by 2-character LCD panel for displaying the time. Then there are a few odds and ends—a 3.3-volt voltage regulator, an LED, and a 5-pin header for programming the device—which all get soldered directly onto the SBC. You need your own wire to hook the SBC up to the LCD. I used 22-gauge rainbow speaker wire. There are also three resistors that attach to the LCD module. Two control contrast and brightness, and one goes between +5 volts and the LCD backlight power.

You solder in the voltage regulator and the LED on the SBC, attach six wires between the SBC and the LCD module, solder on the programming header, and finally run +5 and ground to the SBC and LCD module. The holes are tightly spaced, so if you haven't picked up an iron in a while, you might want to get in some practice first.

Along with the wire, you will also need to provide a source of +5. I was able to find a leftover wall wart with just that voltage, but watch out—some sources claim +5 but really put out far more.

The first time you power up you'll see only the LCD's backlight—the SBC comes with nothing installed on it. To get the board running, you'll need to burn software and data onto the EEPROM and the flash memory, and therein lies the major "gotcha" of this project.

First, you're also going to need to buy a programmer for the SBC. Tuxgraphics

THE RIGHT WORD, ANYWHERE

As a teacher of creative writing, I have one technotrick I couldn't live without—looking up the exact definition of a word on my iPhone while standing at the chalkboard. There are a number of mobile dictionary and thesaurus applications now, but the best ones on a phone are the same as on the Web—Merriam-Webster and Dictionary.com.

The Dictionary.com app is free and comes with a companion thesaurus.

Merriam-Webster's is a pricey US $60 but is based on *Webster's Third New International Dictionary, Unabridged,* the aircraft carrier in a bookstore's fleet of American dictionaries. (*Merriam-Webster's Collegiate Dictionary* is also available, for $25.)

The databases go right on the phone, so there's no time-consuming lookup that depends on a good connection to the Internet cloud. If you love words or need them for your work, these apps are indispensable. —*Steven Cherry*

http://i.word.com
http://dictionary.reference.com/apps/iphone

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

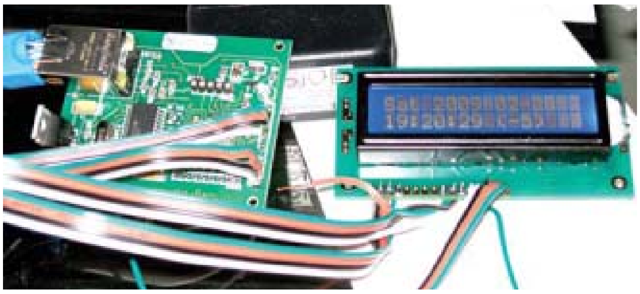

TIME TELLER: Your other do-it-yourself computer projects will know what time it is if you add a single-board computer, such as this one, from Tuxgraphics. It uses the Network Time Protocol to get continual updates from authoritative time servers on the Internet. PHOTOS: JAMES TURNER

sells several; the one I got cost €40.50 ($53.58). Essentially, it's a USB cable with a processor built into it on one end and a 5-pin connector on the other end that plugs into the SBC. The programmer also comes with a LiveCD of a Linux distribution prebuilt to program the SBC.

If you don't feel like booting up Linux every time you want to program the device, the kit also comes with a version of GCC (GNU Compiler Collection) that runs on your regular desktop and compiles for the ATmega168 instruction set, together with the AVR Downloader/Uploader (AVRDUDE), a tool that can communicate with various programming peripherals (including the USB cable I had bought), so that you can program directly on the device. The CD includes precompiled versions of a test program that flashes the LED and displays an OK message on the LCD, a good initial smoke test to ensure you didn't fry anything.

The disk also includes a copy of the NTP application ready to run, but you'll probably need to make some changes—it's preset for a 10.0.0.0 class A subnet, which is ideal for General Electric or MIT but not typically used with a home network. You're going to edit the main.c file, changing the IP address of the device, as well as the IP address of your gateway, the IP address of the NTP server you want to use, and your time zone.

With these edits made, you use the "make" command to build a hex file with the compiled code, and then the avrdude command to write it out to the SBC, which needs to be powered up and connected by the programmer cable. Once installed, the program will try to connect through the gateway you defined to the NTP server. Then it'll start displaying the current time on the LCD, syncing with the server once every hour after that.

For me, the hardest part of the project was the soldering. All told, it took about 3 hours (and several trips to RadioShack) over the course of a week or so to get everything hooked up. Once the programming cable came, it was another 2 hours to get all the software installed and burn the program onto the device.

The final product is a plain black box that tells me what time it is, as long as I have a network connection handy. It also will provide the same data if I point a browser at the IP address of the device. While the project was interesting for its own sake, I'm looking forward to seeing what else I can do with a network-capable SBC in a box, because I can burn onto it whatever software I care to write. —JAMES TURNER

books

PROFESSOR STEWART'S CABINET OF MATHEMATICAL CURIOSITIES

By Ian Stewart; Basic Books, 2008; 256 pp.; US $16.95; ISBN: 978-0-465-01302-9

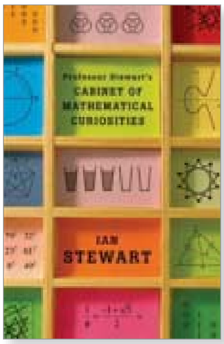

PUZZLES BY THE DRAWERFUL

So why *did* the chicken cross the Möbius strip?

There's a lot to be said for a book that transports you back to your childhood. When I was 11, I was given a subscription to *Scientific American*. Every month, I would turn directly to Martin Gardner's Mathematical Games column, which *Professor Stewart's Cabinet of Mathematical Curiosities* immediately brings to mind.

Open one of the 179 "drawers" in Professor Stewart's cabinet, and you might find just a one-liner ("Why did the chicken cross the Möbius band?") or a seven-page essay on Fermat's last theorem. Many items, like how to find a fake coin in three weighings and Cantor's diagonal argument, are likely to be familiar friends, but plenty of others will be new. And Stewart has put a sheen on some of the oldies—Langton's ant, for example, in which the eponymous insect travels around a checkerboard according to simple rules. I hadn't imagined some of the strategies that might somehow trap it.

I'm not keen on the book's shaggy-dog stories nor on its biographies of famous people (do Americans know who Virginia Wade and Carol Vorderman are?). But those parts of the book are more than offset by a variety of knots, magic hexagons, square wheels, and topology tricks in which you pretend to slice off your finger. Answers (but not, unfortunately, a proof of the Goldbach conjecture) are to be found at the back of the book, and many have URLs for further information.

The book can be devoured in one giant gulp or savored, one curiosity at a time. As for the mathematics, the puzzles require at most a bit of algebra. Many should be within the grasp of an intelligent 11-year-old. —*Clive Feather*

tools & toys

IntelliTouch Eos Wireless
US $250 and up
http://www.eoswireless.com

BIG HOUSE, CHEAP AUDIO

The Eos multiroom sound system isn't slick, but it won't stress your budget

FOR SEVERAL years now I've had Sonos envy. Several of my friends own these high-end multiroom audio systems; when dinner-party conversation lags, we pass the handheld controller around the table to play memorable songs and reminisce. (Okay, a bit strange, but this is Silicon Valley.) Unfortunately, Sonos systems are complicated and costly: Just the starter pack, with the wireless controller and boxes to hook up two rooms, sets you back US $1000, and that's without speakers.

So when I heard that a company named IntelliTouch, in San Diego, would soon offer a low-cost multiroom audio system called Eos, I was excited. "Low" is a relative term, of course. The Eos Wireless starter pack—a base station, a remote speaker unit, an audio cable for hooking up other components, and a remote—lists for $250 but sells for as low as $150. (The controller is built into the speaker.) Additional speakers are about $130 each.

Eos, of course, does not do as much as Sonos. It doesn't connect to online music services without a computer. Nor can it send different audio to different rooms. It is, however, unbelievably easy to set up. I put the base station and four speakers in five different rooms, and the only hard thing was finding available outlets.

The sound quality was comparable to that of most iPod docking systems, but it was better than most when I installed one of the speaker units in the same room as the base station, for an improved stereo effect. (Each speaker has 2.1 surround sound and a subwoofer.) The advertised range is 50 meters; my house isn't quite that big, but a speaker placed outside, plugged into a garage outlet, worked fine.

Eos broadcasts in the 2.4-gigahertz frequency band, right in there with cordless phones, microwave ovens, and Wi-Fi. But it uses a proprietary spread spectrum, frequency-hopping technology that the company says prevents interference—and indeed, I didn't experience any.

The Eos remote controller works only on a line of sight from the base station, so to change a song I have to run downstairs. That makes Eos more like an iPod dock than a true multichannel multiroom system. But the multiroom feature was a huge boon when my family didn't want to interrupt its weekend chores, in and out of the house, for the Obama inauguration concert. I repositioned the four speakers into critical areas, including one out on the front porch and one in the backyard. Eos saved that day.

I do have a few quibbles. The power adapter design is conceptually clever: It looks like it's built into the speaker, so you can plug the speaker directly into the wall, but you can also pop the adapter out of the speaker, revealing a cord you can plug into a less accessible outlet or power strip. Fabulous idea, but the implementation wasn't so successful. One of the power supply slots was jammed shut, so I couldn't use that speaker on a tabletop; it had to stay with the wall outlet. Another speaker had the opposite problem: The spring door was broken, so I could use it only with the cord out. Admittedly, these were press loaner units that had likely taken some abuse.

Also, there was a slight delay between turning on a remote speaker and getting sound out of it. Now, I know there has to be a delay. But every time I turned a speaker on, I found that I'd keep dialing up the volume, thinking that I wasn't hearing the music because it was turned down too low. By the time the signal kicked in, the sound would be so loud that it was painful.

Would I buy Eos? Probably not as a multiroom audio system—not until the remote works without line of sight. Would I buy Eos if I was shopping for an iPod dock? Definitely. Its price is competitive in that market, and being able to move speakers around or easily set up an indoor/outdoor sound system is a great bonus. —TEKLA S. PERRY

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

reflections BY ROBERT W. LUCKY

Wicked Problems

SOME PROBLEMS have such complex social, economic, or organizational interactions that they can't be solved fully. They've become popularly known as "wicked problems."

I'm not fond of these problems, though I've seen my share of them—they seem to be ubiquitous in systems-engineering design. However, I'm fascinated with the name. It reminds me of "fuzzy logic," a brilliant and oxymoronic phrase that juxtaposes an adjective connoting warmth and softness with a noun that implies something cold and mechanical. In "wicked problem," an adjective meaning evil or sinful, usually assigned to humans, is attached to an abstract, inert noun. The name suggests that the problem itself is consciously malicious. It knows that someone is out there working on it, and it is going to stop that person from getting anywhere.

Wicked problems are unsolvable not because they are mathematically difficult but because we can't define them well enough to quantify things. We don't encounter such problems in our engineering education, where all the exercises have solution paths, usually to precise numerical answers. If only the real world were always so obliging! I am reminded of Douglas Adams's famous fictional supercomputer built to answer the one question: "What is the meaning of life?"

MGM/THE KOBAL COLLECTION

This surely is the mother of all wicked problems. As many readers of *The Hitchhiker's Guide to the Galaxy* know, after long consideration the computer returns the answer "42." We sometimes give similar oversimplified solutions to problems that are, well, wicked.

Invoking another evocative name from modern physics, wicked problems demonstrate entanglement: What we do here influences something else way over there in some mysterious way that we don't fully understand. When the ramifications of our design extend beyond our organizational, knowledge, technical, or authority boundaries, we partition the problem and draw a virtual box around the part we control, pretending that no effects propagate beyond this box.

Perhaps all real engineering problems are in some sense wicked, but many critical problems in government may be especially so. Consider airport security. Every time I take off my shoes and walk through that strange freestanding doorway, I think there must be a better way. But we lack the data needed for an engineering solution. We can try to calculate the costs and benefits: We know the salary and equipment expenses, and we can measure the probabilities of missed detection and false alarms. But the true cost of the system should include the lost time, uncertainty, and aggravation of millions of airline passengers, as well as the lost revenue from would-be passengers who are thus discouraged from traveling, and so forth. These costs would have to be weighed against the incalculable cost of failing, even once, to prevent an airline disaster. To make matters worse, we understand only poorly the motivations and plans of potential hijackers. And as in almost all such problems, there is an overriding question—if we didn't spend the money here, where might it be better used? On the other hand, what would be the costs to society if the public at large lost confidence in air safety and stopped flying altogether?

Problems in allocating defense acquisition funds are similarly wicked. For example, what are the costs and benefits of improving the surveillance capabilities of drone aircraft versus investing in satellite technology? Each promises better intelligence, but what is a pound of intelligence worth? Isn't the ultimate objective to win the war? But what war did you have in mind? And what did you mean by "win"? The unanswerable questions ascend into the skies, far above your pay grade, whatever it may be.

Because such problems are wickedly unsolvable, the practice is to set immutable specifications for the constituent subsystems, leaving little opportunity to trade off cost, performance, and scheduling parameters—and leaving equally little room for the analytical and methodical approach that is the essence of engineering.

