IEEE Spectrum

THE MAGAZINE OF TECHNOLOGY INSIDERS

12.11

SPECTRUM.IEEE.ORG

THERE'S A STORM RAGING AROUND YOUR PLANE. AS IT LANDS, WHO DO YOU WANT BEHIND THE CONTROLS— A HUMAN OR A MACHINE?

CLICK HERE TO ACCESS THE DECEMBER ISSUE OF the institute

IEEE


6G

ieee spectrum

volume 48 number 12 international

12.11

38


Name
Dr. Dennis Hong
Job Title
Associate Professor of Mechanical Engineering, Virginia Tech
Area of Expertise
Robotics
LabVIEW Helped Me
Convey and respond to vast amounts of data in real time
Latest Project
Design and prototype a car that can be driven by the blind in just 4 months
NI LabVIEW
LabVIEW makes me better because
CODE REUSE
saves time and effort

ieee spectrum

volume 48 number 12 international

12.11

11

UPDATE

OPINION

DEPARTMENTS

SPECTRUM.IEEE.ORG AVAILABLE 1 DECEMBER

The World's Most Powerful Patent Portfolios, Revealed

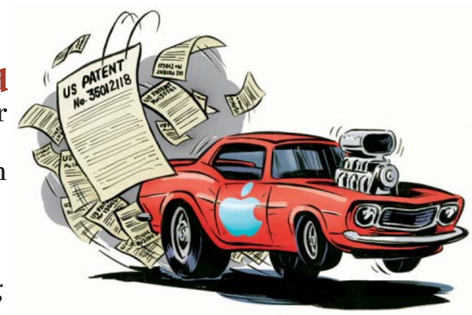

Our annual Patent Power Scorecards, compiled and analyzed by the team at 1790 Analytics, show that U.S. companies like Apple, Yahoo, Microsoft, and Google are inventing their way to dominance. And Chinese companies like Huawei have started to make inroads, even as the flow of patents from Japanese and German companies slows.

THEINSTITUTE.IEEE.ORG AVAILABLE 6 DECEMBER

MAKING THE JUMP INTO GAMES
In our special issue on gaming, learn what it takes to enter this growing industry.

HELP WITH GETTING YOUR GAME ON
A variety of skills are needed to make it as a game developer. IEEE offers several products and services that can help, from continuing education courses and tutorials to webinars and journals.

IEEE STANDARDS ASSOCIATION TAKES ON VIRTUAL WORLDS
These 3-D environments are still considered by some as best suited for games and entertainment. Other people see them as a potential business tool—if certain challenges can be overcome. The IEEE Standards Association recently gathered together virtual-world experts and users to foster cross-industry collaboration and reach a consensus on how best to address these challenges and set standards for the industry.

ONLINE WEBINARS & RESOURCES

AVAILABLE AT http://spectrum.ieee.org/webinar

1 December: Better Simulation of Dielectric Barrier Discharge Plasma Actuators Using VORPAL

15 December: Fluid Structure Interaction Simulation With COMSOL

Electrostatic Actuation With COMSOL Multiphysics

Big Demand Drives Small Solutions–Scalable Baseband Solutions

Accelerating Development of Smart Grid Power Electronics Control Systems With RCP and HIL Techniques

Getting More Benefits From Modeling Vacuum Electronic Devices

Power System Contingency Analysis and Security Constrained Optimal Power Flow in MATLAB

The Emerging Technology Forum Series: Reinventing Infrastructure for High Speed Rail

Plasma Modeling With COMSOL Multiphysics

The IEEE Emerging Technology Forum Series: Silicon Valley's Impact on the Automotive Industry

Requirements Management's Key Role in the Nuclear Industry

Modeling and Simulation of HEV and EV Power Electronics

New product release library:
http://spectrum.ieee.org/static/new-product-release-library

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2011 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 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of American Business Media, the Magazine Publishers of America, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

CLOCKWISE FROM TOP LEFT: MATT MOYER/GETTY IMAGES; JONATHAN CARLSON; EXPERIENCE FIRST DESIGN

back story

DAVID SCHNEIDER

Flying the Robotic Skies

THAT'S SENIOR EDITOR Philip E. Ross standing beside an RQ-4 Global Hawk, a Northrop Grumman unmanned aerial vehicle stationed for the moment at the Naval Air Station Patuxent River, in Maryland.

In this picture, the machine is the only thing that knows how to fly. Sure, it needs a little supervision from ground controllers. But, as Ross freely admits, he couldn't land a kite even if he had a control tower full of pilots talking him through it.

He traveled to Patuxent with his coworker David Schneider, also a senior editor, to see a few robo-planes in action. But it didn't happen that day. These things are very hush-hush. In fact, while snapping this photo, Schneider was told to aim his lens toward the tail of the plane so as to not reveal any secret wiring that might have been visible through an open port on the plane's fuselage.

As Ross reports in this month's cover story, autonomous planes have so transformed the U.S. military that the Air Force is now training more ground controllers than pilots. But what Ross wanted to know was when, or whether, pilotless planes will start carrying people or cargo in civilian air space. After all, if automation can handle the rigors of combat, why can't it fly little air taxis from Podunk to a big-city airport?

Plenty of people are afraid of flying, and many others can fly only with white knuckles and a stiff drink. Ross's editor, upon reading the first draft of the article, declared, "*I'm* never getting in a plane without a pilot." This is the human factor, and it poses the last and highest obstacle to a civilian UAV industry.

Undeterred, Ross has already volunteered for the maiden flight of the first pilotless airliner. Put him down for a ticket. And while you're at it, sell him another one, to the moon. ❑

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 Data is in *IEEE Spectrum,* Vol. 48, no. 12 (INT), December 2011, p. 60, or in *IEEE Spectrum*, Vol. 48, no. 12 (NA), December 2011, p. 76.

ieee spectrum

EDITORIAL

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

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

EDITORIAL DIRECTOR, DIGITAL Harry Goldstein, h.goldstein@ieee.org

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

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

SENIOR ASSOCIATE EDITORS Steven Cherry (Resources), s.cherry@ieee.org; Erico Guizzo, e.guizzo@ieee.org

ASSOCIATE EDITORS Rachel Courtland, r.courtland@ieee.org; Marisa Plumb, m.plumb@ieee.org; Joshua J. Romero (Online), j.j.romero@ieee.org; Eliza Strickland, e.strickland@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 Foster, r.foster@ieee.org; Nancy T. Hantman, n.hantman@ieee.org

INTERN Katie M. Palmer, palmer.k@ieee.org

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

ART & PRODUCTION

SENIOR ART DIRECTOR Mark Montgomery

DEPUTY ART DIRECTOR Angela Howard

ASSISTANT ART DIRECTOR Brandon Palacio

PHOTO EDITOR Randi Silberman Klett

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy

EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia

SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

WEB PRODUCTION COORDINATOR Jacqueline L. Parker

MULTIMEDIA PRODUCTION SPECIALIST Michael Spector

EDITORIAL ADVISORY BOARD

Susan Hassler, *Chair*; Gerard A. Alphonse, Marc T. Apter, Francine D. Berman, Jan Brown, Raffaello D'Andrea, J. Roberto B. De Marca, Hiromichi Fujisawa, Kenneth Y. Goldberg, Susan Hackwood, Bin He, Erik Heijne, Charles H. House, Christopher J. James, Ruby B. Lee, John P. Lewis, Tak Ming Mak, Carmen S. Menoni, David A. Mindell, C. Mohan, Fritz Morgan, Andrew M. Odlyzko, Larry L. Smarr, Harry L. Tredennick III, Sergio Verdú, Jeffrey M. Voas, William Weihl, Kazuo Yano

EDITORIAL CORRESPONDENCE

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

ADVERTISING CORRESPONDENCE

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

REPRINT PERMISSION

LIBRARIES: Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970.
For other copying or republication, contact Business Manager, IEEE Spectrum.


Imagination
Ideas
Innovation
Creating engineers for a borderless planet
Professional Master's Program
Enhance your B.S. degree in ECE and enhance your value to industry.
• Earn a Master's in Electrical and Computer Engineering in 12-18 months.
• Complete an optional project as part of the degree
• Fellowships available for qualified applicants
Ph.D. Program
Create original research in a program ranked in the Top 5 in the U.S.
• Enter the program with either a B.S. or an M.S.
• Study with highly respected faculty with entrepreneurial experience
• Be part of a collaborative team at a university known for its interdisciplinary research.
• Dean's Fellowships available for qualified applicants
Carnegie Mellon
Electrical & Computer ENGINEERING
Applications accepted at:
www.ece.cmu.edu
Women and Minorities are encouraged to apply

contributors

CHRISTIAN CACHIN and **MATTHIAS SCHUNTER**, who wrote "A Cloud You Can Trust" [p. 28], are computer scientists at IBM Research–Zurich. A cryptography expert concerned with cloud security, Cachin likes to start the day with a 5 a.m. row on Lake Zurich. Schunter, technical leader of the European Union–funded TClouds project, prefers bicycling to all other forms of transportation. Both say some of their best ideas about computer security occur when they're in transit.

KENNETH R. FOSTER reviews new versions of Mathematica and Maple [p. 23], descendants of the early math program Macsyma, which he coveted as a physics doctoral student in the late 1960s. "I would have done anything for access to that program," Foster says, slightly in awe of the ubiquity of today's packages. An IEEE Fellow, Foster is a professor of bioengineering at the University of Pennsylvania.

NEIL SAVAGE writes about camera technology that lets you focus *after* you snap in "A New Angle on Imaging" [p. 15]. The method could be used to create 3-D images with a single camera, but Savage has another aspiration: the perfect candid photo. "As you wait for your camera to focus, people stop and they want to pose," he bemoans. Not having to focus would help catch subjects unawares. Based in Lowell, Mass., Savage has also written for *Technology Review* and *Nature*.

VINOD SREEHARSHA, who wrote "Brazil Starts a Chip Industry" [p. 13], is a freelance journalist based in São Paolo, Brazil. Prior to moving there in 2010, he covered politics, economics, technology, and entrepreneurship from Argentina and Venezuela for *The New York Times*, *The Miami Herald*, and *Venturebeat*. Sreeharsha says he relocated to Brazil because it's "a deeply complex nation—the more you learn, the more you realize you don't know."

MICHAEL S. STRANO and **KOUROSH KALANTAR-ZADEH**, who wrote "Nanodynamite" [p. 44], share an interest in extremely small systems. While on sabbatical from the Royal Melbourne Institute of Technology, in Australia, Kalantar-zadeh joined Strano's nanotechnology research group at MIT. The team was working on measuring the acceleration of a chemical reaction along a nanotube when they made the serendipitous discovery that the reaction generated power. Now the two researchers are using their combined expertise in chemistry and nanomaterials to explore this phenomenon.

WARREN TOOMEY teaches at Bond University, in Australia. He was bitten by the Unix bug in 1982 while still in high school, when he spent two weeks at the University of Wollongong learning about computers. There he encountered "an amazing system called Unix." His later discovery of software for simulating the computers that ran the earliest versions of Unix got him collecting vintage copies of this operating system. In "The Strange Birth and Long Life of Unix" [p. 34], he explains why Unix was so influential.

