ieee spectrum

THE MAGAZINE OF TECHNOLOGY INSIDERS

05.09

MEMORY MAKER

ROBERT DENNARD, WINNER OF THIS YEAR'S IEEE MEDAL OF HONOR, FIGURED OUT HOW TO DO DRAM RIGHT

THE 25 MOST REMARKABLE CHIPS EVER

TOUCH SCREENS GET TOUCHY-FEELY

THE NETFLIX CHALLENGE: HELP US RECOMMEND MOVIES, WIN A MILLION DOLLARS

IEEE

Celebrating 125 Years
of Engineering the Future

volume 46 number 5 international

ieee spectrum 05.09

40

20

24

BETTER PICS: Medical ultrasound images will soon be much sharper [above]; a famed 8-bit processor powers a famed cartoon robot [top right]; Netflix's prize adds the fuel of interest to the fire of genius.

COVER: DAVID YELLEN

THIS PAGE: CLOCKWISE FROM LEFT, PAUL VOZDIC/ GETTY IMAGES; "FUTURAMA" TM AND © 2009 TWENTIETH CENTURY FOX FILM CORPORATION. ALL RIGHTS RESERVED; RANDI SILBERMAN

COVER STORY

44 THANKS FOR THE MEMORIES
Robert Dennard's 1968 invention of the one-transistor dynamic random-access memory set the path of computing for the next half century. For this, he will receive the 2009 IEEE Medal of Honor. *By Sally Adee*

24 THE MILLION DOLLAR PROGRAMMING PRIZE
To improve its recommender system, Netflix put a price on its head. *By Robert M. Bell, Jim Bennett, Yehuda Koren & Chris Volinsky*

30 25 MICROCHIPS THAT SHOOK THE WORLD
These chips unleashed groundbreaking technologies and gadgets—and are one reason that engineers don't get out enough. *By Brian R. Santo*

40 NEXT-GEN ULTRASOUND
Technology from the semiconductor industry promises to supplant the transducers now used for medical imaging. *By Butrus T. Khuri-Yakub, Ömer Oralkan & Mario Kupnik*

UPDATE

9 BAIT TO TRACK DATA THIEVES
A honeypot project in Germany finds viruses and tracks down the data they steal. *By Michael Dumiak*

10 CELLPHONE SECURITY THREAT

12 BRIGHT SPOTS IN THE GLOOM

12 ROBOTIC SEALS AID RESEARCH

13 TOUCH SCREENS WITH FEELING

14 THE RFID AS COMPUTER

15 ULTRAVIOLET RADIOS BEAM TO LIFE

16 THE BIG PICTURE
Fabs get the white-glove treatment.

OPINION

7 SPECTRAL LINES
IEEE's 125th-anniversary celebrations reveal the depth and breadth of its members' contributions to technology.

8 FORUM
Readers discuss the state of the U.S. nuclear arsenal and the demise of the business-method patent.

23 REFLECTIONS
Bob Lucky's head is in the clouds—the computing clouds, that is.

DEPARTMENTS

4 BACK STORY
A man vanishes. Can Google find him?

6 CONTRIBUTORS

18 HANDS ON
A high-definition LCD projector can set you back US $2000—or you could build your own. *By James Turner*

20 GEEK LIFE
How an Apple II–era microprocessor ended up in the head of a 30th-century robot. *By Erico Guizzo & David X. Cohen*

21 TOOLS & TOYS
Solve age-old engineering problems with your kids. *By Sherry Sontag*

22 CAREERS
In 2001, the iPod was a big hit, but Sean Adams had a different idea. *By Sherry Sontag*

56 THE DATA
Is the high-tech unemployment glass half full or half empty? Let's look at the numbers. *By Steven Cherry*

volume 46 number 5 international

ieee spectrum 05.09

LEFT: MAYA ENTERTAINMENT; RIGHT: JEFFREY HAMILTON/GETTY IMAGES

WWW.SPECTRUM.IEEE.ORG

AVAILABLE 1 MAY ON SPECTRUM ONLINE

ANIMATION NATION

India's burgeoning animation and visual-effects industry combines two great passions of India's technorati: moviemaking and software composition. Much of the animation and visual-effects work on feature films such as *Madagascar: Escape 2* and *Alvin and the Chipmunks*, TV shows like Nickelodeon's "Back at the Barnyard," and games like *Gears of War* is being done by Indian animation houses. But as these companies start to expand and create original content, their success will hinge on a combination of economic and cultural factors, many of which still remain uncertain and unpredictable, as Suhas Sreedhar discovered during a recent visit to Bangalore.

ONLINE FEATURES:

LEARN ABOUT CHINA'S DTV standard, which takes advantage of advances in wireless communication and information encoding to create a system that works well even in bad weather and can be viewed on the go.

GOT SOMETHING TO SAY? Join the conversation among our editors, bloggers, and readers in our expanded online comments sections.

ALSO ONLINE:

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

WWW.IEEE.ORG/THEINSTITUTE

AVAILABLE 6 MAY ON THE INSTITUTE ONLINE

TECH EXPERTS SPEAK AT IEEE MEDIA EVENT

Get the scoop on the latest technologies in such areas as cancer recognition, brain-machine interfaces, and cognitive computing, as discussed by seven tech leaders at the IEEE Media Roundtable. The theme of the March event, held in celebration of IEEE's 125th anniversary, was "Embracing Human-Technology Interactions."

NETWORKING HUMANITARIAN TECHNOLOGY

The new IEEE Humanitarian Technology Network enables members to connect and collaborate online with others doing similar work in the field.

ACCESSIBILITY CONFERENCE

The first IEEE conference to identify the accessibility challenges arising from the pervasive use of technology will be held in July in Boston. Areas of focus include transportation, collaborative care, and standards.

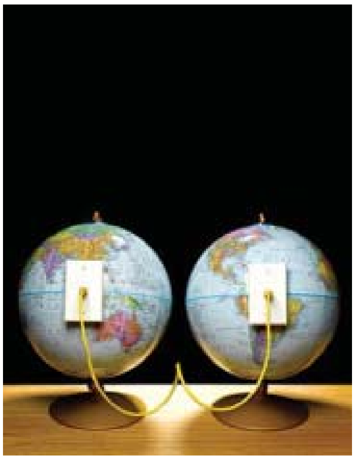

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2009 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine 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

Where in the World Wide Web Is Al Phillips?

LEFT: AL PHILLIPS; RIGHT: LISA SANTO

HIS NAME was Al Phillips, and journalist Brian Santo knew this much about him: He'd been a successful semiconductor executive, having founded Western Digital Corp. in Southern California some 40 years ago. But he seemed to have vanished not long after that.

Phillips was just one of a couple of dozen people Santo had set out to interview for this issue's "25 Microchips That Shook the World." The idea was to get insight about the chips straight from the designers and managers who created them. To reach these people, some long retired, Santo put out calls, sent e-mails, and, of course, searched Google. In fact, he Googled like crazy.

"I wish I could just hook up my brain to Google," sighs Santo [above, right], who started his career at *Electronic News*, a weekly trade newspaper, in 1984 (when, by the way, Google founders Larry Page and Sergey Brin were 11-year-olds still getting allowances from their parents).

After a few days, Santo had gotten hold of all the people on his list. All except one, that is—Al Phillips.

An initial Google search located an Al Phillips in Knoxville, Tenn., but he turned out to be an insurance agent. There was an Al Phillips who was an active member of the Taftville Congregational Church in Connecticut, but he turned out to be a she. There was an Al Phillips who'd recently been a New York Jets cornerback, but no way had he helped architect a chip in the 1970s. Then there was the Canadian wrestler, who was actually better known as Spider Dick, and also the founder of a chain of dry cleaners in Hawaii. Nope, and nope.

Santo kept on Googling, and then he got a break. A letter to the editor of the *Orange County Business Journal*, in California, was signed "Al Phillips, Founder, Western Digital Corp., Newport Coast."

Bingo! From there, Santo turned to the phone book, one of the oldest tools of the trade in journalism—but not a big fat paper one. He used Google's phone book.

Reached at his home in Newport Coast, Phillips [above, left] was only too happy to chat. But first he needed to keep an appointment at a local golf course. Later, when he did have time to talk, he told Santo he was a former Motorola and Rockwell executive who founded Western Digital in 1970. He also mentioned that he'd written an article published in the June 1964 issue of *IEEE Spectrum*, "Monolithic Integrated Circuits," and that at age 80, he's still a busy man. He doesn't play football, but besides golf, he enjoys bicycling and skiing.

Google didn't turn up those details. For some things, you still do need a phone. ❑

CITING ARTICLES IN IEEE SPECTRUM

IEEE Spectrum publishes two editions. In the international edition, the abbreviation INT appears at the foot of each page. The North American edition is identified with the letters NA. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, the first Update page is in *IEEE Spectrum*, Vol. 45, no. 5 (INT), May 2009, p. 9, or in *IEEE Spectrum*, Vol. 45, no. 5 (NA), May 2009, p. 11.

ieee spectrum

EDITORIAL

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

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

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

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

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

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

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

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

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

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

EXECUTIVE PRODUCER, SPECTRUM RADIO Sharon Basco

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

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

INTERN Rosaleen Ortiz, r.ortizdejesus@ieee.org

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

ART & PRODUCTION

SENIOR ART DIRECTOR Mark Montgomery

ASSOCIATE ART DIRECTOR Michael Solita

ASSISTANT ART DIRECTOR Brandon Palacio

PHOTO EDITOR Randi Silberman

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy

EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia

SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

WEB PRODUCTION COORDINATOR Jacqueline L. Parker

MULTIMEDIA PRODUCTION SPECIALIST Michael Spector

EDITORIAL ADVISORY BOARD

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

EDITORIAL CORRESPONDENCE

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

ADVERTISING CORRESPONDENCE

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

contributors

ROBERT M. BELL, CHRIS VOLINSKY, and **YEHUDA KOREN** worked together at AT&T Labs Research, in Florham Park, N.J., where they competed for "The Million Dollar Programming Prize" [p. 24]. **JIM BENNETT,** who was vice president for recommendation systems at Netflix, helped organize the competition to build a better movie recommender. Bell is a principal member of the technical staff at AT&T Labs and a fellow of the American Statistical Association. Volinsky has been director of the labs' statistics research department since 2003. Koren now works for Yahoo Research, in Israel.

BUTRUS T. KHURI-YAKUB, ÖMER ORALKAN, and **MARIO KUPNIK** can all put "Dr." in front of their names, although you can't turn to them for a prescription. Still, the work they describe in "Next-Gen Ultrasound" [p. 40] brings them nearly as often to Stanford University's medical school as to its engineering school. Khuri-Yakub, a professor of electrical engineering, manages technical operations for Stanford's E.L. Ginzton Laboratory, where Oralkan is a senior scientist and Kupnik is a research associate.

ROBERT W. LUCKY ponders the desirability of dispersing all your data into the computing cloud in this month's Reflections column [p. 23]. Lucky, an IEEE Fellow, holds 11 patents and worked for many years at Bell Labs. Before retiring in 2002, he was vice president for applied research at Telcordia Technologies, in Piscataway, N.J.

SHERRY SONTAG is coauthor of the best seller *Blind Man's Bluff: The Untold Story of American Submarine Espionage* (Harper Perennial, 1998). Her interest in digital jukeboxes led her to purchase one made by Slim Devices in 2003; her interest in her two small sons led her to try the children's software game *Time Engineers*. She writes about both in this issue [p. 22 and p. 21].

JAMES TURNER, author of this month's Hands On column [p. 18], says there's more to a do-it-yourself project than just building it. His new LCD digital video projector is sitting idle in his laundry room until he can clear enough wall space to display its capacious images. Turner is a contributing editor for O'Reilly Media and a correspondent for the *Christian Science Monitor*.

DAVID YELLEN, who shot the photos of our Medal of Honor winner [p. 44], became a photographer after working as a musician, fashion designer, and fishing crewman. When not playing with his daughter, he still fishes in Sheepshead Bay, near his Brooklyn, N.Y., home. "Shooting at IBM," he says, "was an adventure," because the Eero Saarinen building is hidden in the woods.

CLOCKWISE FROM TOP RIGHT: CHE GRAHAM; JAMES TURNER; DAVID YELLEN; STEPHEN LACKO

spectral lines

Engineering the Future at IEEE

IEEE INNOVATORS: From left: Roy Want, Krishna Palem, Miguel Nicolelis, Dharmendra Modha, K.J. Ray Liu, Rangachar Kasturi, Katie Hall, IEEE Past President Lew Terman. *PHOTO: IEEE*

Anniversaries are often engulfed by nostalgic memories of past achievements and glory days gone by. Not this one. While commemorating its members' historic engineering accomplishments over the past 125 years, IEEE's 125th-anniversary celebrations—going on throughout 2009 around the globe—focus on how today's engineers are inventing and building technologies that will shape the world's future over the *next* 125 years.

Take, for example, the media roundtable IEEE put on in New York City in March to showcase the work of just seven of its members (you can still view the whole presentation at http://www.ieee125.org/engineering-the-future/media-roundtable.html). It's well worth checking out.

First, Katie Hall, an IEEE Senior Member and chief technology officer of WiTricity, talked about a magnetic resonance–based technology her company is developing that can wirelessly transmit power, from milliwatts to kilowatts, to electronic devices meters away. Rangachar Kasturi, an IEEE Fellow and a professor at the University of South Florida, described his team's work on pattern recognition and its uses in medical image analysis, biometrics, and environmental sensing. K.J. Ray Liu, an IEEE Fellow and a professor at the University of Maryland, College Park, discussed the development of predictive cancer testing—tests that would be able to identify both patients whose cells are transitioning from normal to cancerous *and* the type of cancer that's developing—based on algorithmic analyses of genome-proteome signaling. Dharmendra Modha, an IEEE Senior Member and manager of cognitive computing at IBM's Almaden Research Center, related his team's work on SyNAPSE, a project to build a "brain"—a computer system that can see, feel, and think like a human being.

In another intriguing biomedical engineering development, Miguel Nicolelis, an IEEE Member and professor and codirector of the Center for Neuroengineering at Duke University's Medical Center, talked about a chip his group has built that allows monkeys to control robots across continents with brain waves. Krishna Palem, an IEEE Fellow and a professor at the George Brown School of Engineering at Rice University, reported on his work on low-energy chip technology. He's now using it in a low-power, low-cost "I-Slate" that will give children in impoverished areas—who may have never even seen a teacher—access to learning. Finally, Roy Want, an IEEE Fellow and senior principal engineer at Intel, talked about a mobile solution called dynamic composable computing that could make it possible for you to access any devices or applications you need on the fly—through your cellphone.

Phew! Clayton M. Christensen coined the term *disruptive innovation*, and if some of these developments don't turn out to be disruptive I don't know what will. But then again, sometimes new technologies take their own sweet time to reveal their power and influence. This year's IEEE Medal of Honor winner, Robert Dennard (see "Thanks for the Memories," in this issue), described his own work on one-transistor dynamic RAM like this: "We just didn't imagine how far it would go, how much it would totally change computing."

So if you have students, point them to the roundtable webcast so they can experience the magic of watching people talk passionately about the work they love. If you're a student, watch the webcast and listen carefully. New fields and new professions will be spun out of the work in power, robotics, computing, and biomedical engineering described therein. You will learn, if you haven't already, that a universe of opportunities to invent new technologies and make a difference in the world remain wide open to the curious and adventurous willing to take chances and follow their dreams. —Susan Hassler

For more information about IEEE's 125th-anniversary celebrations, go to http://www.ieee125.org.

forum

LETTERS do not represent opinions of the IEEE. Short, concise letters are preferred. They may be edited for space and clarity. Additional letters are available online in "And More Forum" at http://www.spectrum.ieee.org. 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.

STOCKPILING SECURITY

As a physicist and engineer with over 36 years' experience in nuclear weapons laboratories, I found "What About the Nukes?" [March] well reasoned. At the same time, I have some issues with its conclusions. The ability to design a fundamentally new weapons system without at least a minimal testing program is at best suspect. The answer to this problem as promulgated by one of the creators of the Stockpile Stewardship Program, Victor H. Reis, is to rely on simulation. I was one of the initial laboratory directors of the stockpile stewardship simulation program, the Accelerated Strategic Computing Initiative, and I take great pride in its accomplishments. Nonetheless, I am skeptical that the simulation agenda is sufficiently mature to be relied on as a full alternative to testing. We are better off maintaining designs that we know work than we are fielding untested ones. However, I would argue against halting new developments. The book on nuclear weapons design is far from closed. For that reason alone, the United States would be well advised to maintain a vibrant nuclear weapons research, design, and development program.

William J. Camp
IEEE Member
Cedar Crest, N.M.

I have been an electrical engineer in a process industry for more than 38 years. I don't have any background in nuclear weapons engineering, but I have had exposure to reliability engineering concepts, and I believe that these concepts are quite broad in their application. Your article implies that weapons systems follow a "bathtub" curve, where there is an early period of infant mortality, followed by a long period of stable performance and a period of increasing failures. The end of life for such a system is defined as the point where failures begin increasing. However, modern reliability thinking posits that there are at least six patterns of failure, of which the bathtub curve is only one. In fact, it has been shown that as systems become more complex, two curves predominate. Both are characterized by the lack of a determinate end-of-life stage.

The stockpile should be replaced with more modern equipment, but for reasons other than reliability. Our current nuclear arsenal was designed to deliver large-yield weapons to targets in the Soviet Union. Today the threat from the former Soviet nations is minimal, and we have the ability to precisely target our weapons, obviating the need for such large warheads. Another reason to renew the arsenal is the need to train a new generation of weapons engineers. The maintenance of 50-year-old legacy systems may not be enough to attract the best and brightest. That by itself could lead to failures of the system.

John M. Briggs
IEEE Senior Member
Vancouver, Wash.

How many nuclear warheads does one country need? 5000? 500? 50? 5? Most countries get by with none.

Peter Flanagan
IEEE Member
Melbourne, Australia

PATENT PROTECTION

I enjoyed Steven J. Frank's "The Death of Business-Method Patents" [March]. The *Bilski* decision is a start toward curtailing these insidious bits of patent overreaching, but we have far to go. I also think something more basic is involved, something to do with adjusting to the change from the industrial age to the information age. For example, our present economic troubles are, I would argue, a product of the information age. With the aid of the computer, financial engineers created products and processes that few understood and that led to an artificial, temporary increase in financial assets. Incidentally, many of these financial products and processes were patented.

David S. Holland
IEEE Member
Alexandria, Va.

The author responds: The urgent issue is to not throw the baby out with the bathwater: We don't want industrial-age prejudice to preclude all information-age patents, and if *Bilski* is taken too literally, software will become off-limits almost entirely. If a patent claim requires even a modicum of technology, it should pass muster as patentable subject matter. That doesn't mean a patent will ultimately be awarded, just that the claim is eligible for consideration on the merits.

update

more online at www.spectrum.ieee.org

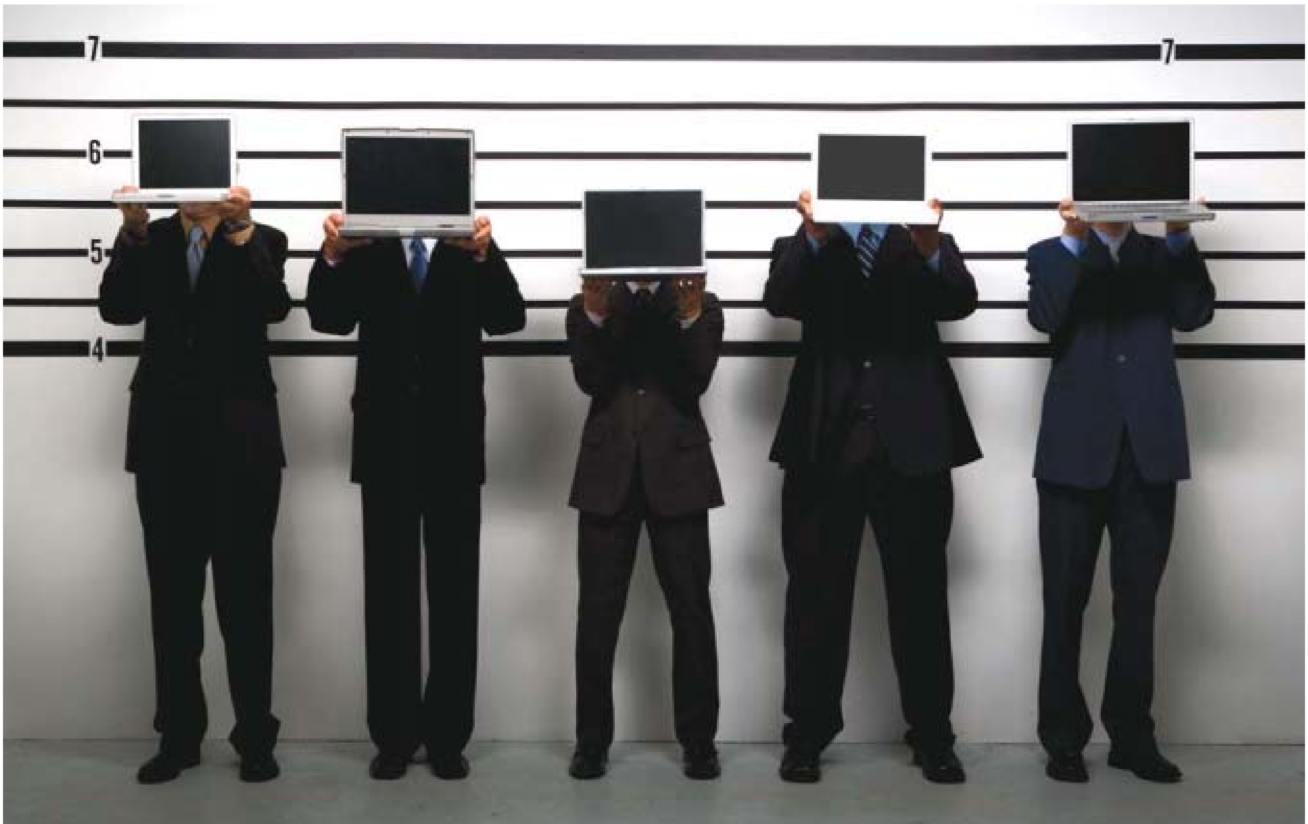

Setting Bait to Track Data Thieves

Germans use unguarded PCs as honeypots

SEVEN COMPUTERS hum through the Rhineland night at the University of Mannheim's Laboratory for Dependable Distributed Systems. All they do is collect bad news and nasty infections from the open Internet.

This is the lab's honeypot network, says Thorsten Holz, a doctoral student at the lab. The honeypots are machines that are walled off from the German university's network but connected to the Internet. By leaving themselves unguarded and pretending to be operated by naive humans, they tirelessly troll for the latest in spam, worms, viral infections, and malware. Then the honeypots execute the bad code and record what happens. Researchers hope that by studying the results they'll get a better understanding of how data is stolen and what happens to it.

Holz's team recently reported the results of a seven-month study that focused on two "families" of the seven or eight current varieties of "keyloggers"—malware that records keystrokes in order to pick up passwords and other sensitive data. The two families were Limbo/Nethell, which uses sham Web sites to infect visitors lured by spam or other social-engineering tricks, and its more sophisticated cousin ZeuS/Zbot, a keylogger that uncoils from an attachment and hides in the user's browser to steal passwords and account information.

The team followed the Mannheim honeypot data to open "drop zones," computers where cybercriminals compile and aggregate their ill-gotten gains before transferring them elsewhere.

In analyzing the drop zones, the team found that:

CATCHING CROOKS: Honeypots can tell security researchers how data is stolen and reveal where cybercriminals hide their stolen goods.
PHOTO: JOHN LUND/SAM DIEPHIUS/GETTY IMAGES

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

“There isn’t a world water crisis—yet.”

WILLIAM COSGROVE, coordinator of the United Nations’ latest report on the state of the world’s water

update

• The drop zones housed 33 gigabytes of account data from 177 000 compromised machines in 175 countries, including 10 775 unique bank-account credentials.

• Facebook is the most popular site for stolen social-network credentials, Windows Live for Web-mail user names and passwords, and eBay for online trading accounts.

• ZeuS can parse account information to read balances (the mean value was US $1700, the average $5225).

Worrisome as the results are, the honeypots themselves are a step forward in understanding black-market data.

2 Days | AVERAGE TIME A COMPUTER REMAINS INFECTED WITH EITHER OF TWO TYPES OF KEYLOGGING MALWARE

Using honeypots to analyze malware came into its own about a decade ago through the work of security pros like Lance Spitzner, author of *Honeypots: Tracking Hackers* (Addison-Wesley, 2003); Marty Roesch, creator of the popular Snort intrusion-prevention and detection software; and Ron Gula (formerly of the National Security Agency), who started working with honeypots in order to defend the big communications networks run by BBN Technologies and GTE.

