
ieee spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
09.09
ieee
RECORDS
Meet the Geeks Behind
THE BEATLES: ROCK BAND
(Written by David Kushner)
STEREO
ST-2047
(ST-X-1-2047)
I
1. MAGICAL MYSTERY TOUR (2:51)
From MIT grad students to media moguls
2. LIKE DREAMERS DO (2:35) Their music video games became a cultural phenomenon
3. REVOLUTION '09 (8:22)
Now they're transforming the Beatles
MFD BY IEEE SPECTRUM MAGAZINE
MARCA REG U.S. PAT NO. 2,631,859
IEEE
Celebrating 125 Years
of Engineering the Future
the institute
CLICK HERE TO ACCESS THE SEPTEMBER ISSUE

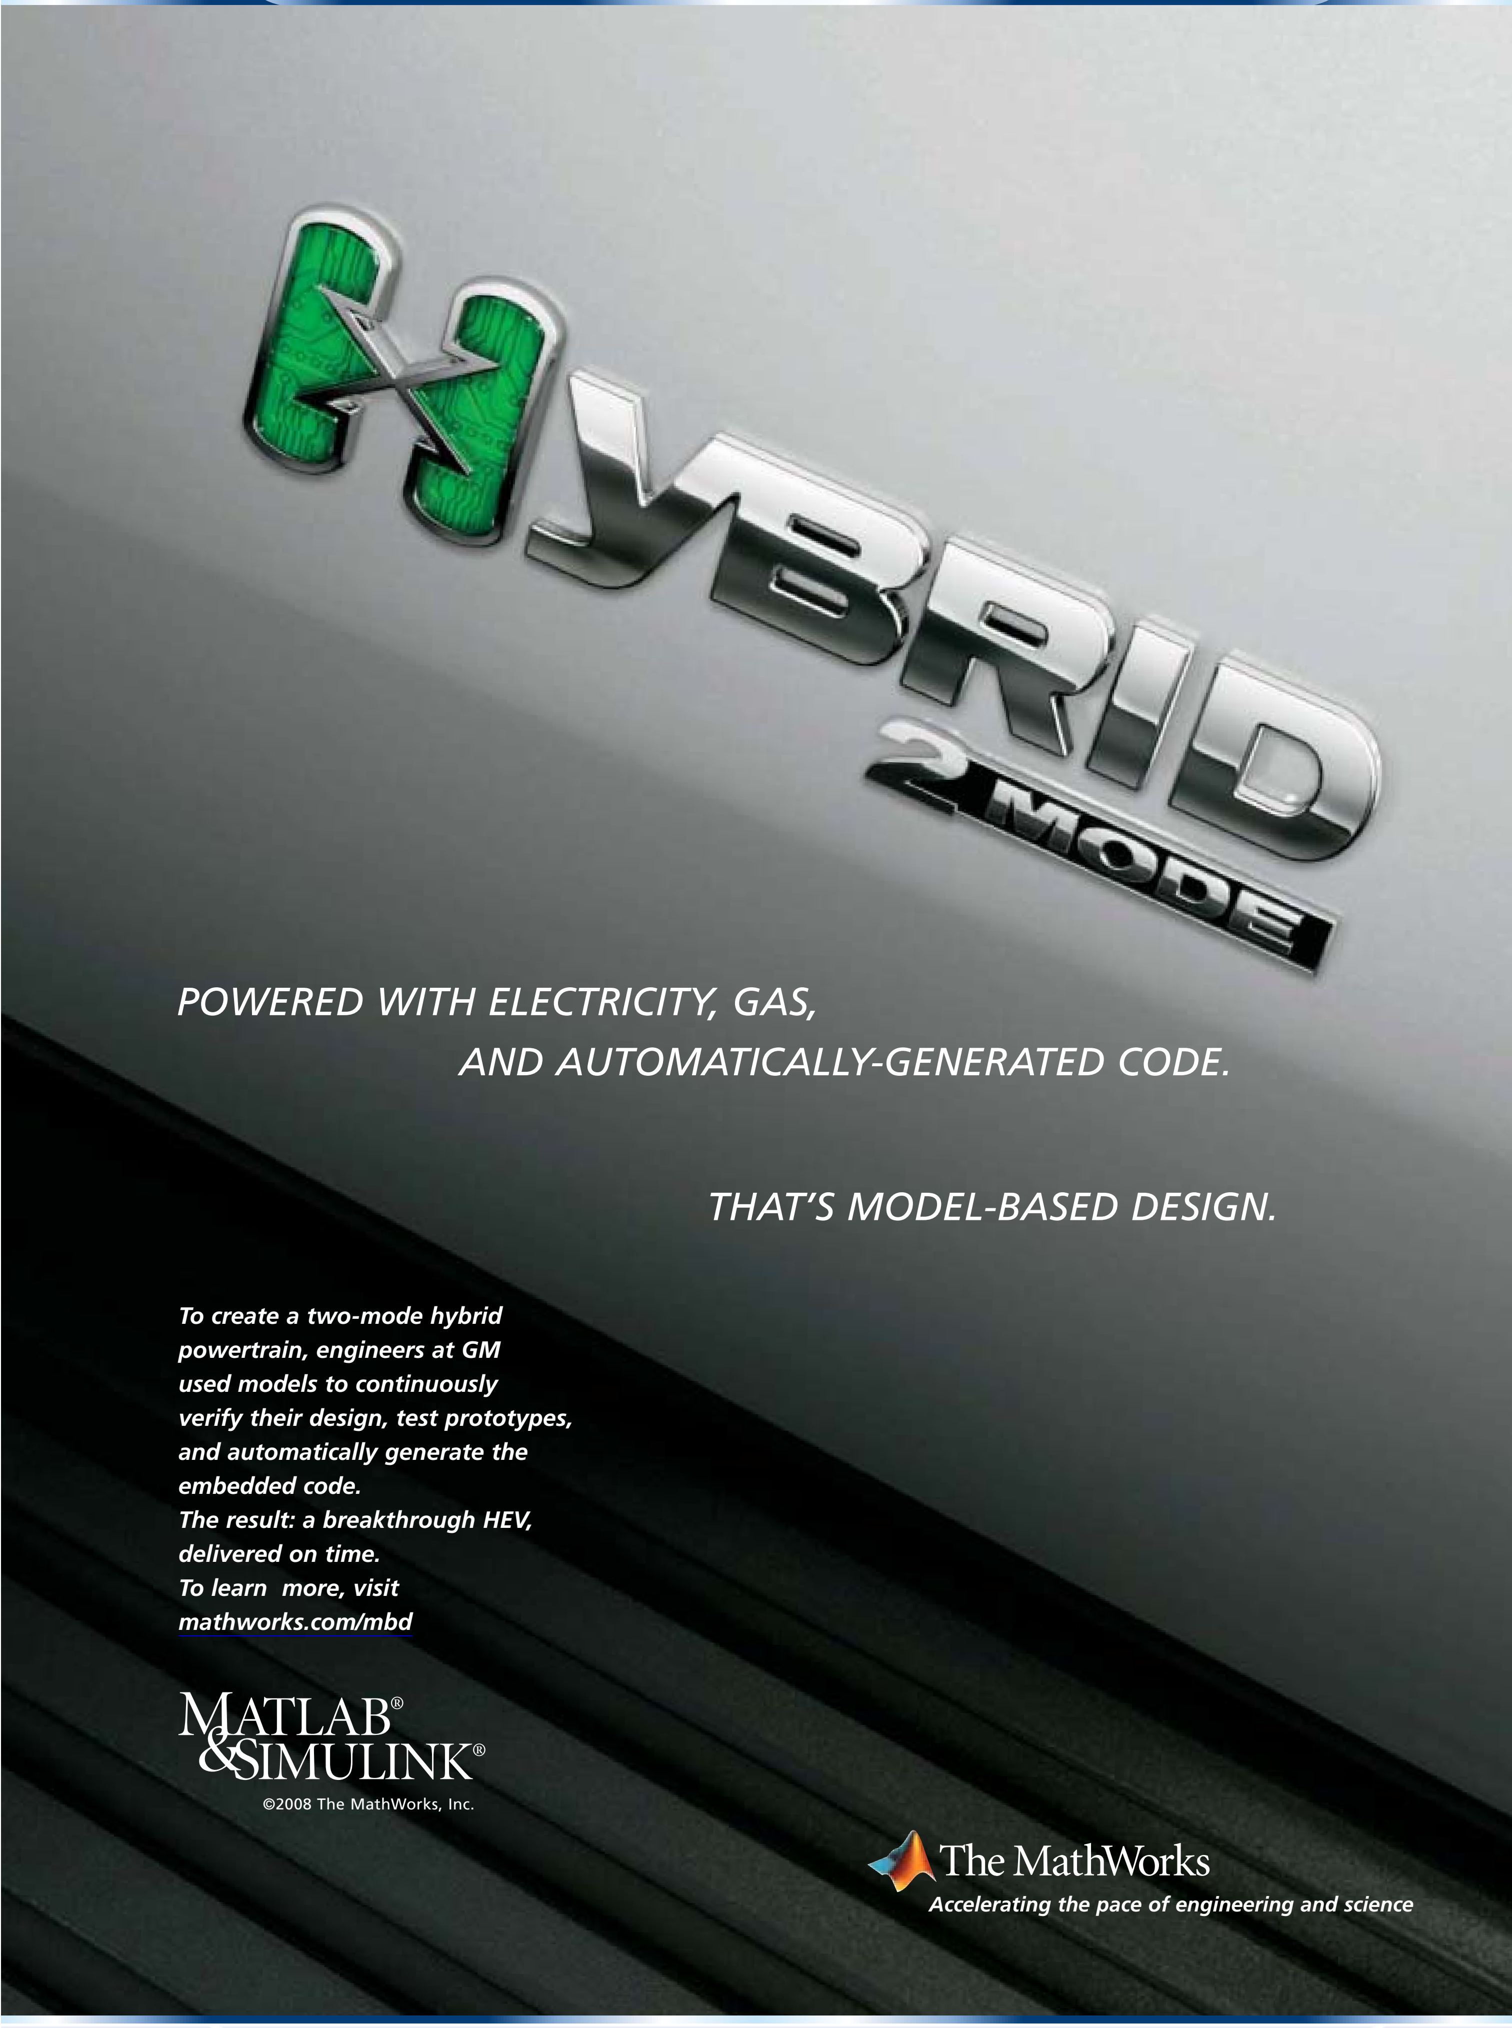
HYBRID
2 MODE
POWERED WITH ELECTRICITY, GAS,
AND AUTOMATICALLY-GENERATED CODE.
THAT'S MODEL-BASED DESIGN.
To create a two-mode hybrid powertrain, engineers at GM used models to continuously verify their design, test prototypes, and automatically generate the embedded code.
The result: a breakthrough HEV, delivered on time.
To learn more, visit mathworks.com/mbd
MATLAB® & SIMULINK®
©2008 The MathWorks, Inc.
The MathWorks
Accelerating the pace of engineering and science

volume 46 number 9 international

ieee spectrum 09.09

26

BANDMATES: Eran Egozy and Alex Rigopulos led the creation of the much anticipated *The Beatles: Rock Band* video game.

COVER: PHOTO-ILLUSTRATION: SEAN MCCABE; HARMONIX FOUNDERS: JOSHUA DALSIMER

THIS PAGE: JOSHUA DALSIMER

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

volume 46 number 9 international

ieee spectrum 09.09

PHOTO: PER-ANDERS PETTERSSON/LIASON/GETTY IMAGES

View exclusive online-only content now by clicking on the headlines

WWW.SPECTRUM.IEEE.ORG

AVAILABLE 1 SEPTEMBER ON IEEE SPECTRUM ONLINE

WIZARDS OF WATER

Water is such a basic human need that it takes real ingenuity to come up with new ways to use, share, and retrieve it. Enter the world's water engineers. In this series of podcasts, we investigate where engineers are helping communities and cities handle acute water challenges. We find carbon nanotubes at work in a stream in Kentucky; talk to a water official who yearns for better flood forecasting in Mozambique [above]; and hear how three women tackled Atlanta's persistent drought—all in pursuit of a more sustainable, water-sensitive future.

NEW SPECTRUM BLOG

INTRODUCING THE NANOCLAST: Based in Madrid, nanotech gadfly Dexter Johnson looks at the big business of making products from the atom up. In his blog, he brings his global perspective on all things nanotech, including scientific breakthroughs, public policy debates, and emerging companies. Watch for his coverage of Nanotech Europe starting 28 September.

ALSO ONLINE

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

TECH INSIDER WEBINARS

SEPTEMBER WEBINAR SCHEDULE

16 SEPTEMBER: Running MATLAB Scripts in the Cloud and on Clusters Using Star-P

17 SEPTEMBER: Robotics and Academic Research

Check out these webinars and more at http://www.spectrum.ieee.org/webinar.

WWW.IEEE.ORG/THEINSTITUTE

AVAILABLE 8 SEPTEMBER ON THE INSTITUTE ONLINE

CANDIDATES TAKE ON THE TOP ISSUES

The IEEE annual election is right around the corner, and it's time to learn where the candidates for 2010 President-Elect stand on important IEEE issues. J. Roberto Boisson de Marca, Moshe Kam, and Joseph V. Lillie answer questions about their top priorities, the recession, green technology, and more.

BOOST YOUR SKILLS AND GET CERTIFIED

In this economic climate, holding onto your job or looking for a new one can be tough. But IEEE can help. Find out how to get an edge on the competition by getting certified in software development and wireless communications.

IEEE HONORS TWO TECH BREAKTHROUGHS

This month IEEE is honoring the ARPANET node at the Stanford Research Institute and the first external pacemaker as IEEE Milestones in Electrical Engineering and Computing. Learn about the history of these achievements.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2009 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at W224-N3322 Duplainville Rd., Pewaukee, WI 53072-4195, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of American Business Media, the Magazine Publishers of America, and the Society of National Association Publications.

back story

ROCKING OUT: Joshua Dalsimer [center] joins Alex Rigopulos [left] and Eran Egozy on stage.

Dropping the Drumsticks

THIS PAST June, *IEEE Spectrum* photo editor Randi Silberman Klett asked photographer Joshua Dalsimer if he'd like to shoot Alex Rigopulos and Eran Egozy, the creators of *Rock Band*, the popular video game that lets players perform songs using instrument-shaped controllers. At the time, Randi couldn't have known how perfect her choice of photographer was. Not only does Dalsimer play the game, he's *in* the game.

Before becoming a photographer, Dalsimer was the drummer for the Mighty Mighty Bosstones, a ska-core band from Boston. The Bosstones are most famous for their 1997 megahit "The Impression That I Get" (refrain: "I've neeeever had to knock on wood"—just YouTube it). But it was another hit, "Where'd You Go?" that made it into *Rock Band*.

"It's really fun to be immortalized in *Rock Band*," says Dalsimer, who now lives in New York City.

On a July afternoon, Dalsimer stepped into the Middle East, a Cambridge, Mass., indie rock club, to photograph Rigopulos and Egozy. Their company, Harmonix Music Systems, is set to release the much anticipated *The Beatles: Rock Band* this month (see article in this issue).

It was a nostalgic homecoming for Dalsimer, who'd performed at the nightclub quite a few times. "Returning there was like my past and present were colliding," he says.

Dalsimer started with the Bosstones in the mid-1980s, when he was 16. The band went on to record several successful albums. He split from the group in 1991 and played in a couple of other bands. In the late 1990s, after taking up photography, he decided to swap his drumsticks for a camera. Did his first career help him with the one that followed?

"They're both in the background, and I like that," he says. "There's also the heavy equipment to carry."

For the shoot at the Middle East, the idea was to create the feel of a live performance, with the Harmonix guys rocking out on stage. Everything looked good, except for one thing—there was no music.

"They had to totally fake it," Dalsimer says. "They did great." ❑

JOSHI RADIN

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. 46, no. 9 (INT), September 2009, p. 11, or in *IEEE Spectrum*, Vol. 46, no. 9 (NA), September 2009, p. 15.

EDITORIAL CORRESPONDENCE

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

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

contributors

ANDREA DI BLAS and **TIM KALDEWEY,** both researchers at Oracle Corp., write about their work using graphics processors to comb through enterprise databases in "Data Monster" [p. 42]. Di Blas, originally from Turin, Italy, was a researcher and lecturer at the University of California, Santa Cruz, when he met Kaldewey. "He was one of my most brilliant students," Di Blas says.

DAVID KUSHNER, an *IEEE Spectrum* contributing editor, envied the faux musicians at the demo he attended for "The Making of *The Beatles: Rock Band*" [p. 26]. "I'm pretty good on the video-game drums," he says. "My favorite song to play in *Rock Band* is 'The Trees' by Rush." Kushner's third book, *Levittown: Two Families, One Tycoon, and the Fight for Civil Rights in America's Legendary Suburb*, was published this year.

MITCHELL LAZARUS is a partner in the Washington, D.C., law firm of Fletcher, Heald & Hildreth. Besides a law degree, he holds two degrees in electrical engineering, plus a doctorate in experimental psychology—excellent training for the work he does in the regulation of new telecommunications technologies. Lazarus helped launch two educational television series and has published five books. He also blogs regularly. But the prose of his that readers are probably most familiar with is the government warning that appears on the packaging of alcoholic beverages, which he penned while helping to pass the U.S. Alcoholic Beverage Labeling Act. Reading "Radio's Regulatory Roadblocks" [p. 38] is not dangerous to pregnant women but should not be attempted while operating heavy machinery.

BABAK A. PARVIZ wakes up every morning and sticks a small piece of polymer in each eye. So it was only a matter of time before this bionanotechnology expert at the University of Washington, in Seattle, imagined contact lenses with built-in circuits and LEDs. "It's really fun to hook things up and see how they might work," he says. In "For Your Eye Only" [p. 32], Parviz previews a contact lens for the 21st century.

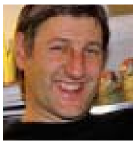

MICHAEL TOMPERT of Raygun Studio, who created the photo-illustration that opens "For Your Eye Only" [p. 32], is originally from Stuttgart, Germany. Before founding Raygun in 2005, he created original art for Apple products like the iPod and iTunes. He can imagine the "bionic" contact lenses in the story coupled with the iPhone, enabling you to scroll through your playlist without taking out your phone.

MICK WIGGINS, who illustrated "Radio's Regulatory Roadblocks" [p. 38], starts with a pencil sketch, which he then scans into his computer. Afterward, he adds color and refines the lines. His illustration for *IEEE Spectrum*'s Technically Speaking column last August was recently chosen from over 5000 entries for inclusion in the 2009 *Communications Arts Illustration Annual*.

CLOCKWISE FROM TOP RIGHT: BABAK A. PARVIZ; ADRIEN TOMPERT; MICK WIGGINS; MITCHELL LAZARUS; BETH KELLY

spectral lines

COLORBLIND/GETTY IMAGES

Working in an Always-On World

THE ATMOSPHERE at this year's *Fortune* Brainstorm Tech meeting, held in Pasadena in July, was unsettled. Unlike previous Brainstorm Tech meetings, which more often than not have been feel-good events, this gathering saw a steady stream of magnates from publishing, advertising, marketing, and technology approach the stage to analyze the storms besetting their industries and predict what the next silver linings will be.

Entertainment and publishing giants like Walt Disney Co., Sony, and News Corp. are struggling to figure out how to make money in a world where people are served up an endless buffet of free news, information, and entertainment, anywhere, all the time. Public relations and marketing firms are scrambling to respond to the Twitter/Facebook effect, which lets companies talk directly to consumers and vice versa, 24 hours a day, without having to go through flacks and marketers.

Telecommunications and computer companies as always are striving to deliver the most flops or baud in the smallest, cheapest, cleverest appliances. But now they're doing it in an environment in which anything less than a killer app has a shelf life of less than a year. Ouch.

Even politicians are chastened. Speaking at the meeting, former U.S. presidential candidate Howard Dean said that while President Obama's use of new Web tools put him in the White House, these same tools have the potential to "put politicians out of business." Citizens, Dean went on to say, may not need politicians to get things done anymore. They can organize for themselves, whenever and wherever they want to.

Into this subdued setting came Ray Kurzweil, inventor, entrepreneur, and technological optimist, to deliver a bracing rendition of his well-known thesis about the exponential advance of semiconductor technologies and their impact on technological innovation. What used to fit in a building, Kurzweil is fond of saying, now fits in your pocket, and what fits in your pocket today will fit inside a blood cell in 25 years. Such advances will continue to fuel a wealth of smart, intercommunicating devices that will allow us to improve our health and our world. They will ultimately lead in the next 20 years to what Kurzweil and his cohorts call the singularity, the moment when machines will achieve consciousness and humans will begin to lead indefinitely long lives.

But what does a daily life look like for the rest of us? What will work be like in 10 or 20 years?

Some say it will be a lot like it is today, only more so. Organizational theorists like Thomas W. Malone, of MIT's Sloan School of Management and author of *The Future of Work*, believe that any job that involves parsing or creating knowledge will be carried out by "e-lancers" who will rarely go to an office. No more sweating in traffic jams, but the already-shrinking divide between work and personal life will continue to disappear.

Others think offices will remain important, but that embedded sensors and intelligent agents, combined with high-powered search technologies, will make some kinds of knowledge work obsolete. Machines will research, collect, sort, update, and weight information. People will decide what to do with it next.

Kurzweil (and no, he's not a game developer) has said that video games "are the harbingers of everything" and thinks that in the next decade or so "full-immersion" virtual reality games will be widely used for business and pleasure. And if computer networks get powerful enough, meetings really could go virtual. "Attendees" would simply project their charts, data—and themselves—as if they were all there for real. What a nice change from the huddle-around-the-speakerphone weekly staff meeting!

This all made me think of a prescient article *IEEE Spectrum* published nearly a decade ago by Professor Philip Agre of the University of California, Los Angeles, called "Welcome to the Always-On World." It described how new communications technologies would push human relationships to evolve in potentially unnerving ways: "The always-on world...is a world of freedom, but it is also a world of anonymous global forces that ceaselessly rearrange all relationships to their liking. We don't understand this world very well, but we will soon have plenty of opportunity to study it firsthand."

Well, our always-on world is now here, and it's clear that it deserves plenty of scrutiny. While the benefits of our wired-up transnational society are many, the "existential downside," as Kurzweil calls it, merits our concern. —SUSAN HASSLER

JAPAN NATIONAL TOURISM ORGANIZATION PRESENTS:

Aesthetic Sciences

A leading Japanese-American Scientist Discusses The Benefits of Holding Meetings in Japan.

FOR IEEE's SPECTRUM'S 125th ANNIVERSARY, Hiroshi Ishii, associate director of MIT's Media Lab, was kind enough to participate in an interview wherein he discusses his work in one of the United States' leading havens of innovation in the field of communications technology. In the interview, Professor Ishii discusses his passion for his work, his Japanese roots, the profound cultural influence Japan had on his ideas, and Japan's role as a pioneer in the field of technology.

The Tokyo-born engineer, renowned for his continual progress in fields ranging from telcom networking to child development, aims to infuse an ethereal outlook towards the engagement of sciences, offering a much more dynamic approach to standard methods and processes. After serving as an associate visiting professor in Toronto, Ishii became the first Japanese faculty member to join the Media Lab team at MIT. His work at Media Lab has garnered significant praise not only in the engineering world but also in the art and design world, earning international accolades by several organizations, including an award from the Computer-Human Interface Academy.

A Curious Nature

Growing up in Japan, Ishii considers the harmonious nature of his upbringing and how it is reflected in his current research. In an interview with Yahoo Japan's Rikunabi-Next, Ishii metaphorically describes how ambient media first made a lasting impression. "We had an abacus that served as a communication tool between my mother and me. The clicking sounds the abacus made when my mother was using it taught me that it was not the right time to ask her to play with me. This notion was used as a concept in Media Lab." Ishii adds, "There was a huge impact on me after I touched an abacus for the first time since I was a child, which made me realize that tangible interfaces, such as an abacus, are direct and straight; opposite of graphical user interfaces such as a Mac."