Wicked problems will never be solved in the conventional sense, but we engineers should do our best to bring this kind of thought to any problem, be it wicked or not. ❑

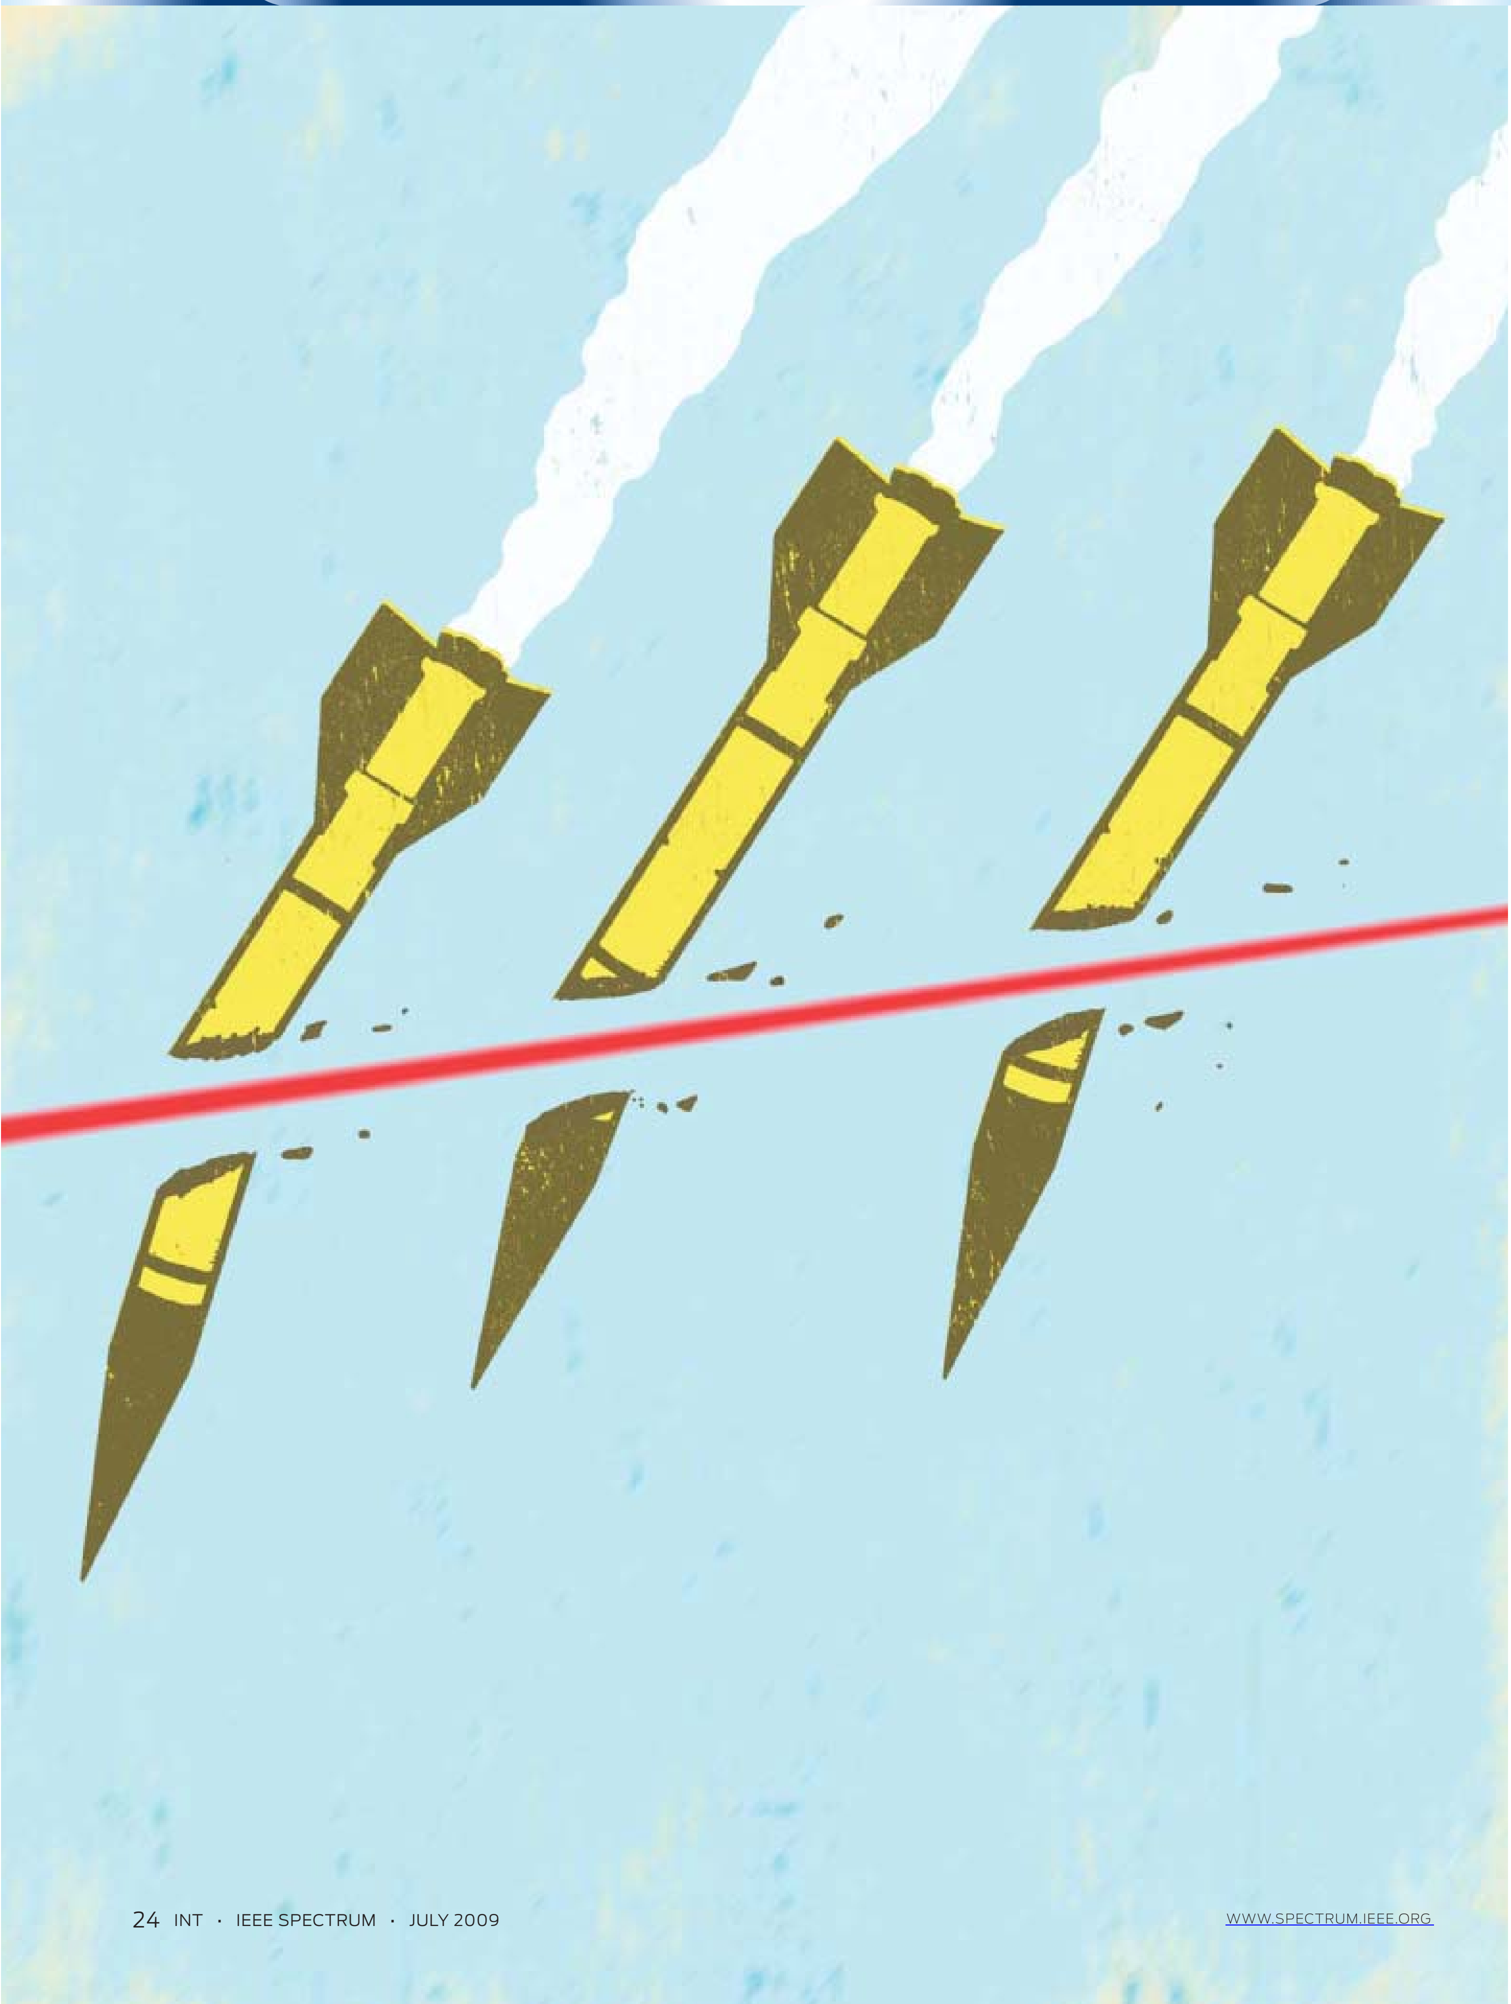

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

RAY GUNS GET REAL

Cheap rockets fired by insurgents are taking a deadly toll in the Middle East. Can a new generation of solid-state lasers blow them out of the sky?

By JEFF HECHT

ROCKETS, MORTARS, and other forms of artillery have a long and grim history on the battlefield. In a conventional war, an army being bombarded by these from afar can respond by firing back at the attacker's battery. But you can't turn the massive firepower of modern armies onto insurgents hiding among civilian populations without courting disaster. Instead of striking the enemy, who run to other hiding spots after firing their weapons, such retaliation would mostly hit civilians.

What the U.S. military dearly wants is a weapon that can defend against such attacks more selectively, shooting down explosive-laden projectiles in the air before they reach their targets. The armament should be easy to field and should strike at the speed of light, but it should not send streams of bullets screaming toward the horizon. In short, the military wants a laser weapon that's small and rugged but powerful enough to ignite explosive payloads on incoming projectiles while they're still a safe distance away.

It's a bold vision for laser defense, bolstered by a dramatic technology demonstration that didn't make Page One: For five solid minutes in March, an electrically powered solid-state laser pumped out 100 kilowatts of infrared light, the first of its kind to make "weapons class."

Each armed service has its own plans for that technology. The U.S. Army and its Israeli allies want truck-mounted lasers to zap short-range rockets on the battlefield or border. The U.S. Air Force wants compact lasers for fighter jets. The Navy wants to defend ships against attacks. And research efforts in China and Russia have been reported as far back as 1995.

And yet, laser weapon R&D is celebrating its 50th birthday this year without much to show for it. In fact, in early April the U.S. Defense Department shelved plans to buy a fleet of 747s to house giant gas-filled antimissile lasers. The old technology was proving too bulky and underpowered to blow North Korean missiles out of the sky without flying within antiaircraft range.

High-energy laser research is at an inflection point. Powered by semiconductors, a new generation of lasers promises new opportunities—and presents a whole new batch of problems.

LASER WEAPONS, like flying cars, have been demonstrated many times, but in the real world their problems have always outweighed their benefits—literally. Weight cripples laser weapons and flying cars alike. Most experimental laser weapons have been so big and heavy that cynical observers have joked that their only conceivable combat use would be to drop them on the enemy.

That's because the size of a laser weapon is inversely related to its efficiency—and laser efficiencies can be pretty dismal. The red helium-neon gas laser long used for classroom demonstrations turned only 0.01 to 0.1 percent of electrical power input into light. The diode lasers used in today's inexpensive laser pointer do much better, converting about 10 percent of the electrical energy they draw from their batteries into light. The rest is lost as heat. This is no big deal for a milliwatt-power laser pointer, because the heat generated is negligible. But it's a thorny problem for a laser weapon. At 10 percent efficiency, it would take 1 megawatt to generate a 100-kilowatt laser beam, leaving 900 kW as heat that must be dissipated somehow.

But that didn't stop the U.S. Missile Defense Agency from building a mega-

DAN PAGE

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

watt laser. To achieve a 1-MW beam with 10 percent efficiency would require a whopping 10 MW of input energy and produce a hefty 9 MW of waste heat. Nevertheless, later this year a beast with such power, called the Airborne Laser (ABL), will be put to the test of blasting dummy nuclear missiles from the sky.

Here's how. ABL is the latest example in a class of high-energy lasers called flowing-gas lasers. They are powered by burning chemical fuels like those that drive rocket engines. Hot molecules in the gas emit a cascade of light emissions, producing a powerful laser beam. Rocket-engine lasers have generated infrared beams that can reach a couple of megawatts for a few seconds at a time. The technology used in ABL can turn more than 20 percent of the combustion energy into laser light in the laboratory, but ABL's efficiency is undisclosed. In such a laser, the exhaust gas carries away the energy left behind as heat.

But so far the US $5 billion ABL can barely squeeze into a Boeing 747. The laser is completely unsuited to the battlefield. It's being designed to destroy long-range missiles rising through the atmosphere a couple of hundred kilometers away, but it's vastly overpowered for the comparatively easy job of hitting slow-moving mortar shells only a kilometer or two away. It would be like shooting deer with a cannon. So in 1996 the U.S. Army and the Israeli Ministry of Defense teamed up to test smaller lasers against mortars and rockets. For that task, they tapped Redondo Beach, Calif.–based aerospace contractor TRW (acquired by Northrop Grumman Corp. in 2002) to build a 100-kilowatt-class flowing-gas laser, a compact version of ABL.

The result, called the Tactical High-Energy Laser (THEL), made laser defense look promising. In 2000, it shot down a short-range Katyusha rocket over the White Sands Missile Range in New Mexico. But by 2004 the United States and Israel agreed THEL wasn't up to the job, ending any further tests.

One problem with THEL was its large footprint. It required several trailer-size containers; outdoor tanks filled with exotic chemical fuels and the "exhausted" chemicals that had fed the reaction (which would need to be protected from enemy fire or sabotage); and a platform-mounted apparatus called a beam director, similar in size and shape to a big sky-scanning searchlight. But the real killer for field commanders was the logistical nightmare they envisioned. Aside from its bulk, characteristic of flowing-gas lasers, THEL would be useless without its special fuels, and it also produced toxic, corrosive hydrogen fluoride gas, which would require special handling. "A chemical laser on a battlefield is more of a hazard than the threat it is trying to mitigate," says John Boness, chief technologist at Textron Systems, a laser-weapons contractor in Wilmington, Mass.

The U.S. Army and Israel had their sights set on something more nimble that could fit on an armored vehicle the size of a motor home and that would be ready to roll when it arrived on the battlefield. John Wachs, chief of the directed-energy division at the Army Space and Missile Defense Command, in Huntsville, Ala., explains the ultimate result of that conclusion: "The Army decided they would prefer solid-state, all-electric devices."