CLOCKWISE FROM TOP LEFT: KENNETH R. FOSTER; LYNNE-DEANNE SMITH; DOMINIC PHILLIPS; KEVIN SAVAGE

Explore the EM simulation universe

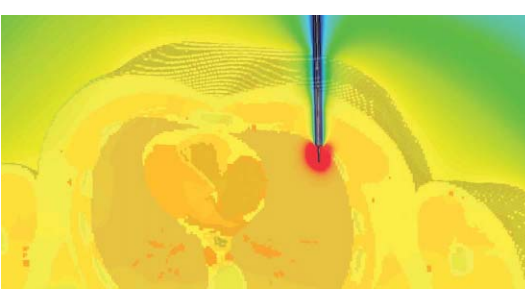

Simulation of cancer treatment by RF thermoablation.

→ Get equipped with leading edge EM technology. CST's tools enable you to characterize, design and optimize electromagnetic devices all before going into the lab or measurement chamber. This can help save substantial costs especially for new or cutting edge products, and also reduce design risk and improve overall performance and profitability.

Involved in biomedical applications? You can read about how CST technology was used to simulate biomedical devices at **www.cst.com/biomed**. If you're more interested in filters, couplers, planar and multilayer structures, we've a wide variety of worked application examples live on our website at **www.cst.com/apps**.

The extensive range of tools integrated in CST STUDIO SUITE enables numerous applications to be analyzed without leaving the familiar CST design environment. This complete technology approach enables unprecedented simulation reliability and additional security through cross verification.

→ Grab the latest in simulation technology. Choose the accuracy and speed offered by CST STUDIO SUITE.

CST

CHANGING THE STANDARDS

CST – COMPUTER SIMULATION TECHNOLOGY | www.cst.com | info@cst.com

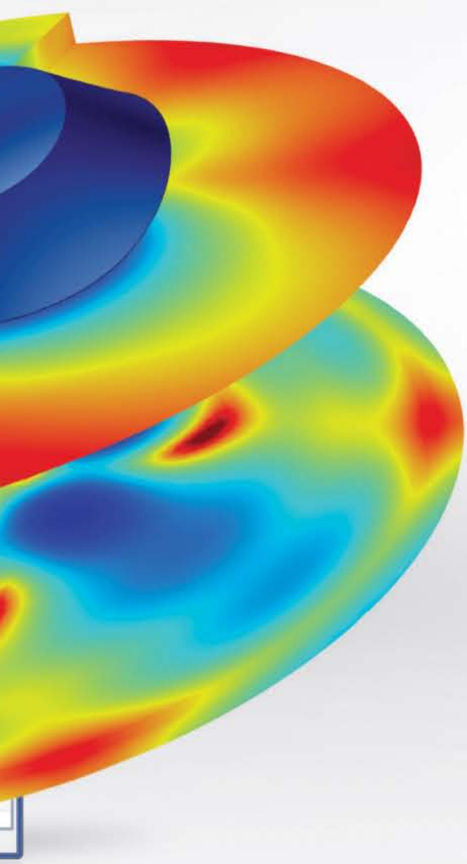

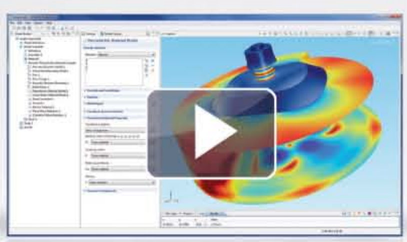

spectral lines

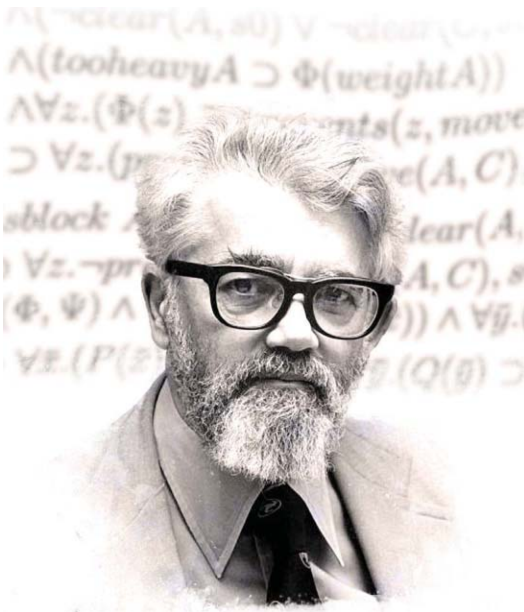

PHOTO: STANFORD UNIVERSITY; PHOTO-MONTAGE: MARK MONTGOMERY

Remembering John McCarthy

THIS PAST October saw the death of John McCarthy, one of the pioneers of computer science and a founder of the field of artificial intelligence (AI), a phrase he is credited with inventing. It capped a sad month that also saw the passing of Apple cofounder Steve Jobs and of Dennis Ritchie, the coinventor of Unix and the C programming language.

John McCarthy was born in Boston in 1927, but he grew up near Caltech, where he got his B.S. in mathematics. He detoured to Princeton for his Ph.D. but ended up at MIT, where he cofounded its artificial-intelligence lab, the world's first, before going on to Stanford in 1962 to found its artificial-intelligence lab.

To hear Steven Cherry's entire interview with Ramanathan Guha, go to http://spectrum.ieee.org/guha1211. See also Katie B. Palmer's blog post at http://spectrum.ieee.org/mccarthy1211.

In between, he found time to invent Lisp, one of the most influential programming languages ever created. McCarthy received computer science's highest honor, the Turing Award, in 1971, the Kyoto Prize in 1988, and the National Medal of Science in 1991. Just this year, he was inducted into the IEEE Intelligent Systems' AI Hall of Fame, along with a fellow AIer, cognitive scientist Marvin Minsky.

McCarthy was probably the first person to seriously consider, in 1955, the question of self-awareness in machines, both as a computer science challenge and as a social issue. He also seems to have been the first to consider computation as a utility in the way that electricity is, an idea that first became time-sharing in the 1960s and lately shows itself in cloud computing [see "A Cloud You Can Trust," in this issue]. When others were focusing their machine-intelligence efforts on chess, McCarthy was working on natural language processing and what we now call robotics. John McCarthy was also an academic, writing dozens of papers in computer science and more than a few in mathematics, all the while overseeing computer science dissertations.

One of John's students, Ramanathan Guha, a former principal scientist at Apple who currently works at Google, had this to say about his thesis advisor when I spoke with him for my weekly *IEEE Spectrum* podcast "Techwise Conversations":

"The set of new things that Lisp introduced into the world of programming is so large that it's almost impossible to think of programming languages without the contributions of Lisp. Everything from conditionals, to recursion, to the idea of mutable data structures—yes, before McCarthy introduced them, programming languages did not even have the 'if..., then...' statement. And pretty much everything—as [programmer] Paul Graham once mentioned—everything that we have been slowly introducing into programming languages since then has been like somebody set up this trust fund of ideas, and we're slowly able to take things off of there. We've still not taken everything out of there. Even things like garbage collection came into mainstream commercial computing only in the mid-'90s. He had it back in the '60s. And he came up with the specification, and the beauty of the specification was that an entire programming language could be specified in half a page, and it was just stunning in its beauty. And I remember [pioneering computer scientist] Alan Kay telling me, if you ever create anything in your life, it needs to have the beauty of Lisp. The idea of beauty in computer science was something that John McCarthy brought into the picture....

"Being a student of John McCarthy was something like being exposed to a supernova at close range. He really, really pushed the boundaries in terms of thinking—people use the cliché 'thinking outside the box.' He didn't understand the word box. *To him there was no problem that was outside his scope."*

—STEVEN CHERRY

Correction

In "Transistor Wars" [November], we state that Intel is putting its new 3-D transistors into production in 2012. In fact, the first chips bearing the devices are already in production and will be commercially available in 2012.

2012
Technology Milestones

YEARS AGO	
03	NASA Ares Rocket
13	Nikon achieves 2.66 megapixel resolution
23	High-Definition Television
33	Compact Disc Audio Player
43	Birth of the Internet
53	Semiconductor Planar Process and Integrated Circuit
63	Binary Automatic Computer
73	Atanasoff-Berry Computer
83	Yosami Radio Transmission
93	Rotary Dial Telephone
103	Shoshone Transmission Line
113	Wireless Telegraph
213	Volta's Electrical Battery






IEEE Advancing Technology for Humanity

www.ieee.org/benefits

Thank you for your membership.

From 2011 IEEE President, Moshe Kam

"Stay engaged with IEEE to make the most of your membership experience. Through my membership in IEEE, I developed a large network of colleagues and peers, learned from and consulted the best in my field, and was offered numerous opportunities to develop new projects and cooperate."

Achieve your goals, regardless of the employment climate.

IEEE Membership offers opportunities in 3 key areas:

1. Finding a job
- Professional networking
- Local activities
- Consultants database

2. Advancing your job
- Career webinars
- IEEE-USA Career Nagivator
- Continuing education
- Mentoring Connection

3. Volunteering
- Learn leadership skills
- Time management
- Improve people skills
- Increase confidence

Turn the page for IEEE Member discounts ⇨

Find your benefits at
www.ieee.org/benefits

IEEE Advancing Technology for Humanity

update

more online at spectrum.ieee.org

DIRTY JOB: Men and children work at a gold mine in Mongbwalu, Democratic Republic of the Congo, in 2005. The mine is controlled by one of the many warring militias in the area. Electronics firms could face bad publicity for using gold from such mines.

PHOTO: MATT MOYER/GETTY IMAGES

Cracking Down on Conflict Minerals

Electronics companies face new rules on minerals mined in war zones

IN THE JUNGLES and mountains of the Democratic Republic of the Congo, battles are raging, part of a 13-year-long civil war. Most of the world has paid little attention to the murder and rape that still dominates life in the DRC's eastern provinces. But U.S. electronics companies like HP, Intel, and Apple recently became deeply interested, thanks to a provision on "conflict minerals" that was slipped into a 2010 financial reform law, the Dodd-Frank Act.

The minerals provision is intended to deprive the Congo's warlords of funds by cutting off sales from the mines they control. It focuses on the ores that produce the "three Ts": tin, tantalum, and tungsten, as well as gold. Public companies that use these metals in their products will be required to investigate their supply chains, determine if they use metals that were mined in the DRC, and disclose their findings to the U.S. Securities and Exchange Commission (SEC), in their annual reports, and on their websites. If its minerals did originate in the DRC, a company must submit a larger report on whether the purchase of these minerals financed or benefited armed groups in that part of Africa. The SEC is expected to issue final rules for implementing the law before the end of the year, and companies are scrambling to get ready.

While the conflict minerals law applies only to companies that are required to file annual reports in the United States, it's expected to have an international impact. Since mineral suppliers sell to electronics companies around the world, any change in operations they make for

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

72% Percentage of respondents to a survey who approved of research into geoengineering or attempts to purposefully manipulate climate. A total of 3105 people—from Canada, the U.K., and the United States—took part.

update

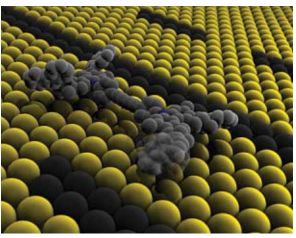

news briefs

Electric Nanocar
Scientists at the University of Groningen, in the Netherlands, and other institutions, have constructed a molecule that looks and seems to move like a four-wheeled car. When zapped with electrons from the tip of a scanning tunneling microscope, the chemical bonds change, rotating the molecule's "wheels" so that it moves in a straight path. It's good only for short trips: Each rotation of the wheels moves it forward less than 1 nanometer.

the U.S. market will have ripple effects elsewhere.