Years later, Ishii traveled all around Japan, staying in youth hostels as he cultivated an appreciation for the arts, architecture, and nature of his home country. Noting that each region of Japan carries a rich and historic heritage, Ishii believes that the aesthetics of Japanese arts influenced his work in a subconscious yet profound way.

A Balance of Symmetry

It should come as no surprise that Ishii would pursue a career that involves balancing the tangible with the virtual. One of his more high profile projects, Tangible Bits, seeks to eliminate the use of hardware as applied to computers, which he considers to be a machine that could visually represent the invisible without necessarily using a mouse and a keyboard as devices of manipulation. In a 2004 article with MIT Spectrum, Ishii notes, "An engineer just makes things work. But the artist asks profound, provocative questions: What feelings does this evoke? How does this relate to the whole? What does it mean? We need to look at the entire picture…Division is dangerous"

In a way, Ishii's work is more poetic than purely scientific. The ability to visually represent the societal impact of a project with his IP Network Design Workbench, for example, evokes not only a practical development but also a certain emotional aspect that inevitably causes users to actually see and feel their influence on the modern populace. The Workbench takes surface computing to a new level, allowing researchers to tactically manipulate every day objects on a special screen that features all pertinent data, right at the users' fingertips. So if an architect, for example, wanted to visually represent and alter his designs, material costs, and measurements in one tangible interface, the Workbench would provide the space. And for the audiophile tech enthusiast, one could look no further than one of Ishii's favorite projects called musicBottles, where ordinary perfume bottles embedded with electro-magnetic tags that sit atop a tangible interface are transformed into musical instruments. The bottles communicate wirelessly with the base as the removal of each top elicits different sounds to create an actual composition, seamlessly fusing fundamental concepts behind art and technology.

A Cultural Influence

The professor's forward-thinking nature combined with his Japanese roots has offered an unprecedented cache around his intellectual circles, an attitude seldom witnessed throughout the tried-and-true practices of the scientific world. His Japanese upbringing is undoubtedly reflected in his research; traditionally, the Japanese culture as a whole has a unique and intensive grasp of the aesthetics of the world and nature. It is this kind of balance between man and nature that directly correlates with his work, creating a perfect harmony of the physical world and technology. "The balance between the advanced technologies and the beauty of tradition makes Japan very unique." Ishii observes. "Japan takes pride in its unique cultural heritage, yet maintains a cosmopolitan approach to innovation and engineering."

There is still an undeniable Japanese influence on Ishii's work. When asked about Japan's current role as pioneers of new technological endeavors, he notes that since the land has limited natural resources, its inhabitants have always had to work hard to create new industries based on strong tradition of arts and engineering. For example, Japan has been at the forefront of the semiconductor, electronics, car, and robotics industries, all of which combine strong senses of engineering and aesthetic design.

Sound meets science with Ishii's Music Bottles project, just one of several fascinating and intuitive projects from MIT Media Labs.

A Nucleus of Commerce

Since several companies in Japan sponsor MIT Media Lab, Ishii returns to his home country three to four times a year to attend conferences. "Too many great international conferences organized by IEEE and ACM are happening in Japan, and I wish I could attend them all", Ishii says. "I often attend those conferences to give keynote speeches, and I always enjoy their hospitality and high standard of professional services. The extensive, efficient public transit system makes it easy to navigate from one venue to the next in large cities. You can easily find the best route and estimated time to travel with trains and subways using online web service."

As media research and development continues, Japan has undoubtedly emerged as a hotbed of technological innovation. "At conferences in Japan one is able to enjoy the fruits of international collaboration and ingenuity while appreciating how local influences have shaped the course of Japanese engineering industries."

WELCOME TO JAPAN

Yōkoso! Japan

CONTACT INFORMATION:

Japan National Tourism Organization
Japan Convention Bureau
Tel: 212-757-5641
E-mail: jcb@jntonyc.org
www.japantravelinfo.com

Tokyo Night Cityscape with Mt. Fuji in the background.

Cherry blossoms at Himeji Castle.

forum

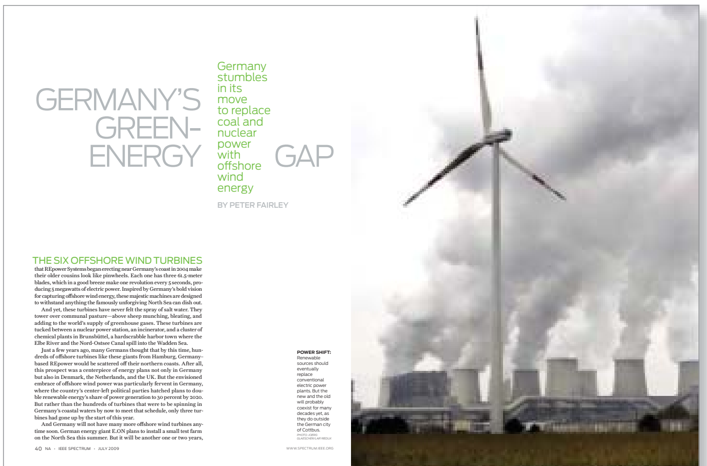

LETTERS do not represent opinions of IEEE. Short, concise letters are preferred. They may be edited for space and clarity. To post your comments online, go to http://www.spectrum.ieee.org. Or write to Forum, *IEEE Spectrum*, 3 Park Ave., 17th Floor, New York, NY 10016-5997, U.S.A.; fax, +1 212 419 7570;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

ENERGY BY THE NUMBERS

PETER FAIRLEY'S comprehensive "Germany's Green-Energy Gap" [July] teaches two lessons. The first is that powering a national electrical grid with green energy other than nuclear is illusory in a densely populated modern industrial country. The second is that politicians, whether they're led by a German physicist or an American lawyer, can't seem to get it right.

Wind turbine output is measured in megawatts. Nuclear plant output is measured in gigawatts. Thus the starting point is 1000 to 1: advantage nuclear. Adding significant power to a grid requires more than the hundreds of turbines cited in the article; it requires thousands upon thousands. And as the article describes, the wind doesn't always blow. The backup is either nuclear or other legacy producers—or candles in the cupboard.

The solar disadvantage is not so straightforward, but existing installations demonstrate that it takes about 10 000 acres of solar panels to match the output of one nuclear plant. Again, it's 1000 to 1: advantage nuclear—when the sun shines. Otherwise, the nuclear advantage is infinite, because there is no solution for large-scale electricity storage for dark nights and cloudy days.

In the German community of Biblis, nuclear reactors generate 60 percent of the electricity. Green Party politicians decreed that the reactors would be shut down in 2012. The article gives the shutdown date as 2022, so apparently some reality has broken out. Are they preparing sites in Hessen for the installation of more than a thousand wind turbines and cutting down all those beloved Hessian trees? Fat chance—and solar would kill more trees than wind.

Someday the people of Germany and other countries—except France, which is 85 percent nuclear—will be in the dark unless politicians learn to do enough arithmetic to see the light.

GEORGE F. STEEG
IEEE Life Member
Potomac Falls, Va.

EXPANDING BROADBAND'S REACH

THE RENEWED interest in broadband over power lines (BPL) because of the availability of US $2.5 billion, plus another $4.7 billion from the recent Recovery and Reinvestment Act, is truly depressing. David Schneider's article, "Is This the Moment for Broadband Over Power Lines?" [Update, July], fails to mention the most drastic problem with BPL: It has the potential to create serious interference and ruin all shortwave and AM broadcast signals within sight of power lines. Radio is dear to the IEEE and is still a big business, so we ought to protest the mistaken rush to BPL by the Washington bureaucracy.

ROBERT BURGER
IEEE Member
Member, American Radio Relay League
Veneta, Ore.

Senior Editor David Schneider responds: The interference issue has indeed cast a long shadow over BPL for years. In particular, the possibility of disrupting amateur radio operators has been most problematic, as was well described in a September 2004 *IEEE Spectrum* article, "Amped Up and Ready to Go" [see http://www.spectrum.ieee.org/semiconductors/materials/amped-up-and-ready-to-go]. While investigating my story, I spoke with a knowledgeable source at the American Radio Relay League. That discussion helped convince me that such interference wasn't going to be a showstopper for BPL.

The BPL equipment now being installed often modulates power lines with frequencies in the range of 30 to 50 megahertz, which avoids amateur bands entirely. Such devices also couldn't interfere with shortwave and AM broadcasts, which are below 30 MHz. Given that the users of radio frequencies between 30 and 50 MHz haven't voiced loud opposition to BPL with the Federal Communications Commission—which has approved this technology—it seems reasonable to conclude that radio interference presents no insurmountable roadblock, at least from a technical standpoint.

CORRECTION

We apologize to the people of the Czech Republic for labeling their country "Austria" on our map in "Germany's Green-Energy Gap" [July].

update

more online at www.spectrum.ieee.org

DESIRABLE: Trolleys like this quaint one running on rail tracks in New Orleans seem a thing of the past. But streetcars could make a comeback, thanks to fuel cells that power each train independently, obviating the need for unsightly power lines.

PHOTO: MATTHEW SEPTIMUS/GETTY IMAGES

Fuel Cells Could Power a Streetcar Revival

Hydrogen power is one way to untether trolleys

Advocates of hydrogen-fuel-cell propulsion have a new target for their technology: trolleys. At the Fifth International Hydrail conference, held 11 and 12 June in Charlotte, N.C., engineers and transit planners concluded that streetcars are an ideal early application for hydrogen propulsion. Traditional trolley cars have drawn their power from catenaries, unsightly overhead electricity supply lines running along city streets. Hydrogen-powered streetcars would eliminate the wires, which otherwise might stand in the way of a streetcar comeback.

At first blush, streetcars may not seem worth devoting much effort to. Many Americans think of them as a quaint anachronism retained by a few cities like New Orleans for nostalgia's sake. But more than a dozen municipalities around the world have restarted and extended trolley car lines because they attract wealthier riders than buses and inspire new, high-density property development. What's more, streetcars—because of the low rolling resistance of steel wheels on rails—require much less energy than the rubber-tired buses that elbowed them aside decades ago. For these and other reasons, says the American Public Transit Association, some 50 cities in the United States are planning new streetcar lines.

Choosing a streetcar design based on hydrogen could save such municipalities millions of dollars per kilometer, say advocates. A 2006 assessment found that roughly one-quarter of the total cost of a planned 16-kilometer rail extension connecting the city and

update

surrounding suburbs would have gone toward installing a catenary system, according to Brian Nadolny, interim project manager for the Charlotte Area Transit System (CATS).

Transit operators are cautiously optimistic about hydrogen-powered trolleys. "If the technology is proven viable, it could revolutionize the industry," says Keith T. Parker, CEO of Via Metropolitan Transit in San Antonio. Parker explains that in addition to saving the installation cost of a catenary system and support poles, the operator would save the maintenance costs.

"The biggest issue with [hydrogen trolleys] is that there isn't one on the ground and running, so municipalities don't have anything to test that would allow us to make a decision about whether to purchase one," says Parker.

Proterra, in Golden, Colo., may be the first to deliver a prototype. It makes a plug-in hydrogen-electric hybrid bus with a battery pack that can be fully recharged in 10 minutes. The bus can travel roughly 50 to 70 km between charges, enough for 2 to 3 hours of operation.

Proterra's system carries more batteries than hydrogen-fuel-cell stacks, because the system can draw most of the energy it needs for a day's travel from the grid, charging periodically at high-voltage links when the bus is in service and via lower-voltage connections when it's out of service. Hill says Proterra hopes to demonstrate to Charlotte's CATS system the potential for turning the bus into a pseudo-trolley by putting it on steel wheels.

"The more widely the issue is discussed now and the sooner a proof-of-concept hydrogen trolley system is formally demonstrated, the smaller the risk that transit systems' investments in overhead trolley planning will be for naught," says H. Stan Thompson, chair of the hydrogen economy advancement team at the Mooresville–South Iredell (N.C.) chamber of commerce.

Of course, hydrogen fuel cells aren't the only way to untether trolleys. Bombardier's German subsidiary has rolled out a system based on contactless inductive transfer of electricity to pickup coils on a streetcar's undercarriage from cables buried beneath the track bed. Shanghai's Sunwin Bus Co. has fielded battery-powered trolleys that quickly draw power every few stops, extending an arm hidden on the top of the car. And Kawasaki Heavy Industry's SWIMO light-rail vehicle travels 10 km or more between stops on a 5-minute charge of its Gigacell batteries.

The French firm Alstom pioneered the wireless streetcar movement with its Citadis tram, which went into service in Bordeaux, France, in 2003 and has cars powered by a segmented ground-level third rail. Only the segment that is completely covered by a train is electrified.

—Willie D. Jones

ALAMY

Britain Mulls Over Digital Radio Transition

The British communications ministry has released a road map calling for the completion of the nation's move from analog FM to the digital audio broadcast (DAB) standard by the end of 2015. But DAB skeptics say that the move will waste money and that consumers just aren't interested. Conversion will require anywhere from 120 to 700 new transmitters, according to Eureca Research analyst Gareth Owen, and there will be 130 million useless FM receivers to toss.

What's more, the British government does not stand to make any money from freeing up bandwidth. Unlike the U.S. 700-megahertz spectrum—which sold for US $20 billion and will be used for new services—the FM space in Britain will simply go to small local commercial and community stations.

For more, see http://www.spectrum.ieee.org/telecom/wireless/britain-mulls-over-digital-radio-transition.

Plastic Surgery 1, Face Recognition 0

In the first test of face-recognition technology vs. cosmetic surgery, face recognition loses

For years, developers of face-recognition algorithms have been battling the effects of awkward poses, facial expressions, and disguises like hats, wigs, and fake moustaches. They've had some success, but they may be meeting their match in plastic surgery.

Systematic studies have tested face-recognition algorithms in a variety of challenging situations—bad lighting, for example—"but none of those conditions had nearly the effect of plastic surgery," says Afzel Noore, a computer science and electrical engineering professor at West Virginia University, in Morgantown. In June, Noore reported the results of the first experimental study to quantify the effect of plastic surgery on face-recognition systems, at the IEEE Computer Society's Computer Vision and Pattern Recognition conference, in Miami. His team of collaborators is based in West Virginia and at the Indraprastha Institute of Information Technology, Delhi, in India.

Using a database containing before-and-after images from 506 plastic surgery patients, Noore and his colleagues tested six of the most widely used face-recognition algorithms. Even in pictures where the subject was facing forward and the lighting was ideal, the best of the algorithms matched a person's pre- and postsurgery images no more than about 40 percent of the time. The researchers found that for local alterations—say, a nose job, getting rid of a double chin, or removing the wrinkles around the eyes—today's systems could make a match roughly one-third of the time. For more global changes like a face-lift, the results were dismal: a match rate of just 2 percent.

"We have to devise systems for security applications knowing that people will aim to circumvent them," says Noore. In particular, researchers must examine a further complication of the plastic surgery problem—the compounding effects of a series of surgeries over time.

Meanwhile, Noore and his coauthors are testing a game-changing hypothesis: that even after plastic surgery, there are features beneath the skin but still observable that remain unchanged. —W.D.J.

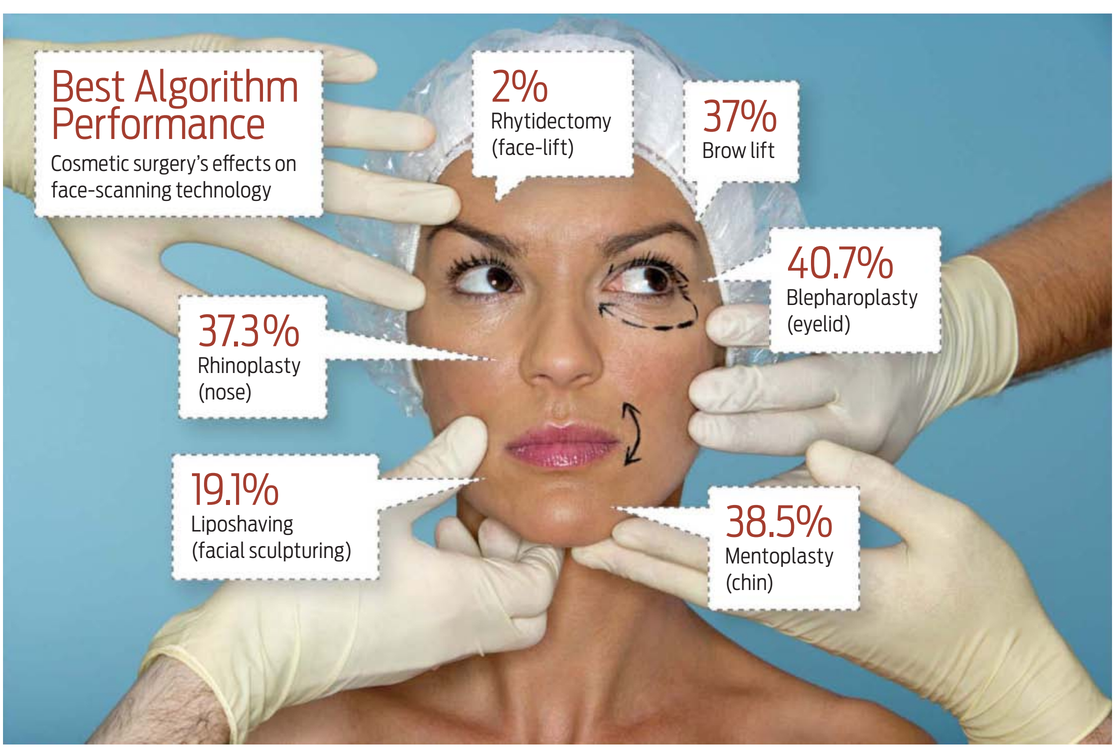

STEPHANIE WOLFSTEINER/GETTY IMAGES

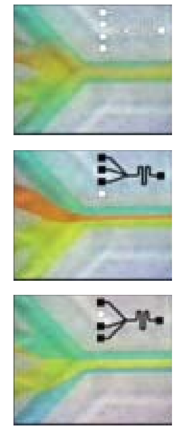

news briefs

Musical Microfluidics

Microfluidic medical diagnostic devices aren't really "micro" at all, because they need bulky pressurized-gas pumping systems to push blood samples and other fluids through their innards. But MEMS masters in Michigan have come up with a microelectro-mechanical solution: resonant cavities that turn musical notes into pressure to precisely control fluid flow. For more, see http://www.spectrum.ieee.org/biomedical/diagnostics/musicpowered-microfluidics.

IMAGE: SEAN LANGELIER/UNIVERSITY OF MICHIGAN

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

update

A High-Tech Oasis in Saudi Arabia

The richest tech university in the Middle East opens this month

On an afternoon in late July, Tony Eastham sat in a half-built office outside Jeddah, in Saudi Arabia. A table, a few chairs still wrapped in plastic, a desk, and some white settees occupied the room. Outside, the whines and growls of construction heralded the rise of a university in the desert.

These facilities are part of the hardware that makes up the King Abdullah University of Science and Technology, an ambitious graduate-level research university with a multibillion-dollar endowment. Known as KAUST, the institution will open its doors to students for the first time this month, and Eastham, the director of laboratories, has quite a few labs left to build. "What's happening here is an experiment, one that is only possible because of the resources available in Saudi Arabia," says Eastham, an IEEE member and until recently an engineering professor at the Hong Kong University of Science and Technology.

The new university is the brainchild of Saudi Arabia's King Abdullah, and its mission is to extend the country's technical prowess beyond the domain of oil. Doing so will be an uphill battle: Pupils in Middle East and North African countries score far below the world average on international math and science tests, and even within the region Saudi Arabian students underperform, according to a 2008 World Bank report on education in the Middle East. This is in spite of the fact that Saudi Arabia's government devoted approximately 30 percent of its annual budget to education through 2003, says the report. Since KAUST's inception in 2006, education spending has almost certainly skyrocketed.

OASIS OF LEARNING: The main entrance to the King Abdullah University of Science and Technology, in Jeddah, is flanked by construction cranes [top]. The photograph below provides a vista of the main campus from the Red Sea. *PHOTOS: JB. PICOULET/S. LOURIÉ*

Once completed, the university will include a top-notch nanofabrication laboratory and a visualization and virtual reality center designed by the University of California, San Diego. The crown jewel is Shaheen, an IBM Blue Gene/P supercomputer capable of 222 teraflops and ranked 14th in the world.

The university's first batch of faculty was recruited by the same people who designed KAUST's course programs—namely, partner universities such as the University of Cambridge, Imperial College London, and the University of California, Berkeley. Stanford University professors patterned the curriculum for applied mathematics after their own and then interviewed the prospective faculty, including David Ketcheson, a newly minted doctor of applied math from the University of Washington. "I never thought I'd graduate and go to Saudi Arabia," Ketcheson says. "But what KAUST has invested in applied math I don't think exists anywhere else in the world."

The inaugural class of 350 students is drawn predominantly from China, Mexico, Saudi Arabia, and the United States. A committee from the Institute for International Education, which also selects students for the Fulbright program, admitted the students. To accommodate the expectations of newcomers thrust into one of the world's most restrictive cultures, the campus and its immediately surrounding community will abide by different rules from the rest of the country. Women will be allowed to drive, and male and female students will study together. Students and faculty will be surrounded by most of the trappings of a Western college town. But a vastly different world lies just outside the university gates, where male guardians oversee many aspects of women's daily life, Internet filtering is extensive, and corporal punishment is a legal tradition. "It's going to be a very bewildering experience for just about everybody who arrives on campus," says David Keyes, KAUST's chair of mathematical and computer sciences and engineering, formerly of Columbia University. —Sandra Upson

121 DAYS The number of days' worth of excess inventory held by photovoltaics manufacturers in the first quarter of 2009, according to iSuppli Corp., in El Segundo, Calif. The industry is at an unprecedented level of oversupply.

New-Breed Browsers Are Harder to Hack

Chrome, Opus Palladianum, and Gazelle will be more secure, with features borrowed from operating systems

JULY WAS a bad month for old-fashioned Web browsers. Troublemakers found ways to infiltrate personal computers by breaking in through Microsoft's Internet Explorer and Mozilla's Firefox.

As security experts issued dire warnings and the companies scrambled to produce software patches, computer scientists at the University of Illinois at Urbana-Champaign quietly put the finishing touches on what many believe is the only solution to a growing number of online security threats: a radically redesigned Web browser.

Samuel King, an assistant professor of computer science who is spearheading the effort, says he and his students plan to release a public version of their new Web browser, Opus Palladianum, or OP, on 1 September. Named after a technique for making mosaics from irregularly shaped material, OP's moniker is a tribute to Mosaic, the original Web browser, also developed on the Urbana-Champaign campus. The development of OP is part of a movement to dramatically improve the security of the Web by making browsers that behave more like operating systems. Microsoft and Google are working along the same lines.

"From a security perspective, browsers are completely broken," King says. The problem with traditional browsers is that the way people use the Web has changed. Instead of just looking up information on static pages coded with HTML, or HyperText Markup Language, people are using the browser to run Web versions of applications that used to reside on a PC, such as e-mail, social networking, and online banking. "I don't think my mother uses anything besides her Web browser," King says.

But browsers weren't built to manage access to applications and sensitive data. And their vulnerability to digital attacks is increasingly attracting everyone from run-of-the-mill mischief makers to sophisticated criminal organizations. Researchers at Sophos, a security company based in Abingdon, England, near Oxford, say they are discovering sabotaged sites every 3.6 seconds, quadruple last year's rate.

TAKING A PAGE: As browsers are increasingly used for sensitive transactions, a security loophole has widened. New designs borrow from operating system features.

King and two of his students began working on OP two years ago with the idea that if they divided a browser into separate subsystems—for instance, the user interface, storage, and networking—they could make it more secure. Communication between the different parts of the program is kept simple and explicit, much as processes are managed in an operating system.

Because the pieces of the browser are carefully kept apart, it becomes much more difficult for bad guys to install software that steals passwords and credit card numbers or sends out spam. It's as if a burglar broke into your home and couldn't get past the front entryway, King says.