LASER TAG: Earlier this year, Northrop Grumman researchers far overshot the 25-kilowatt goal of the Joint High Powered Solid State Laser program by sustaining a 105-kW beam for more than five minutes.
PHOTO: NORTHROP GRUMMAN

SOLID-STATE LASERS are far simpler than their flowing-gas cousins. Small versions have been in the U.S. arsenal since the Vietnam War, but they are just glorified laser pointers: The low-power beams mark targets with invisible infrared spots, which allow smart bombs to home in. Such lasers could probably kill a fly, but nothing bigger.

The first laser ever built was also solid state. It used bright pulses from a flashlamp (which creates an intense white light) to illuminate a stubby ruby rod. The flashes excited chromium atoms in the ruby, which prompted the atoms to release their energy as red photons in a laser beam.

Modern solid-state lasers are powered in the same general way—light from an external source transfers energy to the atoms in a crystal, and these emit laser light. But instead of ruby rods, engineers have developed lasers that use more sophisticated crystals doped with a rare-earth element called neodymium. These simple, durable lasers are used widely in industry and the military.

BEAM ME UP

Competing solid-state laser weapon designs reflect different philosophies

THIN-ZAG TEXTRON FIBER LASER

Textron uses a single oscillator in which the light follows a zigzagging path through six water-cooled slabs, collecting more light than it would if it traveled in a straight line.

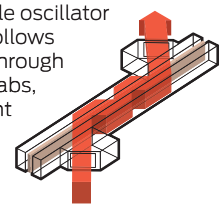

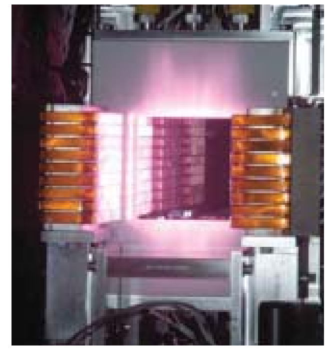

LAWRENCE LIVERMORE HEAT-CAPACITY LASER

Lawrence Livermore's ingenious but elaborate cooling mechanism solved the problem of laser overheating.

NORTHROP GRUMMAN FIRESTRIKE

Each of Northrop Grumman's Firestrike modules generates a 15-kW beam, which can be combined into a 105-kW beam with seven amplifiers.

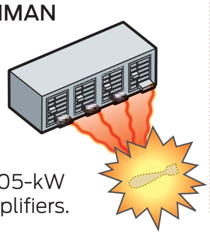

RAYTHEON LASER AREA DEFENSE SYSTEM

Raytheon's LADS is an updated version of its Phalanx Gatling gun and uses a brute force broadband laser to achieve a 50-kW beam.

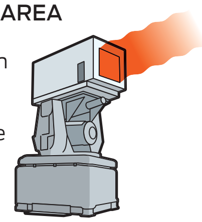

Despite the advances with materials, at most just a few percent of the electrical energy from a flashlamp makes it into a laser beam. Solid-state laser materials don't dissipate heat well, so trying to crank up the laser output too much will warm the crystal rod. Exceeding the strict heat limits of this material even by a negligible amount causes internal stress and degrades the beam quality, which means it won't focus tightly on the target. Add even more heat and the rod can crack or shatter. For decades, this combination of low efficiency and poor heat dissipation made attempts to develop high-energy solid-state lasers seem like a waste of time.

But by the 1980s, the industry began replacing the old flashlamps with semiconductor diode lasers. Both tools for powering solid-state lasers use two-stage processes, with electrical input first converted into light and the light powering the laser. But the diode laser is far more efficient. Like light-emitting diodes, diode lasers generate light when mobile electrons become attached to atoms at the border of two different semiconductor materials. Thanks to strategically placed mirrors, diode lasers convert much more input electricity into light—light that's limited to a narrow range of desirable wavelengths. Contrast that with flashlamps, which give off energy across the entire visible spectrum, the main reason why so much of their light is never absorbed and converted into laser energy.

A diode laser can transform roughly half the input electricity into light, and a solid-state laser can, in turn, convert about half of the energy in that diode laser's output into a high-quality laser beam. So less than a quarter of the input energy emerges in the beam. That's 25 times as much as what you get with a flashlamp. And the improved efficiency lessens the heat-dissipation problem, which should make it easier to construct a battlefield laser.

THEL HAD established that destroying a moving target at a distance of a kilometer or two requires around 100 kW of laser power. That oomph is needed mostly because of the spreading of the laser beam.

Although people often think of laser beams as being pencil thin, beyond a certain distance from the source the beam spreads out conically, like a searchlight. If the beam starts with a diameter of one centimeter, it might well expand to something like 10 centimeters at 200 meters' distance. By the time it hits a target 2 km away, the beam could be a meter across. Extending the target distance from 200 meters for stationary objectives to 2 km for rockets and mortars reduces the laser power per unit area—the critical factor in igniting the explosives in a bomb or rocket—by a factor of 100. And to counter that reduction, you need to boost the power, in turn, by a factor of 100.

Although THEL had gone nowhere, the Army had been hedging its bets all along. In 1997, not long after work started on THEL, the Army tapped Sparta, a defense contractor headquartered in Lake Forest, Calif., to build a Humvee with a turret-mounted solid-state laser to destroy unexploded ordnance exposed on the ground. They named the test system Zeus, after the thunderbolt-wielding king of the Greek gods.

A soldier operating Zeus would use the turret to train a green laser pointer on the target. With the unarmored Humvee parked between 25 and 250 meters away—far enough to keep out of danger—the soldier would then switch on the high-energy laser beam, invisible to the human eye because of its 1-micrometer infrared wavelength. The laser Sparta installed at its Huntsville location, before shipping Zeus to Afghanistan in 2003 for field tests, emitted only a kilowatt—small potatoes by laser-weapon standards. Although the beam Zeus uses can burn exposed skin, it is nowhere near as deadly as an ordinary bullet. But the targets were easy ones, stationary and clearly in the line of sight. Other than the sparkles of dust particles ignited by the heat, the only trace of the beam was the zone it heated on the target, as if it were sunlight focused onto paper. If you were viewing the target with an infrared camera, you'd see a small spot begin to glow as Zeus heated the casing. The steadily brightening spot would grow in size as the laser's heat penetrated deeper, through the case and finally into the explosive payload. The show would end with, well, a bang.

Zeus went on to field trials. At Bagram Air Base in Afghanistan, Zeus destroyed more than 200 rounds of ordnance in six months of field trials, including 51 rounds in one particularly success-

ILLUSTRATIONS: STEVE STANKIEWICZ; PHOTO: LAWRENCE LIVERMORE NATIONAL LAB ful 100-minute span. The Army was impressed: Zeus wasn't powerful enough to destroy buried mines or munitions, but it was small enough to be mobile, and it could take on improvised explosive devices—Iraq's famously deadly IEDs—if they were not buried under sand.

To tackle the problem of scaling up the power of a solid-state laser to compensate for distance, in December 2002 the U.S. military launched a grand challenge called the Joint High Power Solid State Laser (JHPSSL). The Army's goal was a 100-kW electrically powered laser that it could use for distant battlefield targets. The Air Force wanted a similar laser for its F-35 Joint Strike Fighter, and the Navy hoped such a laser could fend off small-boat attacks like the 2000 suicide bombing that killed 17 sailors and seriously damaged the USS *Cole*. Field commanders insisted that the laser not require any special power systems, so it had to run on electricity from the diesel generators that run other front-line equipment.

The first step was a competition to reach 25 kW by late 2005 between four entrants: Northrop Grumman, Textron, Raytheon, and Lawrence Livermore National Laboratory. The 25-kW output beam needed to remain tightly focused for 300 seconds while ensuring that the laser didn't self-destruct. The basic approach was clear—use diode lasers (here acting as what laser physicists call pump diodes) to inject light into thin slabs of laser material containing neodymium atoms, which in turn emit light at 1.06 µm in the near infrared. The pump beam passes through the slab and excites the atoms to produce high-output energy, and the heat can dissipate through the wide top and bottom.

Cooling is a critical issue, because solid-state lasers convert only about 20 percent of their electrical input energy into light output. Again, that may be good by laser standards, but it still means that 80 percent of the input energy winds up as heat—four watts of heat for every watt in the output beam. Much of that heat goes into the laser slab. The results of sending a laser beam through a hot slab are similar to what happens when you park your car in an open lot on a summer day. Just as the rising air currents bend the light over the hood, making it ripple and waver, uneven heating in a slab will break up or spread a laser beam. The farther the beam has to travel in these conditions, the more it will spread and the less damage it can do. And while rocket-engine lasers can blow away their waste

LASER WEAPONS

The path to a deployable laser weapon is littered with canceled projects

THEL

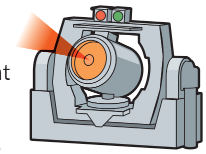

The Tactical High-Energy Laser, a joint project between Israel and the United States, was designed to shoot down short-range rockets.

ZEUS

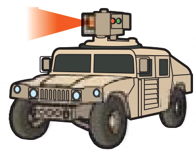

Sparta's Zeus system was a turret-mounted solid-state laser on a Humvee that destroyed unexploded ordnance from a safe distance. It was the first solid-state weapon success story.

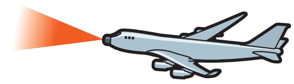

ABL

The Airborne Laser, which is US $4 billion over budget and can barely fit into a retrofitted 747, is designed to destroy long-range missiles a few hundred kilometers away.

heat as exhaust, solid-state lasers need another way to keep from frying.

Lawrence Livermore had already developed a clever but elaborate way to deal with the problem, resulting in what's known as a solid-state heat-capacity laser. Several ceramic laser slabs, 10 cm square and 2 cm thick, are mounted on four rotating wheels. For 10-second intervals, arrays of pump diodes fire 200 pulses per second into the edges of the slabs. During those 10 seconds, the slabs heat from room temperature to 130 °C, at which point the pumping stops for a quarter of a second as fresh slabs are rotated into place and the hot ones are pressed between heat sinks to cool. Then the cycle repeats.

This design was an ingenious solution to the cooling problem, and the Lawrence Livermore laser reached 25 kW. But the lab didn't make the cut for the 100-kW round of JHPSSL. Instead, the Army chose two other designs, from Northrop Grumman and from Textron Systems. Rather than moving the slabs, both designs used liquid coolant and heat exchangers to keep the slab temperatures uniform. The big difference between the Northrop Grumman and Textron systems was in the optical layout.

To appreciate how different these designs were, you need to understand that a laser is an oscillator with photons bouncing back and forth between a pair of mirrors. Those photons stimulate the laser crystal to emit more photons of the same wavelength and phase, which is to say that their electric and magnetic fields oscillate in lockstep. There are two ways to boost laser power: The first is just to build a bigger oscillator with more laser material. The second is to amplify the beam by passing it through a more-excited laser material without mirrors.

Northrop Grumman chose the latter method—amplification. The output of a single oscillator was split between two amplifiers to generate a total of 27 kW for 350 seconds, longer than JHPSSL's stated 300-second goal. For the 100-kW demonstration, Northrop redesigned the amplifiers so that each generated 15 kW and fit onto a table 1 meter square. Reaching 100 kW required splitting the oscillator output among seven amplifiers, then carefully matching the phases of the waves coming out of each in the final beam.

In March of this year Northrop Grumman crossed the finish line, generating a single 105-kW beam for more than 5 minutes and turning 19.3 percent of the input electrical power into output photons. The company's JHPSSL test bed weighed in at 7 metric tons and measured 2 by 2 by 2.7 meters, but Northrop Grumman engineers have also developed a compact version called Firestrike. Its 15-kW amplifier modules fit into field-ready boxes measuring 30.5 by 58.4 by 101.6 cm, and seven of them stack to make a 100-kW weapon that will weigh only a quarter as much as the test bed. That includes power supply and cooling equipment but not the electronics that adjust the output phases and direct the combined beam toward the target.

Textron avoids the problem of combining beams by building one big oscillator. "We have a single beam running through six slab modules in series," says Textron's John Boness, who oversees the company's laser program. The light follows a zigzag path through the water-

ILLUSTRATIONS: STEVE STANKIEWICZ

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

cooled slabs, passing through more laser material and collecting more light than if it had traveled in a straight line.

Meanwhile, DARPA has its own program for a slab-laser weapon, called the High Energy Liquid Laser Area Defense System (HELLADS). Using an arcane design from General Atomics in which some of the light passes through the coolant, developers are gunning for a 150-kW laser weapon weighing only 750 kilograms. At only 5 kilograms per kilowatt, the HELLADS laser module would require only 40 percent of the weight of Northrop Grumman's Firestrike for each kilowatt of output, meaning that HELLADS could actually work aboard a fighter jet.

Yet it's a long way from there to delivering a weapon that can blow enemy rockets out of the sky. The Army must test a 100-kW beam against realistic targets, and developers must build compact versions for field use.

WHY DIDN'T Zeus make the cut for JHPSSL? In Zeus, and in many industrial lasers, the light-emitting material is a special type of optical fiber in which a central core containing light-emitting atoms is surrounded by a layer that guides the pump light along the fiber. This design can convert up to 30 percent of the input power to laser output. Fiber lasers are able to run 24/7 on production lines.

For these reasons, fiber lasers have been getting a lot of attention. Their higher efficiency promises smaller and more portable lasers, and perhaps most important, fibers have larger surface areas than do slabs, making them easier to cool.

Part of the efficiency advantage comes from the fiber structure itself, which increases the fraction of the pump light absorbed. Fiber lasers are more rugged than slabs because they don't have to send their light through discrete optical components, which can drift out of alignment. Their toughness has made them a winner in industry and piqued the interest of DARPA.

But fiber lasers have their own limits. A single fiber laser can't combine high power and high beam quality, and no single-fiber laser can reach 100 kW. The light-guiding cores that transmit the light are small and can carry only a limited amount of light before photon interactions steal energy from the beam. By spreading the power over a wide range of wavelengths, an industrial fiber manufacturer, IPG Photonics Corp., in Oxford, Mass., has reached 6.5 kW with a 10-μm-wide core. Fiber lasers with larger cores can reach higher powers—50 kW—but with poorer beam quality.

However, some experts question whether 100 kW is essential. "It's an urban legend that if you can't get to 100 kilowatts, you won't have utility," says Mike Booen, vice president of advanced missile defense and directed-energy weapons for Raytheon Missile Systems. Even a single industrial fiber laser delivering tens of kilowatts can be of some military use. "In 2006 we blew up two 60-millimeter mortars at a little over 500 meters" with a 20-kilowatt commercial laser, he says.