“What’s amazed me the most is the explosion in the different ways they’re used,” Spitzner says. “There’s client honeypots, [Voice over Internet Protocol] honeypots, Bluetooth honeypots. Ten years ago there were really only one or two choices.” Spitzner founded the Honeynet Project in 1999, and it now has worldwide presence.

Security firms use honeypots to identify new malware they must defend their clients against. But according to Roesch and Yuval Ben-Itzhak, chief technical officer of the Web security company Finjan, the kind of half-year-long, in-depth analysis the Mannheim group did takes too long to be done by most security firms.

The Mannheim lab’s next project is to further automate the honeypots. “We have automated collection, analysis, and monitoring. The next step is to do this on a larger scale,” Holz says. —MICHAEL DUMIAK

The Mobile Infections Threat

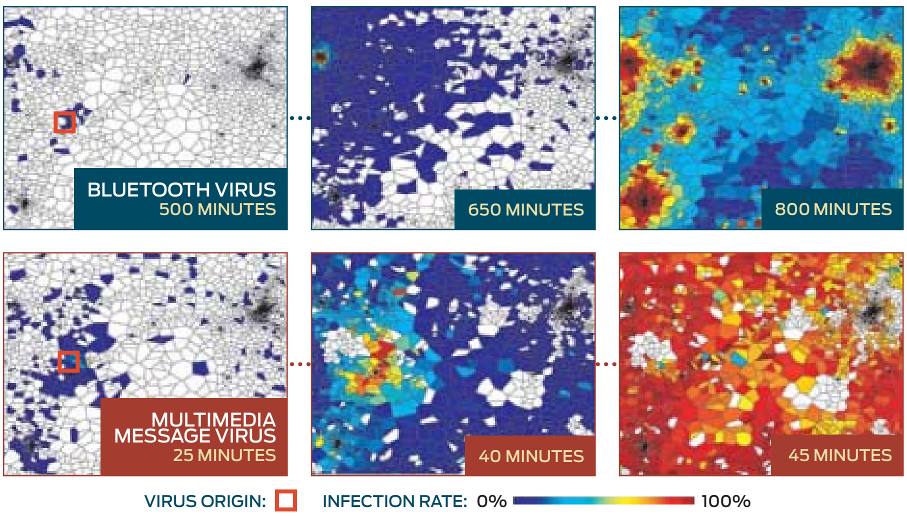

Why haven’t smartphones been overrun by a deluge of viruses the way PCs have? According to researchers at Northeastern University, in Boston, these phones have been safe because no single mobile-phone operating system has achieved a big enough market share for a major virus outbreak to occur. But that could change. The Northeastern group used anonymous call data from 6.2 million mobile-phone subscribers to simulate the epidemiology of a virus outbreak. Like flu and SARS, Bluetooth viruses spread slowly, through close physical proximity [top]. On the other hand, viruses embedded in multimedia messages such as picture messages spread much more rapidly [bottom] but are confined to groups of people who know each other and share the same kind of phone.

CENTER FOR COMPLEX NETWORK RESEARCH/NORTHEASTERN UNIVERSITY

update

Some Bright Spots in the Gloom

THE TECHNOLOGY industry is suffering mightily from the global economic crisis. Worldwide semiconductor revenue is expected to drop 24.1 percent in 2009, to US $194.5 billion, according to Gartner, a technology research firm in Stamford, Conn. Meanwhile, revenue from enterprise software—that is, corporate-scale systems—will be flat at about $222 billion. Still, there are some sectors that will gain, some bright spots in the gloom. —SAMUEL K. MOORE

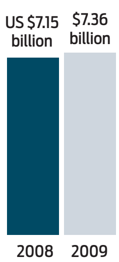

WORLDWIDE LED REVENUE

LEDs: The LCD TV market may be all gloom and doom, but LEDs are increasingly becoming the way to light the screen up. They replace cold-cathode fluorescent lamps, leading to slimmer TVs or screens with an improved contrast ratio. The LCD TV market will gobble up $163 million worth of LEDs in 2009, according to market research firm iSuppli, in El Segundo, Calif. That's more than double last year's figure. By 2012, LCD TV makers will be spending $1.4 billion on them. Manufacturers are using LEDs more often to backlight LCD TVs because of efficiency improvements and the development of lower-cost manufacturing in Taiwan and South Korea, says Jagdish Ribello, iSuppli's LED analyst. Even more gains will surely come, now that engineers have figured out a way to fight an efficiency-sapping phenomenon called LED droop.

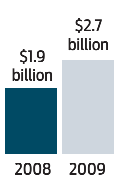

WORLDWIDE VIRTUALIZATION-SOFTWARE REVENUE

Virtualization: In a slump, the drive to do more with less is acute. So virtualization software, which basically lets a company get more out of its computing infrastructure by creating the equivalent of multiple machines per computer, is a natural fit. Gartner says the overall market will grow 43 percent to $2.7 billion in 2009. The largest component of that will be the $1.3 billion worth of virtualization management software Gartner expects firms to sell this year.

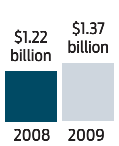

WORLDWIDE MEMS REVENUE

MEMS: Open up 1 in 10 new cellphones shipped in 2008 and you'd find a MEMS accelerometer, according to iSuppli. And mobile phones will continue to drive growth this year—not just for accelerometers but for zoom and auto-focus actuators, pico projectors, gyroscopes, and RF filters, too. Together with consumer items like game controllers and digital cameras, mobiles will push the market up 12 percent in 2009, to nearly $1.4 billion. MEMS manufacturers are making it easier to include accelerometers. For example, in March, STMicroelectronics unveiled an accelerometer for handhelds that can be hooked directly to a mobile's battery without an intervening voltage regulator and that works consistently even as the battery runs down.

Robotic Baby Seal Could Diminish Dementia

It's soft, it's cute, it's cuddly, and it's powered by two 32-bit reduced-instruction-set-computer microprocessors.

Paro may look like a toy, but it's quickly attracting the serious attention of rehabilitation researchers. In Japan, more than 1000 units have been sold to care providers in nursing homes and hospitals, as well as to consumers who want a robotic companion.

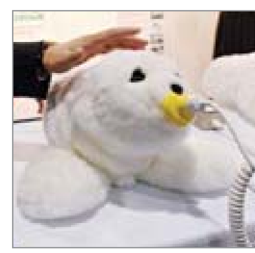

FEED ME! Paro gets its charge from a plug that looks like an infant's pacifier.

PHOTOS: TOP, PICTURE CONTACT/ALAMY; BOTTOM, ROBYN BECK/AFP/GETTY IMAGES

Short-term experiments in Japan and the United States show that Paro can have positive effects on the mental health of some elderly people. Now long-term studies are under way in Europe. The results could lead to specialized versions of Paro to help specific groups of people, such as elderly individuals suffering from dementia or children with autism. For more, see http://www.spectrum.ieee.org/may09/moreparo.

MICHEL TCHEREVKOFF/GETTY IMAGES

Touch Screens With Feeling

Engineers add texture to touch-screen devices

One of the most sought-after new features on mobile devices is the touch screen. But that name is a misnomer, according to a group of researchers at Northwestern University, in Evanston, Ill., who point out the obvious—that these displays offer minimal tactile feedback. But suppose touch screens could touch you back: You could "feel," say, the edges of the buttons on a virtual keypad or the links on a Web page.

Recently, at the IEEE-sponsored World Haptics Conference 2009, in Salt Lake City, Michael A. Peshkin and J. Edward Colgate, mechanical engineering professors and codirectors of Northwestern's Laboratory for Intelligent Mechanical Systems, described their candidate for such an enhanced touch screen—a device they call the Tactile Pattern Display, or TPaD. It can create the illusion of texture on an unadorned piece of glass.

The 25-millimeter-diameter prototype takes advantage of the relatively high coefficient of friction between human skin and glass. "Glass is remarkably 'sticky,'" says Peshkin. "You think of it as being smooth, but the coefficient of friction between glass and your fingertip is about 1." (That's about the same as rubber on dry concrete.) But when the glass is vibrated ultrasonically, a cushion of air forms between your finger and its surface.

The TPaD's 1.6-mm-thick glass layer is set into ultrasonic oscillation by a piezoelectric ceramic disc glued to the glass. The amplitude of the oscillation can be controlled, reducing the coefficient of friction by up to a factor of 10. The higher the amplitude of oscillation, the lower the coefficient of friction.

When you run a finger across the TPaD, "you get a very strong tactile sensation of something being there, like a bump, a dip, or an edge," says Peshkin. To create the feeling that you're rubbing your finger against a file grating, for example, you turn the plate oscillation on for the grooves and off for the ridges.

The entire plate vibrates, so the amount of friction is the same all over the TPaD's surface at any given time. But because the oscillations are modulated as your finger's position changes, the device fools you into thinking that there are varying amounts of friction at different locations. The prototype uses optical sensors to keep track of your finger's position. The friction reduction can be switched on and off so quickly (within about 4 milliseconds on average) that the pitch of virtual bumps or dips can be made far finer than what a fingertip can discern.

The group is currently working on several engineering challenges, including creating larger versions of the TPaD. They've already gotten it to work on screens comparable in size to those on cellphones.

Another technical hurdle is limiting how much power the TPaD draws, which is hugely important in mobile devices, where your power budget is measured in milliwatts. "In principle, the TPaD does not need a lot of power," says Peshkin. "But in practice, you run into issues like parasitic losses because of power being dissipated through things like mountings."

Asked when an online shopper can expect to use a TPaD to tell how a garment feels, Peshkin said, "That's getting into haptic virtual reality, which is a ways down the pike." —Willie D. Jones

news briefs

ELECTRIC TOMATO

A Purdue University food scientist has designed a device that generates a room-temperature plasma field inside a sealed food package. The plasma makes ozone that kills *E. coli* and other unfriendly bacteria. The ozone eventually reverts to ordinary oxygen.

PHOTO: TOM CAMPBELL/PURDUE UNIVERSITY

update

RFID Chips Gain Computing Skills

One way to do long computations with short bursts of power

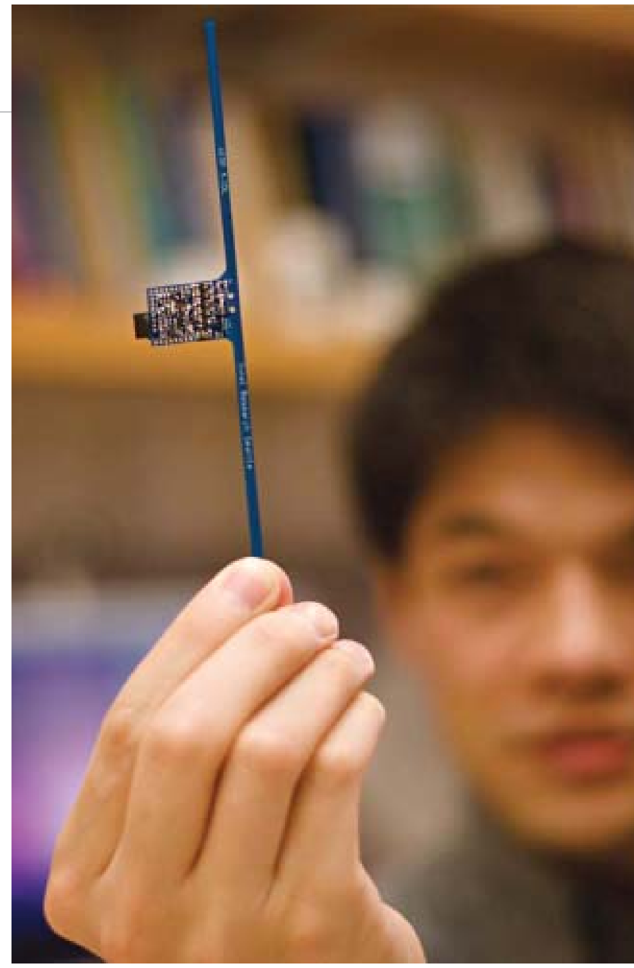

SMARTER THAN IT LOOKS: RFID chips can do complex calculations, thanks to Kevin Fu's software. *PHOTO: BEN RANSFORD*

IN THE 2000 movie *Memento*, the main character tries to solve the mystery of his wife's murder, despite suffering from amnesia that causes his brain to effectively "reboot" every 5 minutes. In the world of computing, "passive" radio-frequency identification (RFID) chips have a similar problem. Dependent for power on infrequent, scavenged RF energy from a reading device, RFID chips may reboot more than once per second and then lie dormant indefinitely, waiting for the next reader to come along.

But computer scientists in Massachusetts are working on software, aptly named Mementos, that could allow an RFID to perform computations that span many power losses and reboots. The software may enable the chip to compute cryptographic protocols, leading to more secure signals. And it might allow RFID chips to be more than just data collectors. They could analyze and possibly take action based on changes to the stress on a "smart" bridge or to trends in a person's vital signs, for instance. Such computational RFIDs could play a role in the transformation of the Internet from a network of computers to a network of things—appliances, cars, smart clothes, and so on.

"We're working on software to make it possible to actually compute, given that our power is going to be disappearing and reappearing," says Kevin Fu, assistant professor of computer science at the University of Massachusetts Amherst.

Mementos does two things: It makes sure that the RFID keeps working toward finishing a computation, and it also keeps the chip in a state such that if it loses power, it can quickly resume work when the power returns. One way the software does that is to have the chip perform energy-intensive tasks, such as writing data to flash memory, only when ample power (more than 2.2 milliwatts) is available.

Ravi Pappu, cofounder of the RFID company ThingMagic, based in Cambridge, Mass., says that the work of Fu's team is very important.

"We have millions of computers everywhere, but computers have been chained to desks and server racks and other kinds of infrastructure," Pappu says. "Rather than being constrained by dragging computers of various shapes and sizes into the real world, could we equip the real world with computing? I think what Kevin and his guys have done is an advance in that direction."

Fu and his colleagues developed their RFID computing software on Intel's prototype Wireless Identification and Sensing Platform (WISP), a postage-stamp-size RFID chip with specs suited for a 1980s home-brew computing enthusiast: a 16-megahertz microprocessor, 512 bytes of RAM, and 8 kilobytes of storage (in the form of flash memory).

The amount of RF power WISP picks up, Fu says, can vary drastically. And Mementos must make programs run on WISP perform under each scenario—sleep (0.2 to 2.5 microwatts); midrange (1.8 to 3.6 mW), which allows it to read from memory and compute; and active (2.2 to 25.2 mW) in which it can even write to flash memory.

"As much as possible we'd like to protect programmers from the underlying problems" of fluctuating power and regular reboots, says Benjamin Ransford, a graduate student in Fu's lab.

Jason Flinn, associate professor of computer science and engineering at the University of Michigan, says he's impressed by Mementos but thinks that it still has a long way to go.

The University of Massachusetts work "asks some very interesting questions and has some preliminary ideas," he says. "But I don't know that they've solved those problems or fully validated it yet."

—MARK ANDERSON

54 BILLION Cubic meters per year of freshwater to be generated by desalination technologies in 2020, according to Lux Research, in Boston. That's about one-third of what flows out of the Nile River.

BAT-SIGNAL: Ultraviolet transmitters beam their signals into the sky [green]. Receivers look to the scattered UV light [gray].
IMAGE: ISTOCKPHOTO

Ultraviolet Radios Beam to Life

Secret military communication scheme from the 1960s is finally practical

The U.S. military has been chasing ultraviolet (UV) communication for decades. Now researchers say radios that communicate using UV light are finally within reach. Working with the Army Research Lab (ARL) in Adelphi, Md., these researchers are mapping out the steps needed to commercialize UV radios. They've reached the last piece of the puzzle: untangling the poorly understood, extraordinarily complex way ultraviolet light scatters. If they can do that, they will have unlocked the secret to a new form of non-line-of-sight communication.

Proposed UV radios communicate in the so-called solar blind portion of the UV-C band—light having wavelengths from 200 to 280 nanometers—which, unlike the sun's UV-A and UV-B rays, is almost completely blotted out by the atmosphere. Near Earth's surface, even a strong UV-C signal would die off within a few kilometers, as individual photons are picked off one by one by oxygen, ozone, and water molecules. But that attenuation also makes UV-C radiation ideal for short-range wireless links, such as in unattended ground sensors. The U.S. military is interested in such short-range communications because they can't be intercepted or jammed outside their intended range. What's more, within its limited range the UV-C band has an inherently high signal-to-noise ratio, enabling the use of very-low-power transmitters, according to ARL scientist Brian Sadler.

In contrast with other optical schemes, which rely on the transmitter sending a signal more or less directly to the receiver, a UV system can take advantage of the signal scattering in the atmosphere. The transmitter beams a modulated signal into the sky in the shape of a cone (think of the Bat-Signal). The receiver is trained on the sky as well, at an overlapping angle. That positioning makes it ideal for sensors in dense urban environments where line-of-sight communication doesn't work. "This goes around corners, through forests, anyplace you can get light," says Russell Dupuis, an electro-optics professor at Georgia Tech.

Early UV-C radio prototypes were far too clunky to make it out of the lab—the transmitter was a massive laser, and the receiver was a bulky vacuum-tube-based photodetector. But thanks to materials-science advances, the new transmitters are tiny, commercially available UV LEDs. The receivers have also shrunk to tiny, solid-state avalanche photodiodes—devices in which a single photon produces an avalanche of electrons. With currently available devices, and under typical operating conditions, a low-power UV-C system with the right overlap could transmit roughly 100 kilobits per second at 10 meters, dropping to less than 10 kb/s at 100 meters, still more than enough for good digital audio.

A major remaining challenge is modeling the behavior of the signal as it scatters randomly in the sky.

"It's not like there's a mirror up there," Sadler says. At the University of California, Riverside, he and electrical engineering professor Zhengyuan Xu experimented with different transmission sources and receivers to characterize the angle at which the transmitter beam and the receiver's field of view should cross. "It's not just about getting the overlap volume as large as possible," Xu adds. "When it's narrower, the signal is sometimes more enhanced." —Sally Adee

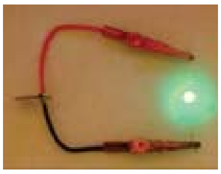

news briefs

BUG BATTERIES
MIT materials scientists say they've used genetically engineered viruses to construct electrodes for lithium-ion batteries. The viruses, which infect bacteria, were designed to grow amorphous iron phosphate along their surfaces and then bond with carbon nanotubes. The result was a nanoscale structure with charge and discharge rates comparable to those of electrodes in state-of-the-art Li-ion batteries.
PHOTO: MIT

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

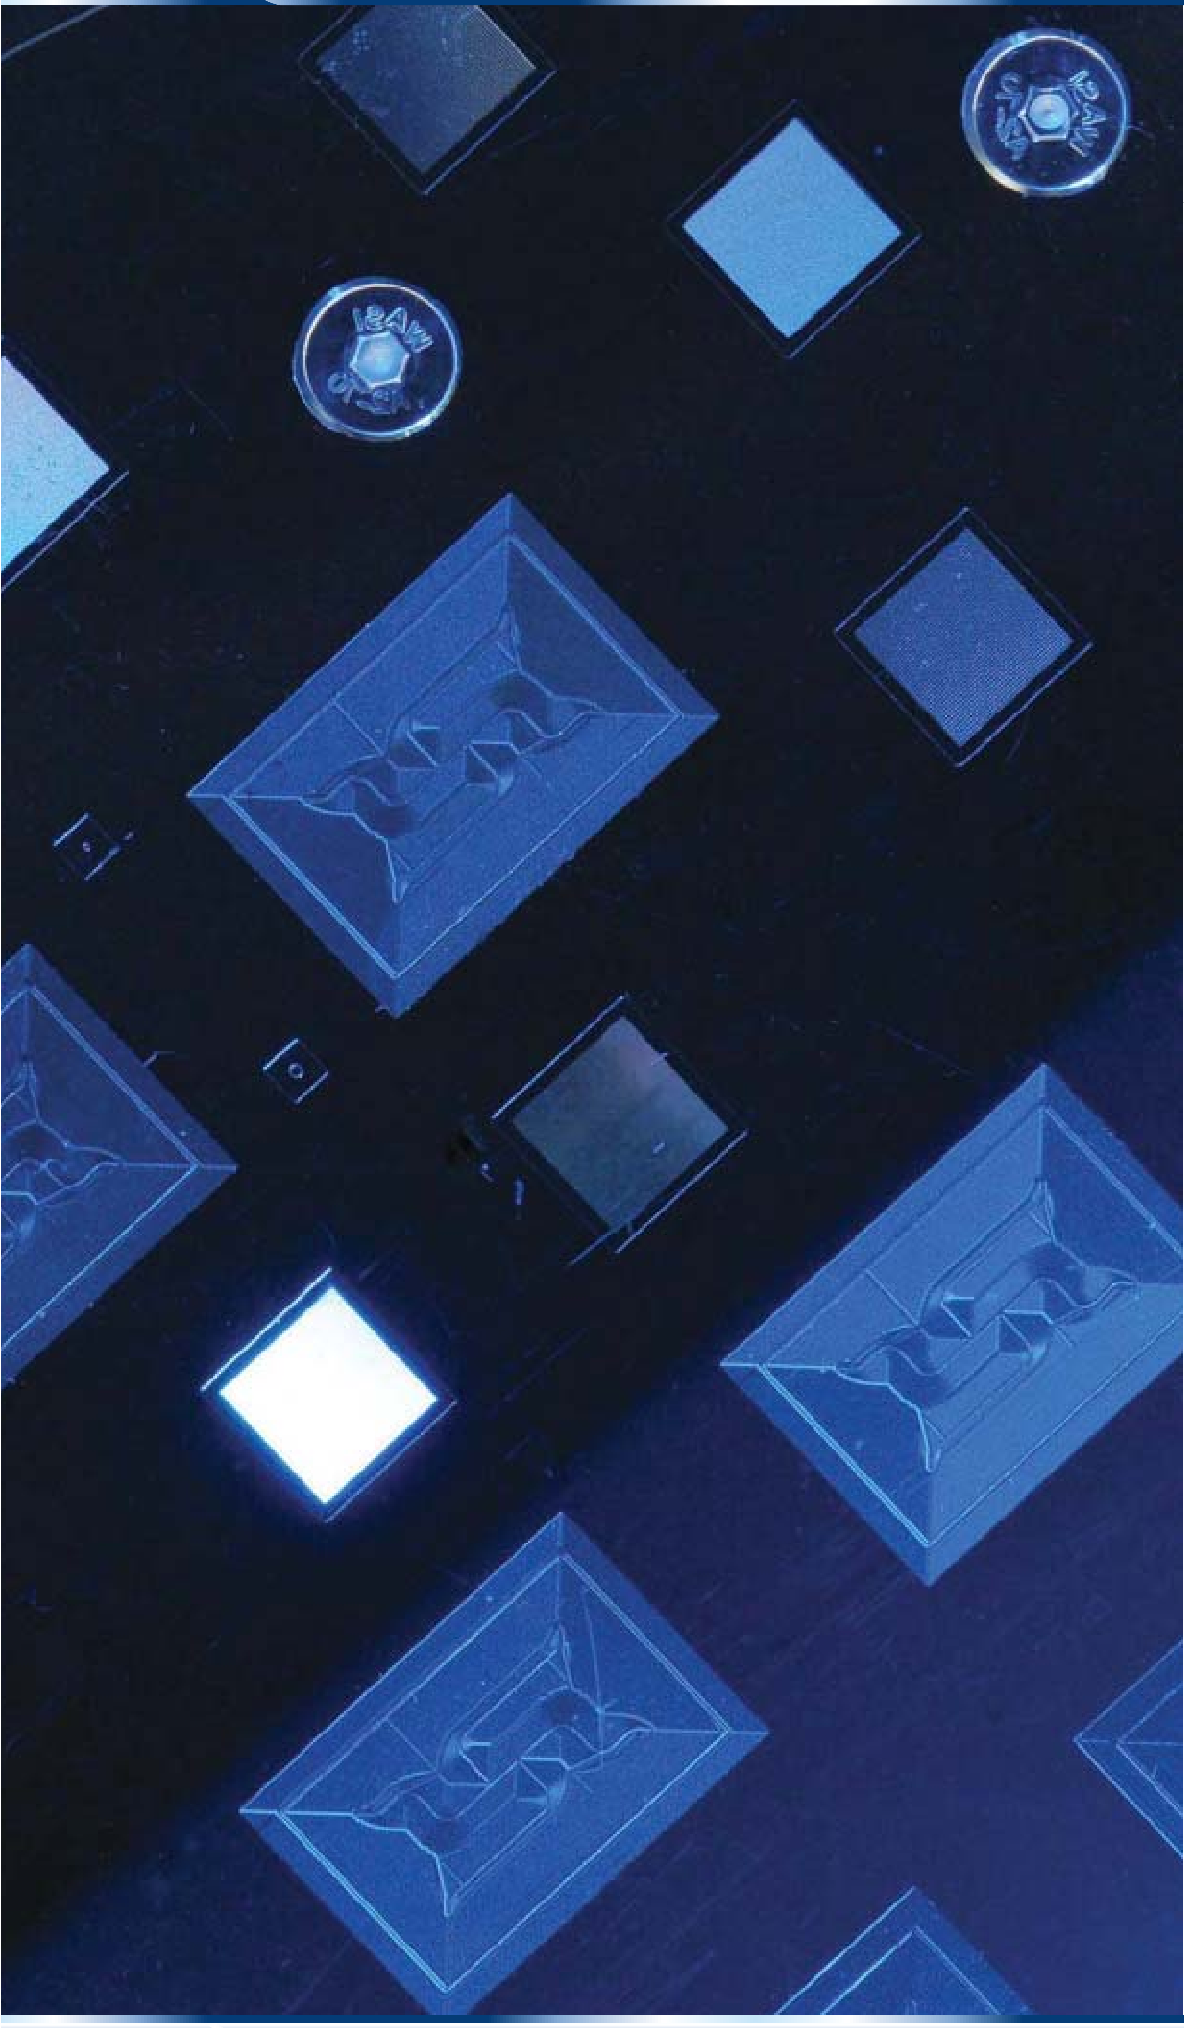

the big picture

A FAB LIFE

No, it's not a butler presenting the lord of the manor with stud and cuff-link options. It's a clean-room technician holding a polymer film disc that's been embossed with dozens of microchips using Jenoptik's microscale thermoforming process. Jenoptik says its hot-stamping technique can cheaply and reliably burn in nanometer-scale features to create chip-scale biomedical laboratories. The fully automated process can yield dozens of different chip configurations, including microscopic scaffolding for cell cultivation and tiny reaction chambers prewired with microelectrodes.