Meanwhile, computer scientists at Google and Microsoft are adopting similar ideas. "We were very much aware of the limitations of the [traditional] browser," says Brian Rakowski, a Google product manager.

Like King and his students, Google engineers concluded that they could get greater security and more reliable performance by separating the browser's functions into different processes. They released the Chrome browser in September 2008 as an open-source project for review. The company is also working on a new, stripped-down operating system, the Chrome OS, designed to support Web browsing.

King says Chrome gave him some significant ideas for how to improve OP, and he also borrowed some ideas for displaying content in the browser from Microsoft's Gazelle project, which like OP and Chrome seeks to build a next-generation browser. The inspiration flowed both ways. King consulted on the Gazelle project and also received some funding from Microsoft Research's Internet Services Research Center.

—ELISE ACKERMAN

news briefs

Say Cheese T-shirts that can snap photos may one day be possible, thanks to the creation of an optical fiber that can detect images. MIT photonics researcher Yoel Fink and his colleagues have created a polymer fiber that can detect the angle, intensity, phase, and wavelength of the light hitting it, information that can be used to re-create a picture of an object. For more, see http://www.spectrum.ieee.org/semiconductors/optoelectronics/optical-fiber-fabric-snaps-pictures.

IMAGE: FINK LAB/MIT

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

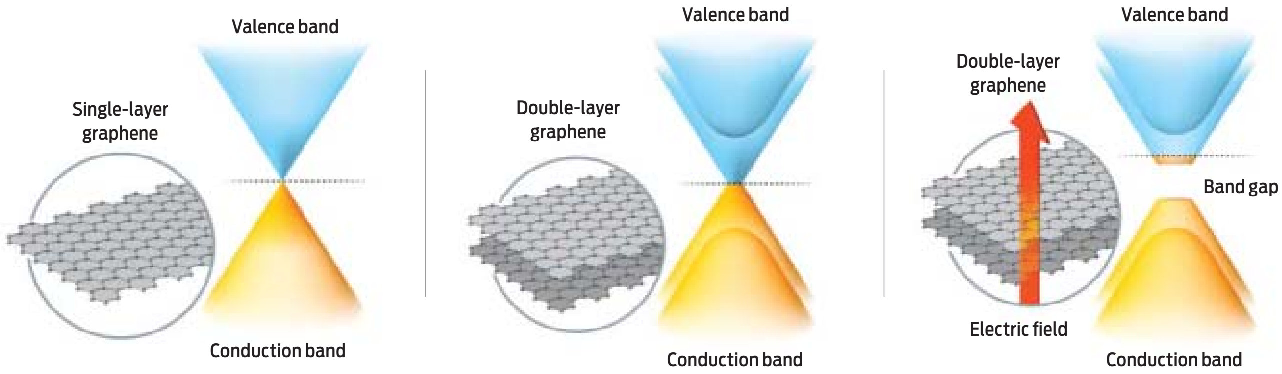

BREAKING SYMMETRY: The existence of an energy gap between the conduction and valence bands of a material is what makes it possible for it to act as a semiconductor. In both single-layer and double-layer graphene [left and middle], the valence and conduction bands meet at a point, with no band gap. The introduction of an electric field perpendicular to the layers [right] creates an asymmetry, which generates a band gap. Though small, the gap is tunable, creating possibilities for new devices.

Graphene Makes Transistors Tunable

But even the "wonder material" has its limits

Graphene, a one-atom-thick layer of carbon atoms, is the strongest material ever tested and more conductive than the purest silver. Since its isolation in 2004, scientists have been dreaming of impossibly tiny graphene transistors. But there's a major challenge: The material has no electronic band gap, the semiconductor property that controls the operations of transistors, lasers, and other solid-state devices.

It's now been demonstrated, however, that graphene can be given a band gap—and not just any band gap, but a tunable one. The ability to modify a device's energy gap could lead to detectors that respond only to a particular wavelength of light or light emitters whose color is controlled.

Researchers at the University of California, Berkeley, and Lawrence Berkeley National Laboratory produced the band gap in a field-effect transistor, or FET, made from two layers of graphene. Though the maximum energy gap is only about a quarter of those typically found in silicon devices, the researchers believe it is big enough to allow for new kinds of semiconductor nanodevices.

Graphene lacks a band gap because of its symmetrical structure—its atoms scatter electrons in such a way that they cancel each other out. But "when we apply an electric field perpendicularly to the graphene, we break the symmetry and create a band gap," explains Feng Wang, head of the Ultra-Fast Nano Optics Group at UC Berkeley and lead author of a recent *Nature* paper describing the work [see diagram, "Breaking Symmetry"]. The symmetry can also be broken by doping one layer of graphene with metal atoms, but such doping is hard to control.

The team built their device by making a bottom gate layer out of silicon, with a thin layer of silicon oxide as an insulator. On top, they placed the two layers of graphene, with gold at each end to act as a source and a drain. Above that they put a layer of sapphire, and above the sapphire they placed a second gate made of platinum. By applying an electrical field at the gates, they were able to create a continuously tunable band gap ranging from 0 to 250 millielectron volts.

That's well short of the 1.1-electron-volt band gap of silicon, but still valuable, says Wang, who believes he'll be able to achieve a maximum band gap of about 300 meV. "I think it can be useful, because it's already a dramatic change, but it will not function exactly like silicon," Wang says. It takes a larger band gap to get a big enough contrast between the high-voltage "on" state of the device and the low-voltage "off" state. "If you need the ultimate off state, this will not be ideal," he concedes.

Indeed, the on/off ratio of such a bilayer graphene transistor would not be high enough to meet the specifications laid out in the International Technology Roadmap for Semiconductors, says Gianluca Fiori, an assistant professor in the information engineering department at the University of Pisa, in Italy. He and his colleague Giuseppe Iannaccone ran computer simulations of bilayer graphene FETs using a supply voltage of 0.5 V between the source and the drain, as called for in the road map, and determined that the band gap could not suppress the flow of current enough to make a sharp distinction between the on state and the off state.

"Maybe you can think of using this kind of material not for high-performance transistors but for ultralow-power devices," Fiori says. "If you think that in the future you can use this kind of transistor in portable electronics, maybe you will only recharge the batteries once per month."

Wang believes that bilayer graphene may be more useful as an optical material than as a competitor to silicon transistors. For instance, it might be used for infrared light sources and sensors for measuring biological molecules, many of which have unique signatures in the infrared. —Neil Savage

EMILY COOPER

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

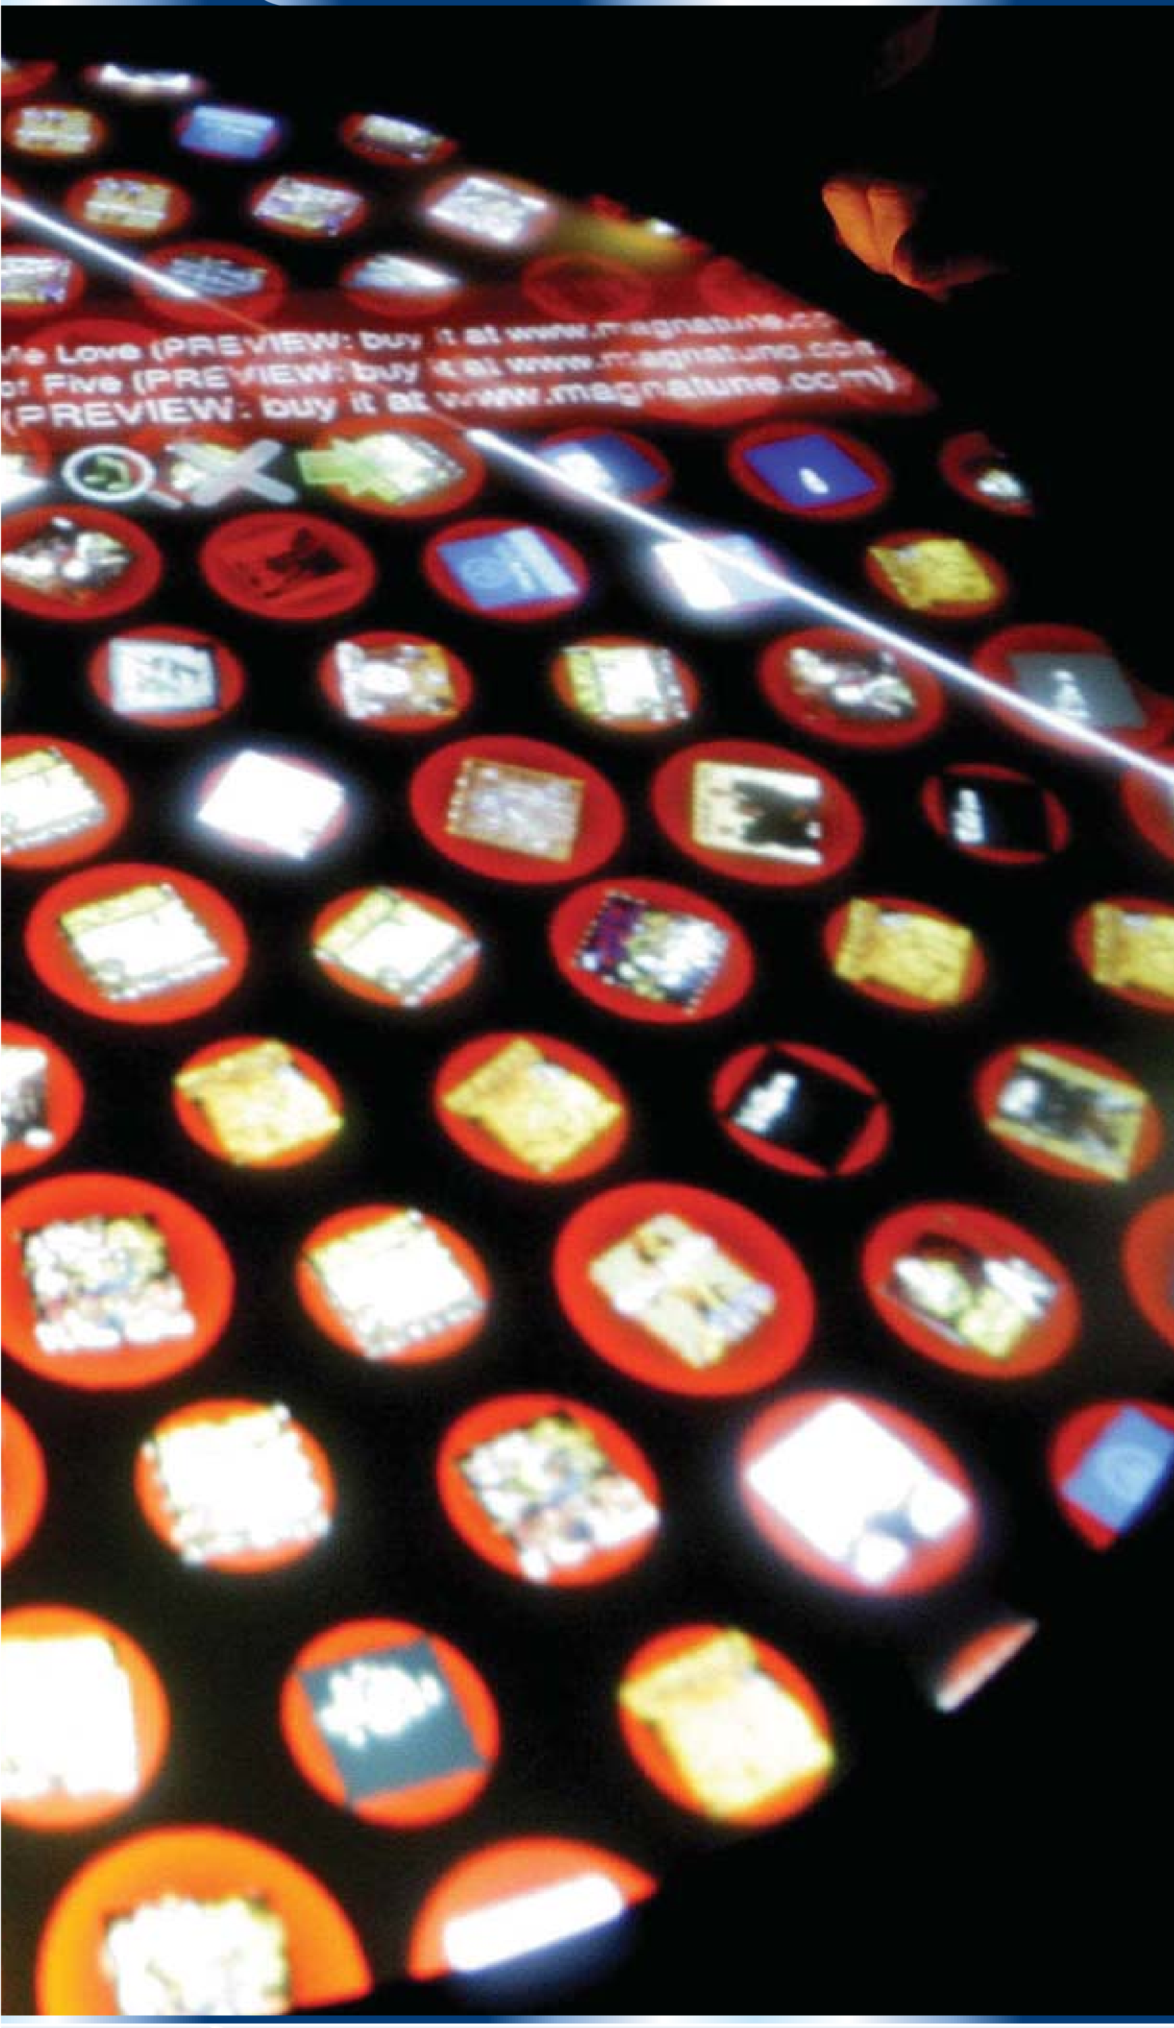

the big picture

Place and Play
Making music ain't what it used to be. Technologists at the Universitat Pompeu Fabra, in Barcelona, have reengineered the music synthesizer, making its guts interchangeable and transforming the electronic instrument into a piece of performance art.

The ReacTable's multitouch display lets users create a custom-made instrument on the fly—capable of, say, the low growl of a trumpet and the distorted whir of an electric guitar—by sliding or rotating blocks, called tangibles. The blocks' surfaces are printed with patterns that the ReacTable is programmed to recognize as parts of a synthesizer, such as voltage-controlled and low-frequency oscillators. A video camera inside the table detects the presence of fingertips and tangibles and prompts the system to light up and generate the sounds suggested by the arrangement of the objects.

PHOTO: JOSEP LAGO/ AFP/GETTY IMAGES

careers

THE iPHONE'S MUSIC MAN

The Ocarina app came from Stanford professor Ge Wang's sonic explorations

PIED PIPER: Ge Wang's iPhone app, which turns the phone into an ancient flute, has been downloaded a million times at US $1 each.
PHOTO: TIMOTHY ARCHIBALD

NOT LONG ago, a kitchen mishap at a Palm Desert, Calif., dinner party caused an hour-long delay between the appetizer and the entrée. Conversation began to lag. Then folks began pulling out their iPhones—not to check their e-mail but to play Ocarina, a US $1 iPhone app that lets users blow into the phone's microphone while pressing on virtual keys to mimic an ocarina, an ancient flutelike instrument.

Who would have thought of blowing into an iPhone to make music? Ge Wang, that's who. Wang, an assistant professor at Stanford University, has made a career of turning computing devices into musical instruments.

Born in China, Wang was 9 when his parents brought him to the United States, where his father was working toward a degree in operations research at Georgia Tech. In time Wang picked up the accordion and guitar. When he matriculated at Duke University, in Durham, N.C., he thought about studying music. His father, however, suggested he pair that pursuit with something more practical.

So Wang majored in computer science and fit in as many music classes as he could. After hearing an electronic music piece by composer Paul Lansky, he finally saw how to merge his two interests. He spent six years as a graduate student at Princeton designing electronic instruments, working with the nascent Princeton Laptop Orchestra and developing ChucK, a programming language optimized for the creation of sound. In 2006, his professor and mentor, Perry Cook, told him it was time to graduate. "It was the sensei saying to the student, 'Go out into the world; you are ready,'" Wang says.

Wang landed at Stanford University's Center for Computer Research in Music and Acoustics as an assistant professor, where he teaches courses like "Compositional Algorithms, Psychoacoustics, and Spatial Processing." He also established the Stanford Laptop Orchestra and a mobile-phone orchestra and continues to work on ChucK. And in 2008 he joined with Ph.D. student Jeff Smith to start Smule, a company that develops iPhone applications.

"When the iPhone came along, I saw this miniature computer that's portable, personal—even intimate—and pervasive," he says. So when Smith came to him, already with investors ready to back the company (the seed round eventually totaled $1.7 million), he couldn't say no.

Smule's first iPhone app, introduced last September, was the Sonic Lighter, a virtual flame. As far as anyone knows, this was the first app to be controlled by the microphone. To blow out the flame, you blow into the mike. The app has another unique feature: global visualization—a map of the world that shows in real time where other people have Sonic Lighters lit. And sometimes that information is surprising.

"For many weeks," Wang said, "there was a point of light inside the Arctic Circle—we have no idea what it is, there is no visible land mass there—but there were six iPhones there using Sonic Lighters." Wang theorizes that they belonged to workers on an oil rig.

The lighter led to Sonic Boom (a virtual firecracker), Sonic Vox (a "real-time voice shifter"), and then, last November, to Ocarina. It works just like a real ocarina: You cover "holes" on the display while blowing into the mike. The sound wave is computed on the fly—the sounds are not prerecorded—using information on the holes covered, the force of the air on the microphone, and the vibrato. Like the Sonic Lighter, Ocarina includes a map for global visualization. And besides seeing where people are playing the Ocarina, you can listen in.

Smule has sold over a million downloads of Ocarina and now has 15 employees. Its latest application, Leaf Trombone World Stage, combines the sound of blowing across a leaf with a slide to change the pitch, and it lets listeners judge users.

Meanwhile, Ocarina continues to serve as an icebreaker for strangers, provide a sound track for holiday parties, generate YouTube videos, and entertain bored dinner guests.

—TEKLA S. PERRY

E-bracadabra

Technology is magical, and now more than ever, magic is technological

It's a bird! It's a plane! It's...*gone*? Illusionist Franz Harary first wowed audiences 15 years ago by making the space shuttle disappear. "There's a kind of glass that becomes reflective when electrically charged, so you're able to produce and vanish objects by turning their reflections on and off," he says.

The late Arthur C. Clarke famously said, "Any sufficiently advanced technology is indistinguishable from magic." For its part, technology has long played a vital role in the world of magic. One of the world's first steam engines served a magic trick when Heron of Alexandria, the first-century scientist and engineer, opened a temple door with a secret counterweight.

Today magic shows up most often as electronics—for example, a deck of cards embedded with RFID tags that transmit their identities to a hidden electronic decoder. But audiences have come to expect a technological explanation, says Harary. "Magic has always had to stay ahead."

Long ago, magicians had a grace period. In the mid-1800s, Jean Eugène Robert-Houdin, the father of modern magic, employed a barely known phenomenon, electromagnetism. He had a giant electromagnet built into the stage floor to hold down a metal chest. Adults couldn't budge it when the magnet was turned on, but children could lift it easily—when the magnet was turned off.

"If we did this trick now, people would assume it involved such a magnet. But at that time, practically no one knew about it," says Richard Kaufman, the editor of *Genii: The Conjurors' Magazine* and author of some two dozen books on magic. With technology a part of everyday life, the challenge becomes more about how well you can use and hide it.

"As soon as magic smells of technology, it's ruined," says veteran illusion designer Jim Steinmeyer. "The art is how it's concealed and what it accomplishes in an unexpected way."

Many magicians see technology as a tool to update old tricks, such as a digital clipboard that transmits to a laptop a shape drawn by an audience member (the laptop is backstage; the magician has to guess the shape). Another effect—where an image of a chosen card shows up as a bruise on the magician's arm—uses a chemical reaction. One chemical is smeared on the arm beforehand, with the second applied by sleight of hand.

"In my experience, the magicians who love 'knuckle busting'—the repetitious training of sleight of hand—are mostly amateurs, more interested in impressing each other than amazing audiences," says mentalist James L. Clark, who edits *Magician Magazine* and has appeared on TV's "Masters of Illusion." "Those who care more about presentation than technique take the path of least resistance."

Yet others embrace science and technology as the act itself. After drumming for the now-defunct alternative rock band the Pixies, former electronic engineer David Lovering took up magic. He appears as the Scientific Phenomenalist, performing scientific and physics-based experiments on stage. Illusionist Jason Latimer, a World Champion of Magic, surprises audiences by hanging hoops and clothes on laser beams.

"Magicians have always looked for new ways to reimagine old effects, says Latimer. "But today technology is giving rise to completely new and original ones." —Susan Karlin

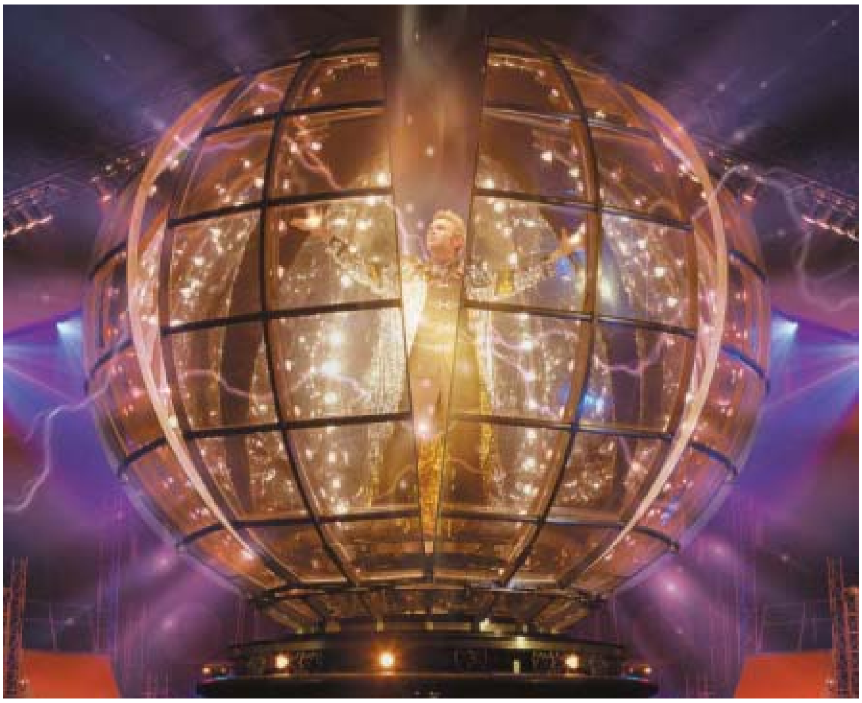

NOW YOU SEE HIM: Magician Franz Harary likes to materialize from inside a robotic sphere.
PHOTO: TAN POW FEI

View this step-by-step how-to audio slideshow now

hands on

EASY RIDER

Convert your bicycle to a human-electric hybrid

WHEN I ride with serious cyclists, they give me endless grief about the plastic milk crate attached to my bike rack. No doubt they're right that it adds considerably to the bike's drag coefficient. Still, it's awfully handy for holding a coat, groceries, or an impulsive garage-sale purchase. So what if the crate detracts from the Lance Armstrong look?

Such purist riders—and the bike shop owners who cater to them—also find other practical modifications anathema. They particularly despise electric motors, which you might use to commute to the office without arriving drenched in sweat, getting your workout on the ride home instead.

PEDAL POWER PLUS: With the help of a battery, a scooter motor, and a second roller chain, the rear wheel of this mountain bike now gets its power from both muscles and electrons.
PHOTO: DAVID SCHNEIDER

Being the practical, crate-toting cyclist I am, I was inspired by a recent encounter with Elise Giddings, an entrepreneur who two years ago helped found Cycle 9 (http://cycle9.com), a bike shop in Carrboro, N.C., devoted to helping people use their bicycles as practical transportation—often by outfitting them with electric-assist motors.