The law doesn't only affect the electronics industry. But the conflict mineral issue has been linked in the public mind to electronics because the three Ts play crucial roles in smartphones, TVs, and laptops. Tin is used in solder and thus found on every circuit board, tantalum is used in capacitors, and tungsten is used in the vibrating motors of many phones.

Electronics companies had been warned that they'd eventually have to account for their use of these minerals. So firms like HP and Intel asked the Electronic Industry Citizenship Coalition (EICC) and the Global e-Sustainability Initiative, two trade groups, to investigate the industry's options.

The groups found that it's extremely difficult to determine the origin of the tantalum used in a certain batch of smartphones. But they also realized that only about 45 smelters worldwide deal with the three Ts, buying the ores from suppliers and turning them into pure metals. After several years of research, the industry groups came up with the Conflict-Free Smelter Program, which is currently in the pilot phase for its first metal, tantalum.

The program asks each smelter to allow an annual independent audit of its mineral procurement process. If the auditors are convinced that no minerals are sourced from the Congo's conflict mines, that smelter is certified as "conflict free," allowing companies to buy its metals without worry. While the program is voluntary, EICC spokeswoman Wendy Dittmer says many smelting firms believe it's in their interest to participate.

"Electronics companies are starting to ask questions all the way down their supply chains," she says. "That certainly makes the buyers of the minerals very interested in being able to talk about their own due diligence."

There are concerns that the law may backfire. By making the reporting requirements more onerous for companies that source minerals from the DRC, the law may reduce demand from all DRC mines, even those that aren't in conflict regions and don't finance armed groups.

These concerns about such a de facto ban led Motorola Solutions to initiate the Solutions for Hope Project, in which Motorola and several other companies formed a relationship with a conflict-free tantalum mine in the DRC's Katanga province.

To establish the program, Michael Loch, Motorola's director of supply-chain corporate responsibility, visited the mine and accompanied a shipment of ore along its export route. "This pilot allows our industry to stay engaged in the area," says Loch. "We didn't want to abandon the region." But he acknowledges that it took a lot of effort to get the process in place for one mine and says it may be difficult to scale up the program.

The pilot programs should provide a framework to make compliance easier. Still, companies around the world are waiting for the SEC's final rules with some anxiety. And there may be some efforts to block the rules' enforcement through U.S. courts. The U.S. Chamber of Commerce, for one, has discussed the possibility of a lawsuit. The chamber disagrees with the SEC's initial compliance cost estimate of US $71 million, saying that costs will instead be counted in the billions of dollars.

One thing is already certain about the SEC rules: There will be no fines for using conflict minerals. Even so, activists think it will have its intended effect, because companies will want to avoid bad publicity.

"For years we have been unknowing consumers of these minerals because companies have turned a blind eye," says Sasha Lezhnev, a policy consultant on conflict minerals with the human rights group Enough. "This will enable consumers to make choices on whether or not to buy products from companies that are sourcing from these mines."

—Eliza Strickland

THE CONGO CONNECTION

The Democratic Republic of the Congo supplies minerals used to produce industrially important elements. Metals derived from "conflict minerals" play critical roles in modern electronics. With the exception of tantalum, the amounts are small compared to the available world supply.

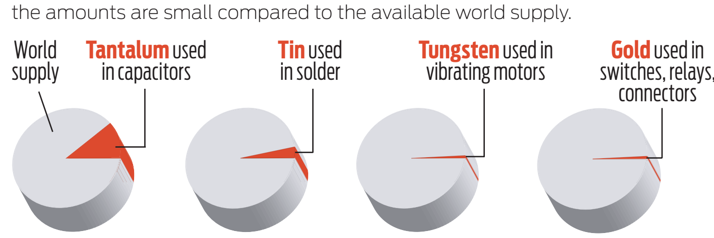

SOURCE: UNITED STATES GEOLOGICAL SURVEY

RANDY WIND AND MARTIN ROELFS

COW COUNTERS: Brazilian-built RFID chips were tested for tracking cows at a ranch in southern Brazil. *PHOTO: CEITEC*

Brazil Starts a Chip Industry

An RFID for cattle is its first semiconductor manufacturing venture

IN THE COMING DECADE, Brazil wants to be known for integrated circuits, not just Ipanema. The country hopes to produce silicon chips alongside soybeans. Its attempt to do so will begin in January, when production is completed on the first batch of locally designed and produced RFID chips at the government-run National Centre for Advanced Electronic Technology, known by its Portuguese acronym, Ceitec.

Working with X-Fab Semiconductor Foundries, based in Erfurt, Germany, Ceitec expects to have approximately 1 million chips ready for market next month. Dubbed the Chip do Boi ("bull chip"), the low-frequency (134.2-kilohertz) RFID series will be used in Brazil's cattle industry to track roaming livestock and gather data on the deforestation they cause in the country's vast Amazon region.

The chip has been tested over the past two years in the predominantly agricultural states of Mato Grosso do Sul and Minas Gerais. Ceitec forecasts a demand for the chip of 1.5 million units next year, with an expected growth of 10 percent or more for the next decade. The cow tracker is the first in a series of planned products that could include chips for digital TV and WiMax systems.

While the technology itself is not cutting edge, this is an important development in a country still highly dependent on exporting commodities, and where nearly all electronics are imported and exorbitantly expensive, says Kevin Suh, an executive familiar with the country's technology sector. Brazil imports all of its semiconductors—US $4.64 billion worth in 2010, according to a study for the Brazilian Development Bank (BNDES).

The push into semiconductors is part of a larger effort to diversify the Brazilian economy, something foreign investors are beginning to notice. Silicon Valley venture capital firms such as Redpoint Ventures and Accel Partners have made multiple investments in Brazilian Internet start-ups recently.

Suh, who is chief operating officer of HT Micron Semiconductors, a joint venture between South Korean semiconductor packager Hana Micron and Brazilian holding company Parit Participações, has closely followed the country's technology efforts. HT Micron currently focuses on packaging, but Suh says that "long-term plans are to become a complete semiconductor company."

Of Ceitec and the Chip do Boi, he says, "it's not a breakthrough innovation, but given the realities of the Brazilian government and its hopes of developing [a technology sector], it can be viewed as a breakthrough innovation." He also believes that looking for niche applications "is good strategy in the beginning" but that in the long term the chip industry will have to diversify.

Brazil began trying to foster a semiconductor industry in 2007 with several measures, such as the National Microelectronic Program, which offered semiconductor firms tax breaks and other incentives. But the effort was to little effect.

The government then created Ceitec in 2008 in Porto Alegre in the southern state of Rio Grande do Sul, home to Brazilian president Dilma Rousseff, who had a hand in choosing its location.

US $38 trillion Global investment needed in energy-supply infrastructure from 2011 to 2035, according to the International Energy Agency.

update

But the firm's development has been marred by management problems (the first CEO was ousted) and by Brazil's infamous government bureaucracy. What's worse, the original silicon-wafer processing equipment, donated by Motorola, was outdated: It was meant for 150-millimeter wafers instead of 200 or 300 mm. Replacing it was a laborious effort. And so the Chip do Boi was set back by over a year.

Ceitec is just one of several government initiatives intended to make Brazil more competitive in technology. Earlier this year the government created the Science Without Borders program, which will provide scholarships for 100 000 Brazilian science and engineering students to study and obtain additional degrees overseas. The government says it will fund 75 000 scholarships and the private sector will fund 25 000.

Another initiative, the government agency Financiadora de Estudos e Projetos (FINEP), provides seed-stage funding to Brazilian start-ups, as does the Fundo Criatec, which is run by BNDES. And these efforts are having some success.

The IC design start-up SiliconReef is one company receiving funding from Fundo Criatec. Founded by Marília Lima and Tiago Lins and incubated at the Recife Center for Advanced Studies and Systems, the company is developing a microchip to improve the efficiency of photovoltaic cells for cellphones and other portable devices. Lima says SiliconReef expects to start selling the chip in the United States and Asia by March 2012.

Whether the efforts of the Brazilian government succeed or not remains an open question. But it will likely need greater participation from the private sector, says Carlos Paz de Araujo, a winner of the IEEE Daniel E. Noble Award. The prolific Brazilian inventor and founder of memory-chip company Symetrix Corp., in Colorado Springs, says he has been trying to stimulate Brazilian semiconductor manufacturing since 2006. In particular, he says that the last few steps in manufacturing Symetrix's Trinion FeRAM chip, including adding the ferroelectric material, can be done in Brazil.

And while Symetrix is working with some semiconductor design houses in Brazil, Paz de Araujo says that he has had a hard time getting Brazilian government officials or Ceitec to understand the opportunities that technologies like FeRAM offer. He says that with some changes in mind-set, "Brazil has an incredible chance to be state of the art." —Vinod Sreeharsha

A Prettier Pylon

There are more than 88 000 electrical transmission pylons in the United Kingdom, and if the country is to meet its renewable energy goals, it's going to need a lot more. So back in May, the Royal Institute of British Architects launched a competition to invite architects and engineers to rethink the pylon. Bystrup Architecture Design and Engineering, of Copenhagen, took the £5000 (about US $8030) prize, beating out 250 rival designs. At 32 meters, the pylon stands 18 meters shorter than today's variety, and at 20 metric tons, it's lighter by a third, too.

BYSTRUP

A New Angle on Imaging

Capturing the direction of light beams can make for after-the-fact focusing

ALL CAMERAS capture the intensity of light as it strikes their imaging chips. Color filters provide a second set of data, sorting the rays into different wavelengths. But new devices—including one produced commercially and others still in the lab—are starting to capture a third piece of information: angle. This allows cameras to go beyond focusing on a single plane to taking images at many different depths of a scene at once. Cameras that capture both intensity and angle will allow for refocusing already-snapped pictures, lensless cameras, and the creation of 3-D images with a single camera, according to researchers in several competing groups.

Alyosha Molnar, assistant professor of electrical and computer engineering at Cornell University, in Ithaca, N.Y., has developed angle-sensitive pixels for CMOS imagers. Each pixel is made up of a photodiode beneath two layers of diffraction grating, one of which is slightly out of alignment with the other. The top grating creates an interference pattern on the grating beneath it. Depending on how that pattern of light and darkness lines up with the second grating, light will either pass through to the photodiode below or be blocked. So whether the diode sees bright light or dim depends on the angle the light is coming from.

Forming an image requires having a detector with pixels in many different orientations and gratings with different amounts of spacing between them. Taken together, the pixels produce a series of measurements that can be processed to create an image with any focus the user chooses. And because the camera can use the same data to produce two images at different depths, it can also generate a 3-D image.

According to Molnar, camera chips with angle-sensitive pixels—known as light-field imagers—should be cheap to produce, because the gratings are inscribed in the layers of wires that are already built into the detector.