IMAGE: JAN-PETER KASPER/DPA/LANDOV

hands on

1

2

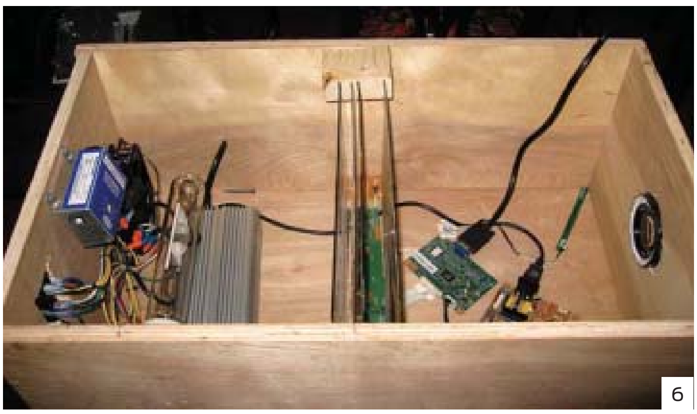

6

7

8

THE PROJECTOR PROJECT

Build your own digital movie projector

A GOOD HIGH-DEFINITION LCD projector can still set you back US $2000, but if your wiring and woodworking skills are up to speed, you can put one together for around $600.

It's not like you have to build it from scratch. At least two companies, DIY for Life and DIY Projector Kits, sell kits and supplies that make this something any engineer could do in a weekend.

The key items are an LCD panel and a pair of Fresnel lenses. A Fresnel lens is thin and light, made out of a piece of plastic into which concentric grooves have been cut, giving it a large aperture and short focal length. (The first Fresnel lenses were invented for lighthouses.) It doesn't do a good job of projecting light, though, so you also need a different kind of compound lens, known as a triplet lens.

Here, the first Fresnel lens beams light from a source through the liquid-crystal display, which I'd stripped from an LCD panel. On the other side a second Fresnel focuses the image to a point on the third lens, which projects the image out onto a plane about 3 meters away.

That's the heart of the project; the rest is just the casing. The bulb gets quite hot, so a Lexan or tempered glass heat shield is placed right before the innermost lens. A fan at the back pulls cooling air down over the LCD panel from a slit in the cover.

I bought nearly everything from DIY Projector Kits: the lenses, heat shield, high-intensity light and ballast, a thermostat to run the fan, some hardware to mount the bulb, and a set of measured plans for a generic 15-inch display-based enclosure. The entire package, including shipping, ran just under $450.

I picked up a refurbished Samsung 15-inch LCD monitor for $124 on eBay. The Samsung had been rated by DIYers as very easy to strip, and sure enough, it took me less than 30 minutes to do it. I overbought on wood, probably using about $50 worth of high-quality 2-centimeter plywood.

The trickiest part was getting the light-path measurements right so that everything lined up. Then there was the adjustable focus ring. The instructions in the kit suggest starting with a 5-cm shower drain. Perhaps drains differ from region to region, but after a lot of carving with a Dremel hand tool and some plastic fumes inhaled, I finally did have a pretty nice adjustable focus.

I mounted the lens with a 1.25-cm vertical offset, to provide ventilation,

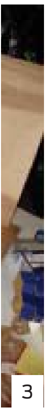

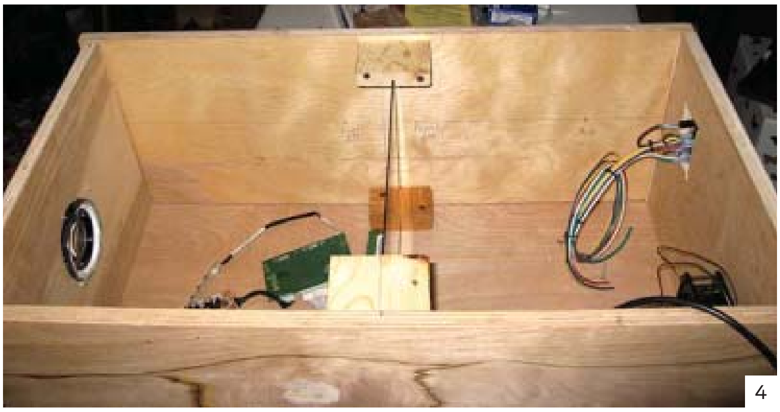

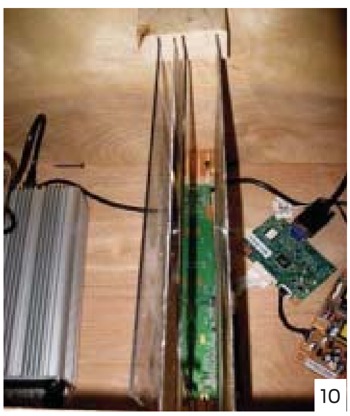

LIGHT SHOW: The heart of the digital movie projector is an LCD panel [1], which first needs to be stripped from its casing [2]. A triplet lens is mounted at one end of what will be the projector's wooden case [3]. The LCD panel is mounted in the middle. Note the lens and bulb positions marked [4, 5]. Here is the projector's interior, seen from the side, front, and rear [6, 7, 8]. When completed, the rear will house the bulb, ballast, and a thermostat for the fan [9], while the center will have the two Fresnel lenses, the LCD panel and control circuitry, and a heat shield [10]. At last! Lights, projector, action [11]!

PHOTOS: JAMES TURNER

and attached the bottom and right-side panels. Next I cut rear holes for the fan and plug/power-switch module and mounted them. Then I attached the rear panel, cutting grooves into two pieces of wood to space the middle lenses, the LCD panel, and the heat shield the correct distance from the front lens. Finally, I wired the electronics and the power to the LCD panel.

I tested the optics before adding the right panel, with more grooves to hold the tops of the lenses and the LCD. While I was installing the LCD panel, I twisted one of the mounts, cracking the display. It still functioned well enough for testing.

Finally, I put the top on, temporarily without a ventilation slit or hinges. I attached my MacBook to the LCD panel input and fired up the lamp. It worked perfectly the first time!

A few words of advice: Experience with optics isn't required but would be helpful. And besides the electronics, which involves little more than basic wiring, you'll need to be comfortable working with wood. Nothing was particularly difficult, but the measured plans leave a lot for you to figure out.

You should also be aware that at the end of the day you'll end up with a large, heavy appliance—a piece of furniture, in fact. The version I built measures 76 by 43 by 28 cm and weighs around 10 kilograms. Also, unless the LCD panel can handle video input, you'll have to drive the display from a computer. There's also no zoom—you control display size by moving the unit.

The project took me about 25 hours over the course of two weeks. I still have to hinge the cover, put in the ventilation slit, and seal all the light leaks. And, of course, I'll need to replace the cracked LCD panel.

During testing, the projector threw a very nice wall-filling image from my MacBook's Digital Visual Interface port. I don't have the equipment to measure the output, but 1500 to 1700 lumens is commonly reported for DIY projectors of this type. By contrast, a high-definition projector from Best Buy will put out around 1200 lm. It will also cost hundreds of dollars the first time you have to replace a bulb, which you'll have to do every 2000 to 4000 hours. The one I used should last around 10 000 hours, and it cost $50. Now if I only had a wall in my house big enough to use it with! An ideal projection area would be around 2.5 meters, measured diagonally. —James Turner

DIY For Life: http://lumenlab.com
DIY Projector Kits: http://diyprojectorkits.com

geek life

THE TRUTH ABOUT BENDER'S BRAIN

David X. Cohen, of "Futurama," reveals how MOS Technology's 6502 processor ended up in the robot's head

On 14 November 1999, an episode of "Futurama," the animated sci-fi comedy series conceived by "The Simpsons" creator Matt Groening, jolted computer geeks with a display of technological acumen absolutely unprecedented in prime-time entertainment. In the episode, "Fry and the Slurm Factory," a character named Professor Farnsworth points his F-ray at the head of the show's famously ill-tempered robot, Bender. It reveals a little rectangle, apparently a chip, labeled "6502."

The 6502 was a beloved—at least by geeks—8-bit microprocessor created by MOS Technology in 1975. It was the chip that the scruffy-bearded, sandal-wearing Steve Wozniak used to build the Apple II in 1977—"The Machine That Changed Everything," as PC World *once put it. It was also used in the Commodore PET, the BBC Micro, and a host of other systems that fomented the personal computer revolution.*

The chip's cameo on "Futurama" rocked the nerdosphere, prompting a burst of commentary in online forums (the nerd equivalent of cocktail party chatter). There was also a description on Wikipedia, in the site's entry about the 6502. But we here at IEEE Spectrum *had a few questions that somehow didn't come up amid all the noise. For starters: Why the 6502? And also: Is it possible that Bender's maker, Mom's Friendly Robot Co., somehow obtained Bender's design from a scruffy-bearded, sandal-wearing computer geek who lived in northern California circa the late 20th century?*

To get to the truth, Associate Editor Erico Guizzo sought the brain behind Bender's brain. That'd be David X. Cohen, executive producer and head writer of "Futurama," who, as it turns out, is quite a geek himself. Here's Cohen's response:

I SPENT A good percentage of my high school years programming the Apple II Plus in 6502 assembly language, so I have fond memories of long nights alone with this chip. My greatest 6502 achievement was a video game I called *Zoid* that was played heavily by me and my father and no one else. Incidentally, *Zoid* incorporated digitized speech (me saying the word "Zoid," slowed down to make it mightier), which was pretty rare at the time. The digital audio for that single syllable used much more memory than the entire program. I tried to sell the game to Broderbund Software, but I knew I was in for bad news when the return letter they sent me started with a misspelling of my name.

From a programming perspective, a more impressive feat was the creation of an actual working computer language ("FLEET") for the Apple II Plus that I developed with two high school friends, David Borden and David Schiminovich. We called ourselves "The Glitchmasters." This language was intended to make it easy to write high-speed graphics programs (that is, video games) for the Apple. None of us knew anything about compilers, yet using no references on the subject and working entirely in 6502 assembly, we somehow wrote a working compiler. It is even more impressive when you consider that virtually no comments appeared in the program—just page after page of assembly language.

In fact, the resulting compiler was extremely good: It was lightning fast, the language was easy to learn and program, and the compiled programs were comparable in speed to anything we would have written directly in assembly language. I believe it would have been an extremely useful product. However, our timing was astoundingly poor. The compiler was completed in 1984, just as the Apple II was fading forever into oblivion and we were heading off to college. Thus our fabulous compiler never really got used for anything.

In retrospect I would say the limitations of the 6502 forced us, against our wills, to be clever and learn the workings of things at a deeper level. For example, we had to write our own efficient subroutines to multiply and divide 16-bit numbers, using only 8-bit addition and subtraction and bit shifting. As another example, it is possible (in fact trivial) in computer graphics to compute the pixels along a line segment from (A, B) to (C, D) without using division or computing the slope. Again it requires cleverness, though. So I think programming

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

the 6502, especially in the days of limited computer memory, was very useful in terms of learning to think creatively and efficiently.

Moving now 15 years into the future to the year 1999, I was working on an early episode of "Futurama." Bender was being X-rayed (actually, "F-rayed"), and we needed to see what was powering his mighty robot brain. Naturally, the 6502.

While I can claim responsibility for the appearance of the 6502 in "Futurama," I was not the most highly trained computer scientist or engineer on the "Futurama" writing staff. I have a master's degree in computer science from the University of California, Berkeley. However, writer Ken Keeler has a Ph.D. from Harvard in applied math, as well as a master's from Stanford in electrical engineering (and yes, in all seriousness, Ken confirms that he does read every issue of *IEEE Spectrum* and occasionally looks at *Transactions on Information Theory*).

No doubt Woz's head still survives in a jar in the year 3000, and somehow it is probably wearing sandals. So it is quite possible that he provided Bender's design to Mom's Friendly Robot Co. in return for some extra fish food in his jar. ❑

tools & toys

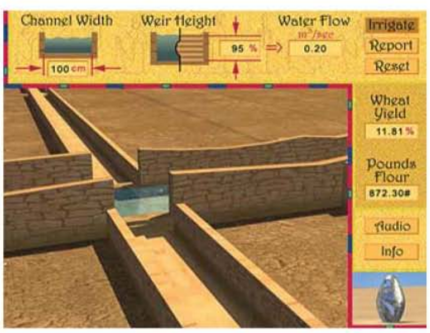

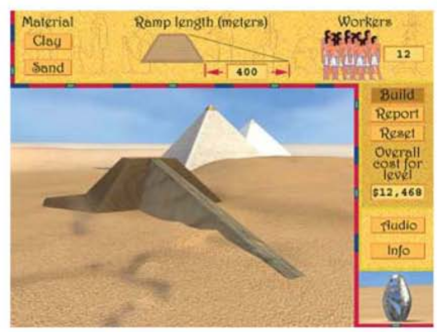

TIME WAITS FOR NO ENGINEER

A new computer game lets kids solve age-old engineering problems

How would you have arranged a radar array to protect London during the Blitz? How do you lower a drawbridge without crashing it, using only medieval technology? How can you build a pyramid out of massive stones using only sand, clay, and ropes? And how do you inveigle kids, 7 to 17, into tackling such questions?

Game designer Ray Shingler has solved all those puzzles with *Time Engineers*, a video game in which players ride a magnetic egg back in time to three eras. In each, players solve two engineering problems. As would-be electrical engineers, they create radar arrays; as mechanical engineers, they design a diesel submarine; and as civil engineers they build those drawbridges and pyramids.

Older kids can click to find mathematical formulas necessary to solve these problems as real engineers do. Coby, my 7-year-old, instead applied trial and error to the basic engineering principles provided. He does have the engineering mind-set though, at one point passionately insisting: "We are in the desert. Sand is cheap!"

I argued for clay as a more cost-effective, sturdier ramp material. It would have less give and require fewer workers to pull stones up. Point one to Mom.

Coby did, however, work out faster than Mom how wide a radar sweep we needed to save London from an air attack. I nailed the number of engines and the speed required to move a submarine through enemy patrols on the first shot. It took us about an hour and a half to play the game through, and we wanted more.

Shingler says *Time Engineers* has been used by hundreds of high school teachers around the country in their classrooms and by thousands of parents—lots of them engineers who want to introduce their kids to the family business.

Sure enough, Coby now thinks engineering is cool, even though *Time Engineers* is not perfect. It's too short, and the animations are dated. But it was fun and well worth both the time to play it and the US $19.95 it costs. Shingler says the money will go toward a flashier, longer version due out next year.

—Sherry Sontag

Time Engineers US $19.95; Software Kids, Valparaiso, Ind.; http://www.timeengineers.com *IMAGES: SOFTWAREKIDS*

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

careers

UNSQUEEZED

Slim Devices, maker of the music-playing Squeezebox, was bought out, but it hasn't sold out

You might expect a garage-born maker of a niche gadget to lose its fans, crash, and burn when it gets taken over by a large corporation. It didn't happen that way, though, to Slim Devices, the tiny company based in Mountain View, Calif., that created the Squeezebox, one of the first and best-reviewed digital-to-analog audio streamers.

The story begins in 2001. Apple's iPod was already selling up a storm, but Sean Adams, a twentysomething geek, had a different idea. He came up with the Slimp3, a high-quality digital-to-analog converter that piped music from an Ethernet connection to home stereos.

Adams assembled the first 100 units himself in his garage and started selling them online for US $249. His creation caught the eye of bloggers, magazine writers, and programmers, who began suggesting changes in code and design—for making a wireless version, for instance, or designing one that could handle just about any compression scheme. Thus was the Squeezebox created, and it, too, was a hit.

One reason for the success was the devotion of the fan base, which served as an unpaid legion of product designers, testers, and even customer service reps. Having them around enabled Adams to keep on developing the product without taking venture capital and hiring engineers.

The next Squeezebox, at $299, came out in 2003, and it still gives more bang for the buck than anything else in the now-expanding market. For audiophiles, the company offered a $2000 version called Transporter in 2006.

The Slim Devices designers wisely decided to keep the gadget dumb. The Squeezebox relies on your personal computer to process the music, meaning that the software for even the very first models can be updated regularly. And because that software is written by a passionate open-source development community, it works well and keeps on getting better.

In October 2006, Logitech bought the company for $20 million, with promises of larger payouts down the road. By now, Slim Devices wasn't just Sean—it was Sean and Dean Blackketter, an early volunteer who bought one of the first 100 devices churned out in Sean's garage, learned Perl so he could work on the software himself, and is now director of engineering. There was also "a marketing guy, a sales guy, a support guy, and a testing guy," Blackketter says.

But the company was still defined by its base of volunteer geeks, techies, and musicians. They had something to say. They had 40 pages of something to say.

"Well, I'm saddened. Once the MBAs, lawyers, and accountants get ahold of Slim, things will change," posted someone with the handle Jasen. "Congrats, Slim Services employees," wrote Bklass. "I can only hope this acquisition made your day, year, lifetime. Personally, my heart sinks on this news."

But high-quality products are still being rolled out, and the open-source, free-labor, this-is-our-gadget crowd is still there, blogging away. Indeed, the fan base recently tested, critiqued, and debugged the Boom, which has a self-contained set of speakers, and the Duet, which adds a freestanding controller to the Squeezebox, letting you play music without turning on your computer.

The Boom was vetted through a secret online forum of Slim Devices regulars, many of whom had already posted how they'd love a Squeezebox with speakers. The secret forum was new. So was a nondisclosure agreement, but the process otherwise remained the same—the fans worked for the chance to own and pick apart a new cool toy.

The takeover was a windfall for the 10 former volunteers—including Bklass—whom Logitech hired to telecommute from places as far-flung as Switzerland, the United Kingdom, Pennsylvania, Texas, and Indianapolis. "Their day job is that they get to do what they did for free," says Blackketter.

Has Squeezebox figured out a way to have the best of both worlds—the passion and free labor of open source and the corporate resources of a major consumer electronics maker? Maybe, but not without some pain. Grumpy volunteers sometimes got into scrapes with paying customers, for example, as they found themselves supporting more and more newbie users—too many for Logitech's official support channels to handle.

Meanwhile, the volunteers look with trepidation to 2010, when a second payout is due. The original Slim Devices crew could receive as much as $69.5 million more, depending on profits. Will they take the money and run? Will it mark the end of Squeezebox as open-source software? The fan base's fears have been filling the message board, enough to pull Adams away from the technical discussions he clearly prefers. With a written "sigh," he reiterated that Slim/Logitech will not go closed-source.

"As for anyone's plans," he wrote, "I don't even plan for next week. Still having fun here. Can't speak for anyone else." —Sherry Sontag

LOGITECH

reflections BY ROBERT W. LUCKY

Cloud Computing

THEY SAY that soon we will do our computing by renting memory and processing on a vast cloud of computers out in the network. They say this is an idea whose time has come. Of course "they" have said "the time has come" about other ideas, and some of those things didn't happen then, or sometimes ever. But this time it's for sure. They say.

I remember the old days of time-shared computing. Big, expensive computers were kept behind big glass walls, tended by shrouded acolytes. For the rest of us, it was "keep your hands off." You rented computation by the second for your dumb terminal.

The miracle of Moore's Law changed all that. Computing equipment has become so cheap, the junk piles are full of it. Who would have ever thought of a terabyte of data just lying around the house? So why should we go back to renting computational power, this time from a network cloud?

Apparently it's precisely because of Moore's Law. Processors and memory are cheap these days—and *really* cheap when you buy, run, and maintain them by the zillions. And the clouds at companies like Google and Amazon are living things, constantly being replaced and upgraded with the latest cheap stuff.

The advantages for us, the cloud's clients, include the ability to add more capacity for peak demand, to flexibly experiment with new services, and to remove unneeded capacity when demand slackens. Within the giant cloud, the laws of probability give the service provider great leverage through statistical multiplexing of all these varying loads. The cloud is also easier to manage—you can install a single software patch to cover all of a company's users, for example.

Those seem like compelling arguments, so what could be wrong with this picture? For one thing, despite the existence of encryption and access-control technologies, many organizations seem reluctant to put their precious, proprietary data out in a public cloud. There is a widely held impression that this might not be a good thing to be doing. Also, there is a practical and administrative problem—once all your petabytes of data are out there in the cloud, can you ever get them back?

While corporate IT organizations ponder the advantages and drawbacks, as an individual user I have a different problem—giving up my autonomy and control. I'll admit from the start that this is a psychological hang-up; nevertheless, it's there. I rather like having my own cycles and bytes; I like being able to tend my own computing garden. Basically, the question is: Do I trust these guys?

Over the years, I've always been both a system designer and a system user, and I've seen that the two sides are not always congruent. The designer feels that the best system is a closed one, where all the features and interfaces are under his control. I don't think that designers have evil intentions or even that they're always trying to make more money for their companies; often the urge to control is simply a part of the quest for engineering perfection.

No, it's not the designers I distrust so much as the businesspeople. I've seen the cellphone that takes pictures that can be sent only to my computer, over the contracted carrier, at a certain cost per picture; the built-in GPS that requires a subscription service; the e-book reader whose crisp electronic page can display only books that come through a particular vendor. All of these may be perfectly good systems, but they don't give me the flexibility I want. So why should I trust the cloud-computing service providers?

There is a counterargument, of course. The cloud advocates say that the open market will create open systems; that openness will give some service providers a competitive advantage. That may already be happening, as there are different offerings that put the user at both the bottom of the stack, where the access is effectively to a bare generic computer, and at the top, where the access is to customized applications.

So is cloud computing an idea whose time has come? It's your call. In the end, that's the way it works. ❑

JOHN HERSEY

The Million Dollar Programming Prize

Netflix's bounty for improving its movie-recommendation software is almost in the bag. Here is one team's account

BY **ROBERT M. BELL, JIM BENNETT, YEHUDA KOREN** & **CHRIS VOLINSKY**

IT'S 7:50 P.M. on 1 October 2007 at AT&T Labs, in Florham Park, N.J., and three of us are frantically hitting the "refresh" buttons on our browsers. We have just submitted our latest entry in the year-old Netflix Prize competition, which offers a grand prize of US $1 million for an algorithm that's 10 percent more accurate than the one Netflix uses to predict customers' movie preferences. Although we have not reached that milestone, we are hoping at least to do better than anyone else has done so far; if we can make it to 8 p.m. with the best score, we'll win a $50 000 Progress Prize. For most of the summer we'd been ahead of our nearest rivals by a comfortable margin, and as recently as 36 hours before this moment, our victory still seemed to be a slam dunk.

The previous day, though, the lead had started to slip away from us. First, the teams then in fifth and sixth places merged, combining their talents to vault into second place, making us nervous enough to submit our best effort, which we had been saving for a rainy day. But before our improved score appeared, we were hit by another combination when our two remaining serious rivals joined forces to tie us. Worse, their entry had come

RANDI SILBERMAN

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

72 seconds before ours, meaning that in the case of a tie, they'd win.

Seventy-two seconds! Could we lose this thing for being a minute too late? Then we realized that there were still 25 hours left, and we still had one more chance. We had to make it count.

We began to receive offers from other competitors to combine our scores with theirs. We politely declined them and planned strategies for our last submission. Sure enough, these bumped up our score by a few hundredths of a percent, at which point we could only wait to see the final score from our newly allied foes.

Refresh...refresh...refresh...