Giddings told me that with a standard hub-motor kit, the conversion might take about an hour. That seemed too easy—and pricey. A good kit (meaning one with decent batteries and a powerful motor) would set me back more than US $1000. So I started scanning the Web in the hope of making the classic DIY trade-off of time for money.

Sure enough, others had ridden this road before me. The design I settled on was shaped by the several generations of electric-bike conversions that Eric Peltzer, a sculptor and inventor in Altadena, Calif., describes at http://www.electricycle.com. It's a mash-up of bicycle, electric scooter, and kart-racing hardware.

My starting point was an old mountain bike. Its frame was beefy enough to stand up to the extra 8 kilograms (17 pounds) or so in hardware, much of it from a 24-volt, 400-watt brushless DC motor. The one I bought, which includes an internal controller, came from Superkids Online and cost a mere $75.

This motor is normally sold for electric-scooter use. In order to retain pedal power, I needed to attach it to the rear wheel in a way that provided adequate speed reduction without interfering with the bike's leg-powered drivetrain. That meant I had to equip the motor with a very small drive sprocket. Searching for one steered me into the world of karting—a sport that involves miniature gasoline-powered race cars that use chains with a smaller pitch than bicycles normally have. I bought a tiny nine-tooth sprocket from Margay Products, of St. Louis, for the motor and a 92-tooth Kevlar-composite sprocket from Precision Karting Technologies, of Wixom, Mich., for the rear wheel. Together they gave my e-bike a reduction ratio of just over 10:1.

Attaching the rear sprocket to the left side of the hub was by far the biggest challenge. I wanted to coast without feeling cogging torque, the drag that a motor's magnets impose as it rotates without power. So I included a freewheel—something almost all bikes have in their rear hubs to make sure that the wheel never turns the pedals.

I started with a bicycle

hub designed to accept a disk-brake rotor on the left side. Using a lathe, I removed much of the metal from the left side of the hub and cut left-hand threads in it to accept an ACS Southpaw freewheel, which is made for bicyclists who prefer to have the chain on the left side. (Believe it or not, for some people this feature is important.) That tactic allowed me to screw a second freewheel on the side opposite the normal freewheel mechanism, which remains in place to serve the standard pedal-powered drivetrain on the right side of the bike. The folks at Cycle 9 sold me a sturdy rim and spokes to go along with the modified hub, and they laced it all together into a wheel. Total cost of the highly modded wheel: about $200.

I got more lathe practice fabricating an aluminum disk to connect the added freewheel with the much larger 92-tooth sprocket, and more still when it came to attaching the nine-tooth drive sprocket to the motor. That involved machining down an adapter I had purchased with the motor to a diameter just 0.001 inch (0.0254 millimeters) bigger than the inner diameter of the drive sprocket. For me, a mere Sunday machinist, that proved a delicate operation: I must have taken dozens of micrometer measurements along the way. I then heated the sprocket with a propane torch, expanding its diameter temporarily, and inserted the adapter, driving it home with a sledgehammer after temperatures equilibrated.

Although lithium-ion battery packs for electric bicycles are readily available, they remain quite expensive. So I went to eBay and found a set of 7.2-volt nickel-metal-hydride modules salvaged from a wrecked Toyota Prius. Although my bike uses only four such modules, for 28.8 volts, I bought extras because the price was right: $20 per module. These modules have a capacity of only 6.5 ampere-hours, but I can keep one pack on the bike and one charging, a welcome luxury.

At this point, the bike lacked only a proper twist-grip throttle control. I had one on order, but I was anxious to test my conversion, so I wired up a plain old 5k potentiometer and lashed it to the handlebars. Big mistake. Accelerating without pedaling was a kick, but when it came time to stop, having to dial down a wobbly potentiometer while simultaneously steering and braking required a dangerously large fraction of my available attention. I did, though, verify that the bike could easily make 20 miles per hour (32 kilometers per hour) on motor power—the legal maximum for electric-assist bikes here in North Carolina. That may seem slow by automotive standards, but for a bike it's really sailing.

The buzz of the motor is audible but not troublingly loud. And when operating under pedal power alone, the bike feels very normal, if a little heavy. But the rapid-fire clicking of the two freewheels reminds me that I'm not riding a standard bicycle and that I can call on electric power whenever I want it. So far I've used the motor only gingerly, limiting the current draw to about 20 A, because I don't want to stress the controller circuitry. That's plenty for the flats and even for modest inclines, so long as the bike is at speed. To launch fast up a hill would require much more current (or lower gearing) than my design provides.

A decent throttle control is expensive, about $50, but a necessity. I also invested about the same amount in a special bag to hold the Prius modules, one that includes a keyed switch so that I can leave the bike chained to a rack without worrying that some prankster will deplete the batteries.

Electrifying a bicycle this way is indeed a labor of love, given that hub motors would accomplish pretty much the same thing in a small fraction of the time, and for not much more money—the bill of materials for my conversion came to roughly $750 when you add in all the bits and pieces required and the associated shipping charges. But I learned a lot, and one day I may revamp the design to attach the motor to the bike's regular drivetrain, which would provide multiple reduction ratios, improving the low-end torque and efficiency.

My only real regret is that I was unable to find a place to attach a milk crate. No matter—my new e-bike could easily haul a whole trailer full of groceries. I would of course look even less like Lance Armstrong, but now at least I can generate as much power on a bike as he can.

—David Schneider

LEFT-HAND DRIVE: An ordinary bicycle hub [top left] was machined to accept a second freewheel mechanism [top right] on the left side. An aluminum adapter connects the added freewheel to a large composite sprocket, which is driven by a 400-watt motor mounted on the frame [bottom left]. Powered by batteries from a Prius, the assembled e-bike takes to the road [bottom right].

PHOTOS: DAVID SCHNEIDER

tools & toys

Instant On

Why can't your computer wake up as quickly as your BlackBerry?

THE COMPUTER inside a BlackBerry or iPhone boots up in seconds, so why can't your laptop?

Some of Microsoft Windows' start-up time is inevitable, says computer architect Peter Glaskowsky, an analyst with the Envisioneering Group, in Seaford, N.Y. "There's too much to do—initialize all of memory, start up the hard disk, assign virtual memory, and more." He says even embedded systems like an iPhone or cable modem, which don't have to look for keyboards and printers, still take 10 to 30 seconds to do these things.

Some of the problem is simply sloppy programming and a philosophy that everything needs to be initialized before anything is available, says Gaurav Banga, chief technical officer of Phoenix Technologies, in Milpitas, Calif., a company that has written BIOS (basic input-output software), going back to the pre-Windows era.

Running a third-party disk defragmentation tool, cleaning out the Windows Registry, and emptying the Recycle Bin and temp files all help speed the boot process, as does updating to the latest service pack. Steve Bass, author of the *TechBite* newsletter, advises savvy users to prune the list of start-up programs, run third-party utilities like TuneXP, or even reload the operating system and start fresh.

Depending on your machine (and your technical confidence), a BIOS update may slice off a few seconds. Banga says Phoenix makes a replacement BIOS that takes only 2 to 4 seconds to load, compared to a typical 10 seconds for the one on a netbook like the Samsung NC10.

None of these fixes, though, get you instantly up and running at an airport gate. Rather than fight with Windows, one increasingly practical option is an "alternative boot environment" (ABE), which gets online much faster, mainly by cranking up only those programs that are needed. ABEs aim to have your browser connected to a known network within 12 seconds.

Current ABEs include Phoenix's HyperSpace and Xandros's Presto, both available as downloads, and DeviceVM's Splashtop, available only when preinstalled by a computer manufacturer. All these are—although they don't tout it—Linux-based environments.

Yet another alternative is to put your computer to sleep rather than shutting it down. But many machines still take too long to wake up and shake hands with the network."The big bottleneck is reinitializing the peripheral devices," says Envisioneering's Glaskowsky. "The device drivers have to wake up the hardware, set up tables in memory, and get ready for normal operation. But we have an example of a mainstream operating system that can come out of sleep in a second or two: Mac OS."

There's hope that Microsoft's next system, Windows 7, will do better. Phoenix's Banga says, "We are working with the kernel team. You might see a Windows 7 that boots in 30 to 40 seconds, instead of the 120 that Vista takes."

Microsoft is also independently trying to reduce boot and resume times for Windows 7. Michael Fortin, who leads a team in the Core Operating System group, recently wrote in a company blog that strategies include initializing drivers in parallel when possible, cutting down on the number of system services being started, and reducing the microprocessor, disk, and memory demands of the services that are absolutely needed.

"To me, the real solution for instantly available computing is not to wake quickly; it's to never sleep," says Glaskowsky, who helped lead Montalvo Systems, a company working on just such an architecture before it sold its assets to Sun Microsystems in April 2008. "You design the system so it's always running the major desktop OS but functioning internally like a cellphone—always on, even if it's a very low operating state."

Or, for checking e-mail and browsing the Web as fast as making a phone call, Glaskowsky says, just use your phone. —DANIEL P. DERN

GREG MABLY

reflections BY ROBERT W. LUCKY

The First Book of Electronics

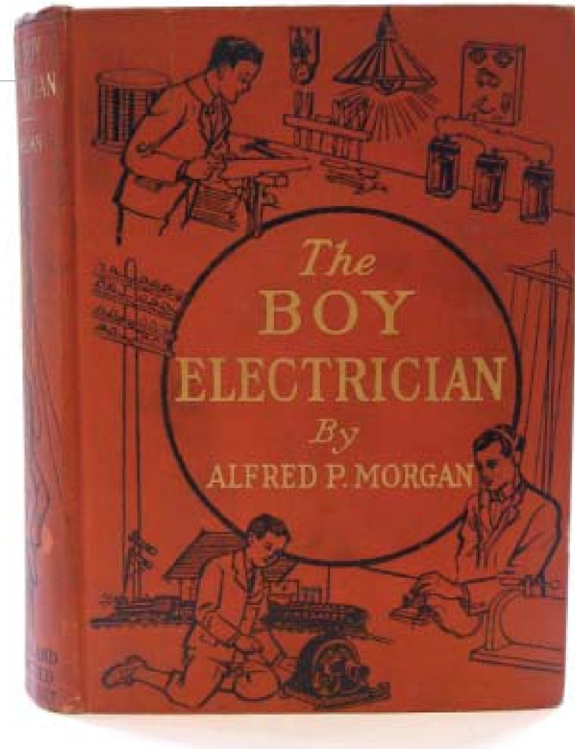

RECENTLY, I was looking for something online, or probably browsing aimlessly, when I happened on a name I hadn't thought of since I was a child: Alfred P. Morgan. Someone had uploaded a digitized version of *The Boy Electrician*. I was instantly swept back more than half a century to my local library. In my mind I saw the familiar metal shelving and the blue-gray binding of my favorite book, also written—and illustrated—by Morgan: *The Boys' First Book of Radio and Electronics*.

I remember how often I went back to that little library, enthralled by Morgan's descriptions in this 1954 book of radios and the like, things that could be built using household parts and a few easily obtained electronic components, like vacuum tubes. Alfred P. Morgan had had a great influence on my childhood, but for all these years, I had forgotten him.

A quick Internet search revealed that Morgan was an electrical engineer with a degree from MIT. He lived in Upper Montclair, N.J., and as a child learned about electricity by visiting with Thomas Edison and Nikola Tesla. He eventually wrote more than 50 books, mostly for children and mostly about electrical things. Perhaps the word *boy* is in many of his titles because the Morgans had three sons. In the 1970s the books were retitled to remove the gender bias.

Many are still available from various publishers, and I found a number of Web sites where people proudly display the devices they have made based on Morgan's descriptions. In fact, you can still buy kits with parts and plans for Morgan's radios. However, I suspect that they are built by older people reliving the past rather than by children dreaming of the future.

Browsing the digitized version of *The Boy Electrician* was like opening a time capsule. How the world has changed since 1913, when Morgan published his hand-drawn illustrations of electronic apparatus! But my reverie soon turned to concern. Are there comparable tomes today that will arouse in children a passion for electronics? Maybe popular children's books today instead have titles like *Boys' and Girls' First Book of Hedge Fund Management*. The world is so very much more complicated today. I worry: Has the wonder that made Morgan's books possible disappeared?

As a child I thought I understood the world of science. I even remember reading and presumably understanding articles in *Scientific American*. Now I confess that I am challenged. Not only is there so much more to try to understand now, but worse, I no longer understand the things I used to think I understood. For example, in Morgan's book there is a drawing of a battery (which you could make) connected to a wire carrying current near a compass needle. Well, the needle deflects, obviously. But now I ponder: Why? How does the needle know that there is a current nearby?

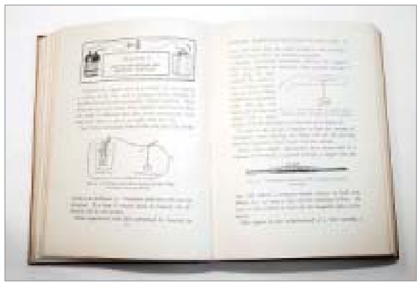

CHILDPROOF: With more than 300 illustrations, *The Boy Electrician* describes projects as complex as a wireless receiving set and as simple as deflecting a compass needle with electric current [shown here].

RANDI SILBERMAN KLETT

Even in college, as I studied electrical engineering, I never worried very deeply about such things. What exactly is an electric or magnetic field? A dangerously shallow knowledge of quantum electrodynamics has clouded the issue, and the more I read the less I understand. I used to think I understood the atomic model—electrons orbiting around neutrons and protons. Now the so-called standard model for particle physics has 12 fundamental particles with the forces between them mediated by another dozen obscure particles, and all of it acting according to an equation that takes multiple pages to print, yet is regarded as incomplete. I wonder if schoolchildren can get excited about something like the Z-gauge boson.

It's not that I want to return to the simplicity of the past. Complexity is a good thing. In complexity lies enormous potential. Imagine how constrained our profession would be if we could use only simple batteries and tubes. But Morgan's book has directions for putting together your own telephone from household parts. Could there be comparable instructions for making your own cellphone? I think not. Much has been gained, but something has also been lost. ❑

TWO OF US: Harmonix founders Alex Rigopulos and Eran Egozy launched a game revolution.

PHOTO-ILLUSTRATION: SEAN McCABE; HARMONIX FOUNDERS: JOSHUA DALSIMER

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

THE STAGE WAS DARK. The curtain, drawn. And the crowd, ready to rock. A hundred lucky fans perched on the edge of their seats waiting for the invitation-only show to begin. Suddenly, the drapes parted as the joyful opening riff of "I Want to Hold Your Hand" filled the room. The spotlight shone on a drum kit bearing the iconic black-and-white logo of the Beatles.

But the performers twisting and shouting on stage were not four young men from Liverpool. Nor were they members of any of the many Beatles tribute bands. The group performing at the Los Angeles Convention Center this past June consisted of six scruffy young geeks. A long-haired coder belted out the verse. A tattooed woman and a cheery guy added the harmonies. A stocky Asian-American played what appeared to be a tinier version of Paul McCartney's familiar Hofner bass. Another guy held a likeness of George Harrison's Gretsch guitar. And a really enthusiastic player smacked the drums.

All were employees of Harmonix Music Systems, a video-game company in Cambridge, Mass., and they weren't actually playing the song; they were demonstrating the most hotly anticipated new game of the year, *The Beatles: Rock Band*, which hits stores worldwide this month.

The title, to be available for the Microsoft Xbox 360, Sony PlayStation 3, and Nintendo Wii, allows gamers to perform along with the Fab Four by singing and playing instruments that work as video-game controllers. The goal is to stay in sync with the music. As a song plays, color-coded dots, representing the musical notes, cascade down the TV screen. The guitarist and bassist must press the corresponding colored buttons on their instruments, and the drummer has to hit the right drum pads. Vocalists must sing on pitch as the lyrics scroll across the top of the screen.

MOP TOPS: Harmonix created CGI tools to fine-tune the gazes and grins of the digital Beatles.

The game, created by Harmonix and published by MTV Games, features 45 career-spanning songs and pixelated Beatles characters performing at locations like the Cavern Club in Liverpool and Shea Stadium in New York City. Paul McCartney and Ringo Starr, as well as Yoko Ono and Olivia Harrison, the widows of John Lennon and George Harrison, helped Harmonix pick the songs and tweak the graphics.

But the game also stands out for all the music technology behind it. To create the separate instrument and vocal tracks the game needs, Harmonix engineers had to use special audio filters to extract every note from the master tapes. They also developed a pitch-evaluating system that can monitor three players singing harmonies—*sun, sun, sun, here it comes*—a key part of the Beatles' music.

Last year, word of the game's impending release generated Beatlemania-like buzz. This is the Beatles' first step into the digital entertainment domain. Apple Corps (not to be confused with the Apple of Macintosh and iPod fame), which manages the rights to the band's original recordings, has long been reluctant about selling songs through online stores. Now the video-game deal might convince Apple Corps that it needs to explore other distribution channels to keep the band relevant.

The new Beatles game is the third and most ambitious title in the *Rock Band* franchise, one of the most successful video games ever made. Since debuting in 2007, *Rock Band* and its sequel, *Rock Band 2*, have earned more than US $1 billion in North America alone. The series is an outgrowth of *Guitar Hero*, the pioneering music game, which has grossed over $2 billion. The genre has become a phenomenon of its own, garnering fans of diverse ages and backgrounds. There are *Rock Band* parties and tournaments, and gamers flock to online forums to share tips and discuss the latest peripherals. It's karaoke for Generation Net.

ALL SCREEN SHOTS: HARMONIX MUSIC SYSTEMS

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

HIT THAT NOTE: Players score points by singing on pitch and hitting the right notes with their instrument-shaped controllers.

"These games are all about tapping into the community," says Stephen Prentice, an Egham, England–based analyst for technology research firm Gartner. "They've made games social."

And they've made real-life stars out of the inventors of *Rock Band* and *Guitar Hero*, Harmonix founders Alex Rigopulos and Eran Egozy. The duo launched the company in 1995 when they were graduate students at MIT. In 2006, after the meteoric success of Harmonix's music games, MTV purchased it for $175 million. Now the company employs 300 people, a mix of musicians and engineers.

"We're trying to let people feel the awesome power of performing on stage," says Egozy, who leads the company's engineering staff.

Rigopulos, who as the CEO spends most of his time on the business side but still helps design the games, adds that Harmonix has been trying to "raise the bar with every game that we release in terms of accessibility and emotional impact." The challenge, he says, has been solving various technical problems along the way. And there have been plenty.

RIGOPULOS AND EGOZY first started experimenting with music and technology as members of the computer-music group at MIT's Media Lab in the early 1990s. Rigopulos, who played in a Balinese gamelan orchestra, and Egozy, a clarinetist, studied hyperinstruments—interactive devices used to explore new forms of musical expression—under the tutelage of MIT scientist and composer Tod Machover. One project outfitted a Levi's jacket with synthesizer keys and speakers. Another transformed drums into game-style controllers that could play an assortment of audio files in real time.

"It was really the work in that group that helped open our minds to the possibility of fundamentally reinventing people's notion of what music was," says Rigopulos, whose 1994 master's thesis, "Growing Music From Seeds: Parametric Generation and Control of Seed-Based Music for Interactive Composition and Performance," explored how gestures might be used to trigger musical events.

Egozy had similar pursuits. In his 1995 master's dissertation, "Deriving Musical Control Features From a Real-Time Timbre Analysis of the Clarinet," he looked at new methods of digital audio analysis, concluding with a prescient remark that such tools would "shape the way we think about and experience the interplay of computer and musician."

When they formed Harmonix, their goal was to develop games that could make people feel like rock stars—in their living rooms. The company's first major project was *Frequency*. The game, released in 2001, introduced a feature that would become a key part of future Harmonix creations: colored "gems" that represented musical notes. The gems rushed across the screen as the song played, and gamers had to hit the corresponding controller buttons. If you missed a gem, its sound would be muted from the song—and you'd lose points.

In 2003, the company released *Karaoke Revolution*. The idea was to create a game that would judge the pitch accuracy of a player's singing. Previous karaoke systems had crude scoring functions, and the MIT grads wanted a technically rigorous solution.

"I went back to my work in academe to see how it could work in games—that was

DECODING THE BEATLES

The Beatles recorded many songs using two- and four-track tape, with multiple instruments often combined on the same track. Harmonix worked with Giles Martin and sound engineers at Abbey Road Studios to isolate instruments onto separate tracks, which could then be converted into game data. Here's the process for "Taxman," from the 1966 album *Revolver*.

STEP 1. "Taxman" was recorded on a four-track tape. Vocals and instruments were kept on different tracks. But the instrument tracks combined multiple sound elements: track one had drums, bass, and guitar; track four had guitar and tambourine. To generate game data, Harmonix needed each instrument on its own track.

STEP 2. Harmonix and Giles Martin's team at Abbey Road used audio engineering tools—including audio forensic software equipped with advanced digital-signal-processing capabilities—to filter sounds with specific characteristics. They applied multiple filters to tracks one and four, splitting the instruments onto separate tracks.

STEP 3. Harmonix converted every note in the instrument and vocal tracks into game data. The data is used to assess the players' singing and performance at the drums, bass, and guitar controllers.

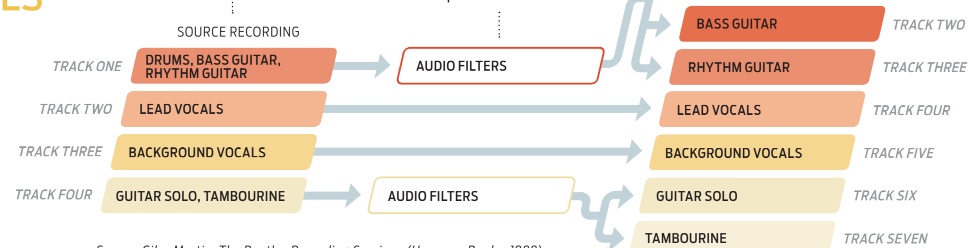

Source: Giles Martin; The Beatles Recording Sessions (Harmony Books, 1988)

pretty cool," says Egozy, who pored over his own work and also a Ph.D. thesis on pitch tracking by an MIT colleague.

His solution was based on a signal-processing algorithm that analyzed the player's voice in real time. The algorithm used a digital filter to discover the frequency of the fundamental note and remove the harmonic overtones. It then compared the filtered signal to a musical instrument digital interface (MIDI) note from the original song. The better the match, the higher the player's score.

In 2005, Rigopulos and Egozy were approached by a company called RedOctane, which had created a small guitar-shaped controller and wanted to build a game around it. The Harmonix founders decided to give it a go. The game was called *Guitar Hero*. The combination of an engaging user interface, innovative controller, and awesome music—including tracks from Jimi Hendrix, the Ramones, and Deep Purple—was a huge success. The game rocked.

A year later, just before Harmonix was acquired by MTV, RedOctane was purchased by Activision, which kept the *Guitar Hero* franchise. Now on its own, Harmonix decided to design new instrument controllers and to develop its most ambitious title yet: *Rock Band*.

"It was important that this thing was a real rock simulation, that it gave you a sense you were a rock star interacting with all your bandmates," says Egozy.

Harmonix spent more than a year perfecting the instruments. The drums posed the biggest challenge. Each of the four drum pads contains a piezoelectric sensor that converts the mechanical vibrations into an electrical signal. The problem was that the four drums sat on the same supporting base, and when the player hit one pad, vibrations propagated and fired the other sensors. Isolating each drum physically would have solved the problem, but it would have made the device unwieldy and costlier to manufacture.

"It might look good on paper, but you have to build it," says Daniel Sussman, Harmonix's director of hardware development, who took long trips to China, where the instruments are produced, to iron out the wrinkles in various designs.