A newer approach, which Molnar is presenting this month at the IEEE International Electron Devices Meeting, would inscribe the gratings in glass on the surface of the detector, generating the same diffraction patterns while blocking less light, thus making the device more sensitive.

Molnar says his approach should be cheaper than a similar technology already hitting the market. Lytro, a start-up in Mountain View, Calif., is taking orders for a consumer camera that allows users to snap pictures first and focus later; the company promises to start shipping them in 2012.

The heart of the device is an array of microlenses that lie over the detector. The microlenses focus light rays from different angles on different pixels in the detector, which yields images at many depths but leads to fairly low resolution. "You lose a lot of resolution, but you get freedom of refocusing afterwards," says Ramesh Raskar, head of the Camera Culture group at the MIT Media Lab, who is familiar with the technology but not involved with Lytro.

Raskar has come up with a different approach. He places a patterned piece of glass between the camera's lens and the image sensor chip. The screen attenuates the light rays entering the lens, with the amount of attenuation depending on the angle of the ray. He says his method is similar to the microlens technique but with no loss in resolution. The Lytro approach, however, requires less computing power.

The three techniques are "redefining the notion of a camera," says Molnar. With the angle-sensitive systems, "you will actually be able to go beyond what a film camera could do."

—NEIL SAVAGE

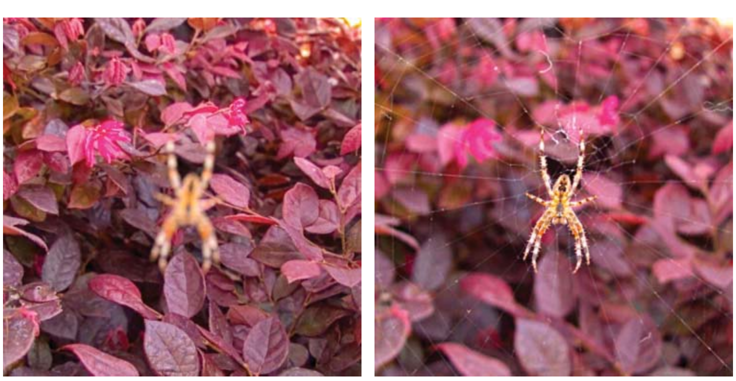

PERFECT PICTURES: Light-field imagers, such as a camera sold by start-up Lytro, use a combination of angle-sensitive pixels and computational trickery to let you focus an image after you shoot it. *PHOTOS: LYTRO*

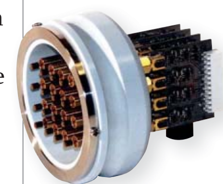

PLESSEY SEMICONDUCTORS

news briefs

Heart Reader
A new kind of electric potential sensor could allow for electrocardiogram readings at a distance, say British engineers. Invented at the University of Sussex and manufactured by Plessey Semiconductors, the electric potential integrated circuit (EPIC) sensor could make monitoring patients in hospitals easier or even form the basis of a brain-machine interface.

spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

update

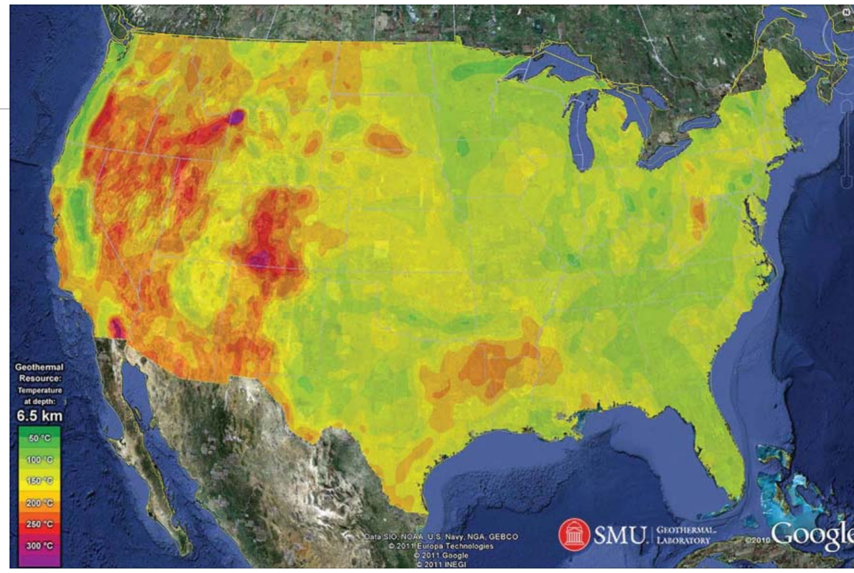

HOT SPOTS: The United States is boiling over with geothermal energy. *IMAGE: SMU GEOTHERMAL LAB*

Geothermal Energy's Promise and Problems

Google-funded research shows U.S. potential is huge, but extraction could cause quakes

GEOTHERMAL ENERGY is touted as one of the few renewable resources that could be used for base-load (round-the-clock) power generation: Earth's heat is always on, and it's not dependent on the vagaries of wind or sun. New research from Southern Methodist University—sponsored by Google's philanthropic arm—suggests massive potential for geothermal power in the United States. But exploiting that resource will be slowed by the cost of the technology—and the fact that it can cause small earthquakes.

Researchers led by David Blackwell at SMU's Geothermal Laboratory set out to update existing maps of the heat beneath our feet, maps that Blackwell says had significant gaps. The researchers doubled the number of locations measured from previous efforts, and by sampling more than 35 000 sites, they found a "technical potential" of almost 3 million megawatts.

"The technical potential is our best estimate of what actually might be extracted," says Blackwell. He says that depths greater than 6.5 kilometers are impractical to access, so his calculation does not take into account a supply of power that's more than 10 times as much at depths up to 10 km.

To put this in perspective, there are only about 3000 MW of installed geothermal capacity in the United States today, and no other country has more. The total installed electricity capacity from all energy sources in the country is right around 1 million megawatts. So if all that geothermal energy were harnessed, it could power the country three times over.

According to Colin Williams, a scientist with the U.S. Geological Survey who has worked on similar geothermal resource estimates, the SMU work is less about new discoveries than about technological optimism. "It's not like they discovered more thermal energy down there, but they're pushing the scenario that you could get more of it out," he says. His own calculations from a 2008 study showed that even the most easily developed geothermal resources could bring 6500 MW online and that more "unconventional" resources represented more than half a million megawatts of potential. That assessment differed from the SMU work in several ways, including stopping at a depth of 6 km instead of 6.5 km, as well as focusing almost entirely on the western United States.

Blackwell also says the most notable improvements over previous estimates are in the East—especially under coal-rich West Virginia. The energy is there, he says, but "the question is, Do we have the will to go ahead and try to really develop it?"

Most existing geothermal projects are at hydrothermal reservoirs where hot water is brought up from below the surface to produce electricity. But the new 3 million MW would almost all require what is known as enhanced geothermal systems (EGS). That technique allows the use of lower-temperature areas by fracturing the rock with high-pressure water, similar to the controversial fracking process in the natural gas industry. And there is ample evidence that EGS produces small earthquakes.

"We know that creating these EGS reservoirs involves making earthquakes. That's just going to happen," Williams says. "The question then becomes, Are we going to be able to control the process of generating the microseismicity so that we don't generate earthquakes that are magnitude 3.5 or 4.0 or something like that?" The U.S. National Academy of Sciences and the National Academy of Engineering have launched an investigation; their report is expected in 2012.

For the moment, cost is also a primary barrier to widespread adoption. USGS's Williams says traditional geothermal electricity is "in the ballpark" in terms of competing with other electricity sources. A 2009 report by the investment bank Credit Suisse quoted a conventional geothermal cost of 3.6 U.S. cents per kilowatt-hour, below the 5.5 cents for coal. EGS is costlier: A 2007 report by consulting firm GeothermEx estimated the best possible cost for EGS systems in the future at 5.4 cents per kilowatt-hour and suggested that the technology won't be truly cost competitive until 2050. "Until EGS is developed on a wide scale, initially it probably wouldn't be competitive," Williams says. "Right now we're looking at sort of slow but steady development." —DAVE LEVITAN