SINCE 1997, WHEN NETFLIX, of Los Gatos, Calif., came up with the idea of sending movie DVDs through the mail to subscribers, its customer base has grown to 10 million. That success stems, in part, from the company's giving quick and easy access to movies. But just as important is the Netflix recommender system, Cinematch, which helps customers get the most out of their memberships. The better the system predicts people's likes and dislikes, the more they'll enjoy the movies they watch and the longer they'll keep up their subscriptions.

As a new subscriber to Netflix, you have several ways to choose movies on the company's Web site. You can browse by genre or search by keyword for a title, actor, or director. After receiving your selections by mail or over the Internet and viewing them, you return to the site to order more movies. At any time, Cinematch lets you rate any movie you've seen by clicking on one to five stars.

As is the case with other recommender systems, such as those of Amazon, the Internet Movie Database, and Pandora, it's in the customer's interest to vote thoughtfully, because doing so helps Netflix figure out his or her tastes. Yet even if the customer declines to offer this feedback, Netflix still notes which movies the subscriber actually orders. After the customer has rated a handful of movies, the algorithm will start recommending titles based on the rating the algorithm predicts the customer will give.

RECOMMENDING MOVIES that customers will enjoy is a complex business. The easy part is gathering the data, which Netflix now accumulates at the rate of some 2 million ratings a day. Much tougher is to find patterns in the data that tell the company which movie a customer would enjoy, if only the person would watch it.

Netflix developed Cinematch to tackle this job using a well-known prediction technique (described below), which it designed to handle the billion or so ratings it had already logged. However, while incorporating so many ratings added to accuracy, it also took a lot of work just to keep up with the increasing scale, let alone to test alternative prediction schemes.

The Netflix researchers were nevertheless curious about the many other techniques for making similar predictions that had been published in the scholarly literature. The problem was that those studies had relied on public data sets containing on the order of a few million ratings, and it would take the small Netflix team a long time to explore how well these alternative techniques worked at a scale a thousand times as large. That is, if they did all the work themselves.

Reed Hastings, the chief executive of Netflix, suggested running a contest. He observed that Netflix had the means, the motive, and the opportunity. And the company had already staged a contest internally between the standard Cinematch system and an alternative algorithm, which did slightly, tantalizingly better.

The Netflix team came up with the basic structure of the contest. They would provide 100 million ratings that 480 000 anonymous customers had given to 17 000 movies. The data set would just fit in the main memory of a typical laptop, allowing almost anyone to compete. Netflix would withhold 3 million of the most recent ratings and ask the contestants to predict them. A

The Neighborhood Model

THE NEAREST-NEIGHBOR METHOD works on the principle that a person tends to give similar ratings to similar movies. Joe likes the three movies on the left, so to make a prediction for him, find users who also liked those movies and see what other movies they liked. Here the three other viewers all liked *Saving Private Ryan,* so that is the top recommendation. Two of them liked *Dune,* so that's ranked second, and so on.

PHOTOS: PARAMOUNT PICTURES

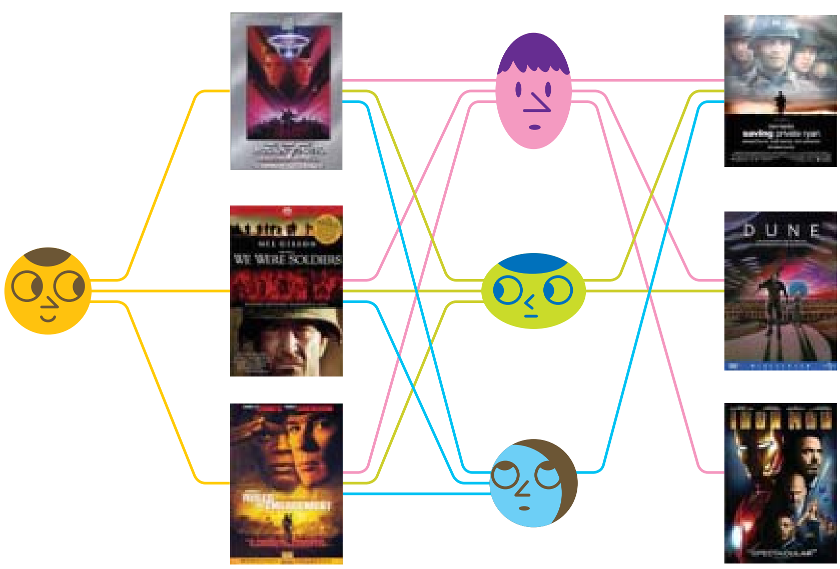

Netflix computer would assess each contestant's 3 million predictions by comparing predictions with actual ratings. The system would use the traditional metric for predictive accuracy, the root-mean squared error (RMSE). The more accurate a set of predictions, the smaller the RMSE will be. The score would then be reported back immediately to the contestant and reflected on an online leaderboard for all to see.

Each such scoring provides valuable information about the hidden ratings—so valuable, in fact, that under certain circumstances it could be used to game the system. The teams were therefore scored once a day, at most. But to help teams estimate how well they might do, Netflix also provided them each with a small representative data set and the score Cinematch had been able to attain for it. Contestants could use that set to test their systems as often as they wanted.

On 2 October 2006, Netflix launched the competition, and within days thousands of teams from hundreds of countries signed up. Within weeks the Netflix Prize Web site was getting hundreds of submissions per day. An online forum was created so that participants could share ideas and techniques, even code. Even more gratifying to Netflix, within months a handful of teams did several percent better than Cinematch. The question then was how much the accuracy would improve in the first year.

LIKE MOST OF THE other top competitors, the three of us at AT&T Labs consulted the rich body of research on ways of solving problems in this domain, known as collaborative filtering.

One of the main areas of collaborative filtering we exploited is the nearest-neighbor approach. A movie's "neighbors" in this context are other movies that tend to be scored most similarly when rated by the same viewer [see illustration, "The Neighborhood Model"]. For example, consider *Saving Private Ryan* (1998), a war movie directed by Steven Spielberg and starring Tom Hanks. Its neighbors may include other war movies, movies directed by Spielberg, or movies starring Tom Hanks. To predict a particular viewer's rating, we would look for the nearest neighbors to *Saving Private Ryan* that the viewer had already seen and rated. For some viewers, it may be easy to find a full allotment of close neighbors; for many others, we may discover only a handful of neighboring movies. Our version of the nearest-neighbor approach predicted ratings using a weighted average of the viewer's previous ratings on up to 50 neighboring movies. (We have since developed a way to use all past ratings, allowing an unlimited number of neighbors.)

A second area of collaborative-filtering research we pursued involves what are known as latent-factor models. These score both a given movie and a given viewer according to a set of factors, themselves inferred from patterns in the ratings given to all the movies by all the viewers [see illustration, "The Latent-Factor Approach"]. Factors for movies may measure comedy versus drama, action versus romance, and orientation to children versus orientation to adults. Because the factors are determined automatically by algorithms, they may correspond to hard-to-describe concepts such as quirkiness, or they may not be interpretable by humans at all. Factors for viewers measure how much the viewer likes movies

The Latent-Factor Approach

A SECOND, COMPLEMENTARY method scores both a given movie and viewer according to latent factors, themselves inferred from the ratings given to all the movies by all the viewers. The factors define a space that at once measures the characteristics of movies and the viewer's interest in those characteristics. Here we would expect the fellow in the southeast corner of the graph to love *Norbit*, to hate *Dreamgirls*, and, perhaps, to rate *Braveheart* about average.

PHOTOS: AMADEUS, *THE SAUL ZAENTZ COMPANY; ALL OTHERS, PARAMOUNT PICTURES*

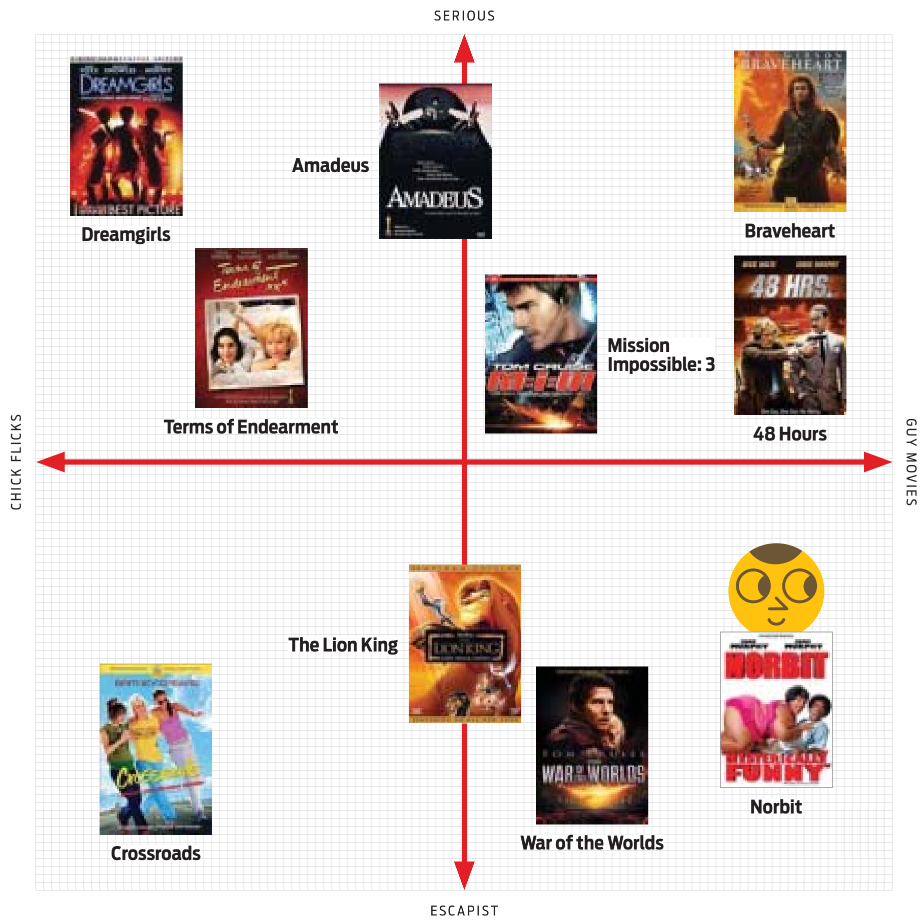

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

that score highly on the corresponding movie factor. Thus, a movie may be classified a comedy, and a viewer may be classified as a comedy lover.

The model may use 20 to 40 such factors to locate each movie and viewer in a multidimensional space. It then predicts a viewer's rating of a movie according to the movie's score on the dimensions that person cares about most. We can put these judgments in quantitative terms by taking the dot (or scalar) product of the locations of the viewer and the movie.

Both collaborative-filtering techniques work even if you don't know a thing about what's in the movies themselves. All you need to care about is how the viewers rate the movies. However, neither approach is a panacea. We found that most nearest-neighbor techniques work best on 50 or fewer neighbors, which means these methods can't exploit all the information a viewer's ratings may contain. Latent-factor models have the opposite weakness: They are bad at detecting strong associations among a few closely related films, such as *The Lord of the Rings* trilogy (2001–2003).

Because these two methods are complementary, we combined them, using many versions of each in what machine-learning experts call an ensemble approach. This allowed us to build systems that were simple and therefore easy to code and fast to run.

What's more, our ensemble approach was robust enough to protect against some of the problems that arise within the system's individual components. Indeed, the solution we had just submitted on 1 October 2007 was a linear combination of 107 separate sets of predictions, using many variations on the above themes and different tuning parameters. Even so, the biggest improvements in accuracy came from relatively few methods. The lesson here is that having lots of ways to skin this particular cat can be useful for gaining the incremental improvements needed to win competitions, but practically speaking, excellent systems can be built using just a few well-selected strategies.

We don't want to give the impression that heaping together a lot of methods was enough to reach the leaderboard. The Netflix Prize data set poses some huge challenges beyond its immense size. For one, there was enormous variation among viewers and among movies. A rating system must be sensitive enough to tease out subtle patterns associated with those few viewers who rated 1500 movies, without overfitting things—that is, expecting prolific raters' patterns to apply to the many more users who rated 15 or fewer movies. It can indeed be hard to make predictions for a viewer who has provided just a handful of ratings. We improved existing ad hoc methods, designed to address this concern rigorously.

Another critical innovation involved focusing on *which* movies a viewer rated, regardless of the scores. The idea is that someone who has rated a lot of fantasy movies will probably like *Lord of the Rings*, even if that person has rated the other movies in the cate-

Want Advice? Try an Expert

IF YOU want recommendations and you'd rather not rely exclusively on your customers, you can always do the daring thing and consult actual experts. That's the idea behind Pandora, a free Internet radio service that employs musicians to rate songs according to a checklist of criteria, such as pace, rhythm, even the voice of the performer.

The Oakland, Calif.–based company must be doing something right. In the three years since it opened shop, Pandora Media has registered 22 million listeners. So far, all of them are in the United States, although the company is negotiating its way back into Europe, which it left after having problems with music licenses there.

About 2 million people listen to the service on a given day, typically while sitting in front of their computers at work or, increasingly, while clutching their iPhones on the commute home, making for an average session of 6 hours. No wonder Pandora streams more data than any other site except YouTube.

Here's how it works. The listener creates a virtual "channel" by selecting a song, artist, or composer. If a song is chosen, the site compares it to its database of 600 000 songs, each rated by one of its musical experts. The site then selects *another* song it deems to be a close relative and keeps on playing such relatives. (Pandora can't give you your first choice because its licensing contracts ban it from playing songs to order.)

When I selected "A Hard Day's Night," by the Beatles, the first song I heard was "She Loves You," by the same band. I listened for a long time before getting my first choice.

You rate a song by clicking on either a thumbs-up or a thumbs-down icon, and the algorithm adjusts its weighting of the musical checklist it uses to select subsequent songs. What's more, a thumbs-down will keep the channel from ever playing the same song again. You have to be careful, because the more thumbs-down you give, the narrower the channel becomes, and in the extreme case you may "thumb yourself into a corner," says Tim Westergren, founder of Pandora.

Even then, he notes, you'd only hobble that one channel. Nothing you do on one channel affects the others, and you may create as many channels as you want.

Westergren says he got the idea for Pandora when he was a young musician working on scores with moviemakers who had very different likes and dislikes. He wanted to find a way to encode those differences in a database he dubbed the Music Genome, paying musicians to do the enormous amount of work.

It may seem strange to use so much manpower as a supplement to computer power, but it makes sense when humans alone can handle the job—a peculiar

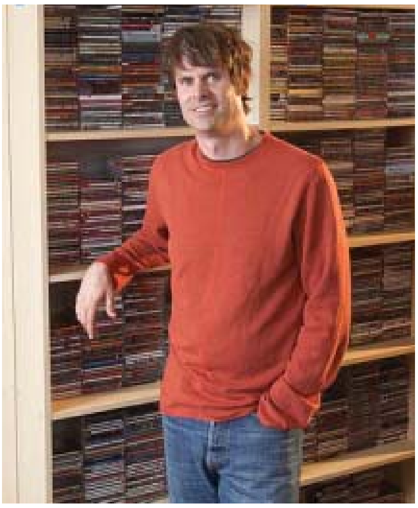

SONG CENTRAL: Tim Westergren and a few CDs waiting to be rated.
PHOTO: RAFAEL FUCHS

gory somewhat low. By replacing numerical scores with a binary who-watched-what score, the data set is transformed from one with mostly missing pieces (the cases in which users don't rate movies) to one that is completely full (using the value 1 when there are ratings and a 0 when there aren't). This approach nicely complemented our other methods.

FINALLY, AT 7:58 on that fateful October evening, all of the scores for the top teams were posted for the last time on the leaderboard, and ours came out highest, with an 8.43 percent improvement on Netflix's algorithm. Our nearest rival scored an 8.38 percent improvement. We didn't do well enough to land a million dollars, but still, we won.

While we've since come very close to the goal of 10 percent, experience has shown that each step forward is harder than the one that came before it, presumably because we've already exploited most of the clues in the data. Nonetheless, we continue to made progress. During 2008 we mined the data for information on how users' behavior changed over time. Later, we joined forces with another team to win the 2008 Progress Prize. Currently we stand at 9.63 percent improvement and are still working hard on the problem.

Now that the confetti has settled, we have a chance to look back on our work and to ask what this experience tells us. First, Netflix has incorporated our discoveries into an improved version of its algorithm, which is now being tested. Second, researchers are benefiting from the data set that the competition made available, and not just because it is orders of magnitude larger than previous data sets. It is also qualitatively better than other data sets, because Netflix gathered the information from paying customers, in a realistic setting. Third, the competition itself recruited many smart people in this line of research.

In any case, the new blood promises to quickly improve the state of the art. Such competitions have invigorated other fields. The various X Prizes that have been offered for advances in genomics, automotive design, and alternative energy have shown an excellent return: By some accounts the recent $10 million Ansari X prize, awarded for suborbital spaceflight, generated $100 million of private investment in space travel.

The competition also validates the concept of collective intelligence popularized in James Surowiecki's 2005 book *The Wisdom of Crowds* (Anchor Books). He argues that the sum of many independent votes is often superior to any one vote, even if made by the greatest expert. For Netflix, the vast number of independent ratings allows for surprisingly good predictions. The power of this collective intelligence is also being harnessed in, for instance, Amazon.com's product recommendations and the collaborative editing of the online encyclopedia, Wikipedia. With the rise of social networks on the Web, we can expect to see and contribute to even more powerful examples in the coming years. ❑

---

field sometimes called artificial artificial intelligence. One example of AAI is setting puzzles, or "captchas," for visitors to a Web site to solve, both to prove that they're human beings and not bots and to perform some useful chore, such as deciphering the blurred letters from a scan of an old book. Other AAI programs lure people to do such work by providing entertainment or, as Amazon's Mechanical Turk does, money.

Westergren got seed money for the Music Genome in 2000, at the very end of the dot-com bubble. When the bubble burst, he and his colleagues labored almost without income for five years before another injection of capital came through. Even now, Westergren says, Pandora is focused solely on growth and so does not turn a profit. It gets most of its revenue from the banner advertisements its site displays to listeners every time they click on something in the site, something they must do from time to time to prevent Pandora from going silent. It also gets a small royalty whenever a listener buys a song by clicking through to a vendor, such as iTunes or Amazon.

One advantage of using experts is that they can categorize songs that are new, by bands that are unknown. They can also provide a way to get at music that fell out of fashion before Internet rating became possible. Such too-new and too-old songs constitute a big part of the "long tail"—the huge inventory of items that each

MUSIC 101: Pandora offers a running commentary on its songs and artists.

sell in very small numbers yet collectively amount to a big part of the online marketplace. Mining that tail is one of the main jobs of any recommender system.

"In book publishing, genres are the equivalent of what we're doing. A brand-new author can say, 'Mine's a historical mystery novel,' and thus put data into the product without having any customer reviews," Westergren says. "But our theory is that it's not good data, not granular enough and not objective."

Tom Conrad, the chief technical officer of Pandora, says that "musical genomes" sometimes turn up connections you'd probably never get with other methods. He cites the '80s pop star Cindy Lauper, who recently recorded a new record that didn't sell in great numbers. "We analyzed it for its genome and found that the record sounds an awful lot like Norah Jones. So we are able to play Lauper's songs when you start a Norah Jones song. There's a Metallica ballad that's musically a nice fit for Indigo Girls. So start an Indigo Girls station and you might get this ballad."

Conrad says that Pandora isn't so proud of its expert-rated system that it can't learn from the collaborative-filtering techniques pioneered at Amazon, Apple, Netflix, and other firms. He contends that the two approaches are complementary.

"We have benefited by peering inside the approaches tried by some of the thinking that went into the Netflix Prize competition, and we've incorporated some of the ideas into our own system," Conrad says. "I'm friendly with the Netflix personalization team; we've talked over the past two years or so. We wanted to have more qualitative information; they wanted more quantitative. Now we both use both. Netflix has human editors who try to capture technical aspects of the movies."

When the two approaches meet, experts will use computers as much as computers use experts. We will have achieved the perfect chimera: a man-machine mind meld. —*Philip E. Ross*

25 MICROCHIPS THAT SHOOK THE WORLD

BY BRIAN R. SANTO

In microchip design, as in life, small things sometimes add up to big things. Dream up a clever microcircuit, get it sculpted in a sliver of silicon, and your little creation may unleash a technological revolution. It happened with the Intel 8088 microprocessor. And the Mostek MK4096 4-kilobit DRAM. And the Texas Instruments TMS32010 digital signal processor.

Among the many great chips that have emerged from fabs during the half-century reign of the integrated circuit, a small group stands out. Their designs proved so cutting-edge, so out of the box, so ahead of their time, that we are left groping for more technology clichés to describe them. Suffice it to say that they gave us the technology that made our brief, otherwise tedious existence in this universe worth living.

We've compiled here a list of 25 ICs that we think deserve the best spot on the mantelpiece of the house that Jack Kilby and Robert Noyce built. Some have become enduring objects of worship among the chiperati: the Signetics 555 timer, for example. Others, such as the Fairchild 741 operational amplifier, became textbook design examples. Some, like Microchip Technology's PIC microcontrollers, have sold billions, and are still doing so. A precious few, like Toshiba's flash memory, created whole new markets. And one, at least, became a geeky reference in popular culture. Question: What processor powers Bender, the alcoholic, chain-smoking, morally reprehensible robot in "Futurama"? Answer: MOS Technology's 6502.

What these chips have in common is that they're part of the reason why engineers don't get out enough.

Of course, lists like this are nothing if not contentious. Some may accuse us of capricious choices and blatant omissions (and, no, it won't be the first time). Why Intel's 8088 microprocessor and not the 4004 (the first) or the 8080 (the famed)? Where's the radiation-hardened military-grade RCA 1802 processor that was the brains of numerous spacecraft?

If you take only one thing away from this introduction, let it be this: Our list is what remained after weeks of raucous debate between the author, his trusted sources, and several editors of *IEEE Spectrum*. We never intended to compile an exhaustive reckoning of every chip that was a commercial success or a major technical advance. Nor could we include chips that were great but so obscure that only the five engineers who designed them would remember them. We focused on chips that proved unique, intriguing, awe-inspiring. We wanted chips of varied types, from both big and small companies, created long ago or more recently. Above all, we sought ICs that had an impact on the lives of lots of people—chips that became part of earth-shaking gadgets, symbolized technological trends, or simply delighted people.

For each chip, we describe how it came about and why it was innovative, with comments from the engineers and executives who architected it. And because we're not the *IEEE Annals of the History of Computing*, we didn't order the 25 chips chronologically or by type or importance; we arbitrarily scattered them on these pages in a way we think makes for a good read. History is messy, after all.

As a bonus, we asked eminent technologists about their favorite chips. Ever wonder which IC has a special place in the hearts of both Gordon Moore, of Intel, *and* Morris Chang, founder of Taiwan Semiconductor Manufacturing Company? (Hint: It's a DRAM chip.)

We also want to know what *you* think. Is there a chip whose absence from our list sent you into paroxysms of rage? Take a few deep breaths, have a nice cup of chamomile tea, and then go to http://www.spectrum.ieee.org/may09/25chips. There you can drop us a line and check out the runners-up that didn't make the list and more favorite picks by other luminaries.

Signetics NE555 Timer (1971)

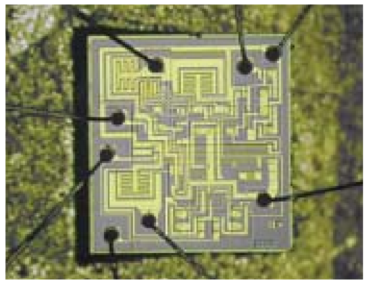

IT WAS THE SUMMER of 1970, and chip designer Hans Camenzind could tell you a thing or two about Chinese restaurants: His small office was squeezed between two of them in downtown Sunnyvale, Calif. Camenzind was working as a consultant to Signetics, a local semiconductor firm. The economy was tanking. He was making less than US $15 000 a year and had a wife and four children at home. He really needed to invent something good.

And so he did. One of the greatest chips of all time, in fact. The 555 was a simple IC that could function as a timer or an oscillator. It would become a best seller in analog semiconductors, winding up in kitchen appliances, toys, spacecraft, and a few thousand other things.

"And it almost didn't get made," recalls Camenzind, who at 75 is still designing chips, albeit nowhere near a Chinese restaurant.

The idea for the 555 came to him when he was working on a kind of system called a phase-locked loop. With some modifications, the circuit could work as a simple timer: You'd trigger it and it would run for a certain period. Simple as it may sound, there was nothing like that around.

At first, Signetics's engineering department rejected the idea. The company was already selling components that customers could use to make timers. That could have been the end of it. But Camenzind insisted. He went to Art Fury, Signetics marketing manager. Fury liked it.

Camenzind spent nearly a year testing breadboard prototypes, drawing the circuit components on paper, and cutting sheets of Rubylith—a masking film. "It was all done by hand, no computer," he says. His final design had 23 transistors, 16 resistors, and 2 diodes.

When the 555 hit the market in 1971, it was a sensation. In 1975 Signetics was absorbed by Philips Semiconductors, now NXP, which says that many billions have been sold. Engineers still use the 555 to create useful electronic modules—as well as less useful things like "Knight Rider"–style lights for car grilles.