Another challenge was generating the gems—the graphic representations of the musical notes—for each song. The set list included tracks of varied styles, with songs from famed groups like Radiohead and Nirvana, but also from indie artists and even bands formed by Harmonix employees. The company tried to automate the process as much as possible, but to ensure greater accuracy, its sound engineers still arranged the gems by hand.

"Fights sometimes break out among our engineers about where to put the gems," Rigopulos jokes.

Rock Band came out in 2007. Again, the result was a tremendous success, with rave reviews from gamers and critics. But even more significant, it opened the door to a deal with the greatest rock band ever.

THE IDEA of a Beatles-based game first came up in 2007 during a chance conversation between MTV executives and George Harrison's son, Dhani Harrison, who mentioned he was a big fan of *Rock Band*. The executives put him in touch with Harmonix, and a discussion with Apple Corps ensued.

But before things went far, Giles Martin, son of Beatles producer George Martin and the current engineering chief for the Beatles' catalog, noticed a major hurdle standing in the way. *Rock Band* requires that each instrument—drums, bass, guitar, and vocals—be laid down onto a separate track. This works fine for a contemporary song recorded on a 48-track system. But roughly the first third of the Beatles' catalog had been engineered using much simpler recording equipment.

Though the Beatles were innovators in the use of two or four tracks, they often put all the instruments on a single track and then used the remaining tracks for ancillary sounds and effects. Apple Corps wanted the game to cover the span of the band's career, so Harmonix would have to find a way to extract the different instruments from those early songs.

"We were thinking, oh crap, this is a daunting challenge," recalls Rigopulos, "but it's a problem worth solving."

Working with Martin, Harmonix first tried to use conventional audio engineering tools to filter specific sounds from the original master tapes, recorded at Abbey Road Studios in the 1960s. But the method required too many steps, and the results often weren't accurate enough.

"You couldn't just set up a single filter to remove the tambourine," Egozy says.

A more hands-on process was necessary. Egozy and his team turned to audio forensic software normally used by law enforcement and in restoration projects. These tools, equipped with advanced digital-signal-processing capabilities, include more controls than those offered by most audio-filtering packages. Myriad audio parameters can be tweaked in order to zero in on specific sound elements.

Egozy says that although sounds with distinct frequency ranges were easily

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

separated using equalizers, the work became much more difficult as the frequencies were intermingled.

"The forensic tools really helped in cases where we had several more full-frequency instruments on the same track," he says. "For example, guitar on the same track as drums, or vocals on the same track as guitar, or all three on a single track."

Fine-tuning the filters was a painstaking process. Harmonix and Giles Martin spent several months going through each song, separating instruments and vocals onto different tracks and then identifying every bit of sound, note by note.

Still, the work was far from over. A big element of the Beatles' musical success was the elaborate vocals—the harmonies that Lennon, McCartney, and Harrison performed in many songs. There couldn't be a Beatles game without harmonies.

The problem now was that previous *Rock Band* games had featured only one microphone and one singer. Harmonix would have to redesign the user interface to fit three vocal lines and also reengineer the way the game evaluates pitch.

Egozy and his team designed a new system, which they called dynamic part assignment, that allows three different microphones to be used and three different vocal lines to be sung simultaneously. As the song unfolds, the vocals are represented on the screen as dashes next to the lyrics—for instance, the middle line of melody in "Day Tripper" appears between the upper and lower harmonic notes.

At every time step in a particular song, the software analyzes each singer's pitch and makes the best educated guess as to what vocal part a player is attempting so it can score each player's singing. If one player is singing McCartney's C and another is singing Lennon's G, for example, the pitch is assessed and assigned automatically to the two players. The software can keep track of the player's pitches even if they switch lines, as Lennon and McCartney sometimes did.

"Players can sing any part without having to decide who is singing what ahead of time," Egozy says. "The matching algorithm makes the whole puzzle work."

Singing the harmonies requires some practice, and so Harmonix made them optional. If singers want to attempt the harmony lines, they get extra points, but if they all sing only the main line, they're not penalized. And with three instruments and three microphones, a living-room Fab Four can actually be a Fab Six.

ONE DAY during the game's development, Harmonix met even greater perfectionists than those on its own staff: the Beatles themselves. McCartney wanted to know why there was a whammy bar—the lever used to vary the string tension—on his Hofner bass (answer: so it could double as a guitar in the game). Starr wanted the Beatles drop-T logo on the bass drum. But Yoko Ono provided perhaps the most moving request: to give the digital Beatles more soul.

Harmonix took a long look at the video-game versions of the band members and realized that Ono was right. Despite all the computer graphics work, there was something dull about the characters.

The problem started with the eyes. In Harmonix's earlier games, the band members would just stare off blankly into space, sometimes appearing aloof on stage. So this time the coders created a special graphics tool—they aptly called it eye tech—that keeps track of each avatar's gaze. It allows animators to fix a character's sight on one of dozens of points, like the back of the venue or the neck of the guitar or even another Beatle.

And then there were the grins—or the lack thereof.

"These guys were smiling all the time," says creative director Josh Randall. "We had to put more joy into the game."

The coders then created another animation tool to allow artists to adjust the Beatles' grins at different points in a song's performance. The small tweaks made a joyful difference.

The final result is a fantastic journey through Beatles history. The game features famous locations where the band played and also imaginary psychedelic landscapes—Harmonix calls them "dreamscapes"—like a mountaintop with a giant sun as a backdrop for "Here Comes the Sun." The game will also include unreleased studio chatter between the quartet and previously unseen archival photos.

All of which should make *The Beatles: Rock Band* another big hit for Harmonix. Critics note, however, that the company's next task may be the most difficult—staying on top of its game.

"The biggest challenge [for Harmonix] is not letting this become a fad," says Edward Woo, an analyst specializing in digital media at Los Angeles–based Wedbush Morgan Securities. "What they have to do is come up with something innovative that gets people to keep buying the games."

Indeed, if you ask the fans, there's no shortage of ideas. Many are already clamoring for a follow-up game that simulates another greatly admired band, such as Led Zeppelin.

"That's the first band everyone asks us about now that we've done the Beatles," Rigopulos admits, but he declined to give more details.

In the meantime, Harmonix is counting on its exclusive deal with the biggest band in music history in order to maintain its edge. In addition to the initial 45 songs, buyers of the game will soon be able to pay to download additional playable tracks, including the album *Abbey Road* in its entirety.

For the Beatles, the embracing of the video-game generation is also a bright move. It's an opportunity to present the band to a younger crowd that could be the next Fab Four fans. It's a crowd that understands the language of cascading multicolored gems and digital avatars.

As McCartney quipped at the Los Angeles event, "Who'd ever thought we'd end up as androids?" ❑

ENGINEERING MUSIC GAMES

Harmonix is responsible for the most advanced and commercially successful music video games ever made. Here are the main innovations in each title the company has released.

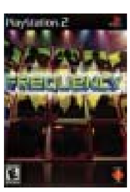

FREQUENCY
2001
Colored gems representing musical notes

AMPLITUDE
2003
Colored-gem lexicon used for music remixing

KARAOKE REVOLUTION
2003
Advanced pitch-tracking system

GUITAR HERO
2005
Plastic guitar used as controller

ROCK BAND
2007
Guitar, bass, drums, and vocals are combined

ROCK BAND 2
2008
Improved drums and guitar controllers

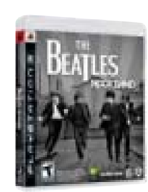

THE BEATLES: ROCK BAND
2009
Pitch tracking of harmonies

FOR YOUR EYE ONLY

A new generation of contact lenses built with very small circuits and LEDs promises bionic eyesight

By Babak A. Parviz

THE HUMAN EYE is a perceptual powerhouse. It can see millions of colors, adjust easily to shifting light conditions, and transmit information to the brain at a rate exceeding that of a high-speed Internet connection.

But why stop there?

In the Terminator movies, Arnold Schwarzenegger's character sees the world with data superimposed on his visual field—virtual captions that enhance the cyborg's scan of a scene. In stories by the science fiction author Vernor Vinge, characters rely on electronic contact lenses, rather than smartphones or brain implants, for seamless access to information that appears right before their eyes.

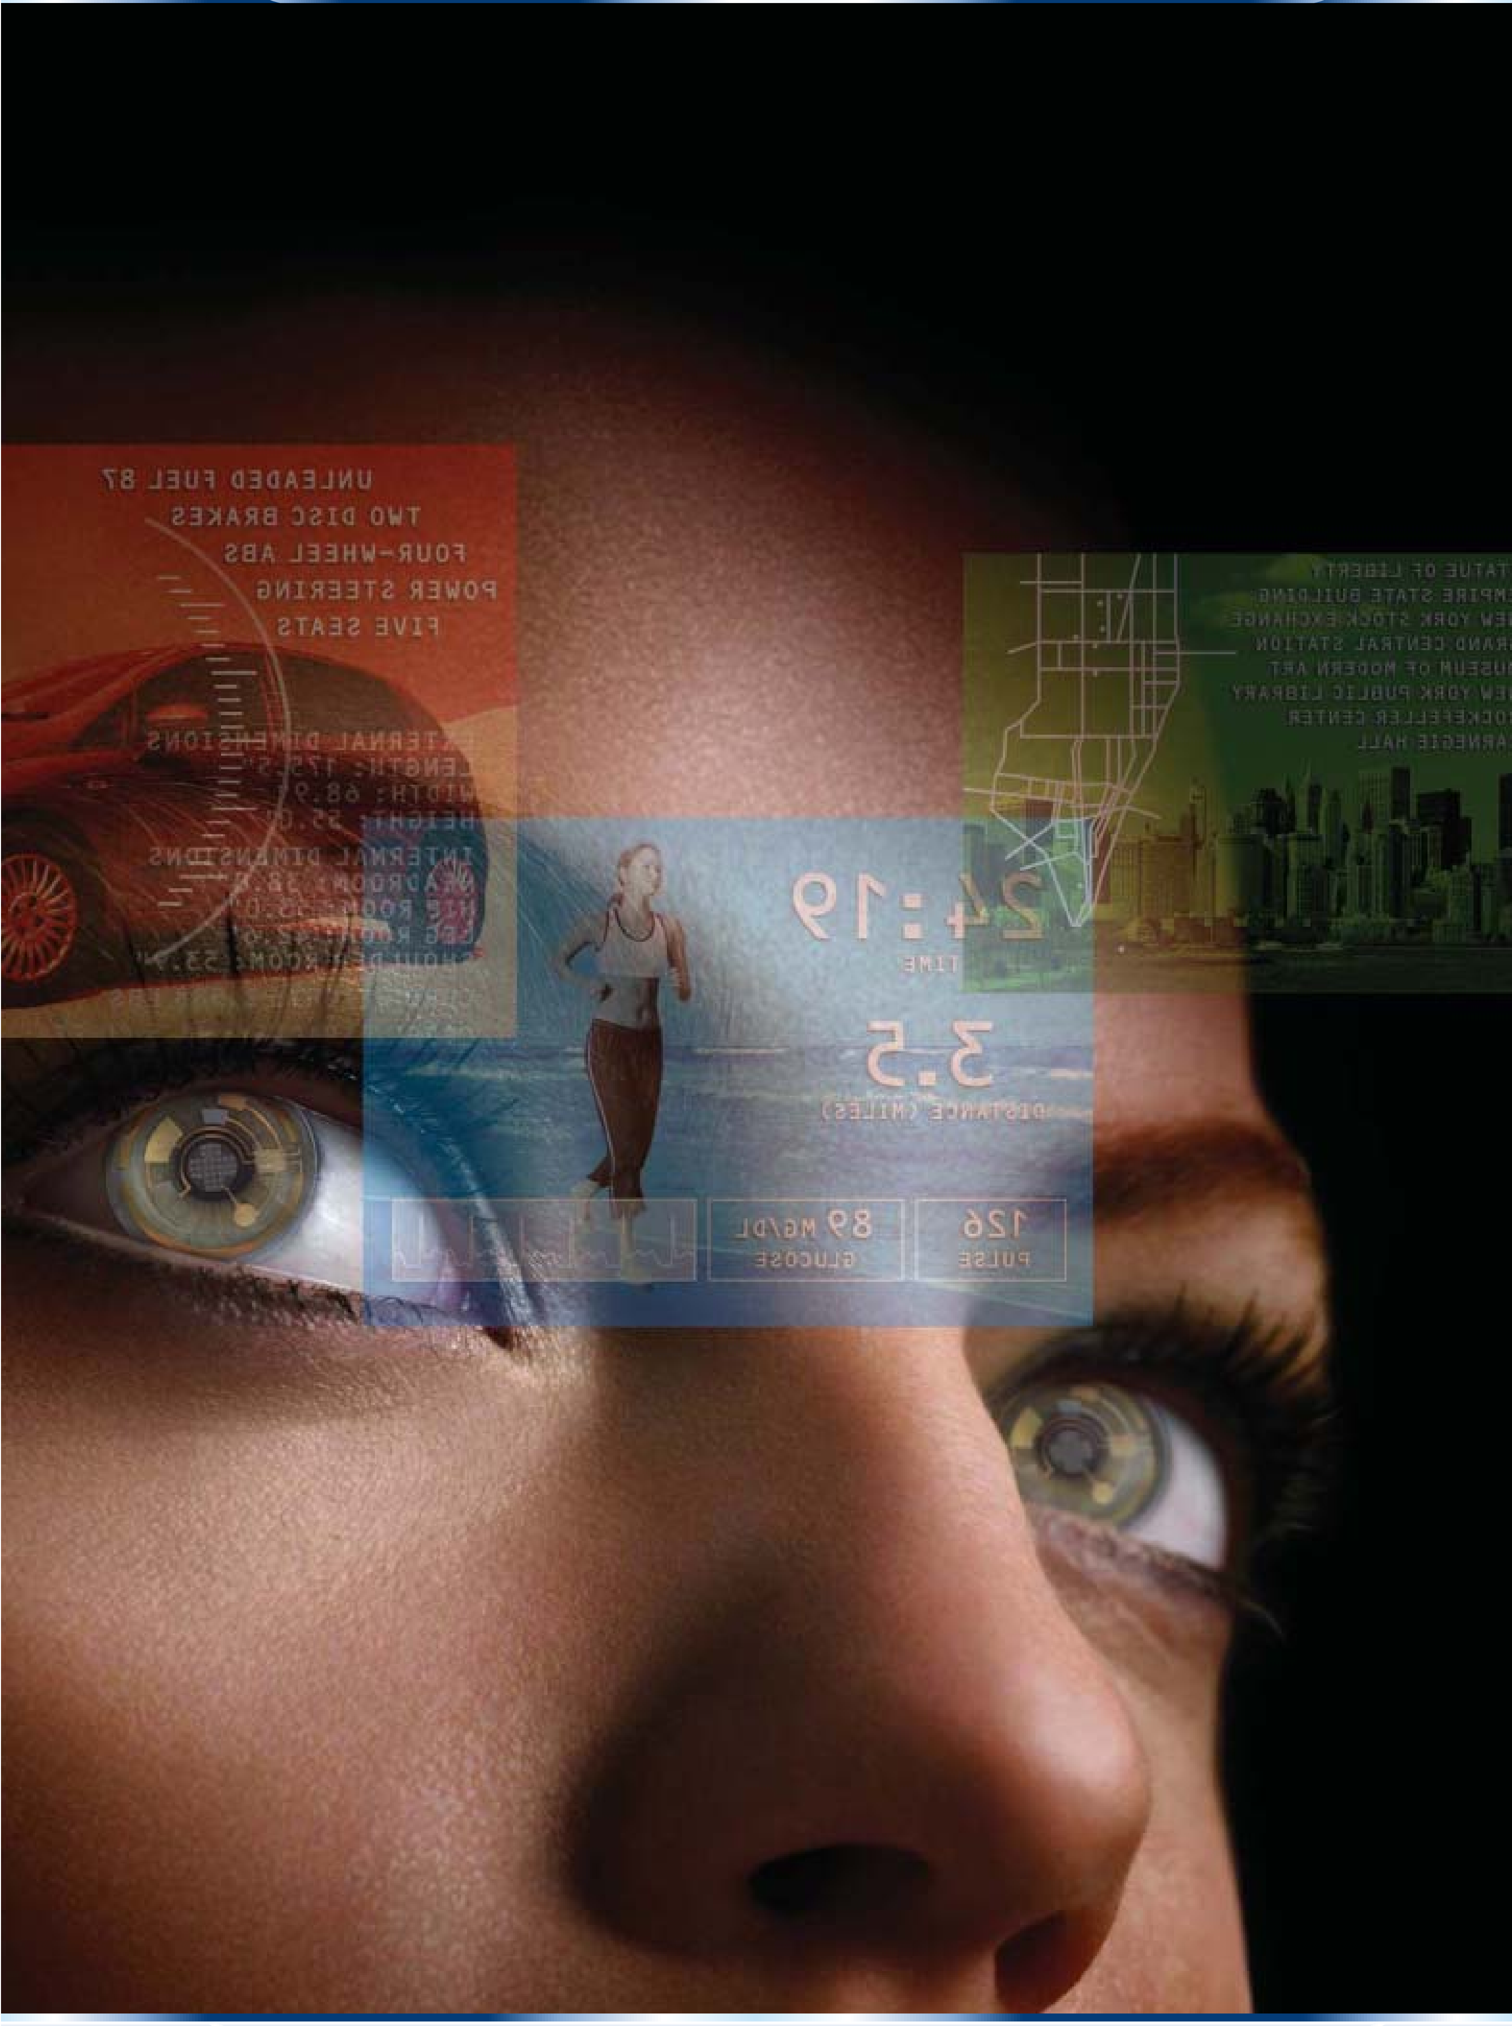
UNLEADED FUEL 87
TWO DISC BRAKES
FOUR-WHEEL ABS
POWER STEERING
FIVE SEATS
24:19
TIME
3.5
DISTANCE (MILES)
89 MG/DL
GLUCOSE
126
PULSE

THESE VISIONS (if I may) might seem far-fetched, but a contact lens with simple built-in electronics is already within reach; in fact, my students and I are already producing such devices in small numbers in my laboratory at the University of Washington, in Seattle. These lenses don't give us the vision of an eagle or the benefit of running subtitles on our surroundings yet. But we have built a lens with one LED, which we've powered wirelessly with RF. What we've done so far barely hints at what will soon be possible with this technology.

Conventional contact lenses are polymers formed in specific shapes to correct faulty vision. To turn such a lens into a functional system, we integrate control circuits, communication circuits, and miniature antennas into the lens using custom-built optoelectronic components. Those components will eventually include hundreds of LEDs, which will form images in front of the eye, such as words, charts, and photographs. Much of the hardware is semitransparent so that wearers can navigate their surroundings without crashing into them or becoming disoriented. In all likelihood, a separate, portable device will relay displayable information to the lens's control circuit, which will operate the optoelectronics in the lens.

These lenses don't need to be very complex to be useful. Even a lens with a single pixel could aid people with impaired hearing or be incorporated as an indicator into computer games. With more colors and resolution, the repertoire could be expanded to include displaying text, translating speech into captions in real time, or offering visual cues from a navigation system. With basic image processing and Internet access, a contact-lens display could unlock whole new worlds of visual information, unfettered by the constraints of a physical display.

Besides visual enhancement, noninvasive monitoring of the wearer's biomarkers and health indicators could be a huge future market. We've built several simple sensors that can detect the concentration of a molecule, such as glucose. Sensors built onto lenses would let diabetic wearers keep tabs on blood-sugar levels without needing to prick a finger. The glucose detectors we're evaluating now are a mere glimmer of what will be possible in the next 5 to 10 years. Contact lenses are worn daily by more than a hundred million people, and they are one of the only disposable, mass-market products that remain in contact, through fluids, with the interior of the body for an extended period of time. When you get a blood test, your doctor is probably measuring many of the same biomarkers that are found in the live cells on the surface of your eye—and in concentrations that correlate closely with the levels in your bloodstream. An appropriately configured contact lens could monitor cholesterol, sodium, and potassium levels, to name a few potential targets. Coupled with a wireless data transmitter, the lens could relay information to medics or nurses instantly, without needles or laboratory chemistry, and with a much lower chance of mix-ups.

Three fundamental challenges stand in the way of building a multipurpose contact lens. First, the processes for making many of the lens's parts and subsystems are incompatible with one another and with the fragile polymer of the lens. To get around this problem, my colleagues and I make all our devices from scratch. To fabricate the components for silicon circuits and LEDs, we use high temperatures and corrosive chemicals, which means we can't manufacture them directly onto a lens. That leads to the second challenge, which is that all the key components of the lens need to be miniaturized and integrated onto about 1.5 square centimeters of a flexible, transparent polymer. We haven't fully solved that problem yet, but we have so far developed our own specialized assembly process, which enables us to integrate several different kinds of components onto a lens. Last but not least, the whole contraption needs to be completely safe for the eye. Take an LED, for example. Most red LEDs are made of aluminum gallium arsenide, which is toxic. So before an LED can go into the eye, it must be enveloped in a biocompatible substance.

So far, besides our glucose monitor, we've been able to batch-fabricate a few other nanoscale biosensors that respond to a target molecule with an electrical signal; we've also made several microscale components, including single-crystal silicon transistors, radio chips, antennas, diffusion resistors, LEDs, and silicon photodetectors. We've constructed all the micrometer-scale metal interconnects necessary to form a circuit on a contact lens. We've also shown that these microcomponents can be integrated through a self-assembly process onto other unconventional substrates, such as thin, flexible transparent plastics or glass. We've fabricated prototype lenses with an LED, a small radio chip, and an antenna, and we've transmitted energy to the lens wirelessly, lighting the LED. To demonstrate that the lenses can be safe, we encapsulated them in a biocompatible polymer and successfully tested them in trials with live rabbits.

SEEING THE LIGHT—LED light—is a reasonable accomplishment. But seeing something useful through the lens is clearly the ultimate goal. Fortunately, the human eye is an extremely sensitive photodetector. At high noon on a cloudless day, lots of light streams through your pupil, and the world appears bright indeed. But the eye doesn't need all that optical power—it can perceive images with only a few microwatts of optical power passing through its lens. An LCD computer screen is similarly wasteful. It sends out a lot of photons, but only a small fraction of them enter your eye and hit the retina to form an image. But when the display is directly over your cornea, every photon generated by the display helps form the image.

PREVIOUS SPREAD: RAYGUN STUDIO

A Twinkle in the Eye

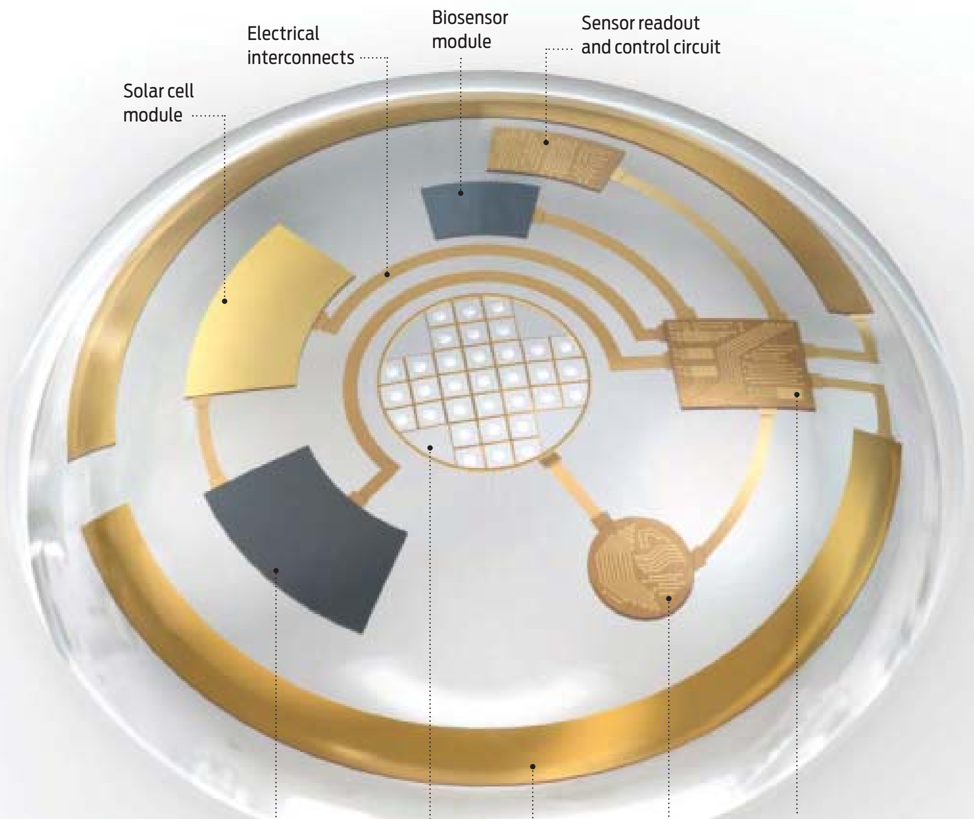

In this lens concept, an antenna at the periphery collects incoming RF energy from a separate portable transmitter. Power-conversion circuitry provides DC power to other parts of the system and sends instructions to the display control circuit. The display, at the center, might consist of LEDs, which would turn on and off, or LCD-like elements, whose transparency would be modulated by the control circuit. An energy-storage module, perhaps a large capacitor, is connected to a solar cell, which could provide a boost to the lens. A biosensor samples the surface of the cornea, performs an analysis, and provides data to the telecommunication module to transmit to an external computer.