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

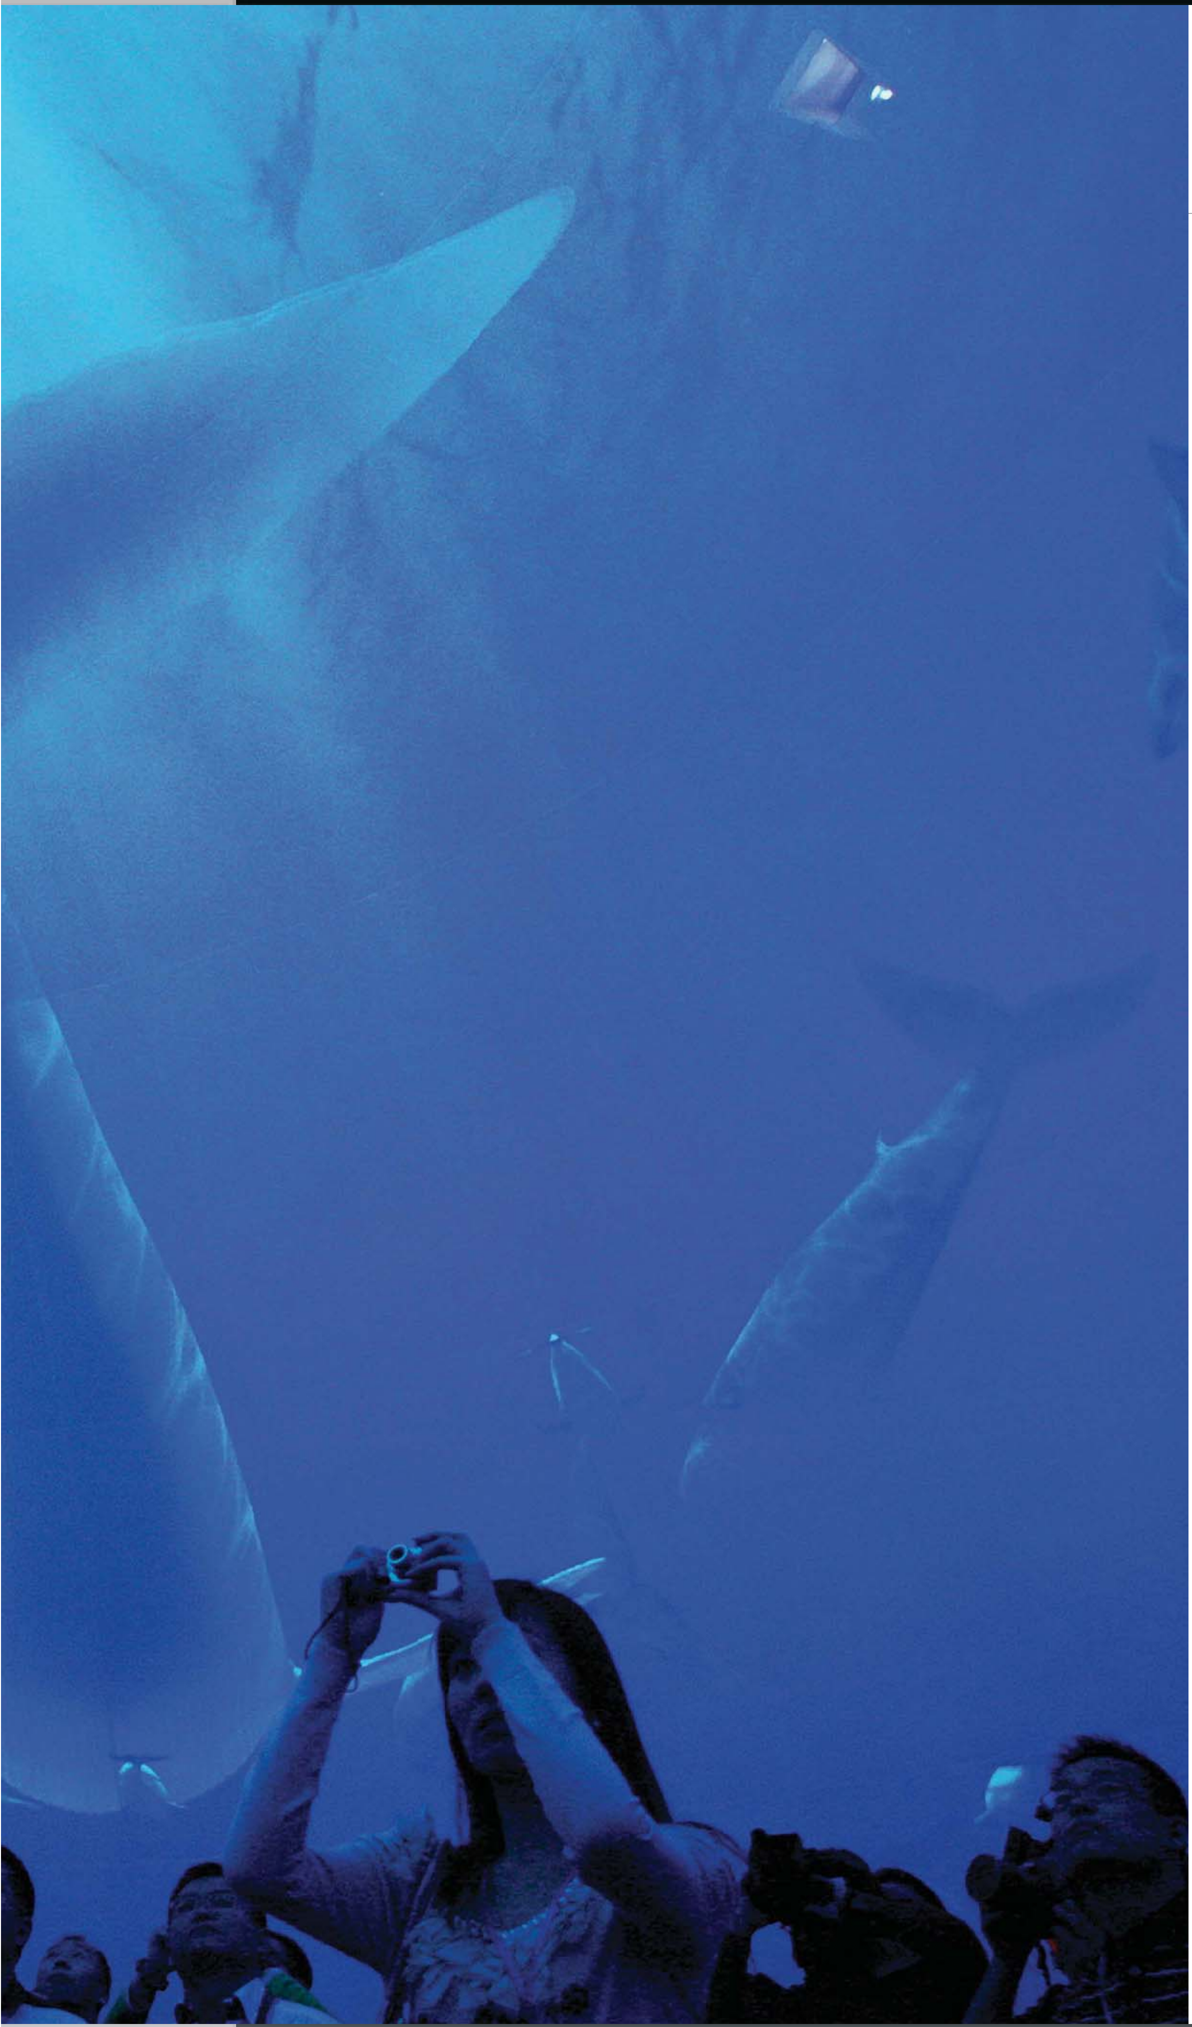

the big picture

CENTER STAGE

Now anyone can swim with the sharks. A film with a deep-sea exploration theme amazes audiences at the Haichang Backdrop Cinema, in Yantai, China. The theater shows films on a 360-degree holographic movie canopy that puts viewers at the center of the action. The fish you can (almost) reach out and touch are the result of 3-D simulation technology combined with an LED lattice and a surround-sound system that adds to the effect.

PHOTO: SHEN JIZHONG/ XINHUA/LANDOV

hands on

Lego Mindstorms-NXT 2.0 Robotics Kit US $240
http://mindstorms.lego.com
PHOTO: LEGO

A DROID FOR ALL SEASONS

These robotics kits can liven up the holidays—and give a kid a taste of robotics engineering

ROBOT PLAYTHINGS have been charming youngsters at least since the first tinplate robots were mass-produced in the World War II era. Today's robotic toys do a lot more than the windup automatons of yore. And if you choose wisely, your gift can also help inculcate some of the rudiments of engineering.

For the smaller tyke, consider an **Ollo Action Kit** (US $30 at Amazon or ThinkGeek). It resembles a plastic Erector Set, but the pieces are joined with snap-together rivets instead of nuts and bolts. Manipulating the rivets still takes fine motor skills, so I wouldn't recommend it for kids younger than about 7. An accompanying booklet shows, with a sequence of very clear pictures, how to first assemble a simple quadruped before moving on to four motorized projects: a windmill, a dog, an insect, and finally a dinosaur.

My 8-year-old son, having assembled countless Lego sets over the years, jumped right to the dinosaur and had it walking around the dining room table in no time. So for a bigger challenge, I presented him with the paragon of robotics gadgetry for children: a Lego Mindstorms set.

First released in 2001, the **Lego Mindstorms Robotics Invention System** got a significant upgrade in 2006, with the introduction of Mindstorms-NXT. Its 577 parts included four different kinds of sensors and three rather sophisticated servomotors. The most recent edition, the Mindstorms-NXT 2.0 Robotics Kit ($240 from Amazon), released in 2009, boosts the overall parts count to 619, contains a slightly different mix of sensors (one ultrasonic range finder, one color sensor, and two contact sensors), and offers such niceties as a built-in Bluetooth radio and the ability to do floating-point calculations.

I'd read about Mindstorms many times, including in this magazine, but I hadn't appreciated how well this system was put together. Hats off to the folks at Lego for producing something that works so well at so many levels.

Any child who can piece together Lego bricks should have little trouble assembling the starter robot, a small tracked vehicle described in the kit's 62-page instruction booklet. The heart of the kit is its computer module, the NXT "brick," which the starter robot holds at a convenient angle for viewing the LCD screen and operating the four buttons. With just that simple user interface, youngsters can quickly get their creations moving and doing various interesting things.

But that's just the starting point. The next step is to install the accompanying software on a computer and start programming your robot du jour using Lego's graphical programming

Ollo Action Kit US $30
http://www.robotis.com
PHOTO: DAVID SCHNEIDER

language, NXT-G. National Instruments helped to develop this user-friendly environment, which is based on NI's well-known LabVIEW (Laboratory Virtual Instrumentation Engineering Workbench) software. With some dragging and dropping, you can quickly develop programs that with the click of a mouse are compiled and then downloaded over USB cable or Bluetooth to the NXT brick.

Lego's NXT-G programming environment is not only easy to master, it's also very good at reporting errors and pointing out the source of the problem (a common one was simply that the brick had shut itself off while I was busy arranging my next program).

Although the graphical programming interface takes a little getting used to and has a few quirks, it provides a budding roboticist with an ideal introduction to programming. My son, whose computer experience was mostly limited to playing games, quickly constructed and programmed a four-wheeled robot that would spin around until it sensed an object in front of it, at which point it would announce "Object detected—good-bye" and then shoot a small wooden ball at whatever had come into its ultrasonic sights. (You'll want to keep an eye on things, especially if there are younger siblings in the house.)

A determined Mindstormer could probably go far using nothing other than NXT-G, but you're not limited to that. The Mindstorms software ecosystem now includes dozens of other ways to program an NXT brick, including Processing, a C-like language and development environment that's also used to program the Arduino family of microcontrollers. You can even control your Mindstorms creation through a Bluetooth link using Microsoft Robotics Developer Studio. These alternate programming environments should be more appealing to technically advanced kids.

And if you're worried that an adolescent might feel he or she has outgrown Lego altogether, there are more substantial robotics platforms, such as the **iRobot Create** ($130 from iRobot).

The company, makers of the Roomba vacuum-cleaning robot and other automated household helpers, produced the Create specifically for people who wanted to experiment with robotics. It's essentially a Roomba without the vacuum cleaner. Instead, the disk-shaped body of the robot contains a large cargo bay. On top are four threaded hard points, so Create owners can easily add their own superstructures.

I decided I would control the Create using Robotics Developer Studio, so I also purchased a Bluetooth adapter module ($63) along with a Bluetooth USB radio ($40), both from iRobot. While waiting for everything to arrive, I downloaded and installed Robotics Developer Studio and set to work figuring out how to simulate a Create, which is one of several robots this software supports out of the gate.

I was using Microsoft's Visual Programming Language (VPL), which comes with Robotics Developer Studio, so I figured this would be easy enough. After all, VPL is another drag-and-drop programming environment—a bit like NXT-G on steroids. Although Robotics Developer Studio comes with various tutorial examples and Microsoft provides copious descriptions of this software online, I found it all less than helpful and ultimately succeeded in simulating a Create only when I stumbled on one blogger's gentle description of how to set this up.

When the package from iRobot arrived, I added the required 12 AA batteries, plugged the Bluetooth adapter module into the Create, and stuck the Bluetooth USB Radio into my laptop (which lacks built-in Bluetooth). After pairing the two radios, I tested the connection using a terminal-emulator program to send numerical commands to

iRobot Create US $130
http://irobot.com/create
PHOTO: DAVID SCHNEIDER

the Create, just to turn on some of its LED indicator lights. That worked straight away, so I thought I'd soon be driving my real-live Create around the house. Alas, I was wrong. Ironing out the wrinkles required some surprisingly large adjustments to the controls, which had been working fine with the simulated Create.

My difficulties learning how to control a Create through Robotics Developer Studio were disconcerting, particularly after having such a pleasurable experience using NXT-G. But some young people might well welcome the challenge of puzzling out Microsoft's complex software.

In any event, the proud new owner of a Create robot can choose other software packages to run it and any number of hardware accessories to embellish its capabilities. Add a Kinect ranging camera and a laptop running the open-source Robotics Operating System, for example, and you could turn a Create into a TurtleBot or BiliBot, two rather advanced mobile robots now gaining traction in the DIY community.