Texas Instruments TMC0281 Speech Synthesizer (1978)

IF IT WEREN'T for the TMC0281, E.T. would've never been able to "phone home." That's because the TMC0281, the first single-chip speech synthesizer, was the heart (or should we say the mouth?) of Texas Instruments' Speak & Spell learning toy. In the Steven Spielberg movie, the flat-headed alien uses it to build his interplanetary communicator. (For the record, E.T. also uses a coat hanger, a coffee can, and a circular saw.)

The TMC0281 conveyed voice using a technique called linear predictive coding; the sound came out as a combination of buzzing, hissing, and popping. It was a surprising solution for something deemed "impossible to do in an integrated circuit," says Gene A. Frantz, one of the four engineers who designed the toy and is still at TI. Variants of the chip were used in Atari arcade games and Chrysler's K-cars. In 2001, TI sold its speech-synthesis chip line to Sensory, which discontinued it in late 2007. But if you ever need to place a long, very-long-distance phone call, you can find Speak & Spell units in excellent condition on eBay for about US $50.

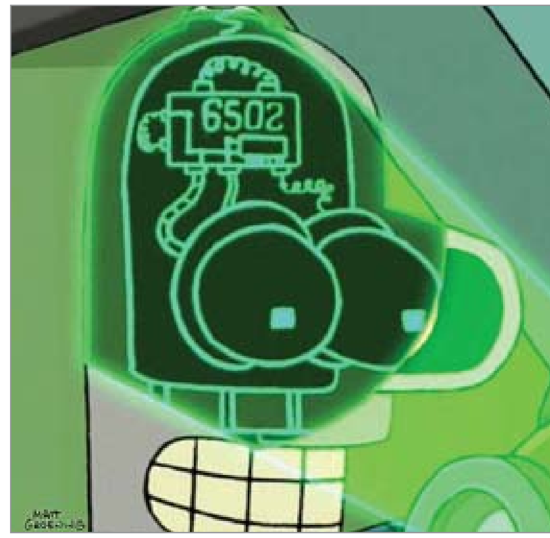

MOS Technology 6502 Microprocessor (1975)

WHEN THE CHUBBY-FACED geek stuck that chip on the computer and booted it up, the universe skipped a beat. The geek was Steve Wozniak, the computer was the Apple I, and the chip was the 6502, an 8-bit microprocessor developed by MOS Technology. The chip went on to become the main brains of ridiculously seminal computers like the Apple II, the Commodore PET, and the BBC Micro, not to mention game systems like the Nintendo and Atari. Chuck Peddle, one of the chip's creators, recalls when they introduced the 6502 at a trade show in 1975. "We had two glass jars filled with chips," he says, "and I had my wife sit there selling them." Hordes showed up. The reason: The 6502 wasn't just faster than its competitors—it was also way cheaper, selling for US $25 while Intel's 8080 and Motorola's 6800 were both fetching nearly $200.

The breakthrough, says Bill Mensch, who created the 6502 with Peddle, was a minimal instruction set combined with a fabrication process that "yielded 10 times as many good chips as the competition." The 6502 almost single-handedly forced the price of processors to drop, helping launch the personal computer revolution. Some embedded systems still use the chip. More interesting perhaps, the 6502 is the electronic brain of Bender, the depraved robot in "Futurama," as revealed in a 1999 episode.

[See "The Truth About Bender's Brain," in this issue, where David X. Cohen, the executive producer and head writer for "Futurama," explains how the choice of the 6502 came about.]

Texas Instruments TMS32010 Digital Signal Processor (1983)

THE BIG STATE OF TEXAS has given us many great things, including the 10-gallon hat, chicken-fried steak, Dr Pepper, and perhaps less prominently, the TMS32010 digital signal processor chip. Created by Texas Instruments, the TMS32010 wasn't the first DSP (that'd be Western Electric's DSP-1, introduced in 1980), but it was surely the fastest. It could compute a multiply operation in 200 nanoseconds, a feat that made engineers all tingly. What's more, it could execute instructions from both on-chip ROM and off-chip RAM, whereas competing chips had only canned DSP functions. "That made program development [for the TMS32010] flexible, just like with microcontrollers and microprocessors," says Wanda Gass, a member of the DSP design team, who is still at TI. At US $500 apiece, the chip sold about 1000 units the first year. Sales eventually ramped up, and the DSP became part of modems, medical devices, and military systems. Oh, and another application: Worlds of Wonder's Julie, a Chucky-style creepy doll that could sing and talk ("Are we making too much noise?"). The chip was the first in a large DSP family that made—and continues to make—TI's fortune.

Microchip Technology PIC 16C84 Microcontroller (1993)

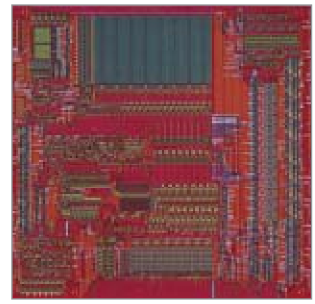

BACK IN THE EARLY 1990s, the huge 8-bit microcontroller universe belonged to one company, the almighty Motorola. Then along came a small contender with a nondescript name, Microchip Technology. Microchip developed the PIC 16C84, which incorporated a type of memory called EEPROM, for electrically erasable programmable read-only memory. It didn't need UV light to be erased, as did its progenitor, EPROM. "Now users could change their code on the fly," says Rod Drake, the chip's lead designer and now a director at Microchip. Even better, the chip cost less than US $5, or a quarter the cost of existing alternatives, most of them from, yes, Motorola. The 16C84 found use in smart cards, remote controls, and wireless car keys. It was the beginning of a line of microcontrollers that became electronics superstars among Fortune 500 companies and weekend hobbyists alike. Some 6 billion have been sold, used in things like industrial controllers, unmanned aerial vehicles, digital pregnancy tests, chip-controlled fireworks, LED jewelry, and a septic-tank monitor named the Turd Alert.

Fairchild Semiconductor µA741 Op-Amp (1968)

OPERATIONAL AMPLIFIERS are the sliced bread of analog design. You can always use some, and you can slap them together with almost anything and get something satisfying. Designers use them to make audio and video preamplifiers, voltage comparators, precision rectifiers, and many other systems that are part of everyday electronics.

In 1963, a 26-year-old engineer named Robert Widlar designed the first monolithic op-amp IC, the µA702, at Fairchild Semiconductor. It sold for US $300 a pop. Widlar followed up with an improved design, the µA709, cutting the cost to $70 and making the chip a huge commercial success. The story goes that the freewheeling Widlar asked for a raise. When he didn't get it, he quit. National Semiconductor was

Intersil ICL8038 Waveform Generator (circa 1983*)

CRITICS SCOFFED at the ICL8038's limited performance and propensity for behaving erratically. The chip, a generator of sine, square, triangular, sawtooth, and pulse waveforms, was indeed a bit temperamental. But engineers soon learned how to use the chip reliably, and the 8038 became a major hit, eventually selling into the hundreds of millions and finding its way into countless applications—like the famed Moog music synthesizers and the "blue boxes" that "phreakers" used to beat the phone companies in the 1980s. The part was so popular the company put out a document titled "Everything You Always Wanted to Know About the ICL8038." Sample question: "Why does connecting pin 7 to pin 8 give the best temperature performance?" Intersil discontinued the 8038 in 2002, but hobbyists still seek it today to make things like homemade function generators and theremins.

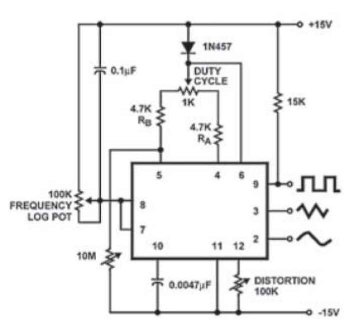

* *Neither Intersil's PR department nor the company's last engineer working with the part knows the precise introduction date. Do you?*

CLOCKWISE FROM TOP LEFT: JANET M. BAKER; MICROCHIP TECHNOLOGY; INTERSIL

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

Western Digital WD1402A UART (1971)

GORDON BELL is famous for launching the PDP series of minicomputers at Digital Equipment Corp. in the 1960s. But he also invented a lesser known but no less significant piece of technology: the universal asynchronous receiver/transmitter, or UART. Bell needed some circuitry to connect a Teletype to a PDP-1, a task that required converting parallel signals into serial signals and vice versa. His implementation used some 50 discrete components. Western Digital, a small company making calculator chips, offered to create a single-chip UART. Western Digital founder Al Phillips still remembers when his vice president of engineering showed him the Rubylith sheets with the design, ready for fabrication. "I looked at it for a minute and spotted an open circuit," Phillips says. "The VP got hysterical." Western Digital introduced the WD1402A around 1971, and other versions soon followed. Now UARTs are widely used in modems, PC peripherals, and other equipment.

Acorn Computers ARM1 Processor (1985)

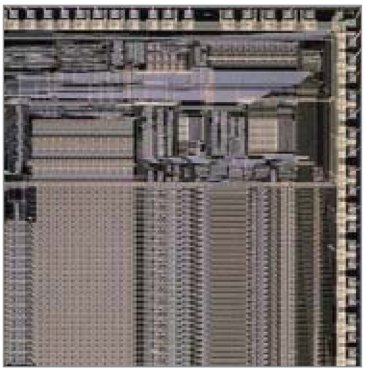

IN THE EARLY 1980s, Acorn Computers was a small company with a big product. The firm, based in Cambridge, England, had sold over 1.5 million BBC Micro desktop computers. It was now time to design a new model, and Acorn engineers decided to create their own 32-bit microprocessor. They called it the Acorn RISC Machine, or ARM. The engineers knew it wouldn't be easy; in fact, they half expected they'd encounter an insurmountable design hurdle and have to scrap the whole project. "The team was so small that every design decision had to favor simplicity—or we'd never finish it!" says codesigner Steve Furber, now a computer engineering professor at the University of Manchester. In the end, the simplicity made all the difference. The ARM was small, low power, and easy to program. Sophie Wilson, who designed the instruction set, still remembers when they first tested the chip on a computer. "We did 'PRINT PI' at the prompt, and it gave the right answer," she says. "We cracked open the bottles of champagne." In 1990, Acorn spun off its ARM division, and the ARM architecture went on to become the dominant 32-bit embedded processor. More than 10 billion ARM cores have been used in all sorts of gadgetry, including one of Apple's most humiliating flops, the Newton handheld, and one of its most glittering successes, the iPhone.

only too happy to hire a guy who was then helping establish the discipline of analog IC design. In 1967, Widlar created an ever better op-amp for National, the LM101.

While Fairchild managers fretted about the sudden competition, over at the company's R&D lab a recent hire, David Fullagar, scrutinized the LM101. He realized that the chip, however brilliant, had a couple of drawbacks. To avoid certain frequency distortions, engineers had to attach an external capacitor to the chip. What's more, the IC's input stage, the so-called front end, was for some chips overly sensitive to noise, because of quality variations in the semiconductors.

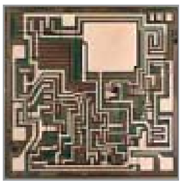

"The front end looked kind of kludgy," he says.

Fullagar embarked on his own design. He stretched the limits of semiconductor manufacturing processes at the time, incorporating a 30-picofarad capacitor into the chip. Now, how to improve the front end? The solution was profoundly simple—"it just came to me, I don't know, driving to Tahoe"—and consisted of a couple of extra transistors. That additional circuitry made the amplification smoother and consistent from chip to chip.

Fullagar took his design to the head of R&D at Fairchild, a guy named Gordon Moore, who sent it to the company's commercial division. The new chip, the µA741, would become *the* standard for op-amps. The IC—and variants created by Fairchild's competitors—have sold in the hundreds of millions. Now, for $300—the price tag of that primordial 702 op-amp—you can get about a thousand of today's 741 chips.

Kodak KAF-1300 Image Sensor (1986)

LAUNCHED IN 1991, the Kodak DCS 100 digital camera cost as much as US $13 000 and required a 5-kilogram external data storage unit that users had to carry on a shoulder strap. The sight of a person lugging the contraption? *Not* a Kodak moment. Still, the camera's electronics—housed inside a Nikon F3 body—included one impressive piece of hardware: a thumbnail-size chip that could capture images at a resolution of 1.3 megapixels, enough for sharp 5-by-7-inch prints. "At the time, 1 megapixel was a magic number," says Eric Stevens, the chip's lead designer, who still works for Kodak. The chip—a true two-phase charge-coupled device—became the basis for future CCD sensors, helping to kick-start the digital photography revolution. What, by the way, was the very first photo made with the KAF-1300? "Uh," says Stevens, "we just pointed the sensor at the wall of the laboratory."

FROM TOP: ACORN COMPUTERS; DAVID FULLAGAR; KODAK

My Favorite CHIP

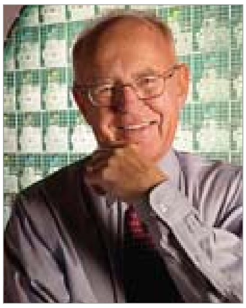

Gordon Moore *cofounder and chairman emeritus of Intel*

There were lots of great chips, but one that will always be dear to me was the **Intel 1103, the first commercial 1024-bit DRAM** [introduced in 1970]. It was the chip that really got Intel over the hump to profitability. It was not the most elegant design, having many of the problems that memory engineers had become familiar with in core memories. But this was comforting to such engineers—it meant that their expertise was not going to be made obsolete by the new technology. Even today, when I look at my digital watch and see 11:03, I cannot help but remember this key product in Intel's history.

Carver Mead *professor of engineering and applied science, Caltech*

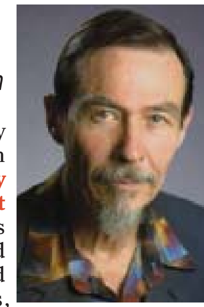

My favorite chip contained only a single transistor, although a remarkable one: a **Shottky barrier gate field-effect transistor** made from GaAs (later called the MESFET and now, using more advanced semiconductor structures, the HEMT). I designed it over Thanksgiving break in 1965. The very high mobility of the III-V materials, together with the absence of minority-carrier storage effects, made these devices far superior for a microwave power-output stage. Despite my efforts to interest American companies, the Japanese were the first to develop these devices. They have made microwave communications in satellites, cellphones, and many other systems possible for many decades.

Lee Felsenstein *computer pioneer*

The humble, ubiquitous, and quite inexpensive **Signetics 555 timer** had a big impact on my career, when I found myself using it to craft various pulse-sequencing circuits, baud-rate oscillators, and ramp generators for my

earliest consulting clients. With a latch triggered by two comparators, a high-current-drive output and a separate open-collector output for discharging capacitors, the 555 offered a wide range of uses with the addition of a few analog components. It suggested that it could provide more than it delivered at times, which helped me refine my understanding of the limits of a chip and the need to consider its operation in the analog domain.

Jeff Hawkins *founder of Palm and Numenta*

My personal favorite—a chip that opened my eyes to what was possible—was the **Intel 2716 EPROM**, vintage late 1970s. The 2716 was nonvolatile, held 2 kilobytes of memory, and, unlike many other chips, used a single 5-volt power supply, so it could be used as storage in a small and practical package. Of course, its big disadvantage was that you needed a UV light to erase it. But with a little imagination you could see that one day it could all be done electrically. In some sense the 2716 is the great-grandparent of today's flash memory chips.

Sophie Vandebroek *chief technology officer, Xerox*

My selection is **Analog Devices' iMEMS accelerometer**, the first commercial chip to significantly integrate MEMS and logic circuitry. Commercialized in the early 1990s, it revolutionized the automotive air-bag industry—and saved lives! Today this type of accelerometer is used in a variety of applications, including the Nintendo Wii and the Apple iPhone. Other types of MEMS chips are more and more present in a broad variety of applications.

Steve Jurvetson *managing director of Draper Fisher Jurvetson*

The **Motorola 68000** was a special one. I built a speaking computer with it. I also wrote a multitasker in assembler. That chip was the workhorse of learning back during my M.S.E.E. days. But wait—other chips come to mind. I also like ZettaCore's first 1-megabyte molecular memory chip, D-Wave's solid-state 128-qubit quantum computer processor, the Pentium (I've got an 8-inch wafer signed by Andy Grove in my office), the Canon EOS 5D 12.8-megapixel CMOS sensor, and the iTV 5-bit asynchronous processor running Forth and a custom OS—whew!

Morris Chang *founder of TSMC*

One of my favorite great microchips is the **Intel 1103, 1-kilobit DRAM**, circa 1970. Reasons: one, huge commercial success; two, started Intel on its way; three, demonstrated the power of MOS technology (versus bipolar); and four, opened up at least another 40 years of life for Moore's Law.

CLOCKWISE FROM TOP LEFT: INTEL; IMPINJ; KHOSLA VENTURES; DAVID DITZEL; LEINO OLE; TSMC; LEE FELSENSTEIN; NUMENTA; XEROX

Vinod Khosla

cofounder of Sun Microsystems and partner at Kleiner, Perkins, Caufield & Byers

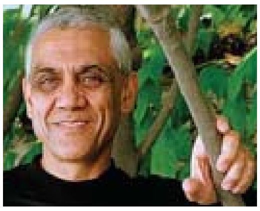

The **Motorola 68010** was my favorite chip. With its power and virtual memory support, it said to the world, "Microprocessors can stand and compete with the big boys—the minis and mainframes."

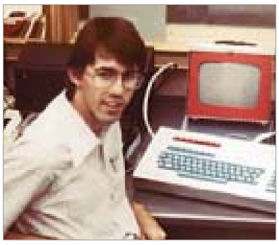

David Ditzel

Intel's chief architect for hybrid parallel computing

My favorite chip: the **6502 from MOS Technology**. This 8-bit microprocessor was used by me and many hobbyists for building our own computers [above, me and my homemade PC in 1977]. I wrote a small operating system that fit in 4 kilobytes, and Tiny BASIC from Tom Pittman fit in 2 KB. The reason it was a great chip was that in the 8-bit era, you could use the first 256 bytes of memory as 128 16-bit pointers to index from, making this machine much easier to program than other 8-bit alternatives. It was a big deal when you had to enter your programs in hexadecimal binary form.

IBM Deep Blue 2 Chess Chip (1997)

ON ONE SIDE of the board, 1.5 kilograms of gray matter. On the other side, 480 chess chips. Humans finally fell to computers in 1997, when IBM's chess-playing computer, Deep Blue, beat the reigning world champion, Garry Kasparov. Each of Deep Blue's chips consisted of 1.5 million transistors arranged into specialized blocks—like a move-generator logic array—as well as some RAM and ROM. Together, the chips could churn through 200 million chess positions per second. That brute-force power, combined with clever game-evaluation functions, gave Deep Blue decisive—Kasparov called them "uncomputerlike"—moves. "They exerted great psychological pressures," recalls Deep Blue's mastermind, Feng-hsiung Hsu, now at Microsoft.

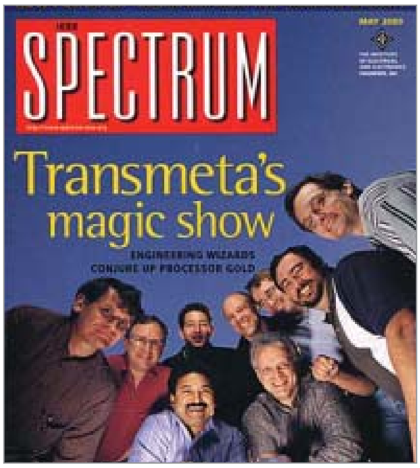

Transmeta Corp. Crusoe Processor (2000)

WITH GREAT POWER come great heat sinks. And short battery life. And crazy electricity consumption. Hence Transmeta's goal of designing a low-power processor that'd put those hogs offered by Intel and AMD to shame. The plan: Software would translate *x*86 instructions on the fly into Crusoe's own machine code, whose higher level of parallelism would save time and power. It was hyped as the greatest thing since sliced silicon, and for a while, it was. "Engineering Wizards Conjure Up Processor Gold" was how *IEEE Spectrum*'s May 2000 cover put it. Crusoe and its successor, Efficeon, "proved that dynamic binary translation was commercially viable," says David Ditzel, Transmeta's cofounder, now at Intel. Unfortunately, he adds, the chips arrived several years before the market for low-power computers took off. In the end, while Transmeta did not deliver on its promises, it did force Intel and AMD—through licenses and lawsuits—to chill out.

Texas Instruments Digital Micromirror Device (1987)

ON 18 JUNE 1999, Larry Hornbeck took his wife, Laura, on a date. They went to watch *Star Wars: Episode 1—The Phantom Menace* at a theater in Burbank, Calif. Not that the graying engineer was an avid Jedi fan. The reason they were there was actually the projector. It used a chip—the digital micromirror device—that Hornbeck had invented at Texas Instruments. The chip uses millions of hinged microscopic mirrors to direct light through a projection lens. The screening was "the first digital exhibition of a major motion picture," says Hornbeck, a TI Fellow. Now movie projectors using this digital light-processing technology—or DLP, as TI branded it—are used in thousands of theaters. It's also used in rear-projection TVs, office projectors, and tiny projectors for cellphones. "To paraphrase Houdini," Hornbeck says, "micromirrors, gentlemen. The effect is created with micromirrors."

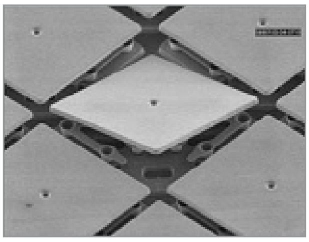

FROM TOP: ADAM NADEL/AP PHOTO; IEEE SPECTRUM; TEXAS INSTRUMENTS

Intel 8088 Microprocessor (1979)

WAS THERE ANY one chip that propelled Intel into the Fortune 500? Intel says there was: the 8088. This was the 16-bit CPU that IBM chose for its original PC line, which went on to dominate the desktop computer market.

In an odd twist of fate, the chip that established what would become known as the *x*86 architecture didn't have a name appended with an "86." The 8088 was basically just a slightly modified 8086, Intel's first 16-bit CPU. Or as Intel engineer Stephen Morse once put it, the 8088 was "a castrated version of the 8086." That's because the new chip's main innovation wasn't exactly a step forward in technical terms: The 8088 processed data in 16-bit words, but it used an 8-bit external data bus.

Intel managers kept the 8088 project under wraps until the 8086 design was mostly complete. "Management didn't want to delay the 8086 by even a day by even telling us they had the 8088 variant in mind," says Peter A. Stoll, a lead engineer for the 8086 project who did some work on the 8088—a "one-day task force to fix a microcode bug that took three days."

It was only after the first functional 8086 came out that Intel shipped the 8086 artwork and documentation to a design unit in Haifa, Israel, where two engineers, Rafi Retter and Dany Star, altered the chip to an 8-bit bus.

The modification proved to be one of Intel's best decisions. The 29 000-transistor 8088 CPU required fewer, less expensive support chips than the 8086 and had "full compatibility with 8-bit hardware, while also providing faster processing and a smooth transition to 16-bit processors," as Intel's Robert Noyce and Ted Hoff wrote in a 1981 article for *IEEE Micro* magazine.

The first PC to use the 8088 was IBM's Model 5150, a monochrome machine that cost US $3000. Now almost all the world's PCs are built around CPUs that can claim the 8088 as an ancestor. Not bad for a castrated chip.

Micronas Semiconductor MAS3507 MP3 Decoder (1997)

BEFORE THE iPOD, there was the Diamond Rio PMP300. Not that you'd remember. Launched in 1998, the PMP300 became an instant hit, but then the hype faded faster than Milli Vanilli. One thing, though, was notable about the player. It carried the MAS3507 MP3 decoder chip—a RISC-based digital signal processor with an instruction set optimized for audio compression and decompression. The chip, developed by Micronas, let the Rio squeeze about a dozen songs onto its flash memory—laughable today but at the time just enough to compete with portable CD players. Quaint, huh? The Rio and its successors paved the way for the iPod, and now you can carry thousands of songs—and all of Milli Vanilli's albums and music videos—in your pocket.

Mostek MK4096 4-Kilobit DRAM (1973)

MOSTEK WASN'T THE FIRST to put out a DRAM. Intel was. But Mostek's 4-kilobit DRAM chip brought about a key innovation, a circuitry trick called address multiplexing, concocted by Mostek cofounder Bob Proebsting. Basically, the chip used the same pins to access the memory's rows and columns by multiplexing the addressing signals. As a result, the chip wouldn't require more pins as memory density increased and could be made for less money. There was just a little compatibility problem. The 4096 used 16 pins, whereas the memories made by Texas Instruments, Intel, and Motorola had 22 pins. What followed was one of the most epic face-offs in DRAM history. With Mostek betting its future on the chip, its executives set out to proselytize customers, partners, the press, and even its staff. Fred K. Beckhusen, who as a recent hire was drafted to test the 4096 devices, recalls when Proebsting and chief executive L.J. Sevin came to his night shift to give a seminar—at 2 a.m. "They boldly predicted that in six months no one would hear or care about 22-pin DRAM," Beckhusen says. They were right. The 4096 and its successors became the dominant DRAM for years.