Indeed, Raytheon has integrated a 50-kW commercial fiber laser into a customized version of its Phalanx Gatling gun. It's a brute-force laser with a lot of power and a mediocre beam. But it's also rugged. Raytheon wanted a laser reliable and robust enough to put in a real war-fighting environment, Booen says, "and not have guys in white coats running around and coaxing it to work." So far, Booen is happy with this laser's operation: "You push a button and it comes on."

Boeing has also tested a kilowatt-class industrial laser in its own multirole weapon system, the Humvee-mounted Avenger. Lee Gutheinz, who directs the high-energy laser program at Boeing Missile Defense Systems, believes that the compact size and inherent ruggedness of fiber lasers makes them attractive for small mobile platforms like a Humvee. But he can't see how to reach the 100-kW level. "Extending the capability of fiber lasers at reasonable performance above tens of kilowatts hurts my head," he says. "I don't want to say it can't be done, but it's very hard."

WILL SOLID-STATE laser weapons be like flying cars that never get off the ground? On one hand, grim headlines continually point out the need for a defense against rockets, artillery, and mortars. And tests have shown that such targets are vulnerable to laser heating. Solid-state lasers are far more practical for battlefield use than scaled-down versions of ABL.

On the other hand, a panel of experts convened by the National Research Council warned the Army last September that high-energy solid-state lasers were not yet ready for battle. And these experts also urged the Army to push beyond 100 kW to demonstrate a 400-kW laser, which they expect to be a more potent weapon, by 2018.

THEL showed that a laser beam can destroy rockets, missiles, and artillery, but success depends on atmospheric effects and target interactions as well as raw power, and those factors depend on wavelength. With solid-state lasers emitting at 1 μm, far from the proven 3.8-μm wavelength of THEL, the Army's upcoming test shots with a 100-kW solid-state laser are therefore essential. But controlled tests alone won't demonstrate military effectiveness. The battlefield is a harsh environment, says Philip Coyle, senior adviser at the Center for Defense Information and a former director of testing for the Pentagon, and no laser systems "are ready for realistic operational testing."

Missile defense may be an appealing idea, but "it is virtually impossible to build defensive laser systems that are cost-effective," says Coyle. He says the key question is whether it is cheaper for the Army to improve its defenses or for the enemy to overwhelm them with more missiles.

Nonetheless, the Pentagon thinks solid-state lasers have a better shot at short-range missiles on the battlefield than the Airborne Laser has for long-range missile defense. In early April U.S. Secretary of Defense Robert Gates announced plans to downshift development of ABL, saying that its "affordability and technology problems" made the program's operational role "highly questionable." ABL will likely not be scrapped entirely but will stay in an indefinite R&D phase. So for that laser weapon at least, the flying-car question seems destined to remain open for a long time. □

TO PROBE FURTHER Read the military fact sheet for the Airborne Laser at http://www.mda.mil/mdalink/pdf/laser.pdf.

An *IEEE Spectrum* feature in 2005 outlined the government's proposed plans for space-based missile defense using directed energy. The article is available at http://spectrum.ieee.org/energy/nuclear/starcrossed. An abbreviated version of the National Research Council's report, "Directed Energy Technology for Countering Rockets, Artillery, and Mortars (RAM)" is available at http://books.nap.edu/catalog.php?record_id=12008.

A Radical

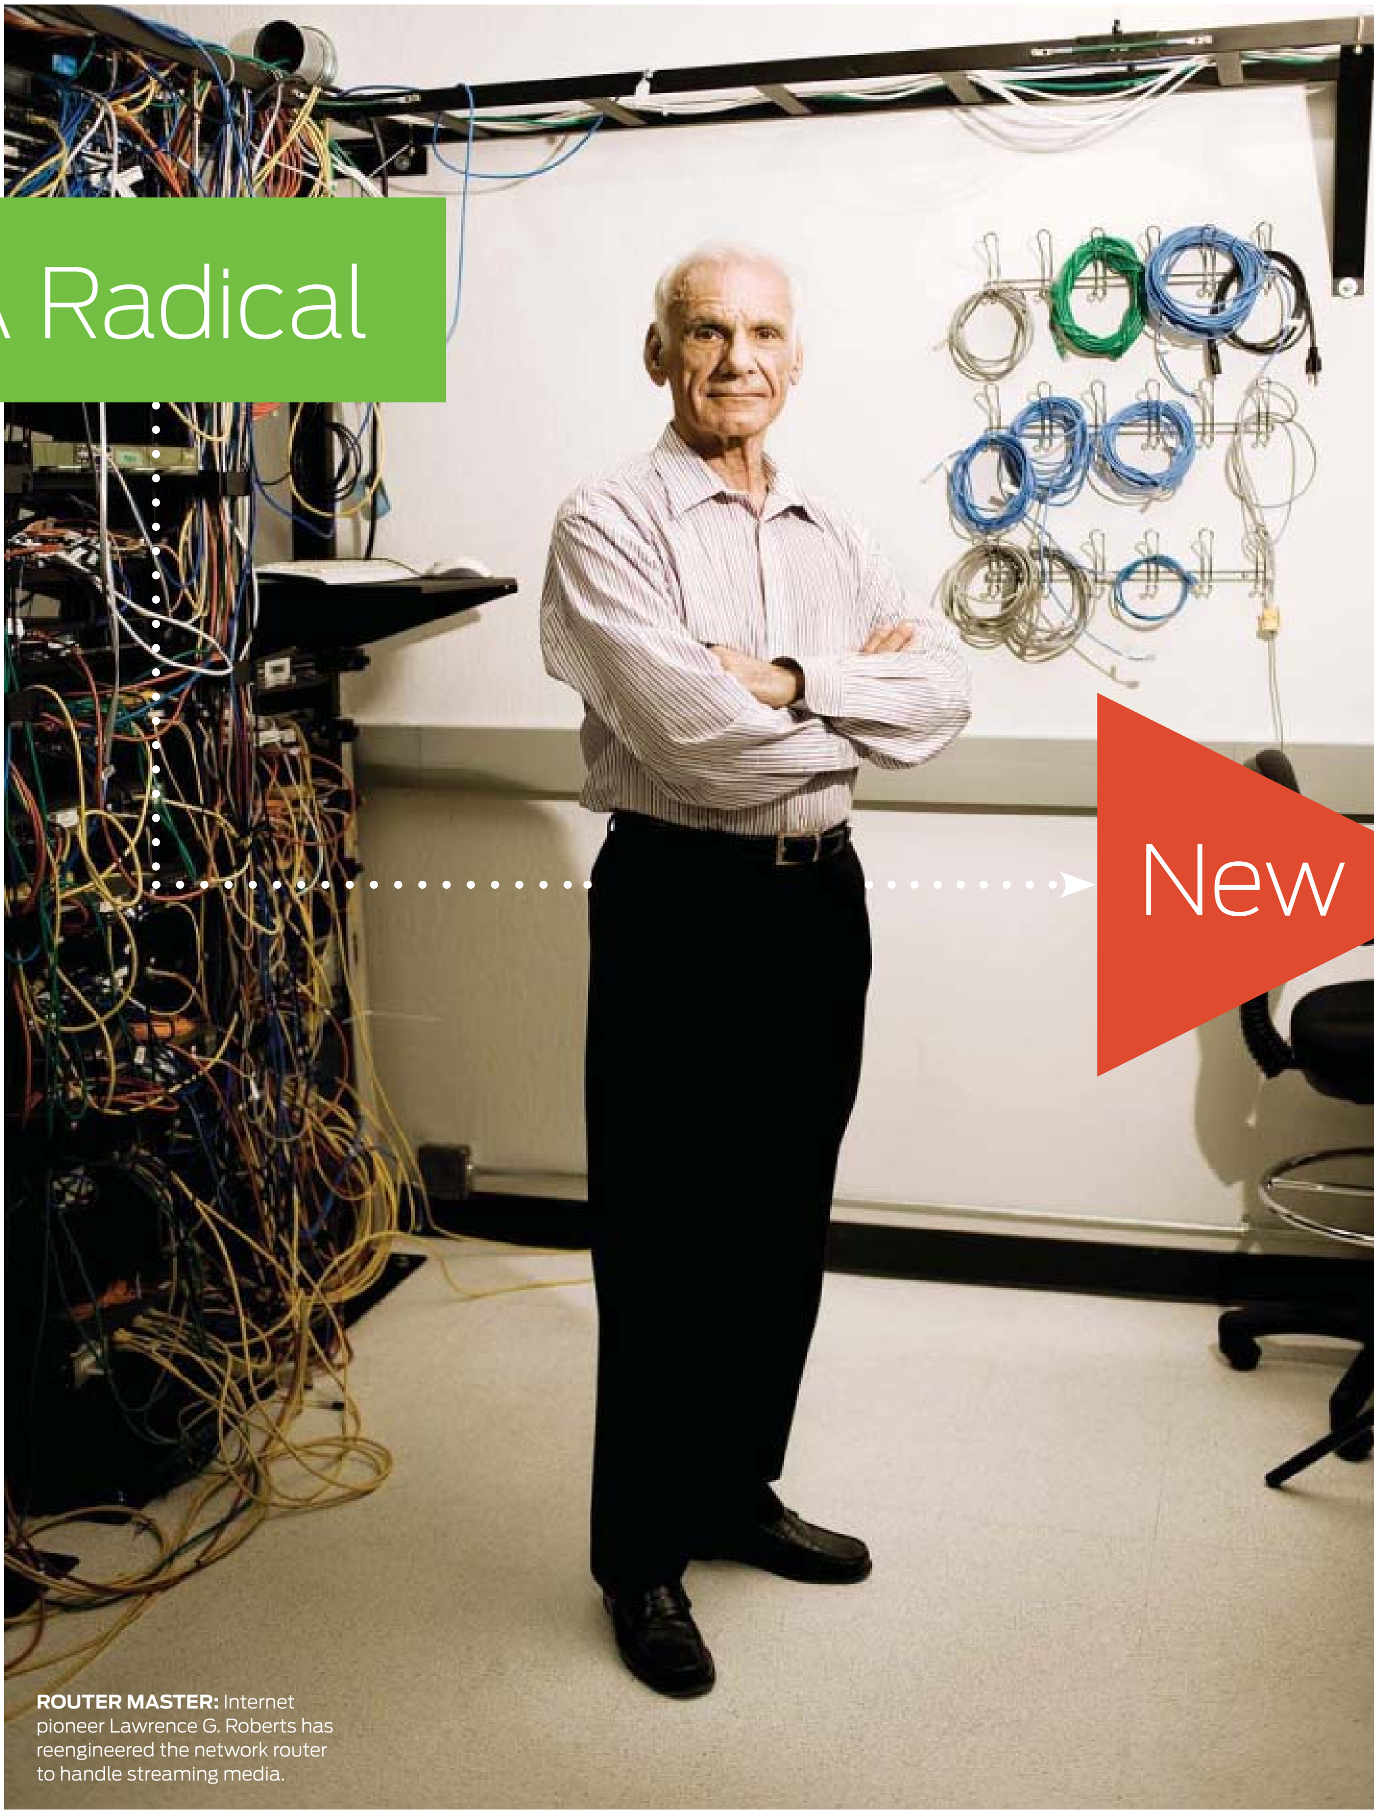

New

ROUTER MASTER: Internet pioneer Lawrence G. Roberts has reengineered the network router to handle streaming media.