One thing is for sure: With some relatively inexpensive hardware to get started, young people can explore robotics at every level, limited only by their imaginations—or perhaps by their supply of AA batteries.

—David Schneider

Prices are those paid at press time. Current prices may vary.

books

Fix That Old Clock Radio!

We don't fix anything anymore, more's the pity. There's an old-world charm, an almost romantic allure, to bringing once-useful gadgets back to life. Repair can save a lot of money, and it's a terrific way to make engineering accessible and attractive to kids.

But trouble-shooting modern gadgets isn't easy. And so, sensibly (if counterintuitively), Michael Jay Geier's entertaining book, *How to Diagnose and Fix Everything Electronic*, pushes an approach that very nearly revels in not knowing the details of how broken devices are supposed to work.

As Geier explains, the days when analog parts drifted out of tune due to wear, weather, or accident are mostly gone, and circuit and system designers have spent the last 30 years working to stamp out the influence of individual component variations. When modern devices fail, they fail hard.

Geier starts with some useful advice about what's probably worth repairing and what isn't. Once you decide to take the plunge, he has a shopping list for your workbench—a soldering iron, desoldering wick, multi-meter, oscilloscope, and so on—but he also reminds you to have stacks of plastic cups and trays on hand to hold all the screws, clips, and other bits from each successive stage of disassembly. He recommends having a digital camera handy, to record how everything was set up before you started fixing it.

The heart of the book is probably its chapter on electronic components: what they're used for, how they're marked, what can kill them, and how to test them outside the context of their circuits. Here's where the difference between building and rebuilding really shows itself. A designer might not care about what marks distinguish capacitors from inductors, resistors, or diodes, but a troubleshooter certainly does, especially nowadays, when consumer electronics are no longer sold with schematics.

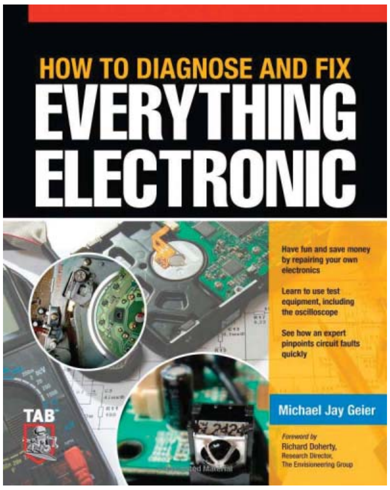

How to Diagnose and Fix Everything Electronic; By Michael Jay Geier; McGraw-Hill/TAB Electronics, 2011; 336 pp.; US $25; ISBN 978-0-07-174422-5

Geier's decades of repair work have given him an arsenal of diagnostic techniques ranging from the obvious—stage-by-stage checking of signals and incoming and outgoing power—to the terrifying: attaching a circuit's power-supply rails to a high-current supply and seeing which components burn up. Through it all, he preaches a healthy skepticism and cautions against declaring a successful diagnosis too early. He gives plenty of examples of dead components that aren't culprits but rather fellow victims of the true cause of failure. (Replacing those components without fixing the real problem will just waste time and money as they burn out again.)

Will I seize on this book as a good reason to get a full set of troubleshooting toys? I don't know. I love the idea of being able to fix all the dead electronic soldiers around the house, but there are also a lot of gadgets that I keep hoping will fail soon so I'll have an excuse to get a shiny new model.

—Paul Wallich

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

tools & toys

Maple 15
US $1245 and up
PHOTO: MAPLESOFT

MATHEMATICA 8 AND MAPLE 15

Two venerable math programs keep up with the times

BORN IN the 1980s, Maple and Mathematica antedate many readers of *IEEE Spectrum*. Over time, both programs have evolved into do-it-all math platforms, with thousands of numerical and symbolic math functions, extensive capabilities for graphics and documents, and much more. With their latest updates, Mathematica 8 and Maple 15, the companies have added many more goodies to already very large and polished packages.

Trying to be all things to all people, Mathematica has long extended its reach far beyond math into document preparation and more. It now incorporates the specialized search engine and database Wolfram|Alpha. Entering "IEEE Spectrum," for example, returns circulation data and some other information for this magazine. The results can be unpredictable: "US Magazine" finds a small town in Arkansas. This same facility allows users to enter commands in natural language (for example, "Plot sin xy"), useful for quick calculations but probably not much more.

Even more significant, in the summer of 2011 Wolfram unveiled its Computable Document Format for creating interactive documents that can be distributed freely. Other users can open them and play around with complex data sets or manipulate math models, thanks to a free browser plug-in that installs a 400-megabyte utility on their computers. But if others want to import data or enter text into the documents, they must purchase additional software.

Maple 15 and its modeling and simulation companion, MapleSim5, have kept a sharper focus on engineering. When speaking to the developers at Maple, I found them particularly proud of the program's new support for multicore processors, parallel programming for a cluster of local computers, and multithread computing, each of which can dramatically speed up large calculations. For example, they described problems involving sparse linear programming that were unfeasible in earlier versions but can be easily calculated in Maple 15.

Wolfram Mathematica 8
US $1095 and up
PHOTO: WOLFRAM RESEARCH

More modest changes (but ones undoubtedly more noticeable to the average user) include a new variable manager that allows users to inspect the definitions of variables and a facility that allows users to view and manipulate tables of data. A new finance package has dozens of commands, including one involving something called a Bermudan-style swaption. Next year's release will predict future stock prices (just kidding).

Mathematica 8 was released November 2010, Maple 15 in April. If you work at an institution with a site license, you've probably already received these updates. If not, take a careful look before buying a pricey upgrade. But for some users, the new capabilities may be deal makers.

Prospective users should examine both updates carefully in light of their own needs; the programs differ in many particulars. In my own experience, many calculations are easily done by following the help menus. But you have to live with these programs for a while to become proficient—not a small investment.

All that aside, both products are major achievements in computer science and have been game changers for many kinds of research. According to *Science Citation Index*, each has been cited by more than 1500 research papers, in such disparate fields as abstract math and marine biology. They're also sophisticated and fun to use, and large user communities exist for both. And unlike us, these venerable products show no signs of slowing down, even as they age and bulk up.

—KENNETH R. FOSTER

http://www.wolfram.com
http://www.maplesoft.com
http://www.wolframalpha.com

geek life

SING A SONG OF CODE

New riffs on two venerable music programming languages

THE 18TH CENTURY'S greatest hacker, Johann Sebastian Bach, composed much of his algorithmically inspired music on the sonic geek's gadget of choice: the adjustable and reconfigurable pipe organ. The gadget of choice for today's musical hackers is the personal computer.

"Bach would be all over this," says David Cottle, associate professor at the University of Utah's School of Music. "A computer is powerful and accurate and ignorant. It'll take your directions as you wouldn't have expected—which might be interesting and fun."

Unlike the pipe organ, the computer needs software in the form of a programming environment. As it happens, one of the two leading musical programming languages, Max, has recently been retooled, while the other, SuperCollider, finally has its own definitive programmer's manual.

Cottle brags that a musical algorithm can manipulate composers' entire canons or styles of music.

"What SuperCollider allows a composer to do is pursue a progressive ethic of asking the question first, then discovering the outcome," he says. "For example, you might create a virtual composer from scratch, entering into a database the probabilities from existing digitized scores of current avant-garde composers. Generating new music based on those probabilities could result in the next step in music evolution. I don't know, but that's the whole point: You can ask that question and SuperCollider will take you there."

Cottle is one of three coeditors of a new edition of *The SuperCollider Book*, a manual for the free and open-source SuperCollider language. It details such procedures as collating input from a computer mouse, tablet devices, and video-game controllers like Nintendo's WiiMote. It also devotes chapters to artificial intelligence applications, like the smart sensing of real-time musical cues that enable an algorithm to "improvise."

But musicians had better be ready to geek up. Halfway into Chapter 1 ("Beginner's Tutorial"), the 680-page book is already nesting commands to generate 12-tone matrices with coding samples that look like outtakes from a C++ manual.

By contrast, Max 6 offers a stable, intuitive, and graphical interface through which a user can quickly set up and start doing interesting things. (If that sounds like Mac versus Linux, it sort of is—Max is available for Windows and the Mac, while SuperCollider runs on those two operating systems and Linux as well.)

Composers and recording artists (including Björk and Radiohead), nonprofessional musicians, and an estimated 1800 universities around the world use Max, says David Zicarelli, CEO of Max's publisher, Cycling '74. Version 6, which launched in November, is built around a virtual canvas on which boxes containing pieces of code (sound generators, effects, video, algorithms) are connected onscreen by patch cords. If you can plug your Gibson into a wah-wah pedal, you can program in Max. But once you do, musical ideas can become connected to anything else that has a digital form.

Curtis Roads, professor of music at the University of California, Santa Barbara, has worked with a group at his university's AlloSphere new media collaborative to use Max to translate MRI brain images into music.

"You can send off agents into a brain generated from real data," Roads says. "These agents come back and sing a song that tells you what the rate of blood flow in that part of the brain is." The resulting sounds can be both interesting compositions and also useful for neurologists seeking subtleties in enormous and complex data sets.

"It's like looking at a beautiful clockwork mechanism," Roads says. "It's so beautifully designed with all these interacting components. There's a fascination of looking at algorithms—and listening to algorithms."

—MARK ANDERSON

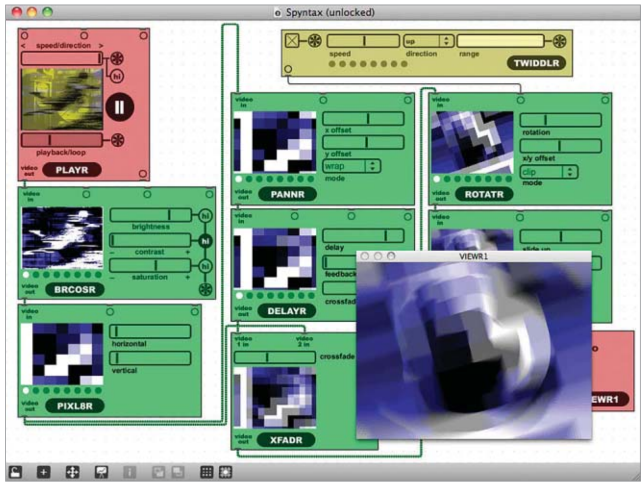

Max 6; $399; Cycling '74
http://cycling74.com/products/max

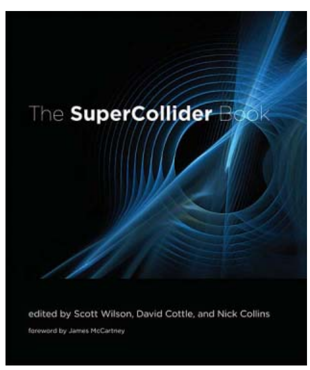

The SuperCollider Book
Edited by Scott Wilson, David Cottle, and Nick Collins; MIT Press, 2011; 680 pp.; US $45; ISBN 978-0-262-23269-2

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

technically speaking BY PAUL McFEDRIES

The Data Gold Rush

Intelligence about baseball had become equated in the public mind with the ability to recite arcane baseball stats. What [baseball statistician Bill] James's wider audience had failed to understand was that the statistics were beside the point. The point was understanding; the point was to make life on earth just a bit more intelligible. —Michael Lewis in *Moneyball* (2003)

ORGANIZATIONS OF all sizes are sitting on mountains of data; what they really need are **knowledge engineers** who can excavate nuggets of valuable information from that data. Earlier this year (in "The Coming Data Deluge," *IEEE Spectrum*, February 2011), I mentioned the concept of *data mining*, which uses sophisticated software and database tools to extract non-obvious patterns, correlations, and useful information from large and complex data sets.