CLOCKWISE FROM TOP LEFT: INTEL CORP.; EIRIK SOLHEIM/EIRIKSO.COM; MOSTEK

Xilinx XC2064 FPGA (1985)

BACK IN THE EARLY 1980s, chip designers tried to get the most out of each and every transistor on their circuits. But then Ross Freeman had a pretty radical idea. He came up with a chip packed with transistors that formed loosely organized logic blocks that in turn could be configured and reconfigured with software. Sometimes a bunch

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

Zilog Z80 Microprocessor (1976)

FEDERICO FAGGIN knew well the kind of money and man-hours it took to market a microprocessor. While at Intel, he had contributed to the designs of two seminal specimens: the primordial 4004, and the 8080, of Altair fame. So when he founded Zilog with former Intel colleague Ralph Ungermann, they decided to start with something simpler: a single-chip microcontroller.

Faggin and Ungermann rented an office in downtown Los Altos, Calif., drafted a business plan, and went looking for venture capital. They ate lunch at a nearby Safeway supermarket—"Camembert cheese and crackers," he recalls.

But the engineers soon realized that the microcontroller market was crowded with very good chips. Even if theirs was better than the others, they'd see only slim profits—and continue lunching on cheese and crackers. Zilog had to aim higher on the food chain, so to speak, and the Z80 microprocessor project was born.

The goal was to outperform the 8080 and also offer full compatibility with 8080 software, to lure customers away from Intel. For months, Faggin, Ungermann, and Masatoshi Shima, another ex-Intel engineer, worked 80-hour weeks hunched over tables, drawing the Z80 circuits. Faggin soon learned that when it comes to microchips, small is beautiful—but it can hurt your eyes.

"By the end I had to get glasses," he says. "I became nearsighted."

The team toiled through 1975 and into 1976. In March of that year, they finally had a prototype chip. The Z80 was a contemporary of MOS Technology's 6502, and like that chip, it stood out not only for its elegant design but also for being dirt cheap (about US $25). Still, getting the product out the door took a lot of convincing. "It was just an intense time," says Faggin, who developed an ulcer as well.

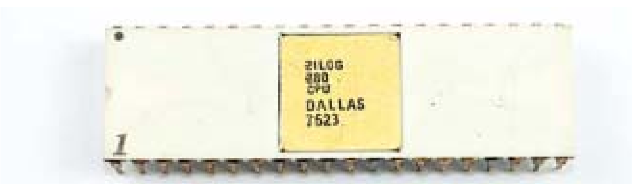

But the sales eventually came through. The Z80 ended up in thousands of products, including the Osborne I (the first portable, or "luggable," computer), and the Radio Shack TRS-80 and MSX home computers, as well as printers, fax machines, photocopiers, modems, and satellites. Zilog still makes the Z80, which is popular in some embedded systems. In a basic configuration today it costs $5.73—not even as much as a cheese-and-crackers lunch.

LEFT: CPU-WORLD.COM; RIGHT: ROBERT GARNER

Sun Microsystems SPARC Processor (1987)

THERE WAS A TIME, long ago (the early 1980s), when people wore neon-colored leg warmers and watched "Dallas," and microprocessor architects sought to increase the complexity of CPU instructions as a way of getting more accomplished in each compute cycle. But then a group at the University of California, Berkeley, always a bastion of counterculture, called for the opposite: Simplify the instruction set, they said, and you'll process instructions at a rate so fast you'll more than compensate for doing less each cycle. The Berkeley group, led by David Patterson, called their approach RISC, for reduced-instruction-set computing.

As an academic study, RISC sounded great. But was it marketable? Sun Microsystems bet on it. In 1984, a small team of Sun engineers set out to develop a 32-bit RISC processor called SPARC (for Scalable Processor Architecture). The idea was to use the chips in a new line of workstations. One day, Scott McNealy, then Sun's CEO, showed up at the SPARC development lab. "He said that SPARC would take Sun from a $500-million-a-year company to a billion-dollar-a-year company," recalls Patterson, a consultant to the SPARC project.

If that weren't pressure enough, many outside Sun had expressed doubt the company could pull it off. Worse still, Sun's marketing team had had a terrifying realization: SPARC spelled backward was...CRAPS! Team members had to swear they would not utter that word to anyone even inside Sun—lest the word get out to archrival MIPS Technologies, which was also exploring the RISC concept.

The first version of the minimalist SPARC consisted of a "20 000-gate-array processor without even integer multiply/divide instructions," says Robert Garner, the lead SPARC architect and now an IBM researcher. Yet, at 10 million instructions per second, it ran about three times as fast as the complex-instruction-set computer (CISC) processors of the day.

Sun would use SPARC to power profitable workstations and servers for years to come. The first SPARC-based product, introduced in 1987, was the Sun-4 line of workstations, which quickly dominated the market and helped propel the company's revenues past the billion-dollar mark—just as McNealy had prophesied.

---

of transistors wouldn't be used—heresy!—but Freeman was betting that Moore's Law would eventually make transistors really cheap. It did. To market the chip, called a field-programmable gate array, or FPGA, Freeman cofounded Xilinx. (Apparently, a weird concept called for a weird company name.) When the company's first product, the XC2064, came out in 1985, employees were given an assignment: They had to draw by hand an example circuit using XC2064's logic blocks, just as Xilinx customers would. Bill Carter, a former chief technology officer, recalls being approached by CEO Bernie Vonderschmitt, who said he was having "a little difficulty doing his homework." Carter was only too happy to help the boss. "There we were," he says, "with paper and colored pencils, working on Bernie's assignment!" Today FPGAs—sold by Xilinx and others—are used in just too many things to list here. Go reconfigure!

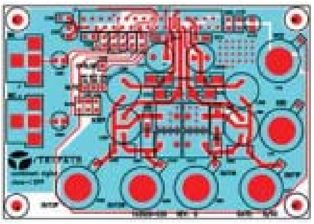

Tripath Technology TA2020 Audio Amplifier (1998)

THERE'S A SUBSET of audiophiles who insist that vacuum tube–based amplifiers produce the best sound and always will. So when some in the audio community claimed that a solid-state class-D amp concocted by a Silicon Valley company called Tripath Technology delivered sound as warm and vibrant as tube amps, it was a big deal. Tripath's trick was to use a 50-megahertz sampling system to drive the amplifier. The company boasted that its TA2020 performed better and cost much less than any comparable solid-state amp. To show off the chip at trade shows, "we'd play that song—that very romantic one from *Titanic*," says Adya Tripathi, Tripath's founder. Like most class-D amps, the 2020 was very power efficient; it didn't require a heat sink and could use a compact package. Tripath's low-end, 15-watt version of the TA2020 sold for US $3 and was used in boom boxes and ministereos. Other versions—the most powerful had a 1000-W output—were used in home theaters, high-end audio systems, and TV sets by Sony, Sharp, Toshiba, and others. Eventually, the big semiconductor companies caught up, creating similar chips and sending Tripath into oblivion. Its chips, however, developed a devoted cult following. Audio-amp kits and products based on the TA2020 are still available from such companies as 41 Hz Audio, Sure Electronics, and Winsome Labs.

Amati Communications Overture ADSL Chip Set (1994)

REMEMBER WHEN DSL came along and you chucked that pathetic 56.6-kilobit-per-second modem into the trash? You and the two-thirds of the world's broadband users who use DSL should thank Amati Communications, a start-up out of Stanford University. In the 1990s, it came up with a DSL modulation approach called discrete multitone, or DMT. It's basically a way of making one phone line look like hundreds of subchannels and improving transmission using an inverse Robin Hood strategy. "Bits are robbed from the poorest channels and given to the wealthiest channels," says John M. Cioffi, a cofounder of Amati and now an engineering professor at Stanford. DMT beat competing approaches—including ones from giants like AT&T—and became a global standard for DSL. In the mid-1990s, Amati's DSL chip set (one analog, two digital) sold in modest quantities, but by 2000, volume had increased to millions. In the early 2000s, sales exceeded 100 million chips per year. Texas Instruments bought Amati in 1997.

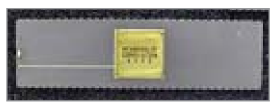

Motorola MC68000 Microprocessor (1979)

MOTOROLA was late to the 16-bit microprocessor party, so it decided to arrive in style. The hybrid 16-bit/32-bit MC68000 packed in 68 000 transistors, more than double the number of Intel's 8086. It had internal 32-bit registers, but a 32-bit bus would have made it prohibitively expensive, so the 68000 used 24-bit address and 16-bit data lines. The 68000 seems to have been the last major processor designed using pencil and paper. "I circulated reduced-size copies of flowcharts, execution-unit resources, decoders, and control logic to other project members," says Nick Tredennick, who designed the 68000's logic. The copies were small and difficult to read, and his bleary-eyed colleagues found a way to make that clear. "One day I came into my office to find a credit-card-size copy of the flowcharts sitting on my desk," Tredennick recalls. The 68000 found its way into all the early Macintosh computers, as well as the Amiga and the Atari ST. Big sales numbers came from embedded applications in laser printers, arcade games, and industrial controllers. But the 68000 was also the subject of one of history's greatest near misses, right up there with Pete Best losing his spot as a drummer for the Beatles. IBM wanted to use the 68000 in its PC line, but the company went with Intel's 8088 because, among other things, the 68000 was still relatively scarce. As one observer later reflected, had Motorola prevailed, the Windows-Intel duopoly known as Wintel might have been Winola instead.

CLOCKWISE FROM TOP LEFT: TRIPATH; PETER CHOW; COMPUTER HISTORY MUSEUM; RAVI BHATNAGAR

Chips & Technologies AT Chip Set (1985)

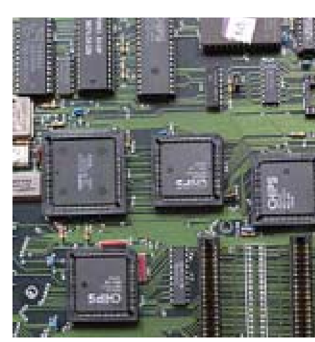

BY 1984, when IBM introduced its 80286 AT line of PCs, the company was already emerging as the clear winner in desktop computers—and it intended to maintain its dominance. But Big Blue's plans were foiled by a tiny company called Chips & Technologies, in San Jose, Calif. C&T developed five chips that duplicated the functionality of the AT motherboard, which used some 100 chips. To make sure the chip set was compatible with the IBM PC, the C&T engineers figured there was just one thing to do. "We had

Computer Cowboys Sh-Boom Processor (1988)

TWO CHIP DESIGNERS walk into a bar. They are Russell H. Fish III and Chuck H. Moore, and the bar is called Sh-Boom. No, this is not the beginning of a joke. It's actually part of a technology tale filled with discord and lawsuits, lots of lawsuits. It all started in 1988 when Fish and Moore created a bizarre processor called Sh-Boom. The chip was so streamlined it could run faster than the clock on the circuit board that drove the rest of the computer. So the two designers found a way to have the processor run its own superfast internal clock while still staying synchronized with the rest of the computer. Sh-Boom was never a commercial success, and after patenting its innovative parts, Moore and Fish moved on. Fish later sold his patent rights to a Carlsbad, Calif.–based firm, Patriot Scientific, which remained a profitless speck of a company until its executives had a revelation: In the years since Sh-Boom's invention, the speed of processors had by far surpassed that of motherboards, and so practically every maker of computers and consumer electronics wound up using a solution just like the one Fish and Moore had patented. *Ka-ching!* Patriot fired a barrage of lawsuits against U.S. and Japanese companies. Whether these companies' chips depend on the Sh-Boom ideas is a matter of controversy. But since 2006, Patriot and Moore have reaped over US $125 million in licensing fees from Intel, AMD, Sony, Olympus, and others. As for the name Sh-Boom, Moore, now at IntellaSys, in Cupertino, Calif., says: "It supposedly derived from the name of a bar where Fish and I drank bourbon and scribbled on napkins. There's little truth in that. But I did like the name he suggested."

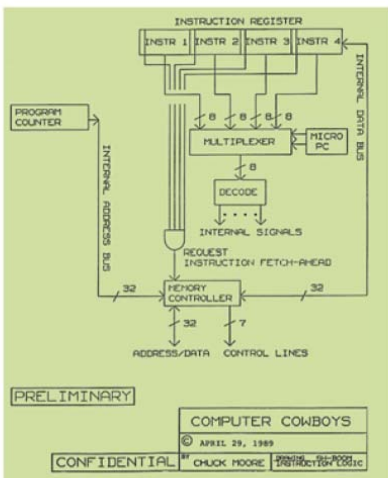

LEFT: CHUCK MOORE; RIGHT: FUJIO MASUOKA

the nerve-racking but admittedly entertaining task of playing games for weeks," says Ravi Bhatnagar, the chip-set lead designer and now a vice president at Altierre Corp., in San Jose, Calif. The C&T chips enabled manufacturers like Taiwan's Acer to make cheaper PCs and launch the invasion of the PC clones. Intel bought C&T in 1997.

Toshiba NAND Flash Memory (1989)

THE SAGA THAT is the invention of flash memory began when a Toshiba factory manager named Fujio Masuoka decided he'd reinvent semiconductor memory. We'll get to that in a minute. First, a bit (groan) of history is in order.

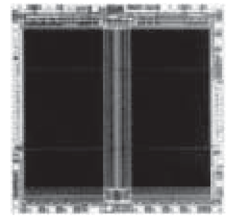

Before flash memory came along, the only way to store what then passed for large amounts of data was to use magnetic tapes, floppy disks, and hard disks. Many companies were trying to create solid-state alternatives, but the choices, such as EPROM (or erasable programmable read-only memory, which required ultraviolet light to erase the data) and EEPROM (the extra *E* stands for "electrically," doing away with the UV) couldn't store much data economically.

Enter Masuoka-san at Toshiba. In 1980, he recruited four engineers to a semisecret project aimed at designing a memory chip that could store lots of data and would be affordable. Their strategy was simple. "We knew the cost of the chip would keep going down as long as transistors shrank in size," says Masuoka, now CTO of Unisantis Electronics, in Tokyo.

Masuoka's team came up with a variation of EEPROM that featured a memory cell consisting of a single transistor. At the time, conventional EEPROM needed two transistors per cell. It was a seemingly small difference that had a huge impact on cost.

In search of a catchy name, they settled on "flash" because of the chip's ultrafast erasing capability. Now, if you're thinking Toshiba rushed the invention into production and watched as the money poured in, you don't know much about how huge corporations typically exploit internal innovations. As it turned out, Masuoka's bosses at Toshiba told him to, well, erase the idea.

He didn't, of course. In 1984 he presented a paper on his memory design at the IEEE International Electron Devices Meeting, in San Francisco. That prompted Intel to begin development of a type of flash memory based on NOR logic gates. In 1988, the company introduced a 256-kilobit chip that found use in vehicles, computers, and other mass-market items, creating a nice new business for Intel.

That's all it took for Toshiba to finally decide to market Masuoka's invention. His flash chip was based on NAND technology, which offered greater storage densities but proved trickier to manufacture. Success came in 1989, when Toshiba's first NAND flash hit the market. And just as Masuoka had predicted, prices kept falling.

Digital photography gave flash a big boost in the late 1990s, and Toshiba became one of the biggest players in a multibillion-dollar market. At the same time, though, Masuoka's relationship with other executives soured, and he quit Toshiba. (He later sued for a share of the vast profits and won a cash payment.)

Now NAND flash is a key piece of every gadget—cellphones, cameras, music players, and of course, the USB drives that techies love to wear around their necks. "Mine has 4 gigabytes," Masuoka says. ❑

With additional reporting by Sally Adee, Erico Guizzo, and Samuel K. Moore

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

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

NEXT-GEN ULTRASOUND

Medical imaging borrows techniques from the microelectronics industry

BY BUTRUS T. KHURI-YAKUB, ÖMER ORALKAN & MARIO KUPNIK

SHARPER IMAGE: Micromachined transducer probes for ultrasound scanners should provide prenatal images that are even sharper than those new parents now get to see. The pictures, though, may never be as crisp as the one in this fanciful photo-illustration.

PHOTO-ILLUSTRATION: PAUL VOZDIC/GETTY IMAGES

ALMOST INVARIABLY, a new baby's photo album begins with a grainy black-and-white picture taken months before birth—a prenatal ultrasound image, which is often detailed enough to inspire comments about the child's resemblance to various members of the family. But jokes about balding uncles notwithstanding, such scans serve a serious purpose and can prove immensely important, as when they allow doctors to diagnose and sometimes even repair a congenital malformation while the baby is still in the womb.

When seeing such an image for the first time, most people are awestruck. How can mere sound waves provide such remarkably clear views? Engineers may well ask something more: How can we give doctors even better ultrasound images? That question has engaged the three of us, along with other members of our Stanford acoustics group, for much of the last decade.

Whereas the signal-processing and image-reconstruction techniques used in medical ultrasonography have made huge advances since this type of imaging became commonplace three decades ago, the business end of the apparatus—the transducer, which converts electrical impulses to sound waves and vice versa—has remained largely unchanged. So we found fertile ground when we began digging for ways to improve those transducers using tools from the microelectronics industry. You will soon find the fruits of those efforts at your local hospital. Indeed, this strategy promises to revolutionize ultrasound imaging within the next few years.

HOW ULTRASOUND imaging works is easy enough to describe, at least in broad strokes. High-frequency (1- to 50-megahertz) sound waves transmitted into the body create reflections when they encounter a change in tissue density or stiffness. These faint echoes are picked up with the same set of transducers used to generate the sound. Or the

A CATHETER'S ULTRASONIC EYES

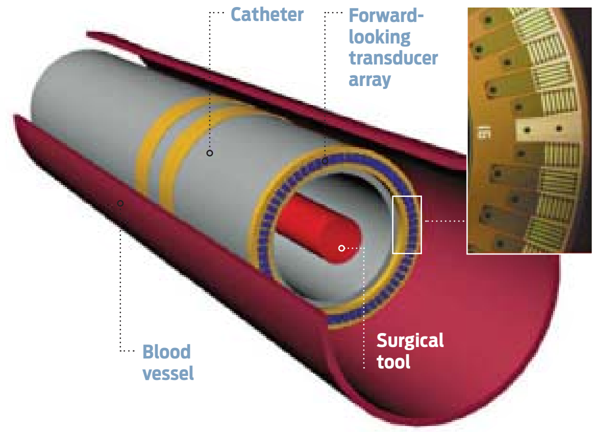

Micromachined capacitive transducers can be easily formed into arrays of various shapes and sizes. Small ring-shaped arrays, for example, can provide forward-looking ultrasonic views for doctors probing an artery with a catheter. Dangerous blockages and certain kinds of heart problems can then be treated using surgical tools that pass through the center of the catheter. The elements of the array depicted in this drawing [blue rectangles] are each composed of many individual transducer units [small brown squares in the inset photograph].

imager may use just a single transducer moved over the body—usually with the aid of much slimy goo, to ensure good acoustic coupling. The resulting electrical signals are then amplified, combined, and displayed as images.

Ultrasonography is valuable for several reasons. For one, it's inexpensive—at least compared with CT (computed tomography) and PET (positron-emission tomography) scanning, or with MRI (magnetic resonance imaging). Also, the low-amplitude ultrasound waves used for imaging do not involve ionizing radiation and are thus harmless to the patient, so repeated scans can be made without worry. And with this technique it is not difficult to get real-time imagery, which doctors may want for such things as guiding a biopsy needle. These virtues make the market for medical ultrasound equipment huge—more than US $5 billion annually, a figure that's only expected to swell in coming years with growing sales of these systems in China and India.

An ultrasound imager has four main parts: the transducer probe, the analog front-end electronics, the digital signal-processing hardware, and the display. Advances in electronics over the years have brought an extraordinary level of refinement to all but the transducer, which means that most of the remaining opportunities for improving system performance lie in the design of this one critical component. In particular, researchers have lately been seeking reliable ways to fashion many individual transducers into compact arrays.

Having a series of transducers laid out in a line—a one-dimensional array—is the simplest example of this strategy. Such transducer arrays are now employed routinely for most forms of ultrasound imaging. Like multielement radio antennas, such arrays can be steered so as to send energy in a narrow, directed beam. Steering an array also works in reverse, allowing it to detect acoustic echoes that come from one particular direction. While a one-dimensional transducer array can be steered and focused within a single plane to make a two-dimensional image, a 2-D array can be steered and focused throughout a volume to make a three-dimensional image—and this can be done in real time.

With this capability, physicians can, for example, follow heart motions in great detail if they want to assess a patient's cardiac functioning. In the not-so-distant future, such ultrasound imaging may even allow robotic surgeons to operate on a beating heart so that patients need not run the risk of having to depend on a heart-lung machine.

In the nearer term, doctors are keen to use small 2-D arrays of tiny ultrasonic transducers to obtain forward-looking images as they probe an artery with a catheter. That would permit them to examine obstructions and map the composition of plaque deposits on vessel walls in three dimensions. What's more, sufficiently small transducers can be arranged in a ring on the end of a catheter, leaving space at the center for an excision device. Such an instrument would allow for simultaneous ultrasound imaging and surgical therapy.

Two-dimensional arrays of ultrasonic transducers would certainly help physicians perform minimally invasive treatments in this way. But making such tiny arrays using traditional transducers is frustratingly difficult. Fortunately, the precise fabrication required can readily be carried out using methods developed by the microelectronics industry, methods that are now routinely used to produce various sorts of microelectromechanical systems, or MEMS.

MEMS fabrication techniques have enabled us to construct something we call a capacitive micromachined ultrasonic transducer. This name, we admit, is an ungainly mouthful, and the acronym we use in our scholarly papers, CMUT, is a bit cryptic to all but a few specialists. Perhaps this is why some of our colleagues in industry refer to this new technology by the more pleasing phrase "silicon ultrasound," which tells you right away what stuff these new transducers are for the most part made of.

BUTRUS T. KHURI-YAKUB

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

NEW VIEW

Medical ultrasound equipment traditionally uses piezoelectric transducers, usually made from lead zirconate titanate. The upcoming generation of capacitive micromachined transducers offers greater bandwidth, which translates into better depth resolution in the images that can be obtained. The improved quality can be seen in a pair of images of the carotid artery and thyroid gland [top panels], which show anatomical detail better when capacitive transducers are used to obtain the image [right]. Such a transducer consists of many individual cells, each containing a silicon membrane separated from the silicon substrate by a thin gap and a combination of silicon dioxide and silicon nitride insulation [bottom].

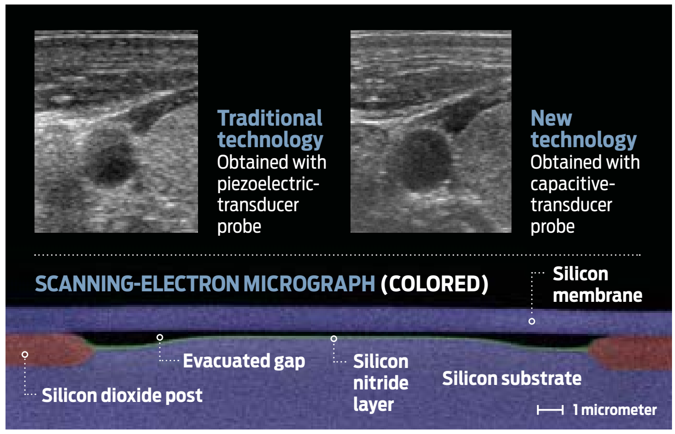

TOP: DAVID M. MILLS/GE (2); BOTTOM: BUTRUS T. KHURI-YAKUB

ULTRASONIC TRANSDUCERS have traditionally been fashioned from piezoelectric materials like quartz or lead zirconate titanate, which many engineers know as just PZT. These are crystals or ceramics that expand or contract in response to an applied voltage. Likewise, piezoelectrics will generate an electric signal in response to being stretched or squeezed, so they can both transmit sound and detect it. This is very old technology, having been invented late in the 19th century, when the brothers Pierre and Jacques Curie demonstrated piezoelectric effects in what for decades remained a laboratory curiosity. The first real application, for sonar, came in 1917.

The procedure used to process a piezoelectric substance into a transducer or an array of transducers relies mostly on age-old manufacturing methods: mixing materials, bonding them, mechanically dicing the resulting assembly, adding wiring—that is, a lot of delicate manual labor. The production of ultrasonic transducer probes, which amounts to a global market of about $1 billion annually, is therefore limited by the many headaches involved in maintaining high yield and good product uniformity in a manufacturing system that depends so much on sharp eyes and steady fingers.

The capacitive transducers we've been pioneering sidestep such issues. By using photolithography and other fabrication techniques of the semiconductor industry, we can make transducer arrays—large or small—with even the most complex geometries, and we can do so very precisely and inexpensively.