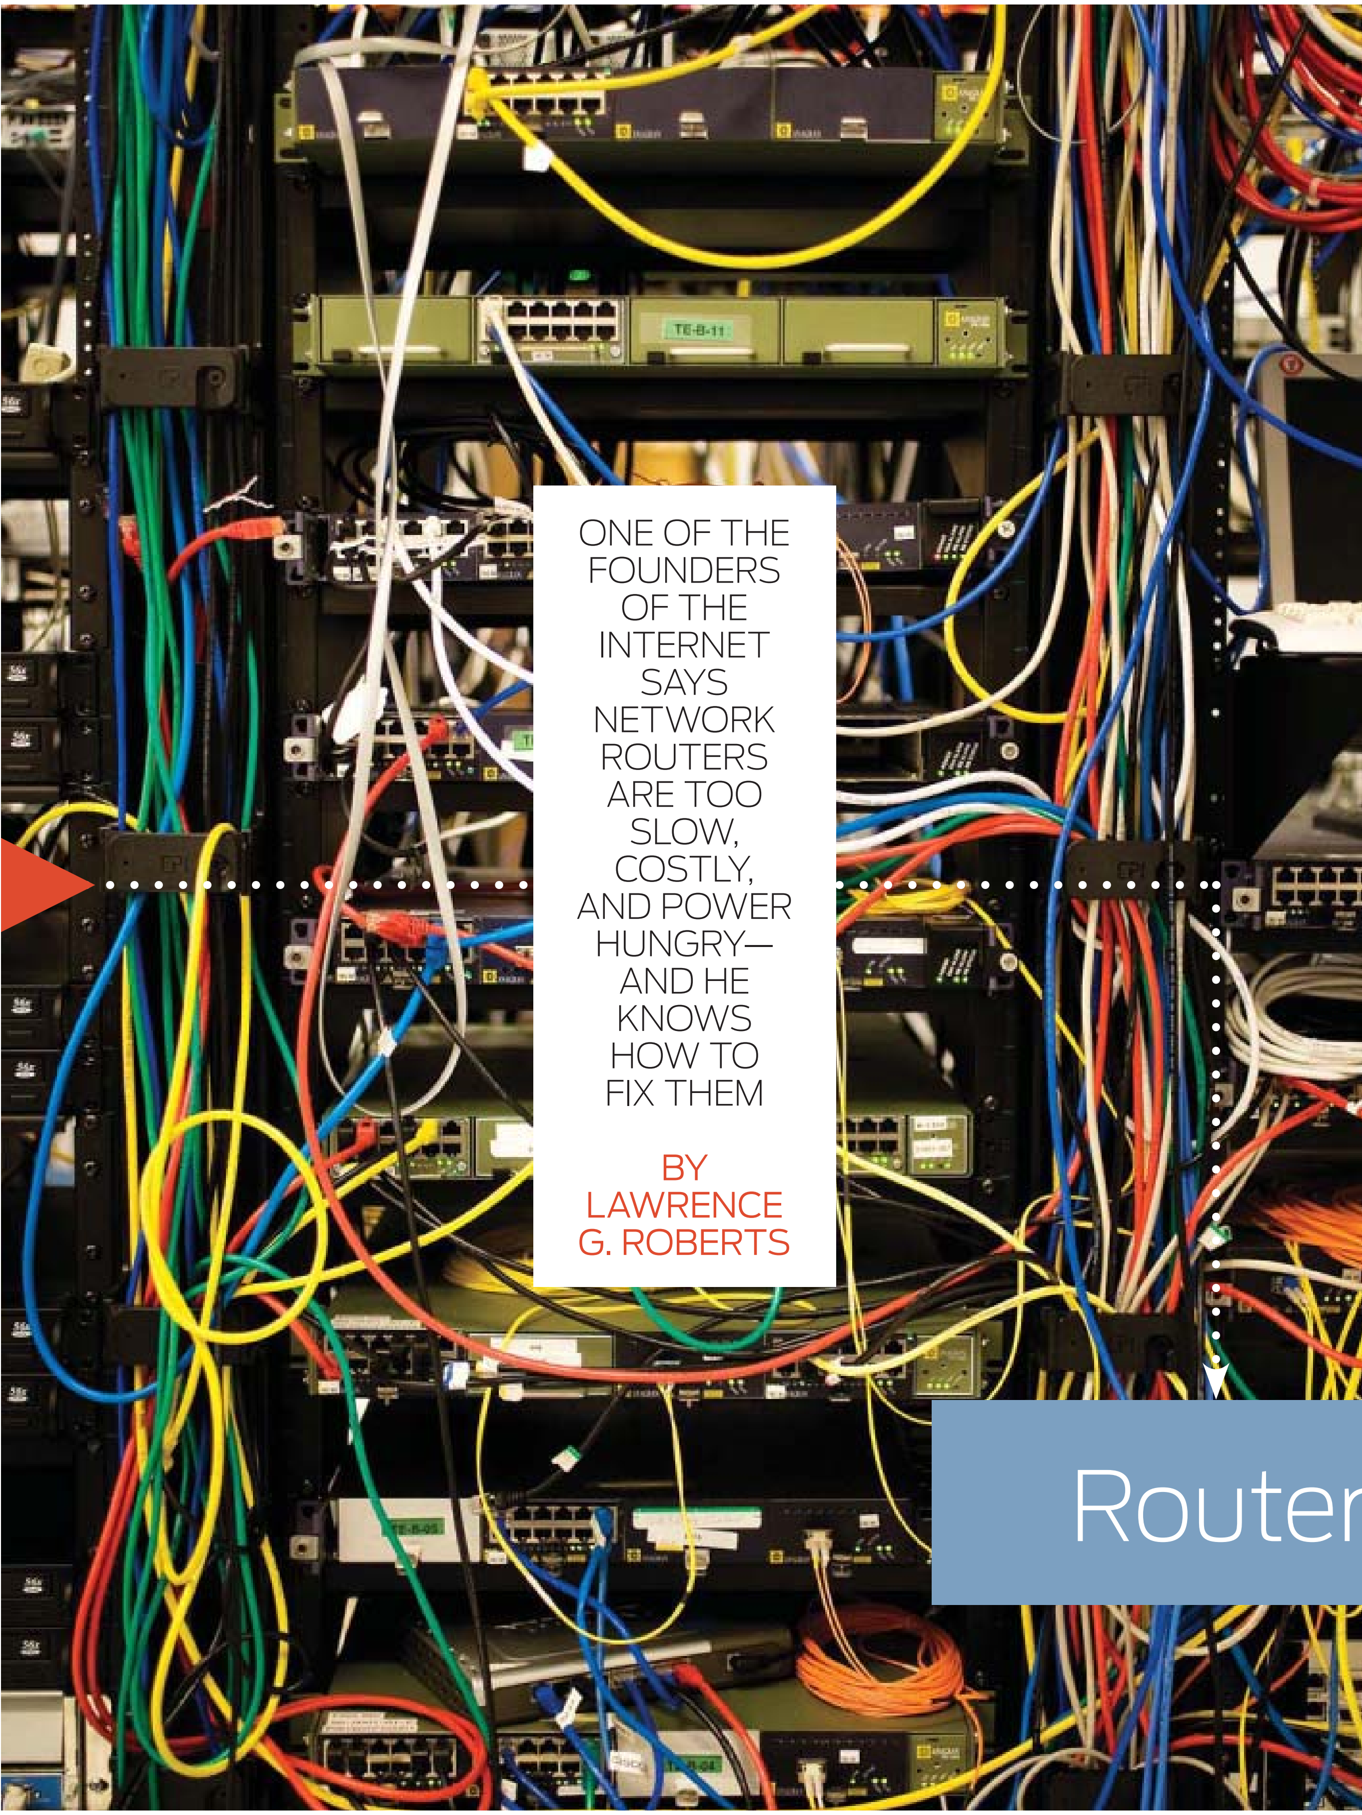

ONE OF THE FOUNDERS OF THE INTERNET SAYS NETWORK ROUTERS ARE TOO SLOW, COSTLY, AND POWER HUNGRY—AND HE KNOWS HOW TO FIX THEM

BY LAWRENCE G. ROBERTS

Router

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

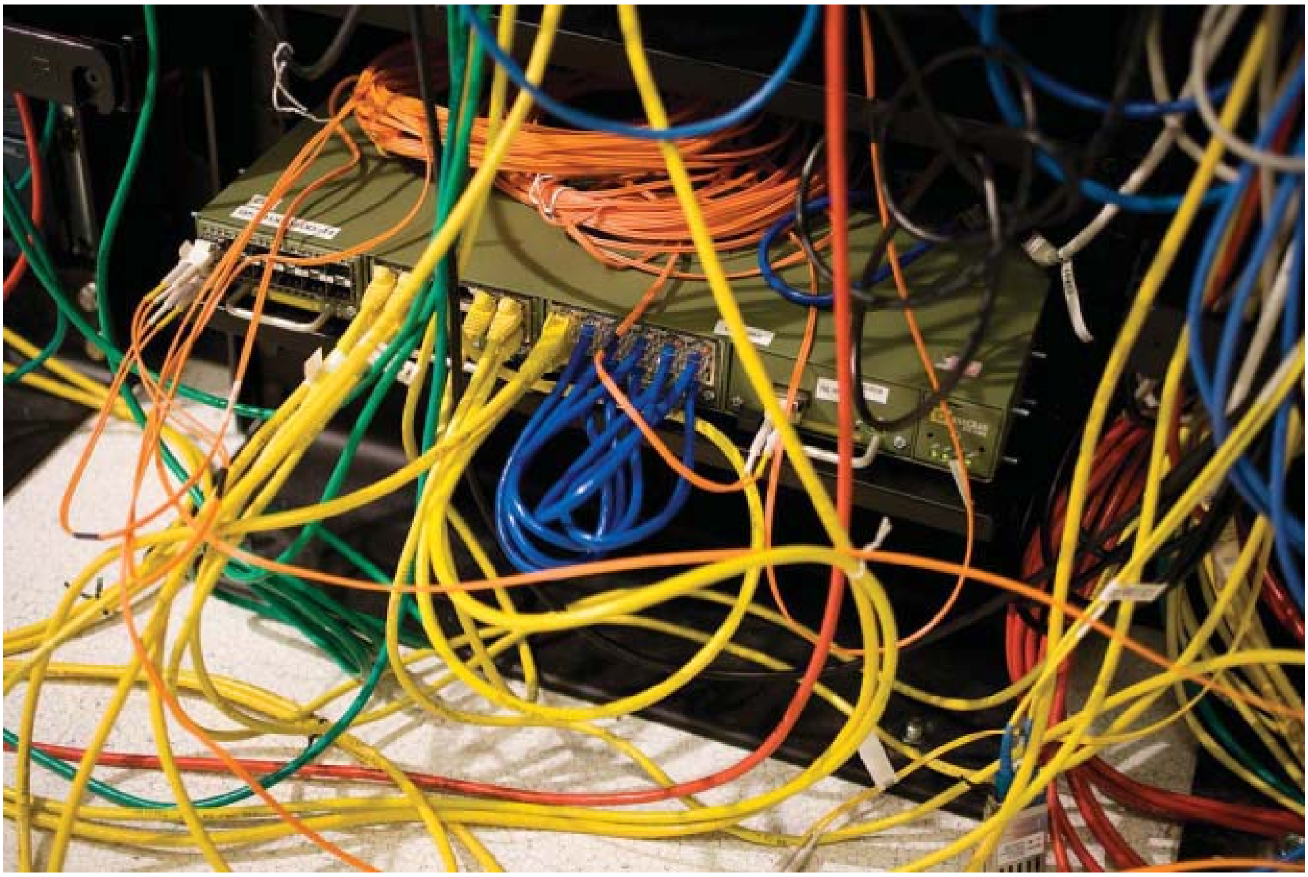

The Internet is broken.

I should know: I designed it. In 1967, I wrote the first plan for the ancestor of today's Internet, the Advanced Research Projects Agency Network, or ARPANET, and then led the team that designed and built it. The main idea was to share the available network infrastructure by sending data as small, independent packets, which, though they might arrive at different times, would still generally make it to their destinations. The small computers that directed the data traffic—I called them Interface Message Processors, or IMPs—evolved into today's routers, and for a long time they've kept up with the Net's phenomenal growth. Until now.

Today Internet traffic is rapidly expanding and also becoming more varied and complex. In particular, we're seeing an explosion in voice and video applications. Millions regularly use Skype to place calls and go to YouTube to share videos. Services like Hulu and Netflix, which let users watch TV shows and movies on their computers, are growing ever more popular. Corporations are embracing videoconferencing and telephony systems based on the Internet Protocol, or IP. What's more, people are now streaming content not only to their PCs but also to iPhones and BlackBerrys, media receivers like the Apple TV, and gaming consoles like Microsoft's Xbox and Sony's PlayStation 3. Communication and entertainment are shifting to the Net.

But this shift is not without its problems. Unlike e-mail and static Web pages, which can handle network hiccups, voice and video deteriorate under transmission delays as short as a few milliseconds. And therein lies the problem with traditional IP packet routers: They can't *guarantee* that a YouTube

ALL PHOTOS: JONATHAN SPRAGUE/REDUX

clip will stream smoothly to a user's computer. They treat the video packets as loose data entities when they ought to treat them as *flows*.

Consider a conventional router receiving two packets that are part of the same video. The router looks at the first packet's destination address and consults a routing table. It then holds the packet in a queue until it can be dispatched. When the router receives the second packet, it repeats those same steps, not "remembering" that it has just processed an earlier piece of the same video. The addition of these small tasks may not look like much, but they can quickly add up, making networks more costly and less flexible.

At this point you might be asking yourself, "But what's the problem, really, if I use things like Skype and YouTube without a hitch?" In fact, you enjoy those services only because the Internet has been grossly overprovisioned. Network operators have deployed mountains of optical communication systems that can handle traffic spikes, but on average these run much below their full capacity. Worse, peer-to-peer (P2P) services, used to download movies and other large files, are eating more and more bandwidth. P2P participants may constitute only 5 percent of the users in some networks, while consuming 75 percent of the bandwidth.

So although users may not perceive the extent of the problem, things are already dire for many Internet service providers and network operators. Keeping up with bandwidth demand has required huge outlays of cash to build an infrastructure that remains underutilized. To put it another way, we've thrown bandwidth at a problem that really requires a computing solution.

With these issues in mind, my colleagues and I at Anagran, a start-up I founded in Sunnyvale, Calif., set out to reinvent the router. We focused on a simple yet powerful idea: If a router can identify the first packet in a flow, it can just prescreen the remaining packets and bypass the routing and queuing stages. This approach would boost throughput, reduce packet loss and delays, allow new capabilities like fairness controls—and while we're at it, save power, size, and cost. We call our approach flow management.

TO UNDERSTAND HOW flow management works, it helps to describe the limitations of current packet routers. In these systems, incoming packets go first to a collection of custom microchips responsible for the routing work. The chips read each packet's destination address and query a routing table. This table determines the packet's next hop as it travels through the network. Then another collection of chips puts the packets into output queues where they await transmission. These two groups of chips—they include application-specific integrated circuits, or ASICs, as well as expensive high-speed memory such as ternary content-addressable memory (TCAM) and static random access memory (SRAM)—consume 80 percent of the power and space in a router.

During periods of peak traffic, a router may be swamped with more packets than it can handle. The router will then pile up more packets in its queue, establishing a buffer that it can discharge when traffic slows down. If the buffer fills up, though, the router will have to discard some packets. The lost packets trigger a control mechanism that tells the originator to slow down its transmission. This self-controlling behavior is a critical feature of the Transmission Control Protocol, or TCP, the primary protocol we rely on with the Internet. It's kept the network stable over decades.

Indeed, during most of my career as a network engineer, I never guessed that the queuing and discarding of packets in routers would create serious problems. More recently, though, as my Anagran colleagues and I scrutinized routers during peak workloads, we spotted two serious problems. First, routers discard packets somewhat randomly, causing some transmissions to stall. Second, the packets that are queued because of momentary overloads experience substantial and nonuniform delays, significantly reducing throughput (TCP throughput is inversely proportional to delay). These two effects hinder traffic for all applications, and some transmissions can take 10 times as long as others to complete.

As I talk to network operators all over the world, I hear one story after another about how the problem is only getting worse. Data traffic has been doubling virtually every year since 1970. Thanks to the development of high-capacity optical systems like dense wave division multiplexing (DWDM), bandwidth cost has been halved every year, so operators don't have to spend more than they did the year before to keep up with the doubling in traffic. On the other hand, routers, as pieces of computing equipment, have followed Moore's Law, and the cost of routing 1 megabit per second has decreased at a slower pace, halving every 1.5 years. Without a major change in router design, this cost discrepancy means that every three years a network operator will have to double its spending on infrastructure expansion.

FLOW MANAGEMENT can solve this capacity crunch. The concept of data flow might be more easily understood in the case of a voice or video stream, but it applies to all traffic over the Internet. Key to our approach is the fact that each packet contains a full identification of the flow it belongs to. This identification, encapsulated by the packet's header according to the Internet Protocol version 4, or IPv4, consists of five values: source address, source port, destination address, destination port, and protocol.