Data mining begins with **data preprocessing**: the gathering of the raw data, which is stored in a **data warehouse** or **data mart**. It continues with **data cleansing**, which removes unrelated or unnecessary data (called **dirty data** or **noise**) and looks for missing information.

As the quote from Michael Lewis suggests, the point of data mining is **knowledge discovery**—the extraction of nonobvious or surprising information hidden in a data set. In data-mining circles, it's axiomatic that the less obvious the knowledge extracted, the more valuable that knowledge is to the organization. Nonobvious patterns represent new opportunities, be it for research, productivity, marketing, or whatever. This is best illustrated by the legendary **diapers and beer** connection, where data miners allegedly noticed that retail sales of diapers and beer would often spike in tandem. Why? Because new dads asked to pick up diapers on the way home from work would also pick up beer. When retailers stocked the products next to one another, sales of both were said to skyrocket.

Another term for finding previously unseen connections in a data set, especially when there are more than two variables, is **pattern mining**, and the quarried patterns are called **association rules**.

Many data sets consist of large amounts of text, such as e-mail, so data-mining projects typically use textual analysis to dredge up connections within that data, a process known as **text mining**. Another promising avenue is **audio mining** (also called **audio indexing**), which is the process of extracting and indexing the words in an audio file and then using that index as data to be otherwise mined. It will come as no surprise that engineers have also come up with ingenious methods for indexing other types of media, including **image mining** and **video mining**. If the data set consists of geographical information, it is called **spatial** (or **geospatial**) **mining**. In this increasingly social world, researchers are turning to **crowd mining**, where they try to unearth useful knowledge from large databases of social information. On a more general level, **Web mining** refers to the harvesting of useful patterns from data sets of Web content, Web usage (such as server logs), and Web structure (such as hyperlinks).

If a data set is just too large to probe efficiently, data miners can often get away with sampling portions of it, a technique variously known as **data dredging**, **data fishing**, or **data snooping**.

Data mining sounds innocent enough on the surface, but privacy advocates warn that it can be used for nonbenign purposes. When Internet service providers and companies such as Google hoard massive data sets that detail the online activities of hundreds of millions of people, **automated data mining** methods can analyze that data to look for patterns of suspicious activity. As computer scientist Jonathan Zittrain has pointed out, "When governments begin to suspect people because of where they were at a certain time, it can get very worrying."

Whether it's a boon or a bane, informative or intrusive, you've seen here that the field of data mining is a rich source of new words and phrases. As I see it, my job here at *IEEE Spectrum* is to sift through the raw material of articles, papers, blogs, and books to uncover new lexical gems and then present them to you in this column. Call it *word mining*. ❑

BRIAN STAUFFER

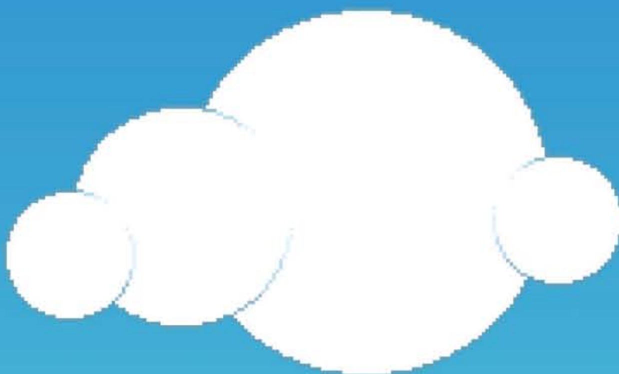

A Cloud You Can Trust

How to ensure that cloud computing's problems—data breaches, leaks, service outages—don't obscure its virtues

by CHRISTIAN CACHIN & MATTHIAS SCHUNTER

This past April, Amazon's Elastic Compute Cloud service crashed during a system upgrade, knocking customers' websites off-line for anywhere from several hours to several days. That same month, hackers broke into the Sony PlayStation Network, exposing the personal information of 77 million people around the world. And in June a software glitch at cloud-storage provider Dropbox temporarily allowed visitors to log in to any of its 25 million customers' accounts using any password—or none at all. As a company blogger drily noted: "This should never have happened."

And yet it did, and it does, with astonishing regularity. The Privacy Rights Clearinghouse has logged 175 data breaches this year in the United States alone, involving more than 13 million records.

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

Such statistics should give you pause, especially if you plan to entrust information that used to exist only on your own computer to an online provider's machines. And yet it's very hard these days to avoid that. Whenever you update your status on Facebook, check your e-mail via Gmail, post your vacation photos on Flickr, or shop, bank, or play games online, you are relying on somebody else's computers to safeguard your stuff. Many businesses, too, are buying into the promise of using computers they don't own or operate, because it gives them affordable and convenient access to computing resources, storage, and networking, as well as sophisticated software and services, that they might not otherwise be able to afford.

Regardless of how exactly they use such Internet-based computing services in "the cloud," these businesses stand to benefit. They gain in particular from the cloud's ability to pool equipment, allowing them to pay only for the resources they use and to scale their operations up or down almost instantaneously. Need more capacity? Just lease it from the burgeoning number of cloud providers, including Amazon, Google, Microsoft, or the company we work for, IBM. Cloud services also provide their customers with detailed metrics that track just how they use their cloud resources. And customers no longer have to wait around for the tech-support guy; their interactions with the cloud provider are almost entirely automated. So rather than being burdened with the expense and effort of procuring and maintaining an in-house computer network, even the smallest business can operate as if it had a world-class IT system.

More and more companies are doing just that. That's why, according to analysts at the technology-research firm Gartner, by next year 20 percent of all businesses will no longer own their own servers. That percentage is likely to grow in the coming years. In short, cloud computing is here to stay.

But this transformation of the IT landscape brings with it some new problems stemming from the very nature of outsourcing and from sharing resources with others. These problems include service disruptions and the inability of cloud providers to accommodate customized networks. But the top concern that businesses have with cloud computing, repeated surveys have found, is security—with good reason. By moving its data and computation to the cloud, a company runs the risk that the cloud-service provider, one of the provider's other customers, or a hacker might inappropriately gain access to sensitive or proprietary information. Customers just have to trust the cloud-service provider to safeguard their data. But unex-

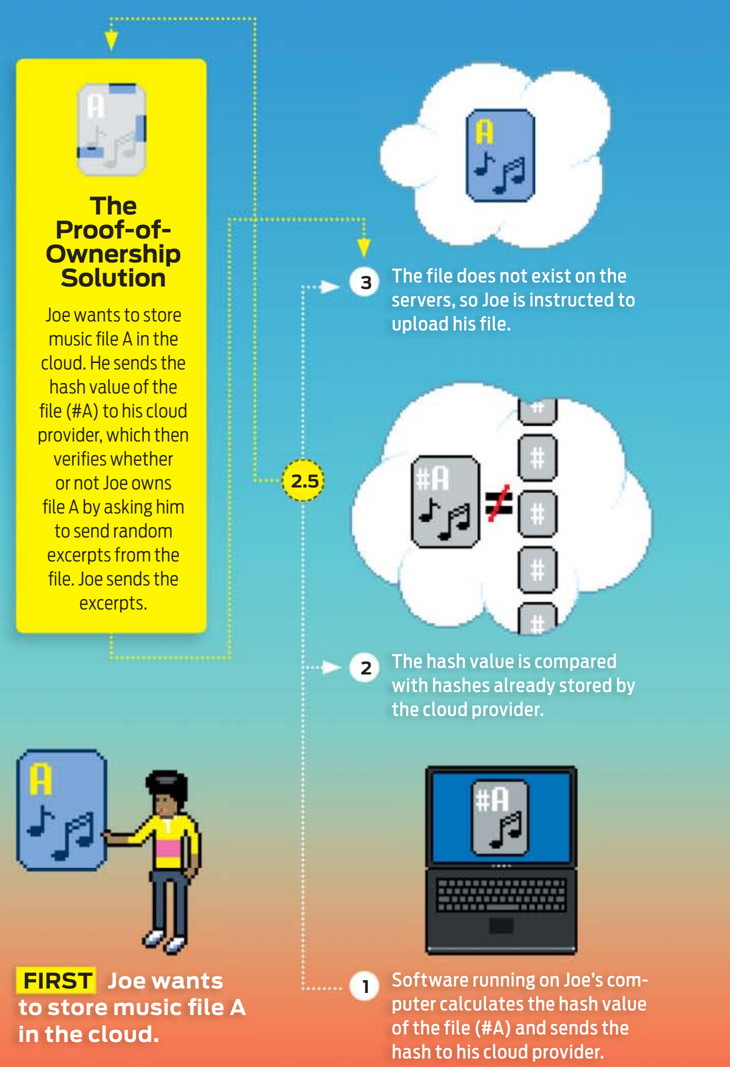

The Problem With Deduplication (and How to Fix It)

MANY CLOUD STORAGE providers use a technique known as deduplication to minimize the amount of data they store. Whenever a user tries to upload a file that has previously been stored, whether by that user or someone else, the cloud provider

ALL ILLUSTRATIONS: QUICKHONEY

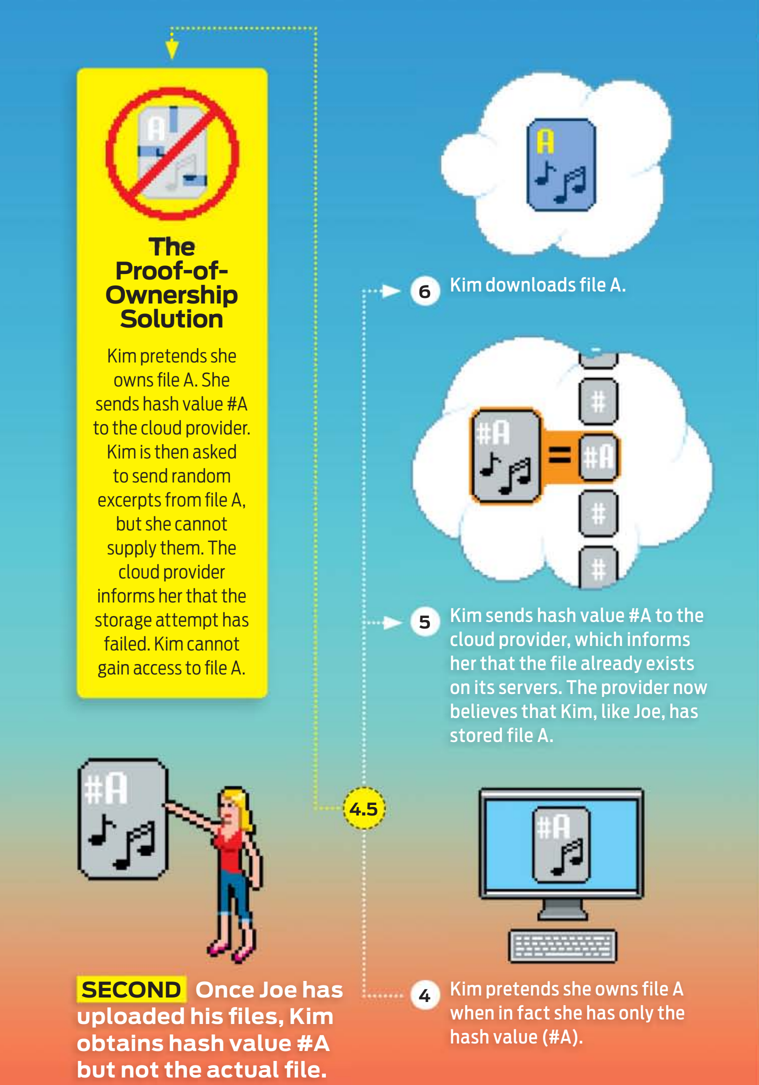

doesn't upload the redundant file; instead, it creates a link between the user's account and the existing file. The example shown here illustrates one problem with deduplication: Cloud customer Kim can surreptitiously download Joe's stored files, thereby exploiting the cloud service for unapproved content distribution. The proof-of-ownership solution, shown in steps 2.5 and 4.5, thwarts Kim's attempt to access Joe's files. The solution was developed by Shai Halevi and his colleagues at IBM Research and Bar-Ilan University, in Israel.