ILLUSTRATION: EMILY COOPER

The beauty of this approach is obvious: With the light coming from a lens on your pupil rather than from an external source, you need much less power to form an image. But how to get light from a lens? We've considered two basic approaches. One option is to build into the lens a display based on an array of LED pixels; we call this an active display. An alternative is to use passive pixels that merely modulate incoming light rather than producing their own. Basically, they construct an image by changing their color and transparency in reaction to a light source. (They're similar to LCDs, in which tiny liquid-crystal "shutters" block or transmit white light through a red, green, or blue filter.) For passive pixels on a functional contact lens, the light source would be the environment. The colors wouldn't be as precise as with a white-backlit LCD, but the images could be quite sharp and finely resolved.

We've mainly pursued the active approach and have produced lenses that can accommodate an 8-by-8 array of LEDs. For now, active pixels are easier to attach to lenses. But using passive pixels would significantly reduce the contact's overall power needs—*if* we can figure out how to make the pixels smaller, higher in contrast, and capable of reacting quickly to external signals.

By now you're probably wondering how a person wearing one of our contact lenses would be able to focus on an image generated on the surface of the eye. After all, a normal and healthy eye cannot focus on objects that are fewer than 10 centimeters from the corneal surface. The LEDs by themselves merely produce a fuzzy splotch of color in the wearer's field of vision. Somehow the image must be pushed away from the cornea. One way to do that is to employ an array of even smaller lenses placed on the surface of the contact lens. Arrays of such microlenses have been used in the past to focus lasers and, in photolithography, to draw patterns of light on a photoresist. On a contact lens, each pixel or small group of pixels would be assigned to a microlens placed between the eye and the pixels. Spacing a pixel and a microlens 360 micrometers apart would be enough to push back the virtual image and let the eye focus on it easily. To the wearer, the image would seem to hang in space about half a meter away, depending on the microlens.

Another way to make sharp images is to use a scanning microlaser or an array of microlasers. Laser beams diverge much less than LED light does, so they would produce a sharper image. A kind of actuated mirror would scan the beams from a red, a green, and a blue laser to generate an image. The resolution of the image would be limited primarily by the narrowness of the beams, and the lasers would obviously have to be extremely small, which would be a substantial challenge. However, using lasers would ensure that the image is in focus at all times and eliminate the need for microlenses.

Whether we use LEDs or lasers for our display, the area available for optoelectronics on the surface of the contact is *really* small: roughly 1.2 millimeters in diameter. The display must also be semitransparent, so that wearers can still see their surroundings. Those are tough but not impossible requirements. The LED chips we've built so far are 300 μm in diameter, and the light-emitting zone on each chip is a 60-μm-wide ring with a radius of 112 μm. We're trying to reduce that by an order of magnitude. Our goal is an array of 3600 10-μm-wide pixels spaced 10 μm apart.

One other difficulty in putting a display on the eye is keeping it from moving around relative to the pupil. Normal contact lenses that correct for astigmatism are weighted on the bottom to maintain a specific orientation, give or take a few degrees. I figure the same technique could keep a display from tilting (unless the wearer blinked too often!).

Like all mobile electronics, these lenses must be powered by suitable sources, but among the options, none are particularly attractive. The space constraints are acute. For example, batteries are hard to miniaturize to this extent, require recharging, and raise the specter of, say, lithium ions floating around in the eye after an accident. A better strategy is gathering inertial power from the environment, by converting ambient vibrations into energy or by receiving solar or RF power. Most inertial power scavenging designs have unacceptably low power output, so we have focused on powering our lenses with solar or RF energy.

Let's assume that 1 square centimeter of lens area is dedicated to power generation, and let's say we devote the space to solar cells. Almost 300 microwatts of incoming power would be available indoors, with potentially much more available outdoors. At a conversion efficiency of 10 percent, these figures would translate to 30 μW of available electrical power, if all the subsystems of the contact lens were run indoors.

Collecting RF energy from a source in the user's pocket would improve the numbers slightly. In this setup, the lens area would hold antennas rather than photovoltaic cells. The antennas' output would be limited by the field strengths permitted at various frequencies. In the microwave bands between 1.5 gigahertz and 100 GHz, the exposure level considered safe for humans is 1 milliwatt per square centimeter. For our prototypes, we have fabricated the first generation of antennas that can transmit in the 900-megahertz to 6-GHz range, and we're working on higher-efficiency versions. So from that one square centimeter of lens real estate, we should be able to extract at least 100 μW, depending on the efficiency of the antenna and the conversion circuit.

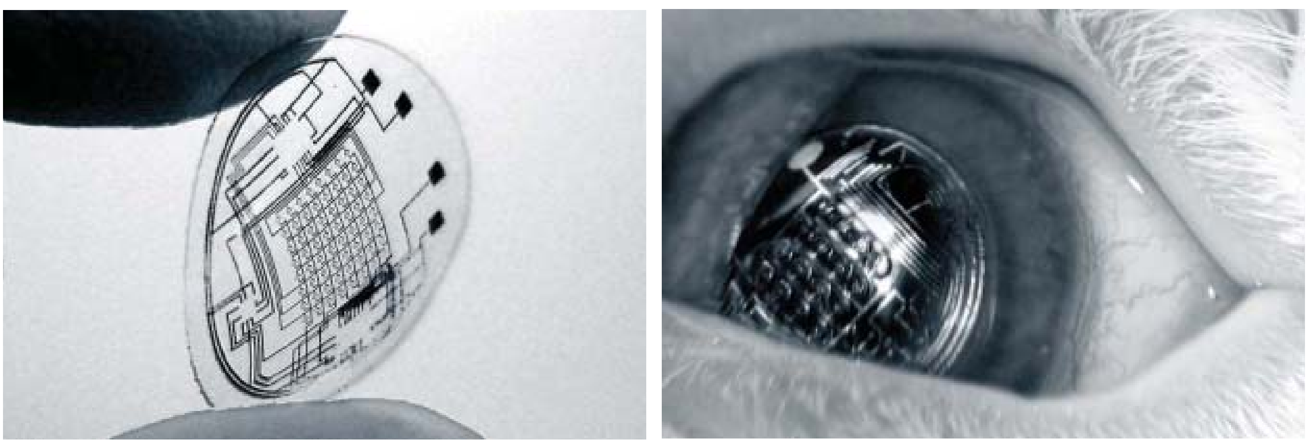

SECOND SIGHT: In recent trials, rabbits wore lenses containing metal circuit structures for 20 minutes at a time with no adverse effects.

PHOTOS: UNIVERSITY OF WASHINGTON

HAVING MADE ALL these subsystems work, the final challenge is making them all fit on the same tiny polymer disc. Recall the pieces that we need to cram onto a lens: metal microstructures to form antennas; compound semiconductors to make optoelectronic devices; advanced complementary metal-oxide-semiconductor silicon circuits for low-power control and RF telecommunication; microelectromechanical system (MEMS) transducers and resonators to tune the frequencies of the RF communication; and surface sensors that are reactive with the biochemical environment.

The semiconductor fabrication processes we'd typically use to make most of these components won't work because they are both thermally and chemically incompatible with the flexible polymer substrate of the contact lens. To get around this problem, we independently fabricate most of the microcomponents on silicon-on-insulator wafers, and we fabricate the

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

LEDs and some of the biosensors on other substrates. Each part has metal interconnects and is etched into a unique shape. The end yield is a collection of powder-fine parts that we then embed in the lens.

We start by preparing the substrate that will hold the microcomponents, a 100-µm-thick slice of polyethylene terephthalate. The substrate has photolithographically defined metal interconnect lines and binding sites. These binding sites are tiny wells, about 10 µm deep, where electrical connections will be made between components and the template. At the bottom of each well is a minuscule pool of a low-melting-point alloy that will later join together two interconnects in what amounts to micrometer-scale soldering.

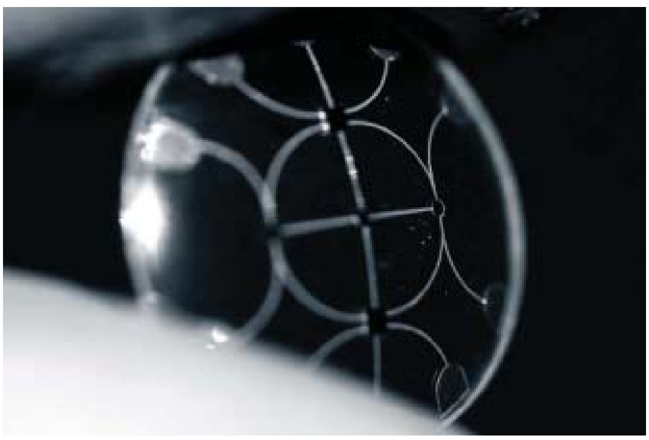

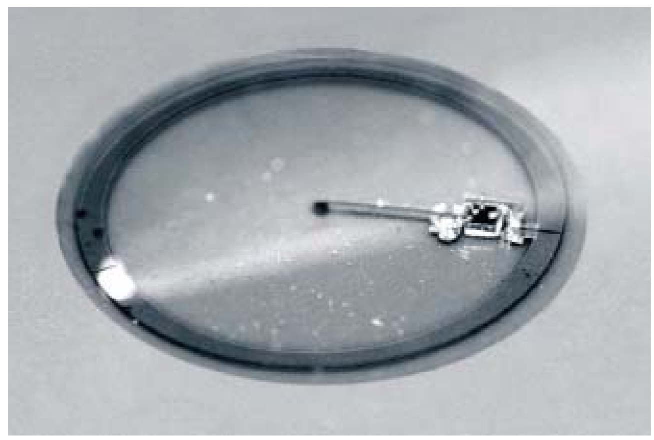

IN FOCUS: One lens prototype [left] has several interconnects, single-crystal silicon components, and compound-semiconductor components embedded within. Another sample lens [right] contains a radio chip, an antenna, and a red LED.

PHOTOS: UNIVERSITY OF WASHINGTON

We then submerge the plastic lens substrate in a liquid medium and flow the collection of microcomponents over it. The binding sites are cut to match the geometries of the individual parts so that a triangular component finds a triangular well, a circular part falls into a circular well, and so on. When a piece falls into its complementary well, a small metal pad on the surface of the component comes in contact with the alloy at the bottom of the well, causing a capillary force that lodges the component in place. After all the parts have found their slots, we drop the temperature to solidify the alloy. This step locks in the mechanical and electrical contact between the components, the interconnects, and the substrate.

The next step is to ensure that all the potentially harmful components that we've just assembled are completely safe and comfortable to wear. The lenses we've been developing resemble existing gas-permeable contacts with small patches of a slightly less breathable material that wraps around the electronic components. We've been encapsulating the functional parts with poly(methyl methacrylate), the polymer used to make earlier generations of contact lenses. Then there's the question of the interaction of heat and light with the eye. Not only must the system's power consumption be very low for the sake of the energy budget, it must also avoid generating enough heat to damage the eye, so the temperature must remain below 45 °C. We have yet to investigate this concern fully, but our preliminary analyses suggest that heat shouldn't be a big problem.

ALL THE BASIC technologies needed to build functional contact lenses are in place. We've tested our first few prototypes on animals, proving that the platform can be safe. What we need to do now is show all the subsystems working together, shrink some of the components even more, and extend the RF power harvesting to higher efficiencies and to distances greater than the few centimeters we have now. We also need to build a companion device that would do all the necessary computing or image processing to truly prove that the system can form images on demand. We're starting with a simple product, a contact lens with a single light source, and we aim to work up to more sophisticated lenses that can superimpose computer-generated high-resolution color graphics on a user's real field of vision.

The true promise of this research is not just the actual system we end up making, whether it's a display, a biosensor, or both. We already see a future in which the humble contact lens becomes a real platform, like the iPhone is today, with lots of developers contributing their ideas and inventions. As far as we're concerned, the possibilities extend as far as the eye can see, and beyond. ❑

The author would like to thank his past and present students and collaborators, especially Brian Otis, Desney Tan, and Tueng Shen, for their contributions to this research.

TO PROBE FURTHER You can find details about the fabrication process using self-assembly in "Self-Assembled Single-Crystal Silicon Circuits on Plastic," by Sean A. Stauth and Babak A. Parviz, in Proceedings of the National Academy of Sciences, *19 September 2006.*

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

HOW FCC RED TAPE SLOWS NEW WIRELESS TECHNOLOGIES— AND WHAT TO DO ABOUT IT

BY MITCHELL LAZARUS

ILLUSTRATIONS BY MICK WIGGINS

RADIO'S REGULATORY ROADBLOCKS

GREAT NEWS! Your team has come up with a new radio technology—one that may have the same impact as Wi-Fi or Bluetooth. Management loves it, funding is in place, patent applications are filed, production is lined up, and marketing is ready to go. This will be huge.

Or maybe not. Your invention could be illegal in the United States. That's an enormous disadvantage in today's global marketplace.

How a new product could be in violation of the law is not hard to understand. Every radio transmitter sold in the United States must comply with technical rules maintained by the Federal Communications Commission (FCC). These rules set limits on power, bandwidth, out-of-band emissions, modulation, and sometimes other properties. They control interference, promote efficient spectrum use, and protect the public from excessive radio-frequency exposure.

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

These rules are needed, and they have teeth. It is a federal violation to import, sell, lease, offer, advertise, ship, or distribute a transmitter (or equipment that includes a transmitter) without first establishing FCC compliance. Nobody goes to jail, but violations can draw large fines. Worse, the FCC will order a noncompliant product off the market. For a start-up built around a single technology, a stop-marketing letter from the FCC can be fatal.

The technical rules that deal with mature products are relatively general. But the FCC tends to regulate newer technologies in much greater detail. The specifics in the rules act like a filter, letting through some kinds of products while blocking or delaying others. Such regulation can, naturally enough, create a barrier to innovation.

The contrast between Bluetooth and Wi-Fi provides a real-life example. Both are unlicensed radio technologies. Their current forms appeared in the late 1990s. They share the same frequency band and are regulated under the same section of FCC rules. The designers of Bluetooth stayed within the technical parameters of the relevant rule, and early Bluetooth products reached the market with no significant holdup from the FCC. But the same rule section delayed some forms of Wi-Fi. The 11-megabit-per-second *b* standard was approved when the FCC staff decided, after a vigorous internal debate, that it complied with the rule as written. Approval of the later, 54-Mb/s *g* standard needed a rule change that took two years. That new rule was more flexible, however, allowing the subsequent, much faster *n* standard to sail through with no fuss at all.

Not surprisingly, the FCC rules are organized around existing technologies. As a consequence, the more innovative a new radio product is, the less likely it is to comply. So, for example, approval of a conventional UHF walkie-talkie takes only a few days. But if the product rests on a novel and creative idea—think of spread spectrum in the 1980s, ultrawideband in the 1990s, TV-band "white space" devices in the 2000s—chances are that it won't reach the U.S. market until the FCC changes or waives the applicable rules.

The need for a rule change or a waiver would be no big deal if it happened quickly. It doesn't. Changes to accommodate new technologies take at least two to three years, and in some cases drag on for four or five years. Amending just one number in a rule can be as slow as adding a whole new category of rules. Even waivers, which are procedurally simpler, need a year or two, sometimes more.

Why does it take so long?

IT WOULD BE EASY just to denounce the FCC for bureaucratic inertia. But the truth is more complicated. You'd be more justified in blaming Congress, the courts—and even the Internet.

A 1946 statute, the Administrative Procedure Act (APA), tells all federal agencies, including the FCC, how to adopt new rules or change old ones. The agency must publish information in advance about the proposed rule and accept comments from the public. This provision lets outsiders explain the problems that a proposed rule might cause and offer suggestions for alternatives. Typically, the agency then issues an order that adopts the rule, possibly altered in light of the received comments. Occasionally, it will announce that it is abandoning the proposal.

For 40 years, this process worked quickly and well. The FCC would issue a concise Notice of Proposed Rulemaking (NPRM), usually just a draft of the rule, with a deadline for comments a month away. To prepare and submit comments, you needed the services of a law firm or the equivalent in-house resources, so only those companies and trade associations that really cared took the trouble. Shortly after the comment deadline, the FCC issued a brief order adopting the new rule, or some variation of it. The whole process typically took six months or less—blindingly fast by federal government standards.

But then, starting in the mid-1980s, several unrelated developments formed a perfect storm of regulatory delay.

COPING WITH THE COMMENTS

In 2008 alone, the FCC fielded more than **23 000 electronically filed comments** on just one technical issue: the use of low-power devices in the TV white-space bands

For one thing, radio technologies began to evolve very quickly, frequently triggering the need to update FCC rules. The year 1985 was a watershed. It brought the first mobile data network (which later served the early BlackBerry devices) and the spread-spectrum rules, the first to authorize high-speed unlicensed digital radio. Ever since, technological changes have arrived at an ever faster pace, which puts a lot of pressure on the FCC to keep the rules current.

A second disruptive element was the emergence of the Web—and in particular the portion of the FCC's Web site that allows the public to file comments electronically. Now anyone with a computer can participate. Unfortunately, a great many do. The result looks like most other examples of Internet discourse: badly spelled, underinformed, and often wildly off-topic postings, with occasional nuggets of solid fact and sound reasoning. A recent rulemaking on unlicensed white-space devices drew 35 000 comments, of which perhaps a few hundred made useful contributions. But somebody at the FCC has to wade through them all.

The ease of submitting comments also encourages frivolous opposition. Companies selling equipment that operates in a particular band sometimes fight a new use of that band on principle, even if their own products would be unaffected. Associations of radio users raise needless alarms and mobilize their members en masse against even very minor threats, just to make the association look important and keep the dues coming in. Competitors of a would-be entrant may oppose a new rule just to slow the arrival of a threatening innovation.

A third major source of delay came from the courts, particularly the U.S. federal appeals court in the District of Columbia. In the course of interpreting the APA procedures over the years, the court has, in effect, added requirements of its own that greatly increase the work of the FCC and its sister agencies.

The court has, for example, held that in an NPRM, providing the text of the rule is not enough. The agency must also give an "accurate picture of the reasoning that [led] to the proposed rule." The court requires, too, that the eventual rule be a "logical outgrowth" of the NPRM, disallowing anything that it would categorize as a "surprise switcheroo."

What's more, the court has stipulated that in adopting a rule the agency must discuss "what major *Continued on page 48*

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

RULEMAKING RULES:
THE U.S. Federal Communications Commission follows a well-established set of procedures when it makes changes to the rules regulating radio technologies. The process is elaborate and can be very time-consuming, as shown by this overview of the many steps involved.