In a way, it's surprising that this approach has taken so long to catch on. After all, condenser microphones are capacitive sound transducers, and they've been common for decades. They change sound into electrical signals using a flexible membrane separated from a solid back plate by a very thin air gap. Both membrane and back plate are conductive, or have conductive electrodes attached to them, so a condenser mic is essentially a parallel-plate capacitor. When sound waves hit the membrane, it vibrates, inducing an oscillatory current from the capacitor when it is biased with a dc voltage.

Condenser microphones can capture sound of superb quality, which is why they are often used in studio recording. In ultrasonics, though, the demands are greater than they are for audio frequencies. For a capacitive design to be as efficient as the existing piezoelectric transducers, the electric field in the gap has to be enormous—hundreds of millions of volts per meter. And when subjected to electric fields of that magnitude, air tends to break down, forming a conductive arc. So if you tried running a normal condenser microphone with a bias voltage high enough to produce such an electric field, you'd soon see sparks fly.

Fortunately, the world works differently at small scales. As you reduce the size of the gap, the electric field required for air to break down increases. So with a sufficiently small gap, you can make a capacitive transducer—essentially, a tiny condenser mic—that supports an immense electric field. Such a transducer can be extremely efficient.

That a small gap can sustain a large electric field has been known for more than a century, but it wasn't until the early 1990s that a few researchers took advantage of this fact and began experimenting with capacitive transducers for ultrasonics. They struggled, though, using mostly conventional machining and plastic-film membranes (in a few cases with micromachined back plates) and were unable to overcome the breakdown issue in their first crude devices. Then in 1994 one of us (Khuri-Yakub) and Matthew I. Haller, who was at that point a graduate student in our research group at Stanford, began to apply micromachining and other MEMS techniques to construct the entire transducer.

At the time, our focus was on equipment for nondestructive testing—looking for cracks in the wings of F-18 fighter jets, to be specific. So we intended these first transducers to be used in air (where they worked surprisingly well). Having previously done a lot of research for the U.S. Navy on sonar, though, we tried out a pair of the new transducers underwater *Continued on page 48*

TO PROBE FURTHER

A detailed summary of the authors' MEMS-fabrication techniques is published in "Capacitive Micromachined Ultrasonic Transducers: Fabrication Technology," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 52, no. 12, December 2005.

The full range of the authors' research in acoustics is described at http://www-kyg.stanford.edu.

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

MEDAL *of* HONOR

Thanks for the Memories

Laptops, workstations, PlayStations, iPhones—they would all be impossible without **Robert Dennard**'s invention of **DRAM**

By SALLY ADEE

Ask Robert Dennard about the invention of DRAM, and he will probably do three things.

First, he will show you the patent notebooks IBM encouraged its inventors to keep, which hold all his ideas about dynamic random-access memory, meticulously dated and witnessed by other people, "to make sure we had proof of our inventions." He stores these pristine notebooks in an armoire under a wall crowded with his awards.

Second, he will spend half an hour showing you how he had the revolutionary idea of substituting a single transistor and a single capacitor for the memory technology then being used—magnetic rings, like miniature Cheerios, each of which stored one bit based on the polarity of its magnetic field. He will draw the circuit diagram for the one-transistor DRAM, including every amplifier, data line, and inverter.

Finally, he will comment on a certain online article that suggests that Intel engineers, rather than Dennard and IBM, should be credited with the invention of DRAM. Intel released a three-transistor DRAM in 1970, three years after Dennard entered the one-transistor DRAM into his patent notebook. The misattribution annoys Dennard to no end: "They asked someone from Intel who worked on the chip, 'Did you invent DRAM?' And he said, 'We don't care about inventions. We care about products.'" Dennard pauses. "A lot of people think Intel invented DRAM, because they were the first to come out with something labeled dynamic RAM," he says. Just about everywhere else, Dennard is credited as the father of DRAM, and for that achievement he is being awarded this year's IEEE Medal of Honor.

The wrestling over who gets credit is no hopelessly irrelevant teapot tempest. Random-access memory inhabits pretty much everything that has electrons coursing through it: your laptop, car, game console, digital camera, and cellphone. The amount of RAM in these devices might even be taken as

DAVID YELLEN

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

a kind of shorthand for their approximate level of performance. That's because ever-increasing memory capacity is one of the key factors driving the evolution of most electronics.

Semiconductor memory is now a large extended family—including EEPROM (electrically erasable programmable read-only memory) and NAND flash—each category with different drawbacks and benefits. But dynamic random-access memory is an important ancestor. "Random access" means what it says: A microprocessor can withdraw any stored "word" (8 bits of data) from this memory in any order.

In Dennard's one-transistor DRAM, each bit of data is stored separately inside its own capacitor. A single transistor controls both reading and writing. A charged capacitor means "1," and an uncharged capacitor means "0." The word *dynamic* in the name derives from the fact that the act of reading the bit discharges it and it must be rewritten back into memory. A capacitor's charge eventually wanes, so the memory must be reinfused with fresh charge several times per second to prevent it from losing information. That fact led one researcher to joke that Dennard had won prestigious awards not for his invention but rather for having the temerity to refer to such a thing as "memory."

Amazingly, in an industry defined by its constant advances and compulsory forward movement, the one-transistor DRAM has endured for 40 years.

In 1958, when Dennard walked into IBM's still-unfinished Thomas J. Watson Research Center for his first day at work, he didn't know exactly how a transistor worked. In those days, not too many engineers actually did. Dennard, fresh out of Pittsburgh's Carnegie Institute of Technology (now Carnegie Mellon University), had just earned a Ph.D. in electrical engineering after completing undergraduate and master's work in EE at Southern Methodist University, in Dallas.

But what he recalls most fondly is his first educational experience, one from a bygone era. "At the National Inventor's Hall of Fame, a bunch of us guys—all very successful—were having a conversation, and we found out that all of us went to one-room schoolhouses," he says. "That was the common denominator."

Growing up in a 5000-person farm community near the Louisiana border of Texas, that's all there was. No Baby Einstein classes for Dennard, no Mozart symphonies on a phonograph. The Depression was just ending; his community hadn't been electrified. "We survived just fine," he says, adding that the secret to his success was that he had a lot of spare time as a child. "I learned everything very slowly and concentrated deeply," he recalls.

In those days, he wasn't interested in science or engineering at all. "I had a crystal radio," he declares, "but I never got that thing to work." What he loved was science fiction; he devoured old anthologies that included authors like Edgar Rice Burroughs and H.G. Wells. "One story really influenced me," he recalls. "It was about probability." The short story, "Inflexible Logic," by Russell Maloney, was published in 1940. To test the theory that patterns would emerge out of randomness, a man assembled six monkeys and set them to typing, to see if they would come up with anything rational or intelligible. After quite a short time, the monkeys began to write some very familiar prose. The man shared the results with his friend, a professor.

"And the monkey was coming up with great stuff, and [the professor] was walking around scratching his head and thinking, It couldn't have happened so soon." Dennard pauses and laughs uproariously. "So he shoots the monkey!"

Science fiction was as close as he got to an interest in science until he took physics classes at SMU, which he attended on a dual academic/band scholarship as a French horn player. He liked his physics classes, particularly the emerging field of semiconductor physics—so much so that he decided to pursue a doctorate in electrical engineering, which was then an interesting discipline that in some ways hadn't quite found itself. "I had some advanced physics courses, solid-state materials, and so forth," he says, "but I still didn't understand exactly how a transistor operated." Armed with his Ph.D., he followed some friends to IBM, which was on a research-scientist hiring binge. He figured he'd stay for a few years. Fifty-one years later, he's still there.

He started as a staff engineer in the applied research group, studying what were then brand-new metal-oxide-semiconductor field-effect transistor (MOSFET) designs and circuit applications. Then one day in the fall of 1966, he attended an internal IBM Research review conference. One project was an attempt to commercialize magnetic core

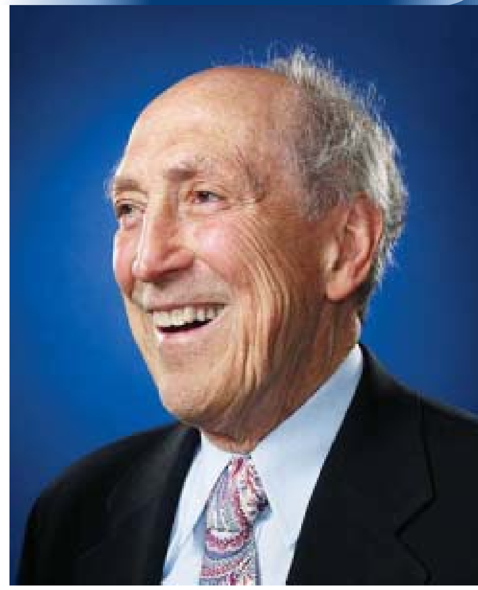

DAVID YELLEN

ROBERT DENNARD

MOST RECENT AWARDS: Charles Stark Draper Prize, IEEE Medal of Honor

DATE OF BIRTH: 5 September 1932

BIRTHPLACE: Carthage, Texas

FAMILY: wife, Jane Bridges, software consultant and teacher; two adult daughters from a previous marriage, each with two children

PETS: Two Scottish terriers, Bonnie and Ferguson

FAVORITE LEISURE ACTIVITY: Scottish country dancing, two nights a week, and choral singing

MANTRA: "Attitude is everything."

CURRENT TITLE: IBM Fellow

FAVORITE MOVIE: "I don't watch movies. They're too loud. The last one I loved was *La Ronde*, which I saw in graduate school 55 years ago."

memory, the standard at the time. The magnetic rings were strung together on a mesh of wire, forming a grid perhaps 30 centimeters on a side. "The truth is, it probably wouldn't have worked," Dennard says. "But it looked good. It was still big, but they had put lots of bits in there."

Dennard went home that night wondering if he could replace the magnetic ring with a small capacitor to store charge. So for the next couple of months he worked on the problem every day and every night. "The first thing I did was put a transistor in series with a capacitor. Then you could write the charge into the capacitor and turn it off." But how to read it? After months of hair pulling, Dennard was seized with his great eureka moment: a single field-effect transistor and data line could both read and write the charge stored in the capacitor. So in 1967 he detailed his invention in his standard-issue IBM patent notebook, and so was born the one-transistor dynamic random-access memory.

DRAM, like almost all great inventions, has many fathers. Three years before Dennard drew his circuit diagram in his notebook, fellow IBMers Arnold Farber and Eugene Schlig had created a memory cell with two transistors and two resistors. A year later, in 1965, IBM researchers refined that idea into a 16-bit monolithic memory array. Also that year, J.D. Schmidt developed a semiconductor random-access memory, but he used six MOSFETs per memory cell, inflating both the footprint and power consumption. Dennard's patent for a one-transistor DRAM was awarded in 1968, but IBM didn't turn it into a product. Instead it shipped computers that used six-transistor SRAM, a technology the company considered less risky.

Then, in 1970, Intel released the first commercially available DRAM memory chip, the three-transistor 1103, which could store 1024 bits [see "25 Microchips That Shook the World," in this issue]. But that's not what people mean when they speak of DRAM today, Dennard insists. "That's why they don't call it 'one-transistor DRAM,'" he says. "It's just DRAM."

Dennard also conceived the scaling theory of MOSFETs, which predicted that the speed of any chip would increase in direct proportion to the decrease in size of its transistors. This theory is commonly—and erroneously—folded into Moore's Law, which actually predicted only the continuing size decreases, not the associated performance increases.

"Bob Dennard was the person who correlated scaling with performance, and it's as important as DRAM," insists Juri Matisoo, who worked on magnetic memory at IBM in the 1960s before Dennard wiped out the competition with semiconductor memory. Matisoo went on to become vice president of technology at the Semiconductor Industry Association. "Moore was projecting the timescale; the IBM people described how to actually do it."

The sprawling T.J. Watson research complex in Yorktown Heights, N.Y., is 61 kilometers north of New York City, but it borrows the city's grid layout, with 40 numbered alleys on each of its three levels. You can't get lost. The Watson campus architecture, finished in 1961, rejected the caste system of corporate ambition: No offices have windows. Instead, architect Eero Saarinen crafted an enormous communal corridor with a three-story wall of windows overlooking rolling hills and leafy greenery in the summer, bare branches under ice and snow in winter. That bucolic view is available to Nobel Prize winners and postdocs in equal measure.

Dennard's window view is a bit of a cheat. He doesn't actually have a window. But his door does open directly onto the magnificent corridor; his is one of only three offices with that luxury.

For all the spartan egalitarianism, Saarinen designed the offices in a bright spectrum of cheery colors. Dennard's office is as brash and upbeat as a Piet Mondrian print. Big color blocks of built-in filing cabinets cover an entire wall. Dennard spent the majority of his career in a blue office buried in the center of the building; the first-among-equals office he has now is a happy lime green.

The IEEE Edison Medal hanging on the wall behind his desk squeezes in next to a row of IBM awards, which in turn rub elbows with a Lemelson-MIT Lifetime Achievement Award. On the armoire beneath, haphazard stacks of plaques suggest that at some point Dennard gave up the 40-year jigsaw puzzle of fitting all the honors onto a single wall.

Soon Dennard will need to reorganize again to make room for his 2009 additions: the Charles Stark Draper Prize—an annual US $500 000 award conferred by the National Academy of Engineering—and the IEEE Medal of Honor. "They're not making it any easier," he says, laughing, as he examines his favorite, the heavy bronze National Medal of Technology awarded to him in 1988 by President Ronald Reagan.

"We just didn't imagine how far it would go," he says, of the one-transistor cell, "how much it would totally change computing."

The wall opposite the awards is almost completely filled by a chalkboard that hasn't been erased in months, or maybe years. Its runes are of different sizes, with some squeezed into the spaces between previous scribblings. There is a small patch of equations with signs for high-*k* metal dielectrics. In another corner, barely visible under some fresher chalkings, is an equation for measuring capacitance. Dennard preserves them all as artifacts of the part he likes most about his role at IBM, which is mentoring incoming employees.

"It's not official mentoring; it's more like being a professor at a university," he explains. "I work with the new people. I work with them on projects, helping define them, monitor progress, and develop the people. Some of them like it. Others want to stay well away," he laughs.

Ghavam Shahidi, who is also an IBM Fellow, says he benefited from Dennard's perspective when he started at the company as a postdoc 20 years ago. "I knew of him for years before I came to IBM," he says. "I only knew of his accomplishments, and that was very intimidating at first. I thought, Here's the man who invented DRAM. This guy is famous." But Shahidi says Dennard was so approachable, down-to-earth, and humble that the impression did not last. "He was not the way I imagined him at all."

Shahidi, who is credited with the development of silicon-on-insulator semiconductor technology at IBM, says many of his epiphanies were born out of long talks with Dennard. "He's great to sit down with and just throw out ideas. He has a broad perspective that he applies to narrow problems."

Dennard applies his perspective liberally, including to the recurring "the end is near" refrain that plagues the semiconductor industry. "The first paper warning of the end of scaling was published by RCA before I even got into this business," he grouses. "You can always find a reason things can't be done."

"That's the thing about the future," he exclaims. "It's totally unexpected. It's been the same for 50 years—we could never see anything more than three years down the road." ❑

Next-Gen Ultrasound

Continued from page 43

for kicks. What we saw knocked us off our chairs. When immersed, the new transducers displayed phenomenal bandwidth, much better than piezoelectrics. Actually, they showed this stellar performance for all of 15 minutes or so; then they stopped functioning altogether.

After a certain amount of cursing and head scratching, we figured out the reason. In our initial designs, the gap between the membrane and the substrate was left open to the outside environment. That's fine for use in air, but when these transducers were immersed, water slowly made its way into the gap, ruining their ability to operate. But it didn't take us long to figure out how to seal these cavities. And eventually we devised ways to eliminate the air inside altogether.

Further work also showed why these capacitive transducers have greater bandwidth than piezoelectrics. The difference arises because a piezoelectric transducer is by nature a highly tuned device, like the pendulum of a clock. At its particular resonant frequency, a piezoelectric transducer undergoes high-amplitude oscillations, even with very little forcing, but at other frequencies, it barely moves at all—which is to say that it has very limited bandwidth.

A capacitive transducer also has a distinct resonant frequency, but only when it's operating in air. When it's immersed in water—or coupled to biological tissues, which are much like water in their acoustic properties—the situation becomes very different. Because the vibrating membrane has so little mass, its movements become highly damped by the watery medium it touches. The same thing happens if you place a pendulum under water. It'll no longer oscillate at its normal resonant frequency, but it can still swing back and forth at the frequency you're using to drive it. This effect, then, lets a single transducer work well over a broad swath of the ultrasonic spectrum.

That's important because it means that the transducer is able to emit and detect the many different frequencies that are contained in a short ultrasonic pulse. The shorter the pulse you use to probe the patient's body, of course, the better the depth resolution in the resulting image. And improved resolution is, after all, just what the doctor ordered.

WE CONSTRUCT one of our transducers by connecting many small units in parallel. Each contains a thin membrane, separated from an underlying substrate by a tiny gap. In our latest designs, the membrane is made of silicon, possibly covered with a metal electrode. Silicon is desirable for several reasons. One is that it has good mechanical properties—it doesn't fatigue, for example—and as long as it is thin, it will flex sufficiently. The substrate is silicon as well, doped with a sprinkling of other atoms to make it highly conductive.

We've developed different recipes for making these devices over the years, but the best scheme we've found uses two different wafers: a garden-variety sili-

---

IEEE

Candidates in 2009 Election

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

Individuals petitioning to be placed on the ballot may do so. Please refer to IEEE Policy 13.7.3. at http://www.ieee.org/web/aboutus/whatis/policies/p13-7.html for specific requirements.

To ensure ballot packages are delivered to the proper mailing address, please visit http://www.ieee.org/go/my_account and update your member profile if necessary.

For more information on IEEE elections and candidates, please visit http://www.ieee.org/elections or contact Carrie Loh, IEEE Corporate Activities, 445 Hoes Lane, Piscataway, NJ 08854, U.S.A.; telephone: +1 732 562 3934; fax +1 732 981 1621; e-mail: c.loh@ieee.org.

Ballot packages will be mailed to all IEEE members eligible to vote on or before 1 August.

IEEE PRESIDENT-ELECT, 2010
J. Roberto Boisson de Marca
Moshe Kam
Joseph V. Lillie

DIVISION II DELEGATE-ELECT/DIRECTOR-ELECT, 2010
Caio A. Ferreira
J. Keith Nelson
Brian C. Wadell

DIVISION IV DELEGATE-ELECT/DIRECTOR-ELECT, 2010
Peter N. Clout
Elya B. Joffe

DIVISION VI DELEGATE-ELECT/DIRECTOR-ELECT, 2010
Gerard H. Gaynor
Jeffrey M. Voas

DIVISION VIII DELEGATE-ELECT/DIRECTOR-ELECT, 2010
Susan K. (Kathy) Land
James W. Moore

DIVISION X DELEGATE-ELECT/DIRECTOR-ELECT, 2010
Vincenzo Piuri
Daniel S. Yeung

REGION 1 DELEGATE-ELECT/DIRECTOR-ELECT, 2010–2011
Peter Alan Eckstein
Albert J. Reinhart

REGION 3 DELEGATE-ELECT/DIRECTOR-ELECT, 2010–2011
Eric S. Ackerman
David G. Green

REGION 5 DELEGATE-ELECT/DIRECTOR-ELECT, 2010–2011
James A. Jefferies
Richard A. Painter

REGION 7 DELEGATE-ELECT/DIRECTOR-ELECT, 2010–2011
Keith B. Brown
Maike Luiken
Marcelo O. Mota

REGION 9 DELEGATE-ELECT/DIRECTOR-ELECT, 2010–2011
Gustavo A. Giannattasio
Norberto M. Lerendegui

STANDARDS ASSOCIATION PRESIDENT-ELECT, 2010
S. Mark Halpin
Steve M. Mills

STANDARDS ASSOCIATION BOARD OF GOVERNORS MEMBER-AT-LARGE, 2010–2011
Stanley L. Moyer

STANDARDS ASSOCIATION BOARD OF GOVERNORS MEMBER-AT-LARGE, 2010–2011
Andrew L. Drozd
T.W. (Ted) Olsen

TECHNICAL ACTIVITIES VICE PRESIDENT-ELECT, 2010
Thomas G. Habetler
Donna L. Hudson

IEEE-USA PRESIDENT-ELECT, 2010
James M. Howard
Ronald G. Jensen

IEEE-USA MEMBER-AT-LARGE, 2010–2011
Winnfort J. Myles
Mauro J. Togneri

con wafer for the substrate and one that's slightly more exotic for the membrane, something known in the semiconductor industry as silicon-on-insulator. The two are bonded together using nothing more than a modest amount of pressure and heat. This two-wafer approach permits us to add the membrane after the pockets that serve as the gaps are already formed, so we can sculpt the membrane and substrate as we wish—they don't have to be just flat planes.

Building transducers from silicon makes it a snap to connect them with the front-end electronics of an ultrasound imager. Although it's possible to fabricate a transducer directly over the associated electronic components on the very same silicon wafer, doing so creates a number of troublesome complications. The better tactic, we've found, is to bond the finished transducer array to a separate wafer containing the electronic circuitry.

Connecting each transducer element may be tricky for tightly packed 2-D arrays, because there isn't much free real estate on the front surface to route a lot of electrical leads. But here again the microelectronics industry has a good solution: Make the connections to the electronics by burrowing down through the transducer substrate and creating vertical conductive channels, which are known in the trade as through-silicon vias.

GIVEN THE many wonderful things we've said about them, you might think that capacitive micromachined ultrasonic transducers would already be in use in medical imaging equipment. Many of the companies that make these systems have indeed embraced this technology, but it hasn't yet reached vendors' shelves. Most of the remaining technical issues are minor, though. Some stem from the electric fields these transducers must contain.

Although there is no chance of arcing across the evacuated gaps, the enormous electric fields can stress the insulating layers to the breakdown point. And even without that, these large fields can inject static electric charge into those layers, which reduces the electric field in the gap, making it necessary to keep adjusting the dc bias field to compensate.

Another challenge with capacitive transducers is that they do not respond as linearly to drive voltages as PZT transducers do. Nonlinearity of the transducer becomes an issue when an ultrasound system is used to image the nonlinear response of biological tissues. Fortunately, there are ways to circumvent this problem, such as purposefully distorting the drive signal to compensate for the nonlinearity of the transducer.

We—along with a slew of engineers at Canon, General Electric, LG Electronics, National Semiconductor, Siemens, and elsewhere—are working to solve these nagging problems and to confront the many other practical realities you have to deal with in any new product. That'll take some time, but it's clear to us that there are no showstoppers here.

It won't be long before this new breed of transducers arrives at hospitals all over the world. So expect those first baby pictures you're shown, among other sorts of ultrasound images, soon to become even more stunning. ❑

Research Group Leader positions

at the International Iberian Nanotechnology Laboratory (INL)

The International Iberian Nanotechnology Laboratory, a recently formed international research organization registered in the UN, is seeking strongly motivated Research Group Leaders, both at Senior level (Principal Investigator, Associate or Full professor level), or at Junior level (tenured track, assistant professor level), to join its new facility in Braga, North of Portugal (www.inl.int). INL central lab facilities are presently being built (€100 M investment for an expected research community of around 400 people at full operation), and will open in late 2009.

INL is recruiting scientists in the following research areas:

- **Nanomedecine:** drug delivery systems, molecular diagnosis systems, cell therapy and tissue engineering.
- **Nanotechnology:** applied to food industry, food safety and environment control.
- **Nanomanipulation:** molecular devices, using biomolecules as building blocks for nanodevices.
- **Nanoelectronics:** Nanofluidics, CNTs, molecular electronics, spintronics, nanophotonics, NEMS, and other nanotechnologies used to build nanodevices and system platforms to support the previous research topics.

Candidates with outstanding CVs in these and related areas will be considered. INL welcomes applicants with previous industrial laboratory experience, and an interdisciplinary research track. INL will offer an exciting, and highly competitive research environment. The remuneration scheme is in line with those offered by other international organizations (IO). Group leaders will be offered substantial starting funds (both for capital equipment and research personnel) to help them jump start their research activities.

Senior Level:
Principal Investigator, Associate or Full Professor Level

Junior Level:
Tenured track positions, Assistant Professor Level

INL facilities will open in late 2009, and are located in Braga (150,000 inhabitants with a high quality and attractive living environment), 30min drive from Porto International Sá Caneiro Airport, 30min drive from the North Atlantic coast, and about 45min drive to the Portuguese-Spanish border and Geres National Park.

Interested applicants should submit a cover letter, curriculum vitae, research statement, and two reference letters (Junior level candidates) to our recruitment website www.inl.int

INL INTERNATIONAL IBERIAN NANOTECHNOLOGY LABORATORY

www.inl.int

THE UNIVERSITY OF HONG KONG

香 港 大 學

Founded in 1911, The University of Hong Kong is committed to the highest international standards of excellence in teaching and research, and has been at the international forefront of academic scholarship for many years. Of a number of recent indicators of the University's performance, one is its ranking at 26 among the top 200 universities in the world by the UK's *Times Higher Education Supplement*. The University has a comprehensive range of study programmes and research disciplines, with 20,000 undergraduate and postgraduate students from 50 countries, and a complement of 1,200 academic members of staff, many of whom are internationally renowned.