All packets that are part of the same flow carry the same five-value identification. So in flow management, you have to effectively process—or route—only the first packet. You'd then take the routing parameters that apply to that first packet and store them in a hash table, a data structure that allows for fast lookup. When a new packet comes in, you'd check if its iden-

FLOW CONTROL: The Anagran FR-1000 can be plugged into existing networks and can manage up to 4 million simultaneous flows.

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

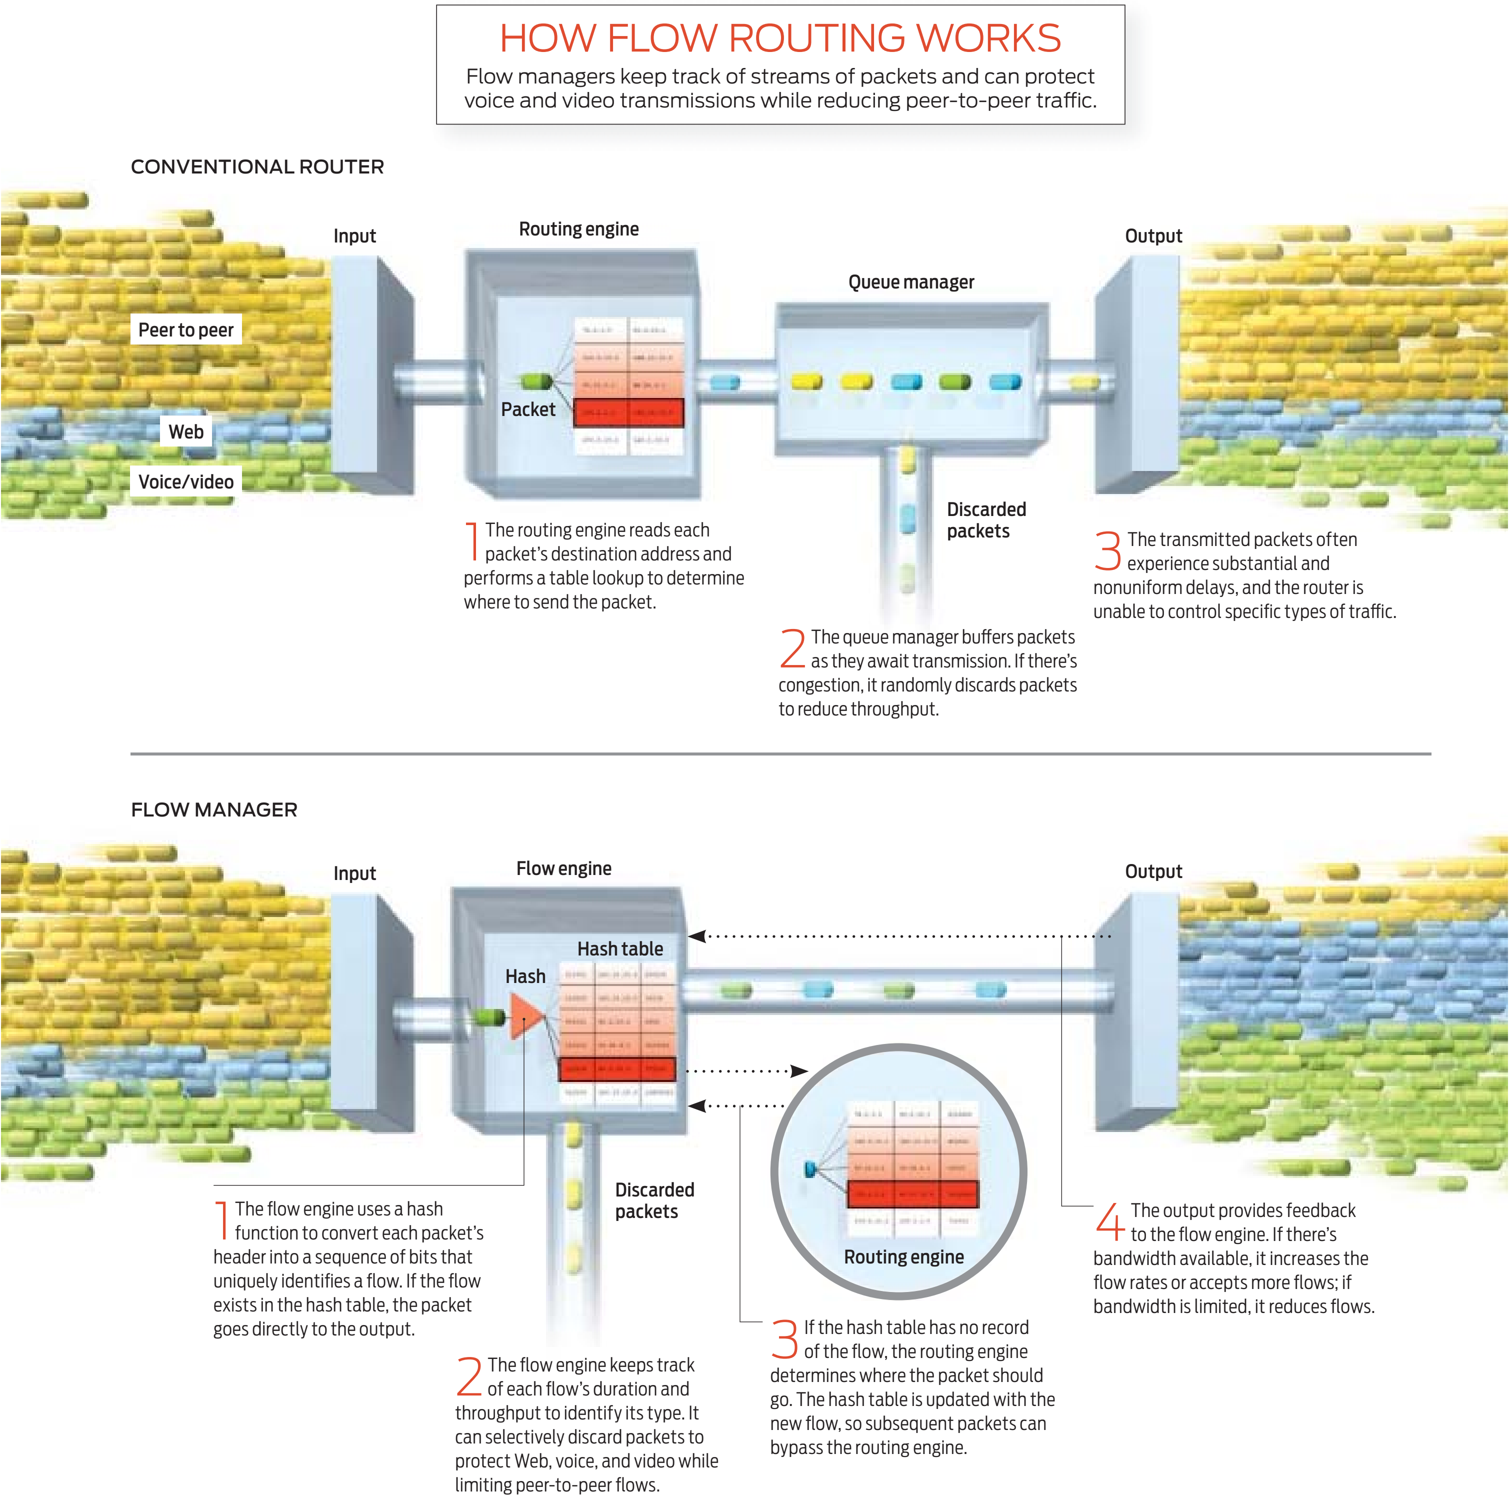

tification is in the hash, and if it is, that means the new packet is part of a flow you've already routed. You'd then quickly dispatch—the more accurate term is "switch"—the packet straight to an output port, thus saving time and power.

If traffic gets too heavy, you'll still have to discard packets. The big advantage is that now you can do it intelligently. By monitoring the packets as they're coming in, you can track in real time the duration, throughput, bytes transferred, average packet size, and other metrics of every flow. For example, if a flow has a steady throughput, which is the case with voice and video, you can avoid discarding such packets, protecting these stream-based transmissions. For other types of traffic, such as Web browsing, you can selectively discard just enough packets to achieve specific rates without stalling those transmissions.

This capability is especially convenient for managing network overload due to P2P traffic. Conventionally, P2P is filtered out using a technique called deep packet inspection, or DPI, which looks at the data portion of all packets. With flow man-

EMILY COOPER

agement, you can detect P2P because it relies on many long-duration flows per user. Then, without peeking into the packets' data, you can limit their transmission to rates you deem fair.

Since the early days of the ARPANET, I've always thought that routers should manage flows rather than individual packets. Why hasn't it been done before? The reason is that memory chips were too expensive until not long ago. You need lots of memory to store the hash table with routing parameters of each flow. (A 1 gigabit-per-second data trunk often carries about 100 000 flows.) If you were to keep a flow table on one IMP of 40 years ago, you'd spend US $1 million in memory. But about a decade ago, as memory cost kept falling, it started to make sense economically to design flow-management equipment.

In 1999, I founded Caspian Networks to develop large terabit flow routers, which I planned to sell to the carriers that maintain the Internet's core infrastructure. That market, however, proved hard to crack—the carriers seem satisfied with overprovisioning, as well as techniques like traffic caching and compression, which ameliorate congestion without addressing the roots of the problem. In early 2004, I decided to leave Caspian and start Anagran, focusing on smaller flow-management equipment to solve the overload and fairness problems. We designed the equipment to operate at the edge of networks, the point where an Internet service provider aggregates traffic from its broadband subscribers or where a corporate network connects to the outside world. Virtually all network overload occurs at the edge.

ANAGRAN'S FLOW MANAGER, the FR-1000, can replace routers and DPI systems or may simply be added to existing networks. It supports up to 4 million simultaneous flows—a combined 80 Gb/s in throughput. Its hardware consists of inexpensive, off-the-shelf components as opposed to ASICs, which increase development costs. We implemented our flow-routing algorithms in a field-programmable gate array, or FPGA, and the router's memory consists of standard high-speed DRAM. The FR-1000 sells in different models, starting at less than $30 000.

Like a regular router, the FR-1000 has input and output ports. But the similarities end there. Recall that in a traditional router the routing and queuing chips consume 80 percent of the power and space. By routing only the first packet of a flow, the FR-1000's chips do much less work, consuming about 1 percent of the power that a conventional router requires.

Even more significant, the FR-1000 does away entirely with the queuing chips. During congestion, it adjusts each flow rate at its input instead. If an incoming flow has a rate deemed too high, the equipment discards a single packet to signal the transmission to slow down. And rather than just delaying or dropping packets as in regular routers, in the FR-1000 the output provides feedback to the input. If there's bandwidth available, the equipment increases the flow rates or accepts more flows at the input; if bandwidth is scarce, the router reduces flow rates or discards packets.

By eliminating power-hungry circuitry, the FR-1000 consumes about 300 watts, or one-fifth the total power of a comparable router, and occupies one unit in a standard rack, a tenth of the space that other routers fill. We estimate that the equipment allows network operators to reduce their operating costs per gigabit per second by a factor of 10.

Measurements of the FR-1000 in our laboratories and by customers showed that networks equipped with the flow manager were able to carry many more streams of voice and video without quality degradation.

Another important capability we tested was whether the equipment could maintain quality of transmissions during congestion. The test involved a 100-Mb/s data trunk using a conventional router and another that included the Anagran flow manager. We progressively added TCP flows and measured the time required to load a specific Web page. The conventional router began to discard packets once traffic filled the trunk's capacity, and the time to load the Web page increased exponentially as we kept adding flows. The Anagran flow manager was able to control the rate of the flows, slowing them down to accommodate new ones, and the load time increased only linearly. The result: At 1000 flows, the flow manager delivered the page in about 15 seconds, whereas the conventional router required nearly 65 seconds.

Another capability we tested was fairness controls. Currently, P2P applications consume an excessive amount of bandwidth, because they use multiple flows per user—from 10 to even 1000. But services like cloud computing, which rely on Web applications constantly accessing servers that store and process data, are likely to expand the problem. We conducted measurements at a U.S. university whose wireless network was overwhelmed by P2P traffic, with a small fraction of users consuming up to 70 percent of the bandwidth. Early attempts to solve the problem using DPI systems didn't work, because P2P applications often encrypt packets, making them hard to recognize. The Anagran equipment was able to detect P2P by watching the number and duration of flows per user. And instead of simply shutting down the P2P connections, the flow manager adjusted their throughputs to a desired level. Once the fairness controls were active, P2P traffic shrank to less than 2 percent of the capacity.

The upshot is that directing traffic in terms of flows rather than individual packets improves the utilization of networks. By eliminating the excessive delays and random packet losses typical of traditional routers, flow management fills communication links with more data and protects voice and video streams. And it does all that without requiring changes to the time-tested TCP/IP protocol.

So is the Internet really broken? Okay, maybe that was an exaggeration. But the 40-year-old router sure needs an overhaul. I should know. ❑

TO PROBE FURTHER For more technical details on flow management and the Anagran FR-1000, visit http://www.anagran.com. See also the white papers "Flow Rate Management" and "TCP Rate Control With IFD [Intelligent Flow Delivery]," available at Lawrence G. Roberts's Web site: http://www.packet.cc.

For more on Internet routing, visit the Web sites of the Internet Research Task Force's Routing Research Group at http://www.irtf.org/charter?gtype=rg&group=rrg and the Internet Engineering Task Force's Routing Area at http://www.ietf.org/html.charters/wg-dir.html.

GERMANY'S GREEN-ENERGY GAP

Germany stumbles in its move to replace coal and nuclear power with offshore wind energy

BY PETER FAIRLEY

THE SIX OFFSHORE WIND TURBINES

that REpower Systems began erecting near Germany's coast in 2004 make their older cousins look like pinwheels. Each one has three 61.5-meter blades, which in a good breeze make one revolution every 5 seconds, producing 5 megawatts of electric power. Inspired by Germany's bold vision for capturing offshore wind energy, these majestic machines are designed to withstand anything the famously unforgiving North Sea can dish out.

And yet, these turbines have never felt the spray of salt water. They tower over communal pasture—above sheep munching, bleating, and adding to the world's supply of greenhouse gases. These turbines are tucked between a nuclear power station, an incinerator, and a cluster of chemical plants in Brunsbüttel, a hardscrabble harbor town where the Elbe River and the Nord-Ostsee Canal spill into the Wadden Sea.