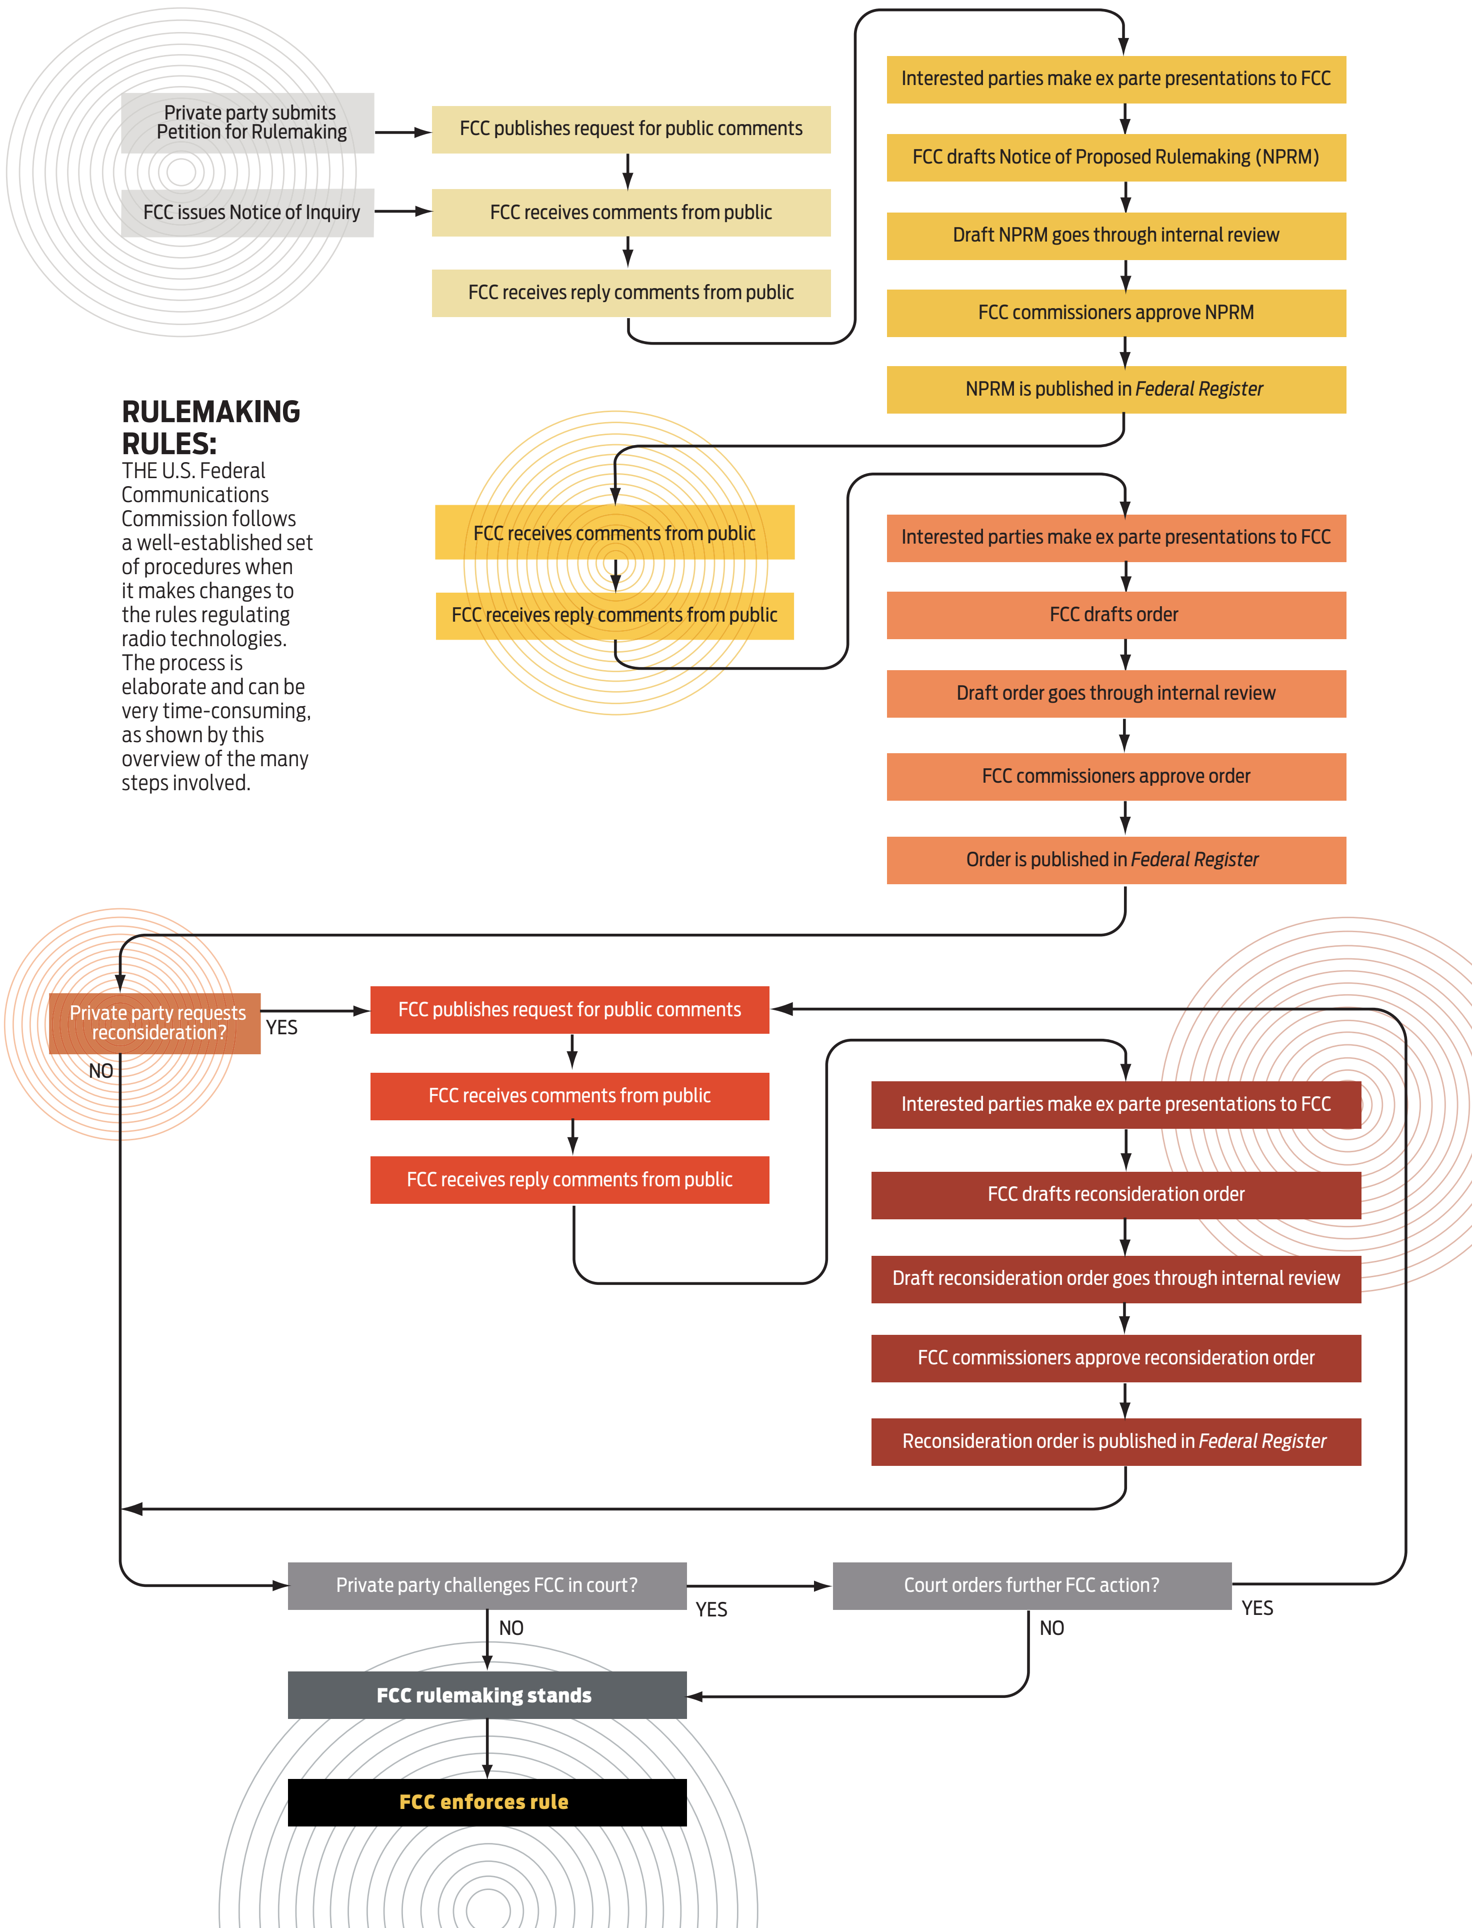

DATA MONSTER

Why graphics processors will transform database processing

BY ANDREA DI BLAS + TIM KALDEWEY

The history of technology is full of breakthroughs in one field that wound up working wonders in a related one. The 300B vacuum tube, introduced by Western Electric in 1937 to amplify telephone signals, found a far more enduring use as a high-fidelity audio amplifier. The atomic clocks first used in the 1960s by the U.S. military to track Sputnik and later to validate Albert Einstein's relativity theories are now the basis of the Global Positioning System. And of course, the magnetron, invented in the 1920s at General Electric and used in radars during World War II, later found itself repurposed as the basis for the microwave oven.

Now add another tech crossover: The graphics coprocessor, invented in the 1970s to churn through voluminous and repetitive calculations and render smooth and realistic-looking images on computer screens, can now chew on large-scale databases.

Database processing is a cornerstone of computing, and it is a market that last year generated approximately US $27 billion, according to technology analysis firm Forrester Research, in Cambridge, Mass. The firm projects that this number—which includes new database licenses, technical support, and consulting—will grow to $32 billion

SUPERBROTHERS

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

by 2013. Every time you bid on an eBay auction, search for a movie on Netflix, look for a Kindle title on Amazon, or do a Google search, massive database applications spring into action, delving into huge quantities of data spread across tens of thousands of machines.

This radical new task for graphics chips evolved from their role as the engine of computer games. So what does sifting enterprise-class databases have in common with rendering virtual monsters in a game? Both require handling huge amounts of data: Realistic-looking virtual monsters require generating millions of pixels every second, while searching large databases involves accessing millions of records per second. So why not take the same hardware that accelerates virtual monsters and put it to work on real-world applications, like the databases that are a large part of our daily lives—more so than pixel monsters?

For the past few years, first at the University of California, Santa Cruz, and now at Oracle, we have been looking for ways to leverage the power of these graphics processors, known as graphics processing units (GPUs). These special-purpose chips are designed to be paired with a central processing unit (CPU) for applications like games and scientific visualization, which demand high graphics performance. GPUs are the progeny of the old video cards, which did nothing but display memory contents on the screen. They ease the computational burden on the CPU by handling the calculations and other simple, highly repetitive operations necessary for rendering the lines, polygons, and surfaces of a full-motion graphics scene. For the price of a low-end computer, high-end graphics cards condense into a single PC card the processing power that just 10 years ago required a supercomputer.

A GPU can deliver hundreds of billions of operations per second—some GPUs more than a teraflop, or a trillion operations per second—while requiring only slightly more electrical power and cooling than a CPU. For the same levels of power and cooling, a GPU can deliver 20 to 30 times as much total computational power. That's a lot less power per calculation.

Take the Nvidia GeForce GTX 285 graphics card, for example. For the price of a quad-core CPU ($300) and 1.5 times the power consumption (200 watts) you get a processor that can rip through 1 trillion mathematical operations per second, about 20 times as fast as a CPU. And because you can plug up to four GPUs into a single server, the GPUs could be used to retrofit existing systems. Four GPUs in place of a single CPU would mean an 80-fold increase in performance for just six times the power consumption. So, with fewer machines, you get less heat and a lower air-conditioning bill.

Captivated by such numbers, researchers have already begun harnessing the power of server-scale GPU computing. In mid-2007, Nvidia unveiled a GPU-based supercomputer called the Tesla platform, which can crunch the numbers that riddle molecular biology, medical diagnostics, and geology. Newer Tesla systems cram up to four GPUs—nearly 1000 cores—into a single 4.45-centimeter-high rack-mount server capable of teraflops speeds. As of June 2009, the Tsubame cluster at the University of Tokyo, made up of 170 Tesla servers, ranked 41st on the biannual list of top 500 supercomputers. It achieved a peak performance of 77 teraflops on the Linpack benchmark, which emulates scientific applications that solve systems of linear equations.

For obvious reasons, we are particularly interested in any advance that promises significant performance gains in database processing. We have developed a parallel search algorithm that exploits the unique architecture of GPUs and has performed very well so far in our experiments.

It's a start, but there's more to be done before the technique is ready for prime time. The sheer performance of these chips, coupled with widespread R&D efforts and the commercial interest in higher performance at lower cost—and that includes power, cooling, and data-center space—suggests that GPUs could soon become a major force outside the graphics arena.

Until the mid-1990s, mainstream video cards were responsible only for displaying the contents of a specific memory region on the screen. Later, the graphics processor took over more and more of the CPU's tasks. But early GPUs were not programmable and could perform only very basic drawing operations.

With the growing demand for more-realistic computer games (the major force driving GPU evolution), the numbers started getting very interesting. A $400 off-the-shelf graphics card could compete with a top-of-the-line custom image-generator board that sold for $200 000 only a few years earlier. With all this power at their fingertips, developers naturally began to wonder: What else could you do with it? What we all really wanted to do was off-load the heavy lifting of compute-intensive tasks from the CPU to the GPU.

However, CPUs and GPUs have radically different architectures, so they require radically different programming approaches. From the programmer's perspective, a GPU is a parallel machine, whereas the CPU is a serial machine. A CPU executes one instruction after another, and each instruction does one thing—for instance, adding the contents of two memory locations or switching the contents of two memory locations. A GPU, on the other hand, can add many pairs of numbers at the same time, possibly all of them together at once. That's because it has many processors that are tightly linked to one another.

Why the different architectures? GPUs were built exclusively to plow through the myriad calculations that specify which pixels get lighted on the screen, and how. In computer graphics, scenes are made up of triangles. Each of those triangles is described by three vectors. (A vector is a set of coordinates that represents a point in three-dimensional space.) Each of those vectors consists of four values: x, y, z, and w (the w is just an algebraic factor to express a point

PARALLEL PLAY

How a GPU deals with information

BEFORE CUDA: To represent the address at right, a GPU had to express each element as a set of four numbers [see above right]. These numbers define the triangles that make up animated objects. Only then could the GPU process the information.

Originally, GPUs were designed to perform a handful of very simple operations, such as coloring pixels and computing lighting effects, while CPUs still had to do most of the drawing work. Nowadays, though, it's the GPU that does most of the work.

The three-dimensional object at right, a vehicle of the sort found in many video games, begins as a wireframe model. The wireframe is composed of many triangles. Starting from only a very simple description, the GPU draws the triangles, then colors and converts them into pixels for display—over and over and over again.

Like a CPU, a GPU also executes a sequence of instructions. But each instruction is carried out by more than 100 internal processors—called processing elements—simultaneously, each of them operating on a different piece of data.

Let's say a processor is presented with a group of 64 vectors (such as one of the three sets of numbers at right) that it must add to a second set of 64 vectors. A CPU would add these in serial order, adding the first value in the first *x-y-z-w* group to the first value in the second group, adding the second value in the first group to the second value in the second group, and so on. A GPU can add all 64 vectors to the other 64 vectors all at once. That's because GPUs can churn through each string of numbers as though they were single numbers. In other words, a task that takes a CPU up to 64 x 4 = 256 instructions will require a single GPU instruction, because vectors can be added with a single instruction. And with GPUs having hundreds of processing elements at their disposal, the 64-vector additions can be done in parallel.

using so-called homogeneous coordinates, which are much more convenient for geometrical transformations). So one triangle, with its three vertices, will be represented by $\{(x_0, y_0, z_0, w_0), (x_1, y_1, z_1, w_1), \text{and } (x_2, y_2, z_2, w_2)\}$. Add to that the fact that each pixel of the triangle must be colored. The color of a pixel is also described by a vector of four values: red, green, blue, and the percentage of transparency. Now consider that a typical graphics scene has millions of triangles and that a piece of animation has at least 24 scenes per second. Before the ascendance of GPUs, to draw a triangle in 3-D space, the CPU would have to compute exactly what the triangle looked like, where the three vertices were with respect to the viewer, what pixels were visible, what part of the background was visible, the colors of all the pixels, and so on. It took too long.

Although a GPU has a limited vocabulary, its level of parallelism is tremendous. It can handle tens, hundreds, even thousands of these triangles at once without breaking stride.

Early on, researchers had the bright idea of taking all that raw power and putting it to work on applications other than graphics. In 2002, Mark Harris, now a computer graphics researcher with Nvidia, coined the term GPGPU, for "general-purpose computation on graphics-processing units."

So what's the catch? Well, let's start with the obvious one. GPUs are designed to do one thing extremely well—and only one: draw triangles on computer screens. After all, they were primarily designed to cater to the video-game market. Until recently, GPU programming was considered one of the "black arts" of computer science: It was esoteric and difficult, the realm of just a few researchers and hackers. It required a two-step mapping process. First, a programmer had to take the problem's data and map it to graphics data that the GPU could handle. Second, he had to figure out a way to express the operations required to solve the problem in terms the GPU could understand. That is, he had to essentially lie to the GPU, representing the computation as a triangle-drawing operation, the only thing a GPU could understand and carry out [see illustration, "Parallel Play"].

To do anything with those vertices, you had to use a combination of programming tools specifically built for graphics, such as OpenGL, Microsoft DirectX, Cg, or the jerry-built assembly languages of individual GPUs.

The exercise might have been fascinating for us geeks (and make no mistake, it was). But it was also time-consuming—far from commercialization and far from being more than just a riveting diversion. Then came CUDA. CUDA changed everything.

A visionary named John Nickolls, at Nvidia, recognized the huge potential of GPGPU. So in 2007, the Compute Unified Device Architecture was unleashed

SUPERBROTHERS

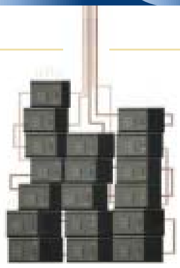

P-ARY SEARCH

Taking advantage of parallelism

To avoid having to look at every entry for every query, a database contains many sorted lists, called indexes. Each processor is assigned an equal part of an index and looks up the first entry in its assigned section. Comparing results—looking for the alphabetically highest part with a first entry alphabetically lower than the search key—is not much of a problem for either a small number of people or processors. But what if there are more than 100 processors?

The communication necessary to coordinate the task increases exponentially with the number of processors. Consider the cat herding you would have to do to manage 100 people—your job would be a lot tougher than theirs.

A better strategy is to let each processor do two comparisons—the first and last entry of its section. Because only one of them will find a result, no further coordination is required. For the GPU, this strategy plays out well. Each of its processing elements (PEs) looks up the first and last entry of its section at the same time, and then waits while one processor either narrows the search scope or returns the result.

BINARY SEARCH

Searching for "Jane Doe" using a traditional binary search takes 10 steps.

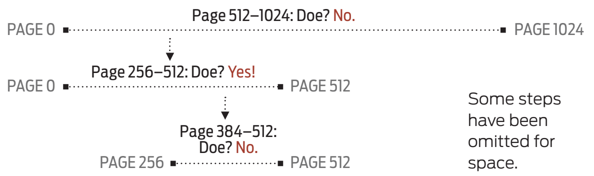

***P*-ARY SEARCH**

Searching for "Jane Doe" with four processing elements now takes 5 steps.

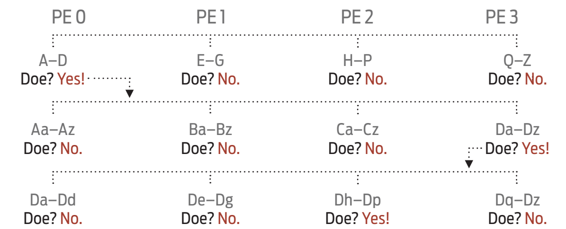

on the world. It divided the history of general-purpose processing on GPUs into two epochs: BC (before CUDA) and AC (after CUDA).

With CUDA, the GPU is much easier to program for general-purpose applications. Now it can understand instructions beyond drawing and digest data beyond coordinates and colors. For example, a programmer can now store the name "Jane Doe" as a simple sequence of characters {J,a,n,e,...} instead of as floating-point numbers.

Nvidia isn't the only company with a programming environment aimed at making GPGPU more accessible. Another major GPU manufacturer, ATI Technologies (now AMD Graphics Products Group), introduced a general-purpose GPU computing technology called Stream SDK. Both Nvidia and AMD Graphics, together with IBM, Apple, Intel, and others, are now putting their bets on OpenCL, a parallel programming language that will let developers write programs for CPUs, GPUs, and other processors in a uniform way. OpenCL is the logical evolution of CUDA.

To understand why anyone would want to put GPUs to work on databases, you've got to first grasp the demands of modern database applications. Searching large-scale databases like Google or eBay usually means finding a small group of specific entries from among tens of millions of possibilities. A processor has to do this in milliseconds at most. And it has to perform this feat not for just one user but for hundreds or thousands of them, all searching the same database for different things at the same time.

Let's say you're looking on eBay for a replacement sound reproducer for a 1906 Edison windup phonograph. The sound reproducer converts the grooves on those old wax cylinders into sound. The specific part you're looking for is a 4-minute reproducer specifically suited to the hard plastic celluloid cylinders that were manufactured later in the century.

You type your query into eBay's search bar. That begins a chain of events. First, the form—the translator between the arcane code behind an application and what you see on your screen—collects your search keywords and sends them to the server.

The server application then checks those terms—"Edison sound reproducer"—against the approximately 20 million listings in eBay's currently open auctions to match some or all of your search terms to entries in the database. Within this cacophony, it finds 18 results that match.

Next, eBay's software displays these 18 entries on your screen. For each entry, it needs to load the item name, the associated pictures, the number of bids, and how many days the auction has left. And it has to cobble together all this information in under 2 seconds from multiple sources. Once all these different elements have been assembled as a Web page, the program loads them all into your browser.

But those 2 seconds must include not only the time it takes to search the database but also the time to retrieve and sort the results, load the pictures, and account for network latency, the network bandwidth, the connection speed, and the user's possibly slow PC. That means the time it took to access the wellspring of eBay data is only a minuscule fraction of the 2 seconds you waited to find out how many sound reproducers were on offer.

In fact, most of the time you spend waiting for a request is taken up by the transfer of data from the server application across various networks to your PC. Add up all those factors and the "reasonable response time" allowed to actually search eBay's databases ends up being on the order of a few milliseconds—at most.

Consider also that at any given moment, Google, Amazon, and eBay have to service huge volumes of requests. While you're looking for your sound reproducer, 200 people are looking for a purple sweater, 3000 are looking for deals on diamond rings, 6000 are combing through used textbooks, and so on. To give you some idea of the scope of that undertaking: With every query, eBay's software must comb through up to 24 million listings, Amazon probably sells hundreds of thousands of items per day, and Google's software must sift through the world's estimated 182 million registered domain names.

Disks—such as those used by these firms' servers—are too slow to deliver millisecond responses for more than a few requests at a time. Instead of waiting

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

for hundreds or thousands of slow disks, wouldn't it be better to just keep the entire database in the much faster main memory? That's exactly what enterprise-class databases try to do: They chop up and distribute data across multiple servers specifically so they can fit the most frequently accessed pieces into memory. This guarantees subsecond response times even in heavy traffic, such as on Black Friday, when an online retailer might sell millions of items in one day.

But that scheme carries a substantial cost: Large server farms require large amounts of energy, cooling, and space. As the volume of online information and the user base keep skyrocketing—and as a result, the servers' workloads—companies have no choice but to keep increasing the number and performance of their servers. GPUs should be able to get the same job done while using a lot less electricity.

The enterprise software field has another reason to be chasing GPUs now. Enterprise data is all the information companies need to run their businesses. For eBay, that's every listing, user profile, feedback, and comment for its nearly 100 million users worldwide. Enterprise data have been growing at a slower rate than the number of transistors on microchips (see Moore's Law), so computer memory is growing faster than the amount of enterprise data.

The implications of that fact are tremendous: It is now possible to handle huge databases in main memory, which means that the data can be pulled in microseconds, as opposed to the several milliseconds needed for disk access. The actual time required for disk access is somewhere between 4 and 10 milliseconds, which might seem good enough, but a million accesses would add up to several hours. One application that can potentially benefit from GPUs is geospatial databases, whose data sets are similar to graphics data. These were actually the first to adapt GPUs for complex operations that require combining results from multiple data sets. The bottom line: When database performance is no longer limited by the response times of sluggish disk drives, a processor like a GPU can really shine.

SURPRISINGLY, SEARCH applications have not yet leveraged the massive parallelism of GPUs. We have developed a parallel search algorithm that optimally exploits the GPU by applying a "divide and conquer" strategy to speed up individual searches. We call our baby *p*-ary search, where *p* is the number of processors. To understand how it works, first consider the way you look up numbers in a phone book. Say you're looking for the phone number of somebody named Godot. Instead of flipping through every page, you could open the book in the middle to see whether the first entry on the current page comes alphabetically before or after *Godot*. Based on the result, you choose the first half or the second half of the book and continue this strategy. With every iteration you cut the search range in half, until you're down to a single page. With this strategy, finding the page containing *Godot* in an imaginary telephone book with 1024 pages will require 10 iterations. In computing terms, this algorithm is called binary search, which is possible, of course, because the phone book data you are searching is already sorted.

You could speed up the process by conscripting three friends, tearing the phone book apart into four equal pieces, and assigning each person to search his or her own piece. If everyone is using binary search, this would cut the search time from 10 steps down to 8. Of course, only one of you will find *Godot*, while the other three waste their time searching in vain.

If you have a spare phone book, there is a faster alternative. Each person looks at the first name of his section. The person whose section's first entry (for example, *Beckett*) alphabetically precedes *Godot* holds the relevant part. The other three sections are discarded. The relevant section is again torn into four pieces and handed out. This process continues until each person ends up with a single page and someone finds *Godot*.

This algorithm converges significantly faster on the solution than the previous one. Each iteration takes up the same amount of time—one lookup, but performed by four people in parallel. With each iteration the search range is reduced to one-fourth instead of one-half. Now searching the 1024-page phone book requires only five steps: 1024 to 256 to 64 to 16 to 4 to 1 [see diagram, "*P*-ary Search"].

The astute reader will notice that the number of steps was reduced by only 50 percent, even with plenty of help. While this appears to be a small gain—and at considerably high cost—it becomes a lifesaver when you have to search tens of millions of entries in a fraction of a second. If the resources are available—and today's high-end GPUs have more than 200 cores—you can throw even more man power at *p*-ary search to improve responsiveness. Using 32 people (or processors), for example, you can search 1024 entries in two steps.

With *p*-ary search a GPU can outperform a similarly priced CPU by up to 200 percent for very large workloads: Thousands of searches over data sets with millions of entries are not an unusual scenario for large Web applications. Only then does the GPU reach its peak performance, pumping out more than 6 million search results per second. To do a single search of a 1024-page phone book, there is no reason to bother with a GPU.

That's because the GPU resides on a PC card. Any data, including search queries, need to be copied via the relatively slow PC-card bus before the GPU can get its hands on them. Applications with high computational intensity, such as physical simulations, do not suffer from these limitations. However, most database applications have a relatively low arithmetic intensity but require handling large amounts of data.

As graphics-card memory sizes increase at a rapid rate—up to 4 gigabytes on the newest models—using the video memory as a data cache reduces traffic on the PC-card bus, thus improving performance. Newer generations of GPUs will also enable asynchronous data transfers to and from main memory, which does not reduce memory traffic but allows the GPU to return results as they become available, thus improving response time.