Associate Professor/Assistant Professor in Microwave/Electromagnetics (Ref.: RF-2008/2009-639)

Applications are invited for appointment as Associate Professor/Assistant Professor in Microwave/Electromagnetics in the Department of Electrical and Electronic Engineering, from September 1, 2009 or as soon as possible thereafter, on a three-year fixed-term basis, with the possibility of renewal. Appointee who demonstrates satisfactory performance will be considered for tenure during the second three-year contract.

The Department offers B.Eng., M.Sc., M.Phil. and Ph.D. programmes. The B.Eng. programme comprises Electronic and Communications Engineering, Information Engineering, Electrical Engineering and Computer Engineering (jointly with the Department of Computer Science). The Department consists of 40 full-time teaching staff and has excellent computing resources, well-equipped teaching and research facilities and supports. Information about the Department can be obtained at http://www.eee.hku.hk/.

Applicants should possess a Ph.D. degree in Electrical/Electronic/Computer Engineering or a closely related field, and an excellent research record. Applicants with higher qualifications and experience will be appointed at a higher level.

Applicants should indicate clearly which field and level (with reference number) they wish to be considered for.

Annual salaries will be in the following ranges (subject to review from time to time at the entire discretion of the University):

Associate Professor: HK$661,980 – 1,023,720 (approximately US$1 = HK$7.8)
Assistant Professor : HK$504,480 – 779,640

The appointment will attract a contract-end gratuity and University contribution to a retirement benefits scheme, totalling up to 15% of basic salary. At current rates, salaries tax does not exceed 15% of gross income. The appointment carries leave, and medical/dental benefits. Housing benefits will be provided as applicable.

Further particulars and application forms (152/708) can be obtained at https://www.hku.hk/apptunit/; or from the Appointments Unit (Senior), Human Resource Section, Registry, The University of Hong Kong, Hong Kong (**fax**: (852) 2540 6735 or 2559 2058; **e-mail**: senrappt@hku.hk). **Applications will be accepted until the position is filled. Candidates who are not contacted within 4 months of the date of their applications may consider their applications unsuccessful.**

The University is an equal opportunity employer and is committed to a No-Smoking Policy

DEAN OF THE SCHOOL OF INFORMATION SCIENCE AND ENGINEERING

Fudan University, Shanghai, China

Fudan University is marching on its way towards a world class institute of high education. With high quality faculty and academic strengths, the School of Information Science and Engineering is embracing new development opportunities and challenges. For speeding up its progress to improve the research and educational performance, the school is inviting qualified scholars and administrators to apply for the position of Dean who is supposed to have made outstanding academic accomplishments in the information science, engineering and related fields.

DUTIES AND RESPONSIBILITIES:

- Recruit and retain high quality faculty, staff and students in the school.
- Reform and optimize the structure of disciplines based on the principle to stimulate new and scientifically orientated development in the information arena.
- Enhance international cooperation and establish strategic and sustainable partnerships with world class universities to step up the international influence of the school.
- Significantly improve the research and educational performance of the school.
- Make effort to acquire national and/or international research projects through competition.

QUALIFICATIONS:

- Demonstrated experience and success in leading a major department or large laboratory; high-level administration experience in international academic organizations will be favorable.
- Carried out pioneered national and/or international research projects with great success in one or multiple fields of information science and engineering.
- In the last six years, a publication record in prestigious international journals in the information science and engineering fields as the corresponding author, and patents; or given invited talks with publication of influential papers at top international academic conferences; or world advanced and innovative research achievements.
- Fluency in English or Chinese regardless of citizenship.
- 55 years of age or under.

SALARY AND TERM:

- The appointment will be made on a five-year fixed-term basis.
- The successful candidate should work in full-time position (no less than 9 months per annum).
- Salary will be commensurate with experiences and qualifications.

APPLICANTS SHOULD FORWARD: curriculum vitae; copies of advanced degrees; certification of present employment; papers and patents published in the past six years; addresses, telephone numbers, and e-mail addresses of three to five scholars for reference and consultation; letter of interest addressing vision objectives, goals and qualifications.

TERM OF APPLICATION: from May 1 to June 10, 2009

FOR MORE INFORMATION, PLEASE GO TO SCHOOL'S WEBSITE:
http://it.fudan.edu.cn

PLEASE DIRECT APPLICATIONS AND / OR INQUIRIES TO:
Mr. GE Qinghua
Talent Recruitment Office, Fudan University
Email: ***yj@fudan.ac.cn***
Tel: 86-21-65642953
or
Prof. QU Xinping
School of Information Science and Engineering, Fudan University
Email: ***xpqu@fudan.edu.cn***
Fax: 86-21-65643159
Tel: 86-21-65643196
86-21-65643561

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 aspires to become a leading global university, excelling in research and professional education. The University is committed to nurturing and developing students' talent and creating applicable knowledge in order to support social and economic advancement. Within the next five years, the University will employ another **200 scholars** in various disciplines including **science**, **engineering**, **business**, **social sciences**, **humanities**, **law**, **creative media**, **energy**, **environment**, and **biomedical & veterinary sciences**.

Applications and nominations are invited for :

Head of Department of Manufacturing Engineering & Engineering Management *[Ref. A/578/36]*

The Department of Manufacturing Engineering and Engineering Management (MEEM) is one of the eight departments under the University's College of Science and Engineering. The College has a strong team of academic faculty, and among them are many National Academy Members/Fellows of Royal Societies and IEEE Fellows.

There are 30 faculty members and nearly 100 research and support staff in the MEEM Department. The undergraduate and taught master's degree programmes offered by the Department are in the areas of Manufacturing Engineering, Mechatronic Engineering & Automation, and Industrial Engineering & Engineering Management. MPhil, PhD, and Engineering Doctorate are also offered. The Department has a strong teaching and research performance, and close ties with the local industries.

Qualifications for Appointment

Candidates should have strong academic and professional qualifications, demonstrated capability for academic leadership and management at an appropriate level in higher education institutions, long-term vision for the Department's development, and excellent interpersonal and communication skills. Candidates should also be able to foster and strengthen external research support from various organizations.

Information and Application

City University of Hong Kong is one of the eight tertiary institutions funded by the Government of the Hong Kong Special Administrative Region. 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) 2788 9334/email : *hrklau@cityu.edu.hk*]. Detailed information of the Department is available at *http://www.cityu.edu.hk/meem.* Please send the nomination or application with a current curriculum vitae to the Human Resources Office by **31 May 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 among the world's top 150 universities according to *The Times Higher Education Supplement 2008* survey.
http://www.cityu.edu.hk

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

THE HONG KONG
POLYTECHNIC UNIVERSITY
香港理工大學

The Hong Kong Polytechnic University is the largest government-funded tertiary institution in Hong Kong, with a total student headcount of about 28,090, of which 14,260 are full-time students, 10,050 are part-time students, and 3,780 are mixed-mode students. It offers programmes at Doctorate, Master's, Bachelor's degrees and Higher Diploma levels. The University has 27 academic departments and units grouped under six faculties, as well as 2 independent schools and 2 independent research institutes. It has a full-time academic staff strength of around 1,300. The total consolidated expenditure budget of the University is in excess of HK$4 billion per year.

DEPARTMENT OF BUILDING SERVICES ENGINEERING

Assistant Professor in Electrical Services

The Department of Building Services Engineering is one of the constituent departments of the Faculty of Construction and Land Use. It offers a full range of programmes leading to the awards of Doctor of Philosophy, Master of Philosophy, Master of Science and Bachelor of Engineering. It has an undergraduate student body of over 400 full-time students, over 200 part-time students and a postgraduate student body of over 350. BSE currently has a total full-time academic staff establishment of 32 engaging in a wide range of active research areas including but not limited to HVAC, Building Energy Studies, Building Environment, Fire and Safety Engineering, Electrical and Lighting Systems and Facility Management. The University and the Department highly encourage a multi-cultural environment concerning both staff and student mix. For more information regarding the Department, please visit the departmental homepage at http://www.bse.polyu.edu.hk.

We are looking for an enthusiastic, established and suitably qualified academic to further advance our excellent reputation through research, teaching and publications.

The appointee will be required to (a) undertake teaching up to master's degree level in Electrical Services and related subjects; (b) be an active researcher in terms of quality publications, external research grants and applied research; (c) supervise undergraduate and postgraduate research/design projects; and (d) provide leadership in developing a programme, a subject discipline, or a research programme, etc.

Applicants should (a) have a PhD degree in an appropriate discipline related to Electrical Engineering; (b) have a good publication record in refereed journals and good potential in bidding for research grants; (c) preferably have professional qualification in a recognised and relevant professional body; and (d) have several years of research experience in an electrical discipline, preferably related to application in buildings.

Remuneration and Conditions of Service
Salary offered will be commensurate with qualifications and experience. Initial appointment will be made on a fixed-term gratuity-bearing contract. Re-engagement thereafter is subject to mutual agreement. Remuneration package will be highly competitive. Applicants should state their current and expected salary in the application.

Application
Please submit application form via email to hrstaff@polyu.edu.hk; by fax at (852) 2764 3374; or by mail to **Human Resources Office, 13/F, Li Ka Shing Tower, The Hong Kong Polytechnic University, Hung Hom, K0owloon, Hong Kong**. If you would like to provide a separate curriculum vitae, please still complete the application form which will help speed up the recruitment process. Application forms can be obtained via the above channels or downloaded from http://www.polyu.edu.hk/hro/job.htm. **Recruitment will continue until the position is filled**. Details of the University's Personal Information Collection Statement for recruitment can be found at http://www.polyu.edu.hk/hro/jobpics.htm.

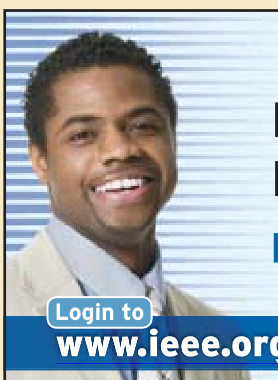

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

the data

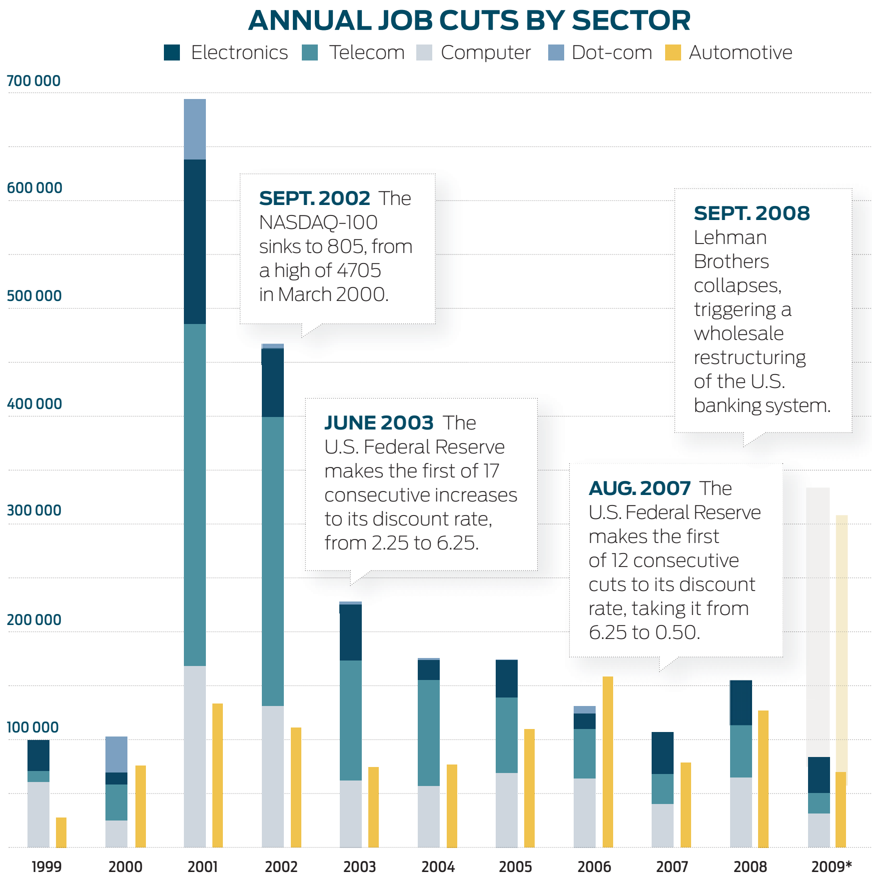

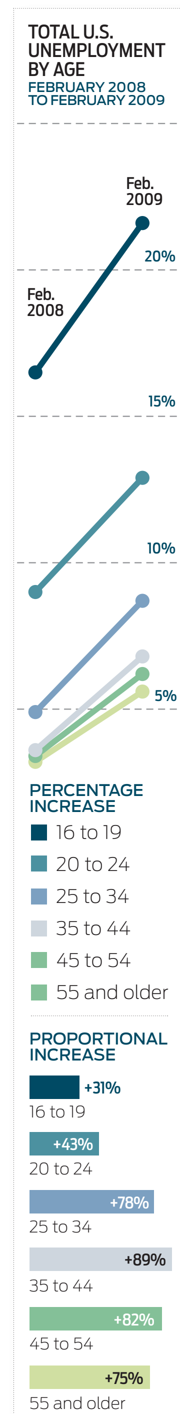

High-Tech Unemployment Is Up, But Not Way Up

UNEMPLOYMENT IN the United States is rising as the current recession cuts deeper than past downturns. Some key technology sectors are certainly feeling a share of the pain, but through 2008, their situation wasn't all that bad when compared with unemployment in previous recessions or even in other economic sectors.

In the first quarter of this year, however, high-tech unemployment spiked dramatically compared to 2008. If extrapolated for the full year, 2009's totals would be the highest in six years, although still only half that for 2001.

Last year's job losses for computer, electronics, and telecommunications, taken together, were up 50 percent compared with the year before—but 2007's numbers were unusually low, according to data provided to *IEEE Spectrum* by Challenger, Gray & Christmas, an employment consulting firm in Chicago. The 156 000 jobs lost in 2008 were a far cry from 2001's high-water mark of 696 000. Compare that to the automotive industry, where the job market is heading over a cliff.

Of course, these figures on job losses don't give a full picture of unemployment, because they don't take into account the new positions being filled. Still, Challenger spokesperson James Pedderson says, the data offer a good picture of how things will be going in the next few months.

Pedderson also notes that unemployment doesn't hit all age groups equally. The overall U.S. unemployment rate has almost doubled for workers 35 to 44, but it is still far higher among the young. Pedderson expects the new crop of college graduates to feel the most pain, with both internships and regular jobs hard to come by. —*Steven Cherry*

*Q1 2009 figures are actual; 2009 totals are projected. SOURCES: Challenger, Gray & Christmas, New York Federal Reserve Bank, http://www.newyorkfed.org/markets/statistics/dlyrates/fedrate.html, Yahoo Finance, http://tinyurl.com/cboecl

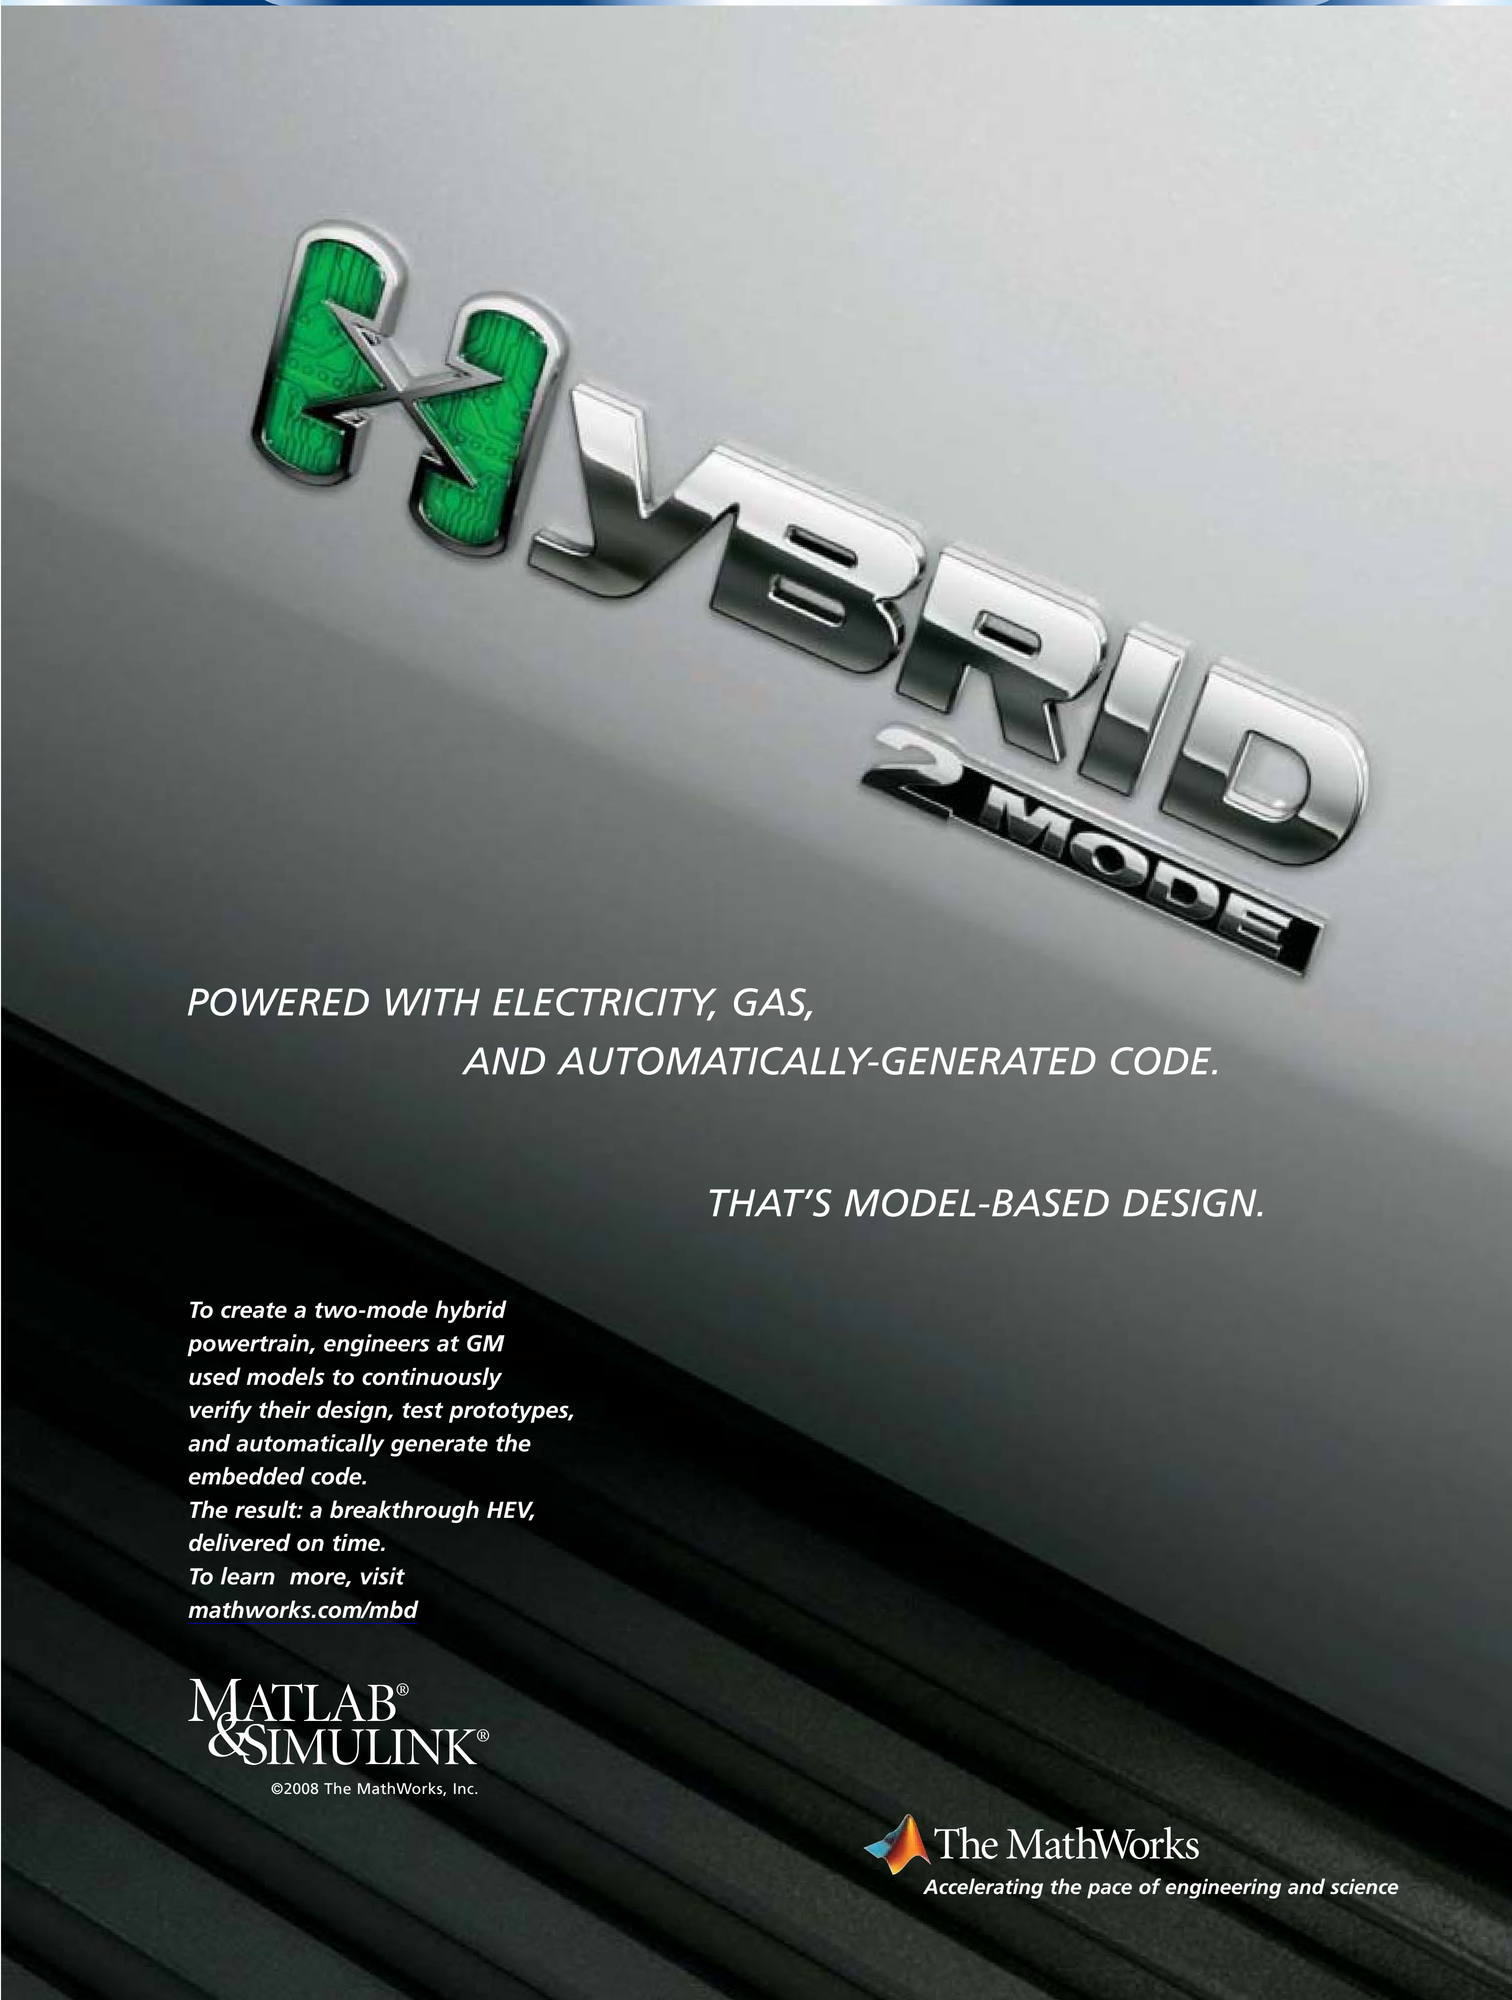
HYBRID
2 MODE
POWERED WITH ELECTRICITY, GAS,
AND AUTOMATICALLY-GENERATED CODE.
THAT'S MODEL-BASED DESIGN.
To create a two-mode hybrid powertrain, engineers at GM used models to continuously verify their design, test prototypes, and automatically generate the embedded code.
The result: a breakthrough 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