Just a few years ago, many Germans thought that by this time, hundreds of offshore turbines like these giants from Hamburg, Germany–based REpower would be scattered off their northern coasts. After all, this prospect was a centerpiece of energy plans not only in Germany but also in Denmark, the Netherlands, and the UK. But the envisioned embrace of offshore wind power was particularly fervent in Germany, where the country's center-left political parties hatched plans to double renewable energy's share of power generation to 30 percent by 2020. But rather than the hundreds of turbines that were to be spinning in Germany's coastal waters by now to meet that schedule, only three turbines had gone up by the start of this year.

And Germany will not have many more offshore wind turbines anytime soon. German energy giant E.ON plans to install a small test farm on the North Sea this summer. But it will be another one or two years,

POWER SHIFT: Renewable sources should eventually replace conventional electric power plants. But the new and the old will probably coexist for many decades yet, as they do outside the German city of Cottbus.
PHOTO: JOERG GLAESCHER/LAIF/REDUX

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

at least, before big offshore wind farms are feeding the German grid. “Germany lost a lot of momentum,” says Eduard Sala de Vedruna, a senior analyst tracking wind energy for the consultancy Emerging Energy Research, of Cambridge, Mass., and Barcelona. “The offshore projects are at a quite immature stage.”

The idea that Germany is playing catch-up with Europe’s most promising strategy for renewable energy is jarring. This is Germany, after all, the country that 11 years ago put the Green Party in government, decided to phase out nuclear power, and pushed wind energy and photovoltaics to grid scale. Today Germany’s installed wind-turbine capacity of 24 gigawatts ranks second only to that of the United States (which has 25 GW). But despite the promises, greenhouse-gas emissions there haven’t plummeted. Rather, they have gone down only slightly since 2000. Germany, it seems, has lost its groove.

The result is a turnabout that would have seemed preposterous even six months ago: “Everyone in the environmental community is looking to the U.S. now,” says Elias Perabo, who codirects a campaign against the use of coal for Germany’s Berlin-based Climate Alliance.

The dearth of offshore wind turbines is just one of several signs of a slowdown in the country’s two-decade-old transition to renewable energy. Germany’s balkanized power grid, split between east and west when the country was divided and not yet fully knit back together, remains ill adapted to the variable flows from renewable energy. And Germany is readying a new generation of coal-fired power plants—including three proposed for Brunsbüttel.

The story of how Germany lost the lead in the transition to greener sources of energy contains a complex blend of backlash, environmental conflict, and competing commercial interests. It is a cautionary tale, showing in particular that public consensus about the urgency of combating climate change is just a first step in delivering a renewable-energy system.

NO COUNTRY HAS PUSHED renewable energy harder than Germany has. And much of that impetus came from one development: disenchantment with nuclear energy, which supplies about a quarter of the country’s electrical needs.

Public opinion turned abruptly against nuclear power in 1986, after the Chernobyl accident in Ukraine sent radioactive fallout over northern Europe and made West Germans uneasy about their own reactors. Popular concern after Chernobyl froze construction of additional reactors and fueled calls from the political left to scrap the nation’s existing nuclear plants. The chancellor at the time, Helmut Kohl, refused to abandon this source of carbon-free electricity, declaring climate change to be Germany’s top environmental challenge.

In this way, Kohl forged a political consensus for reducing greenhouse-gas emissions. But so far it is renewable energy, not nuclear, that has reaped the benefits. In 1990, the German government passed its path-breaking Electricity Feed-in Law, compelling utilities to buy all the power that renewable sources on their grid could generate—and at premium prices. The Feed-in Law thus set off a wind-power boom.

In 2001 that boom boomeranged on nuclear energy under Kohl’s successor, Gerhard Schröder. His Green Party–Social Democrat coalition cited wind energy as proof that Germany had an alternative to dirty coal and Russian natural gas in replacing nuclear power. Schröder’s government passed legislation to shut down all of the country’s reactors by 2022.

For that to happen, though, offshore wind power would be key. Germans, like most people, love the idea of wind power, but not all of them like the idea of having their landscapes marred by 130-meter-tall wind turbines. What is more, thanks largely to the 1990 law, most of the sites on land best suited to wind generation were already occupied. So installing turbines in their offshore territorial waters seemed like the best way around these obstacles. And because winds are in general stronger offshore than onshore, planting turbines far out in the sea promised twice as many hours of peak generation for each megawatt of installed capacity (assuming that offshore equipment functions reliably over time).

With these virtues in mind, the government passed its Renewable Energy Act in 2000, extending the favorable tariffs to wind farms in Germany’s North Sea and Baltic waters. By 2002—the year in which annual installations on land peaked at 3240 MW—developers had filed 29 proposals for offshore farms that together would have had a generating capacity of 63 GW, which was equal to half of Germany’s entire installed capacity at the time. Germany’s ministry for the environment (its Bundesministerium für Umwelt, Naturschutz und Reaktorsicherheit, or BMU) forecast that 500 MW of offshore wind would be operating by 2006 and that an additional 2500 MW would come on line by 2010.

Then the plans crashed headlong into political reality. Almost immediately, conservationists and marine ecologists questioned proposed incursions into near-shore areas where millions of migratory birds breed and feed. The BMU handled that challenge by studying it carefully and then, in 2005, designating permissible zones for wind development that were far from shore and in deep water.

As the UK, Ireland, the Netherlands, Denmark, and Sweden pressed forward with pioneering wind farms installed in water less than 20 meters deep and within 15 kilometers of the shore, Germany’s maritime authority offered developers 20- to 40-meter waters, located for the most part 40 km or more from the coast. That raised the cost and technical risk of German projects. Earlier, the German government had mandated a tariff of at most 9.1 euro cents (13 U.S. cents) per kilowatt-hour for offshore wind-generated electricity, no more than its neighbors were offering, despite the higher costs and risks.

Boosting the tariff to match the challenge faced opposition from the Big Four utilities that dominate Germany’s power sector: E.ON, RWE (formerly called Rheinisch-Westfälisches Elektrizitätswerk), Swedish power giant Vattenfall, and Électricité de France–owned EnBW (Energie Baden-Württemberg). Saddled with purchasing rising levels of wind power at top rates, these companies were pressing Berlin to scrap the special tariffs being offered.

IT WAS THE REVIVAL of Kohl’s center-right Christian Democratic Union party under Chancellor Angela Merkel that delivered the concessions needed to kick-start the offshore-wind industry. In 2006 Merkel’s government—a coalition that also included the Social Democrats and the Christian Social Union—made power-grid operators responsible for running cables to offshore farms. That shaved about one-fifth off the average cost of a project. And last year Merkel improved the revenue side of the ledger, boosting the offshore tariff to €0.15/kWh (US $0.21/kWh).

Slow but sure change is, well, in the wind. In a world that’s putting a price on carbon, Germany’s Big Four power giants are warming to the commercial potential of renewable energy. Over the past few years they have bought into offshore wind by

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

Germany's Wind-Energy Realität

PLACING WIND turbines many kilometers out to sea eases people's concerns about the appearance of such large structures—an often-voiced objection to the many wind-farm installations in Germany [see graph, far right]. Moving facilities offshore also boosts the amount of electricity that a 5-megawatt turbine can generate [see graph, right]. But Germany lags behind many of its neighbors in the construction of offshore wind farms [see map]. The pace of offshore wind development appears too slow to fill the energy gap that will ensue if Germany carries out its plans to decommission all its nuclear reactors, which still provide about a quarter of electric-power generation for that country [bottom].

EXPECTED ENERGY YIELDS
(MILLIONS kWh/YEAR)

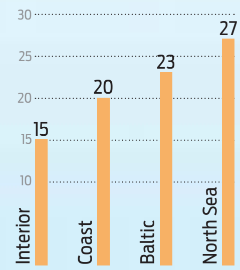

ANNUAL INSTALLED WIND-ENERGY CAPACITY
(HUNDREDS OF MEGAWATTS)

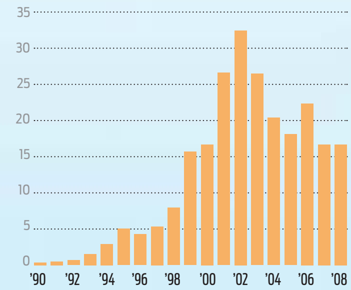

ELECTRICITY GENERATION
(BY FUEL IN TWh)

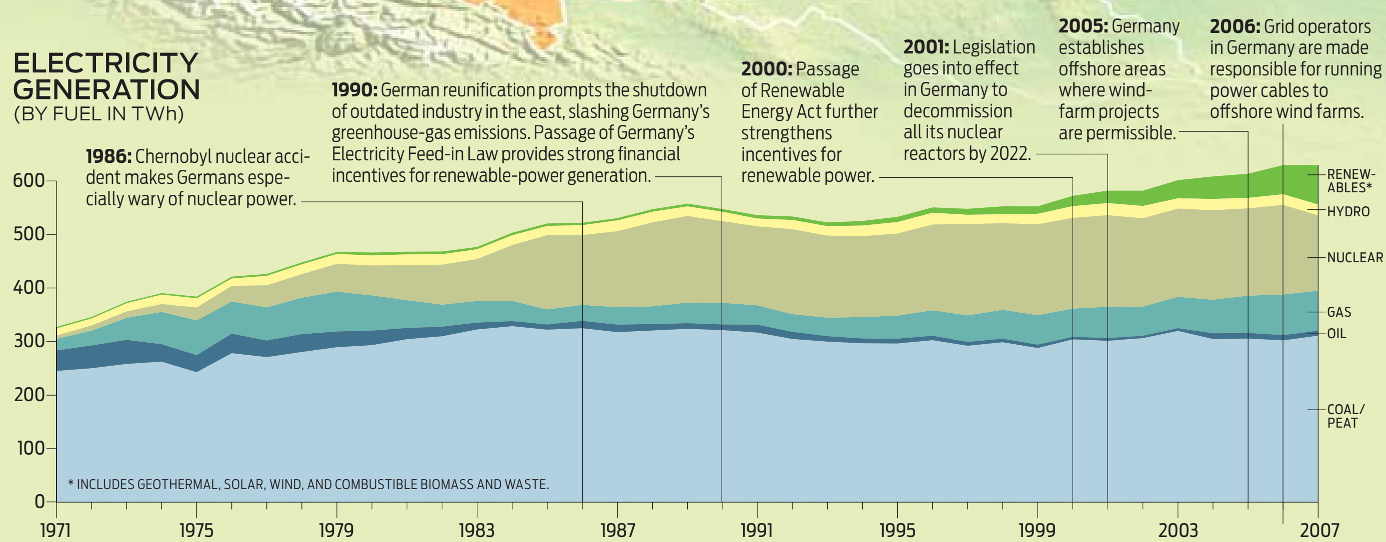

BRYAN CHRISTIE DESIGN

SOURCES: Yields—German Wind Energy Association; capacity—Deutsches Windenergie-Institut; map—European Wind Energy Association; generation—International Energy Agency

MODELING THE FUTURE: German Chancellor Angela Merkel's government has given wind developers a much-needed boost.
PHOTO: CHRISTIAN CHARISIUS/REUTERS/LANDOV

acquiring projects from wind developers. Norbert Giese, director of REpower's offshore business unit, says this shift is critical because most wind developers cannot raise the €1.2 billion (US $1.7 billion) to €1.4 billion (US $2 billion) needed to install a commercial-scale offshore farm.

Turbines have been ordered for more than 900 megawatts' worth of installations off Germany's northern coast. Of that, 60 MW will come from E.ON's €180-million (US $254 million) Alpha Ventus project, a turbine cluster being installed over the next few months 45 km off the North Sea island of Borkum. E.ON will plant six of REpower's behemoth 5-MW turbines and six more turbines of the same capacity from REpower's competitor Multibrid in water 28 to 32 meters deep. Sven Utermöhlen, regional director and CEO for E.ON Climate and Renewables Central Europe, calls Alpha Ventus a trial to gain logistical experience with these water depths and offshore distances.

They are getting their money's worth: Unusually rough seas last August scuttled an attempt to put down 700-metric-ton steel tripods on which to erect the first six turbines. No surprise then that E.ON—which operates globally—will focus its offshore investments elsewhere in Europe until at least 2011 or 2012. "We'd rather proceed sequentially from shore," Utermöhlen says.

RWE has firm plans to install 30 5-MW turbines in 2011 at the Nordsee I wind farm, its first German offshore project. Many more may follow—RWE took an option with REpower for an additional 220 turbines. But there's no guarantee. Martin Skiba, offshore-wind director for RWE's renewable-energy subsidiary, says he's not sure that €0.15/kWh will cover the cost of expanding Nordsee and other projects in Germany.

Even if these projects get built, offshore wind will generate a lot less energy in 2020 than Germany had hoped for. Just a few years ago, the German Energy Agency (Deutsche Energie-Agentur, or DENA) was projecting 20.4 GW of wind power by 2020, but lately the BMU has cut that forecast to just 10 GW. And even that estimate appears optimistic. "These 10 gigawatts are not going to be installed by 2020. That's a fact," says Emerging Energy Research's de Vedruna. He puts the figure at 8.4 GW. If every turbine ran full out, they would together deliver less than a quarter of the 149 billion kWh generated by Germany's nuclear reactors last year.

Germany hopes to make up the offshore shortfall by rejuvenating the onshore market for wind power. Tariff revisions last year added a bonus of half a euro cent per kilowatt-hour for repowering wind farms—swapping in multimegawatt machines in place of older, smaller turbines. That can double or triple output from a wind farm, but only if neighbors consent. That's a big "if," considering that growing resistance to the sight of giant turbines helped to drive Germany's wind developers offshore in the first place.

WHAT HAPPENS WHEN the wind doesn't blow? Even over the North Sea, the breeze sometimes abates, just as it does frequently enough on land. There was a day in January, for example, when a stalled high-pressure system becalmed many of Germany's wind turbines, and just 113 MW flowed from the country's 24 GW of installed capacity. Pushing renewable energy to an ever larger share of power generation means contending

IN 2006 MERKEL'S GOVERNMENT MADE POWER-GRID OPERATORS RESPONSIBLE FOR RUNNING CABLES TO OFFSHORE FARMS

with the problem of backing up this fickle source of energy. Here, too, Germany's ideas are far bolder than its actions.

The challenge is steep. DENA estimates that in 2020 Germany will be able to count on enough renewable generation at any moment to cover just one-eighth of projected peak power demand, even though its wind turbines, photovoltaics, and biomass-fired power plants may constitute more than one-third of installed power capacity.