MORE-POWERFUL databases and more-realistic video games aside, GPUs per se do not enable anything radically different from what can already be done with today's CPUs. However, they may very well be the key to an epochal change. GPUs are democratizing supercomputing the way the PC democratized computing, making an enormous amount of computational power—previously the exclusive domain of government agencies, research institutes, and large companies—available to the masses.

For example, it is possible that somebody out there has the right algorithm to design a human-equivalent machine intelligence but could not afford enough computational power to make it work—until now. After all, what we humans call "experience" in artificial intelligence is a database. ❑

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

Radio's Regulatory Roadblocks

Continued from page 41

issues of policy were ventilated...and why the agency reacted to them as it did." And the adopting order must "respond in a reasoned manner to the comments received,... explain how the agency resolved any significant problems raised,...[and] show how that resolution led...to the ultimate rule."

These requirements and others like them add length and complexity to both the NPRM and the order adopting a rule. A recent order authorizing white-space devices took up 130 pages of single-spaced text, much of it addressing these court-imposed requirements. Even the NPRM (issued 52 months earlier!) ran 38 pages. The time needed to draft, edit, and review these documents is a major source of rulemaking delays.

A fourth change comes from the enormous commercial success of radio-based devices, which amounts to many tens of billions of dollars annually. This kind of money raises the stakes for success or failure in the regulatory process. A rule change allowing a company's technology or (almost as good) denying that of a competitor translates quickly into revenues. This gives plenty of incentive to spend time and money at the FCC lobbying for a favorable result or challenging an unwanted outcome afterward.

There is one last reason for the slow pace of regulatory change. With telecommunications being key to the global information economy, the FCC has been thrust to the forefront of national policymaking. This transformation has a downside. When the FCC staff is deep into matters of national political importance—the transition to digital television, say, or a US $20 billion spectrum auction—issues like a mere technical rule change tend to get pushed aside.

Before suggesting ways to accelerate the process, I should give a few more details on how the system works. A rulemaking gets started in one of two ways. The FCC might issue a Notice of Inquiry, which seeks general views on how best to regulate a new technology. Or a private party might submit a Petition for Rulemaking, which asks for a change or a new rule. Either way, the FCC sets a deadline for public comment, usually a month or two away, and a date for a second round of input, called reply comments, a month or so later. Even after those dates, interested persons can still make written or oral presentations to the FCC staff, called ex parte submissions (Latin for "one sided"). The staff studies the comments, reply comments, and ex parte material. In most cases the result is an NPRM that explains the background and proposes specific rules. After a staff member drafts the NPRM, it goes through a complex internal review process and eventually lands on the desks of the five FCC commissioners. Once they have signed off, the NPRM is released to the public. All of that takes about a year.

The NPRM sets dates for still more comments and reply comments, after which interested parties can again make ex parte presentations. A given rulemaking might spark scores of trips to the FCC and hundreds of written filings. This in itself is a major cause of delay. Every hour that FCC staffers spend meeting with outsiders is an hour not spent generating the paperwork to adopt the rules. And every ex parte filing (some run to hundreds of pages) must be read, analyzed, and evaluated. Worse, each ex parte submission fires up the opponents to chime in with their views. The staff sometimes uses these in-person meetings to probe each side's needs and test possible compromises. But the process still drags out the decision making and bulks up the resulting order.

Once a draft order is ready, it too, like the NPRM, must go through multiple levels of internal review. The FCC's engineers, for example, scrutinize it to ensure that the rules will work in practice. And the lawyers want to be sure that it can withstand an appeal. Eventually, the five commissioners sign off, and the new rules appear in the *Federal Register*. From NPRM to publication of the new rules usually takes at least one, sometimes two, and occasionally three or more years.

But it's not over yet, because any party has the right to ask the FCC for "reconsideration." The white-space ruling of 2008, for example, attracted 19 of these requests. That sets off yet another round of comments and reply comments, another series of ex parte presentations, and another carefully drafted and extensively reviewed order a year

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

or two later. The rules usually stay in effect in the meantime. Reconsideration rarely changes much, but it does divert staff from other work. It also postpones certainty in the final outcome, which can make investors nervous.

Finally, both the original order and the reconsideration order are subject to challenge in court. Such cases used to be uncommon because they are expensive to bring and almost never invalidate the contested rule. At worst, in a technical proceeding, the court might order the FCC to take a second look, while leaving the rule in place. But lately, more companies and organizations have been rolling the dice with court appeals, despite the low likelihood of success.

In 2004, for example, the FCC allowed power companies to provide broadband service over their electric wires. Amateur radio operators, afraid that such operations would interfere with their own communications, opposed the proceeding at every stage. After the FCC denied their request for reconsideration, the amateurs filed an appeal in court. The resulting 2008 court ruling asked the FCC for certain clarifications, but it let the power companies continue to deploy and use the technology in the meantime. These cases remind the FCC staff that they must write every document to anticipate and address the objections an opponent might raise in court. That, of course, takes a lot of time.

A different and somewhat faster approach than changing the rules may work when a product fails to comply but still satisfies the purpose of the FCC's regulations. Perhaps its radiated power is above the maximum allowed, but the intended application is deep inside mines, where it is unlikely to cause interference. Or the new radio equipment might use a band that is ordinarily off limits, but with a modulation that makes it invisible to other users. Such cases are good candidates for a waiver—a ruling that allows the device to be sold despite its noncompliance.

The FCC usually makes waiver requests public, invites comments and reply comments, and carries out the ex parte process. Waivers of technical rules typically take one or two years. Like rulemakings, they are subject to reconsideration and court appeals. Unlike a rulemaking, the waiver applies only to the company that requested it. But the FCC routinely grants identical "me too" waivers to others who ask.

With the needed rule or waiver in place, there is still one step to go, sometimes a big one. Every mobile, portable, and unlicensed transmitter sold in the United States must be certified by the FCC or an authorized company acting on its behalf. With a mature product line, this requires just a few days or weeks. But the first few certifications of a new technology take longer, sometimes three to six months.

HOW MIGHT THE PROCESS BE sped up? Congress could trim back the court-imposed requirements. The FCC could help, too, by making its rules as general as possible, thus reducing the need for frequent changes. But even as things stand, if your company is promoting a new radio technology, you can improve its chances for approval—and get that approval faster—by keeping in mind a few broad guidelines.

First, plan to share the airwaves. Vacant spectrum is scarce, and it is generally auctioned at very high prices. Most new technologies must cohabit with old ones. So choose a part of the spectrum that avoids sensitive or important applications like GPS, radio astronomy, search and rescue, aeronautical, and so forth. And unless requesting a waiver, pick a band already allocated for the same general purpose as your application: fixed, mobile, radar, satellite, whatever. Changing the FCC's allocation for a particular frequency band takes many years. Note also that parts of the spectrum that are shared between the FCC and federal users are available but subject to extra delays. Bands reserved for federal users are barred to rulemakings and inhospitable to waivers.

At the same time, tailor your system to minimize trouble for existing receivers. In a narrowband environment, for example, a system using very short pulses is less likely to affect those receivers adversely than would a system that uses continuous modulation. Keep power levels, bandwidths, out-of-band emissions, and duty cycles as low as is practical. In the end, you should be able to show that your equipment does not significantly increase the risk of harmful interference to others. If you cannot, your system is unlikely to be approved.

DAWDLING FOR A DECADE

The FCC proceeding on ultrawideband rules, which began 11 years ago, is now in its **third round of reconsideration petitions** and is **not over yet**

Expect vigorous opposition in auctioned, amateur, and satellite-downlink bands, even if your application poses little or no realistic threat of interference. If the parties already using the slice of spectrum you need make up a small group, consider negotiating limitations on deployment or operation of your device to obtain their consent. Their approval greatly improves your chances with the FCC.

Think through whether a grant of the requested rule or waiver might authorize other systems that cause more interference than yours. This will be an important internal question at the FCC. Also, show how the public will benefit from your system. A good way to do this is to accumulate written support from potential users. But tell the truth, and deal with any downsides openly. (Remember, lying to the FCC is a federal crime—up to five years in an orange jumpsuit.)

As with so many aspects of life, it's best to keep things simple. Minimize the number of FCC rules to be amended or waived. Doing so may require some extra care in setting up the request, but it is well worth it.

Above all, plan ahead. Start working through FCC issues early in the design process. Build a compliant device if possible. If not, minimize the required departure from the rules, and go to the FCC as soon as the transmitter specifications are firm—but not before. Changing a request in midcourse not only restarts the whole rigmarole from the beginning, it also incurs ill will at the FCC.

The process eventually does work, most of the time. True, it might seem unfair for your company to have to foot the bills and endure the frustrations when competitors can then exploit the outcome for free. But someone else did just that for each of the new radio-based technologies introduced over the past three decades. Consider it the price of progress. ❑

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

香港城市大學
City University of Hong Kong

大隱隱於市 學研研出塵
A World-class University
Celebrating 25th Anniversary

Worldwide Search for Talent

City University of Hong Kong is a dynamic, fast developing university distinguished by scholarship in research and professional education. As a publicly funded institution, the University is committed to nurturing and developing students' talent and creating applicable knowledge to support social and economic advancement. It is a forward looking university with over 20,000 students and more than 3,000 full-time faculty and staff from diverse international backgrounds and ethnicities. Currently, the University has six Colleges/Schools. Within the next five years, the University aims to recruit **200 more scholars** in various disciplines from all over the world, including **science**, **engineering**, **business**, **social sciences**, **humanities**, **law, creative media, energy**, **environment**, and **biomedical & veterinary sciences**.

Applications and nominations are invited for:

Dean of College of Science and Engineering
[Ref. A/592/36]

The College of Science and Engineering is the University's largest academic unit and intellectual core. It comprises eight constituent units: Biology and Chemistry; Building and Construction; Computer Science; Electronic Engineering; Mathematics; Manufacturing Engineering and Engineering Management; Physics and Materials Science; and Building Science and Technology. The College vision is to win international recognition as one of the best science and engineering colleges in the Asia-Pacific region. With seven Academicians, eleven IEEE Fellows, and many promising young scholars, the College has a very strong research profile, pioneering in many of the world's most cutting-edge discoveries. The College offers about 60 taught programmes at various levels, including associate degree to master's degree, and research degree programmes at MPhil and PhD levels.

Qualifications for Appointment
Candidates must be a distinguished scholar and researcher of international standing, with strong research and publication records and the capability to provide academic leadership in both undergraduate and postgraduate programmes. Substantial relevant experience in a senior academic position and evidence of strong management skills appropriate to working in an international environment are highly desirable. The successful candidate should also have a strategic vision to lead the College to new heights in academic achievement and an ambitious drive to foster the highest standards of research, scholarship and education.

Salary and Conditions of Service
Remuneration package will be commensurate with the appointee's credentials and experience. Excellent fringe benefits include gratuity, leave, medical and dental schemes, and relocation assistance where applicable. The appointee will be offered the deanship and a chair professorship simultaneously on a fixed-term contract of three years.

Application and Information
Further information on the post and the University is available at http://www.cityu.edu.hk, or from the Human Resources Office, City University of Hong Kong, Tat Chee Avenue, Kowloon, Hong Kong [Fax : (852) 2788 1154 or (852) 3442 0311/email : deancse@cityu.edu.hk]. Please send the nomination or application with a current curriculum vitae and ask three academic referees to send their references directly to Ms. Rita FUNG of the Human Resources Office by **30 September 2009**.

Please quote the reference of the post in the application and on the envelope. The University reserves the right to consider late applications and nominations, and not to fill the position. Personal data provided by applicants will be used for recruitment and other employment-related purposes.

City University of Hong Kong was ranked the 18th in Asia according to *The Times Higher Education 2009* survey.
http://www.cityu.edu.hk

The Petroleum Institute
Abu Dhabi, United Arab Emirates

Institution: The Petroleum Institute (PI) in Abu Dhabi, United Arab Emirates was created in 2001 with the goal of establishing itself as a recognized institution in engineering education and research in areas of significance to the oil and gas and the broader energy industries. The PI's sponsors include the Abu Dhabi National Oil Company and four major other international oil companies, namely BP, Shell, Jodco, and Total. The Institute is affiliated with and has collaborative programs in place with the Colorado School of Mines, the University of Maryland at College Park, The University of Minnesota, and Leoben and Linz Universities (Austria). For more information, please refer to the PI website: **www.pi.ac.ae**.

FACULTY POSITIONS
- ELECTRICAL ENGINEERING

The Electrical Engineering Department at the PI
is seeking applications for the following positions:

Chaired Professor, Distinguished Professor
Professor, Associate Professor, Assistant Professor
Senior Research Associate, Research Associate

Applicants with research interests and experience in one or more of the following areas: instrumentation and measurements, smart sensors technology, condition monitoring, power quality, power systems, and with interest in applications in the Oil/Gas industry are especially encouraged to apply.

Program faculty will be expected to teach undergraduate and graduate courses, develop an active research program, and to engage in professional and institutional service activities. Opportunities to interact with PI industrial stakeholders and other local industries will be a key feature in the development of a research program.

Interested candidates should submit all materials online:
www.pi.ac.ae/jobs
Review of applications will begin immediately
and will continue until successful candidates are selected
Only short-listed applicants will be notified.

KOÇ UNIVERSITY
DEAN • College of Engineering

KOÇ UNIVERSITY invites applications and nominations for the position of the Dean of the College of Engineering with the appointment effective as early as January 2010. Koç University is a private, nonprofit institution of higher education, founded in 1993. Koç University is committed to the pursuit of excellence in both research and teaching. Its aim is to provide world class education and research opportunities to a high quality group of students by distinguished faculty. The medium of instruction is English. More detailed information about Koç University can be found at the web site: **http://www.ku.edu.tr**

The College of Engineering offers B.S. degree programs in Chemical & Biological Engineering, Computer Engineering, Electrical & Electronics Engineering, Industrial Engineering and Mechanical Engineering. In addition the Graduate School offers M.S. and Ph.D. degrees.

The new Dean is expected to have a distinguished research and publication record in one of the programs covered by the college or related fields. The candidate must provide leadership in teaching and research programs of the college, and foster relations with the academic community and business world. In addition the Dean is expected to have demonstrated ability for organizational and interpersonal skills.

The compensation package is competitive. All information on candidates will be kept confidential. Review of applications will start immediately and continue until the position is filled. Candidates should submit their letter, a list of references, curriculum vitae and further inquiries to:

Professor Yaman Arkun
Chemical and Biological Engineering
Chair of Dean Search Committee
Koç University
Rumeli Feneri Yolu
34450 Sarıyer, Istanbul, Turkey
Phone: 90-212-338 1313 ***E-mail:* *yarkun@ku.edu.tr***

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE | INRIA

Open Positions at INRIA

For Faculty and Research Scientists

INRIA is a French public research institute in Computer Sciences and Applied Mathematics. It is an outstanding and highly visible scientific organization, a major player in European research and development programs.

INRIA has eight research centers in Paris, Bordeaux, Grenoble, Lille, Nancy, Nice – Sophia Antipolis, Rennes and Saclay. These centers host more than 170 groups in partnership with universities and other research organizations. INRIA focuses the activity of over 1200 researchers and faculty members, 1200 PhD students and 500 post-docs and engineers, on fundamental research at the best international level, as well as on development and transfer activities in the following areas:

- Modeling, simulation and optimization of complex dynamic systems
- Formal methods in programming secure and reliable computing systems
- Networks and ubiquitous information, computation and communication systems
- Vision and human-computer interaction modalities, virtual worlds and robotics
- Computational Engineering, Computational Sciences and Computational Medicine

In 2010, INRIA will be opening several positions in its 8 research centers across France:

- Junior and senior level positions,
- Tenured and tenure-track positions,
- Research and joint faculty positions with universities

These positions cover all the above research areas.

INRIA centers provide outstanding scientific environments and excellent working conditions. The institute welcomes applications from all nationalities. It offers competitive salaries and social benefit programs. French schooling, medical coverage and social programs are highly regarded. Visa and working permits for the applicant and the spouse will be provided.

Calendar and detailed application information at:
http://www.inria.fr/travailler/index.en.html

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Faculty Position in Distributed Electrical Systems

at Ecole polytechnique fédérale de Lausanne (EPFL)

The Institute of Electrical Engineering at EPFL invites applications for a **Professor** in the area of **Distributed Electrical Systems.** Qualified dossiers at all levels will be considered.

A strong expertise is requested in the broad fields of classic and non-conventional Energy Generation, Conversion and Storage as well as in Distribution Networks and Technologies. Research areas of interest include, but are not limited to: (a) integration/optimization of distributed energy production and storage technologies, (b) large-scale integration of new loads such as electric transportation and heat pumps, (c) ubiquitous deployment of Smart Grid technologies, (d) broad application of innovative ICT technologies for monitoring and control, (e) interfacing with multi-energy networks.

The successful candidate is expected to initiate independent, creative research programs and actively participate in undergraduate and graduate teaching.

Significant start-up resources and state-of-the-art research infrastructure will be available. Salaries and benefits are internationally competitive.

Applications should include a cover letter with a statement of motivation, curriculum vitae, list of publications and patents, concise statement of research and teaching interests, and the names and addresses of 6 references. Applications must be uploaded in PDF format to the web site **http://iel-search09.epfl.ch**

Candidate evaluation will begin on **December, 1st 2009.**

Enquiries may be addressed to:
Prof. Giovanni De Micheli
Search Committee Chair
INF 341, Station 14, EPFL
CH-1015 Lausanne, Switzerland
E-mail: hiring.iel@epfl.ch

For additional information on EPFL, please consult the web sites **http://www.epfl.ch, http://sti.epfl.ch** and **http://iel.epfl.ch.** Additional information can be garnered from the EPFL Energy Center **(http://energycenter.epfl.ch)** under the direction of Prof. Hans B. Püttgen.

EPFL aims to increase the presence of women amongst its faculty, and qualified female candidates are strongly encouraged to apply.

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

the data

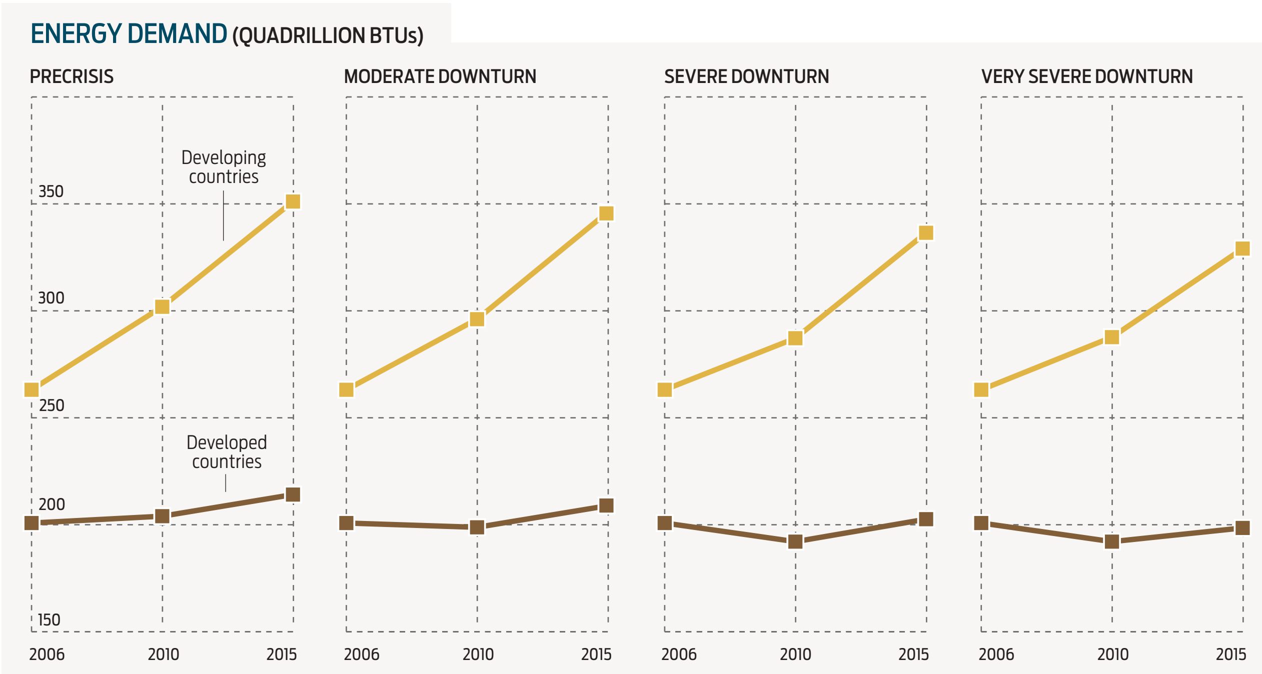

When Will Energy Prices Recover?

JUST AS you'd expect, the global recession has lowered energy demand—and prices, kicking crude oil down from an all-time peak of US $147.27 per barrel in July 2008 to a low of $33.87 in December 2008.

How long will demand languish? According to the McKinsey Global Institute, it depends on how bad the recession gets and whether you live in the developed or developing world.

McKinsey's report, published in March, developed three scenarios based on different economic assumptions. The "moderate" scenario is based on economists' consensus forecast for the overall shape of the global recession. The "severe" scenario assumes that the recession will be deeper than economists expect and the "very severe" that it will be much deeper. McKinsey included these last two scenarios because economists have had to revise their forecasts downward every month since mid-2008.

In the developing world, energy demand remains strong in all three scenarios. Even in a very severe economic downturn, it rises 9 percent from 2006 to 2010 and an additional 16 percent by 2015.

In the developed world, however, even a moderate downturn reduces 2010 energy demand to below 2006 levels, and the very severe scenario keeps it down there all the way through 2015.

The same pattern held for oil specifically. McKinsey predicts that global demand will continue to rise in the developing world even in the face of a very severe downturn. In the developed world, however, even a moderate downturn will keep 2015 demand for oil 1.4 percent below 2006 levels, and in a very severe downturn, 6.7 percent less. —*Steven Cherry*

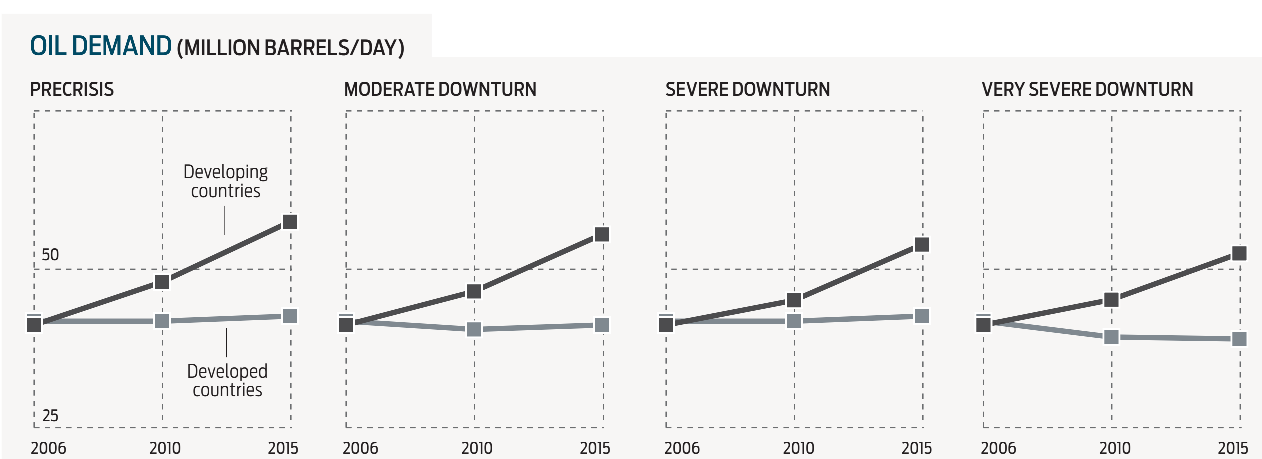

Source: "Averting the Next Energy Crisis: The Demand Challenge," McKinsey Global Institute