In the long run, a more robust grid might be able to import enough renewable energy to cover such shortfalls. German analysts have been in the vanguard of modeling the backup benefits of a trans-European supergrid, in which high-voltage DC

(HVDC) transmission lines would provide access to the renewable resources available at any given moment across a broad area of Europe, the Middle East, and North Africa. Energy-systems consultant Gregor Czisch, for example, has shown that Europe could meet all its power needs with an HVDC supergrid sharing northern Europe's offshore wind energy, Scandinavia's hydropower, and North African solar energy.

The European Union embraced that concept last year. However, the idea of exchanging tens of gigawatts of electric power over thousands of kilometers looks positively hallucinatory given the state of transmission systems in Germany today, where wind-turbine installation has far outpaced expansion of the grid.

At present, weak connections between the regional grids controlled by the Big Four's transmission subsidiaries hinder the distribution of wind power, heavily concentrated in northern Germany, to the rest of the country. Already, transmission bottlenecks require grid operators to idle some wind farms in especially blustery conditions. A 2005 study by DENA led to a consensus plan for upgrading the grid, designating 850 km of high-tension AC lines to be built to integrate wind power. But a combination of tight credit and community opposition has stalled some of these high-priority projects, frustrating renewable-energy advocates. "The grid extensions should have started 10 years ago," says Hermann Albers, president of the German Wind Energy Association.

Czisch, meanwhile, complains that Germany isn't facing the future and building the HVDC cables required to provide a backup supply of renewable power from elsewhere. The root problem, says Czisch, is the Big Four, whose power plants face greater competition with every megawatt of added transmission. "There must be an independent state-owned organization that can calculate what the grids should look like in the future," he says.

The good news here is that pressure from European Union competition authorities forced the Big Four to put their transmission subsidiaries up for sale. The bad news is that until these assets are sold, those companies have even less incentive to invest in better transmission lines.

AT THE MOMENT, however, DENA has more pressing priorities. One of them is dealing with the spreading opposition in Germany to coal-fired power, which has caused a handful of projects to be stalled, blocked, or dropped. Last year DENA determined that cancellation of coal projects could leave Germany without sufficient conventional energy capacity to back up its wind and solar power. This "power gap," says DENA's CEO, Stephan Kohler, could reach 12 GW by 2020.

The gap could even be wider, says Kohler, if energy-efficiency gains slide. For example, the DENA's analysis assumes that Germany will reduce power demand by 8 percent by 2020. "This is a realistic scenario," says Kohler. But then again, "last year power demand increased 0.9 percent," he concedes. "That's not good."

Climate activists call Kohler's position disingenuous. For one thing, the construction of new coal plants directly contradicts the government's energy-efficiency plans, because only a few facilities will capture their waste heat for district heating. A climate scorecard of G8 countries prepared last year by environmental consultancy Ecofys concluded that the build-out of coal-fired plants would "lock Germany into a high level of carbon intensiveness for the next 40 years."

Two years ago Germany vowed that by 2020 it would cut greenhouse-gas emissions by 40 percent from its 1990 level, inspiring its European Union partners to agree to a combined reduction of 20 to 30 percent if other countries, such as the United States, contribute their share. But Ecofys, confirming most other analyses, concluded that an increased reliance on coal and an emissions-heavy automotive sector meant that Germany was "lagging behind its aspirations."

The only hope for neutralizing the greenhouse impact of the coal surge—carbon sequestration—is at least a decade away from implementation. Companies such as Vattenfall are engineering new technologies to reduce the punishing costs of capturing and storing CO_2, but they know they will have to convince the same communities that reject wind turbines and power lines in their backyards to accept the presence of CO_2 underneath them.

POSTPONING THE NUCLEAR phaseout would take the steam out of Germany's coal surge and erase DENA's projected power gap. Many expect Chancellor Merkel to do just that if her Christian Democratic Union party wins a majority in federal elections this fall. It would be a bitter pill for German antinuclear activists. But most of their fellow citizens seem ready to compromise.

The German Physical Society, a scientific body that has been preoccupied with the dangers of climate change since the early 1980s, has pleaded for continued use of nuclear energy. In 2005, the society assessed Germany's energy options and calculated that only a nuclear extension would enable it to meet its 2020 goals for CO_2 reduction. The society predicted then, correctly, that the supply of offshore wind energy would fall well short of the government's projections. It concludes in its report (and has repeatedly confirmed) that no other measure even comes close to making up for the increased carbon emissions from coal that the nuclear phaseout would stimulate: "Whilst originally we had hoped to have sufficient leeway to compensate for the loss of CO_2-free electricity derived from nuclear power, today we are forced to realize that such an equation will not balance out."

"I think I could live with that," agrees Brunsbüttel-area resident Stephan Klose when asked about keeping nuclear alive. Klose lives 12 km south of Brunsbüttel's cutting-edge wind turbines, its proposed coal-fired power plants, and its 33-year-old nuclear reactor, which has two more years of operation left before its scheduled shutdown. He likes renewable energy and fiercely opposes more reliance on coal.

Shutting down reactors, reasons Klose, will not resolve challenges posed by nuclear power, such as the long-term management of spent fuel rods. "We have the nuclear waste anyhow," says Klose. But last year he joined a local group working against coal, because he believes that coal-fired generators in Brunsbüttel will pollute the area with mercury and contribute to global climate change. The climate concern has a particular resonance in this marshy region where sinking dikes stand between communities and the swiftly rising sea levels that these people fear.

As for carbon capture and storage neutralizing the climate risk, Klose says no way, raising the specter of sequestered CO_2 escaping to the surface and causing mass asphyxiation. "If there's a leak and you have a 1- to 2-meter-high level of CO_2, every animal, every human being within this zone will die," says Klose. "I think you can't take that risk." ❑

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

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

the data

A Telecom Diet Rich in Fiber

UNDERSEA CABLE construction is booming, with 15 cables laid last year, as many as in the peak year of 2001. Even more cables will be laid this year, according to a recent report from TeleGeography Research, an analytical firm based in Washington, D.C., whose parent company is in Exeter, England.

Prior to the late 1990s, subsea cables primarily carried telegraph and voice traffic. Then the dot-com and telecommunications manias fueled unprecedented construction under the sea as well as on land. In 2001 US $13.5 billion was spent. But with every boom comes a bust; only about $2 billion was spent during four lean years from 2004 to 2007.

In the dot-com boom, spending was speculative and therefore through the roof. This time around, carriers and wholesalers are adding capacity only where new demand had soaked up the earlier glut. Spending will be lower because most new cables will cover shorter, regional runs.

Even during the lean years, global demand never slaked, growing at an average compound annual rate of 54 percent in the years from 2002 to 2008. In Latin America and the Caribbean, traffic grew the most, at more than 75 percent. —*Steven Cherry*

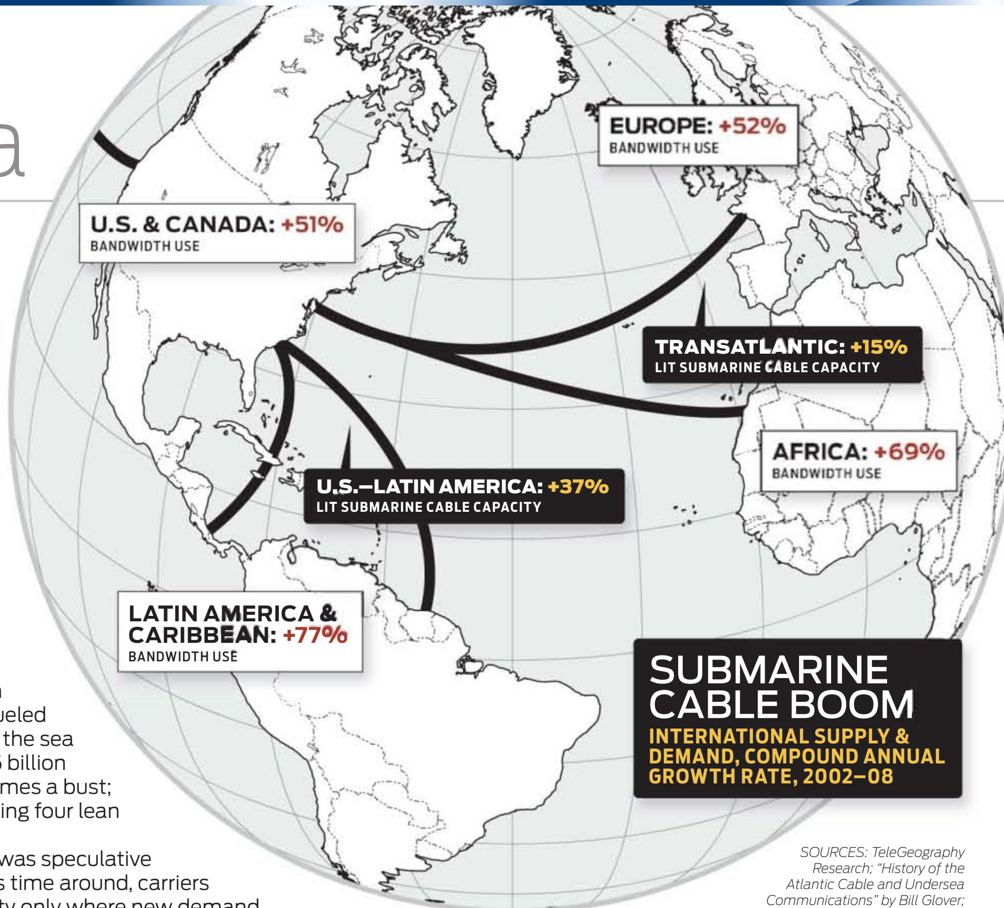

SOURCES: TeleGeography Research; "History of the Atlantic Cable and Undersea Communications" by Bill Glover; Wikipedia

CONSTRUCTION COST OF SUBMARINE CABLES
(US $, BILLIONS, YEAR CABLE ENTERED SERVICE)

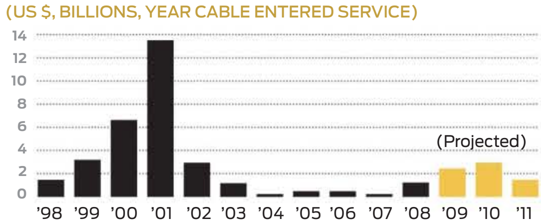

NUMBER OF NEW SUBMARINE CABLES ENTERING SERVICE

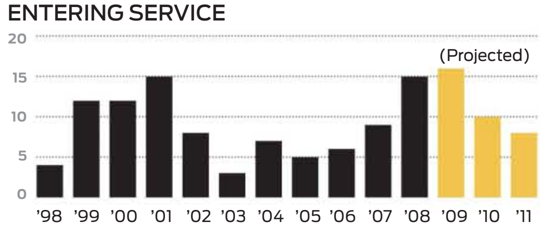

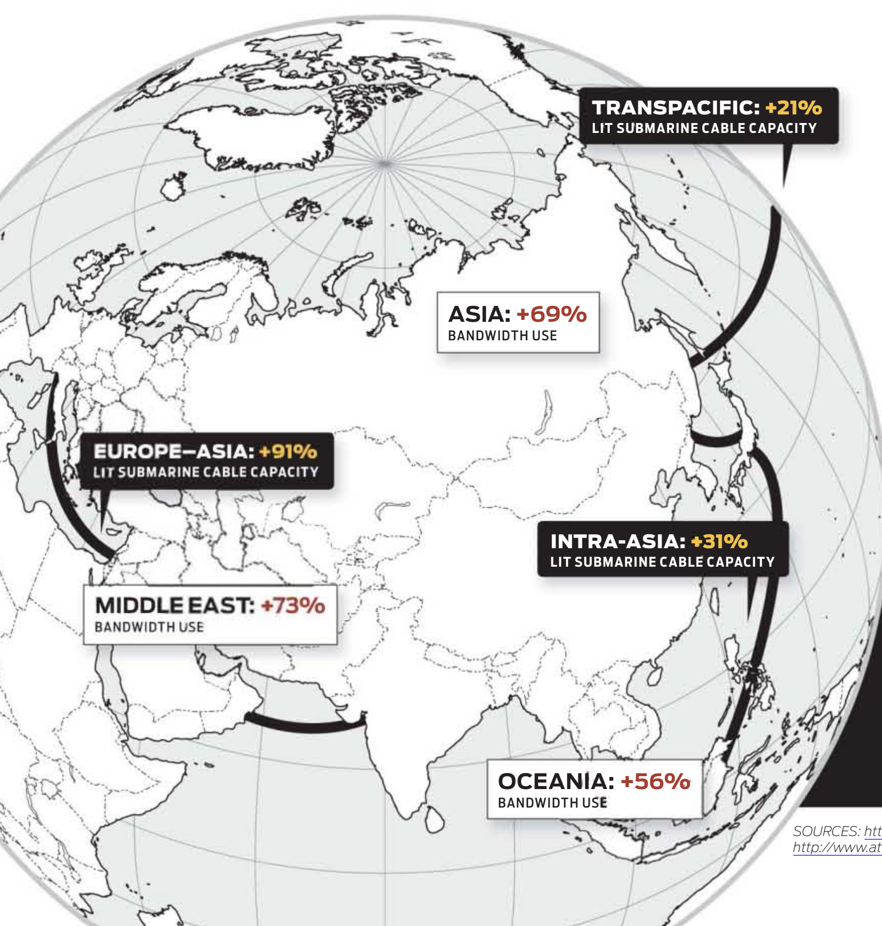

150 Years Under the Sea

1851: First successful commercial telegraph cable laid, England to France (copper wire, gutta-percha, tarred yarn, tarred hemp, and iron wires).

1858: First transatlantic cable.

1902–1906: Cables cross the Pacific.

1956: TAT-1, the first submarine transatlantic telephone cable (coaxial cable, polyethylene insulation, vacuum-tube repeaters).

1988: TAT-8, the first fiber-optic Atlantic cable. Capacity: 40 000 voice circuits.

2001: VSNL Transatlantic laid a pair of multiterabit cables under the Atlantic. Total bandwidth: 5.12 terabits per second.

2010: Pacific Unity will lay multiple fibers. Total capacity: 7.68 terabits per second.

SOURCES: http://www.telegeography.com; http://www.atlantic-cable.com/Cables/CableTimeLine/index.htm