pected things can and do happen, even when you're dealing with well-established and presumably well-run companies. So it's no wonder that many IT managers remain jittery. If businesses are going to reap the full benefits of cloud computing, cloud providers will need to do much more to address security concerns. Here's an overview of how we think they could start.

ALTHOUGH "CLOUD COMPUTING" is the current buzz phrase, the concept has been around for half a century. In 1961, the late John McCarthy, an artificial-intelligence pioneer, proposed a different term for what is essentially the same thing: utility computing. "Computing may someday be organized as a public utility, just as the telephone system is a public utility," McCarthy said. "The computer utility could become the basis of a new and important industry."

But at the time, and for several decades afterward, computer hardware and software weren't up to the task. Only in the past few years, with the advent of high-bandwidth networking, Web-based applications, and powerful and cheap server technology, has McCarthy's vision finally been realized.

In many ways, cloud computing is just another form of outsourcing. But traditional outsourcing arrangements—contract manufacturing, say—come with legal, organizational, and technical controls. Cloud computing hasn't yet developed such protections. Nobody even agrees on what the best practices should be. Typical cloud-service agreements guarantee only that the provider will make its "best effort" to deliver services. Rarely do these providers pay any penalties should their services suffer an outage, breach, or failure.

What's more, many cloud providers do not consider security a top priority, according to a report released in April by CA Technologies and the Ponemon Institute. The study, which surveyed 103 providers in the United States and 24 in Europe, found that the majority "do not consider computing security as one of their most important responsibilities and do not believe their products or services substantially protect and secure the confidential or sensitive information of their customers." They also stated that it is "their customer's responsibility to secure the cloud" and that "their systems and applications are not always evaluated for security threats prior to deployment to customers."

If you want a higher level of security, you may be hard-pressed to find a provider that will customize its services to satisfy your concerns. Right now, cloud providers favor one-size-fits-all services: By offering a single and fully standardized cloud service, they can maximize economies of scale and thus lower costs. The downside, however, is that the cloud services they provide meet only the most basic requirements for security. That's fine for many of their customers, to be sure, but certainly not for all. Today's cloud is like the Model T Ford circa 1914: You can have any color cloud you want, as long as it's black.

Ironically, the greatest risk in cloud computing stems from its greatest advantage: resource sharing. Let's say you've developed an online game, and rather than buy your own servers, you lease computing time from a service like Amazon Elastic Compute Cloud, also known as EC2. That way, if your game becomes an overnight hit, you won't have to worry about thousands of players crashing your servers.

To run your code, you create what are called virtual machine images on the Amazon EC2 servers. Each VMI is the software equivalent of a stand-alone computer running its own operating system. In addition to specifying the number and type of VMIs you want, you can select where you'd like each one to reside. Amazon lets you choose among six geographic regions, each having one or more data centers. You can even spread out your virtual machines within a region by putting them in different availability zones. Each VMI will then be assigned to a physical server in a data center and will remain there as long as it is active.

Amazon promises that your virtual machines will be kept "virtually isolated" from those of other customers. But virtual isolation is not good enough. Research by Thomas Ristenpart and his colleagues at the University of California, San Diego, and MIT showed that a determined outsider stood a very good chance of putting his virtual machine onto the same server as another customer's VMI and then launching attacks from it.

Being located on the same server would give the attacker access to information about the target, such as cache usage and data traffic rates. And what could be done with that kind of information? Let's say the EC2 server is running encryption for one customer's VMI. To allow its CPU to run more efficiently, the server does a lot of caching; in this case, information about the encryption key might be available through the cache. And because the server cache is shared by the customers, the attacker might be able to access information about the key and thereby gain entry to the other customer's encrypted data.

More recently, Sven Bugiel and his colleagues at the Darmstadt Research Center for Advanced Security, in Germany, looked at a practice that allows Amazon cloud customers to publish their VMIs for others to use. Of the 1100 Amazon Machine Images (AMIs) the researchers looked at, about 30 percent contained private data that the creators had unintentionally published. These data included cryptographic keys, passwords, and security certificates, which attackers could extract and then use to gain illegal access to services that were built around the AMIs. And while both of these studies looked at Amazon's cloud services, such vulnerabilities aren't unique to them—they probably also exist on popular cloud services such as Microsoft's Azure and the Rackspace Cloud.

Data-storage clouds are also vulnerable. One class of attacks exploits a space-saving technique known as data deduplication. Many files that people upload to the cloud end up being duplicates—identical copies of software user manuals, say, or MP3s of Lady Gaga's "Telephone." Deduplication allows a data-storage cloud to keep only one copy of each file. Any time a customer attempts to upload a file, the contents of the file are first compared with other stored files. The new file is uploaded only if it doesn't already exist in the cloud; otherwise the customer's account is linked to the stored file.

In a paper they presented at the Usenix Security Symposium in August, Martin Mulazzani and his colleagues at SBA Research, in Vienna, described several ways in which deduplication could be used to access files uploaded to Dropbox. One way to do this involves hash values, which are short, unique digests assigned by an algorithm to a stored file. When a customer attempts to upload a file, the Dropbox software running on his computer first calculates the hash values of the chunks of data in the file. The hash values are then compared with hash values already stored by Dropbox. If the file does not yet exist, it gets uploaded.

But if customer Joe inadvertently shares the hash values for his stored files with customer Kim, or if Kim steals them from Joe, she will be able to freely download those files without Joe's permission or even his knowledge. Mulazzani and his colleagues worked with Dropbox to plug this and other security holes they had identified before going public with them.

CLEARLY, ISOLATING CUSTOMERS from one another is, or should be, a major concern for cloud providers. No data from one customer should be exposed to any other, nor should one customer's behavior affect another. With traditional outsourcing, isolation is achieved by maintaining dedicated physical infrastructure—separate production lines at a contract manufacturer, for instance—for each customer and by wiping clean all shared computers (such as workstations storing customer designs) before reuse.

To plug known security holes, cloud providers sometimes offer add-on services. For instance, Amazon Virtual Private Cloud allows the customer to specify a set of virtual machines that may communicate only through an encrypted virtual private network.

What Is Cloud Computing?

There are many informal understandings of what cloud computing is and isn't. Those of us in the field tend to use the definition outlined by the U.S. National Institute of Standards and Technology, which describes it as services that are accessed over a standardized network and that satisfy the following characteristics:

On-demand self-service

A customer can unilaterally order computing capabilities, such as server time and network storage, as needed and automatically, without any human interaction with the cloud provider.

Broad network access

Capabilities are available over the Internet and can be accessed from various platforms, such as mobile phones, laptops, and PDAs.

Resource pooling

The cloud provider's resources are pooled to serve multiple customers, with physical and virtual resources—including storage, processing, memory, network bandwidth, and virtual machines—dynamically assigned and reassigned according to customer demand. The customer generally has no control over or knowledge of the exact location of the provided resources but may be able to specify location at a higher level (such as country, state, or data center).

Rapid elasticity

To quickly scale up or down, capabilities can be rapidly ordered or released, in some cases automatically. To the customer, the capabilities available may appear to be unlimited and can be purchased in any quantity at any time.

Measured service

Cloud systems control and optimize resource use by measuring how each customer uses their services (such as storage, processing, and bandwidth). Monitoring, controlling, and reporting resource usage provides transparency for the provider and its customers.

EC2 also allows its users to define security groups, which operate like firewalls to control the incoming connections to a virtual machine.

There are a few products and services on the market aimed at enhancing cloud security. IBM's Websphere Cast Iron and Cisco IronPort, for instance, provide secure online messaging. Informatica Corp. offers businesses a way to protect sensitive information by masking it as realistic-looking but nonsensitive data. Eventually, though, cloud security systems will need to be fully automated, so that customers can detect, analyze, and respond to their own security issues, rather than rely on the cloud provider's staff for support and troubleshooting.

Other cloud security offerings include e-mail scanning, which checks for malicious code embedded in messages, and identity-management services, which control users' access to resources and automate related tasks, like resetting passwords. The trick, though, is for cloud providers to integrate the customer's own security measures with their operations.

That turns out to be hard to do, according to research by Burton S. Kaliski Jr. and Wayne Pauley of EMC Corp., in Hopkinton, Mass. They argue that the very features that define cloud computing—automated transactions, resource pooling, and so on—make traditional security assessments difficult. Indeed, they noted in a paper last year, these features offer opportunities for security breaches. For example, cloud providers continuously monitor and measure activity within their networks, to better allocate shared resources and keep their costs down. But the collection of that metering data itself opens up a security hole, say Kaliski and Pauley. A devious customer could, for example, infer behavioral patterns of other customers by analyzing her own usage.

DESPITE THESE AND OTHER vulnerabilities, cloud customers can do quite a bit to boost their security. Businesses typically operate their in-house data networks according to the principle of least privilege, conferring to a given user only those privileges needed to do his or her job. A bank teller, for instance, doesn't need full access to the bank's mainframes. The same rule should apply when using a cloud service. Though that may seem obvious, many companies fail to take this simple step.

Users can also do things to confirm the integrity of the cloud infrastructure they're using. Amazon CloudWatch allows EC2 users to do real-time monitoring of their CPU utilization, data transfers, and disk usage. The CloudAudit working group, a volunteer effort whose members include big cloud operators like Google, Microsoft, and Rackspace, is also exploring methods for monitoring the cloud's performance. In the future, trusted computing technology could make it possible for a customer to verify that the code running remotely in the cloud matches certain guarantees made by the cloud provider or attested to by a third-party auditor.

Of course, the strongest protection you can give to the information you send off premises is to encrypt it. But it isn't possible to encrypt everything: Data used in remote computations, for instance, cannot be encrypted easily.

IS THERE SUCH A THING as a totally secure cloud? No. But we, along with many other cloud-security researchers around the world, are constantly striving toward that goal.

One such effort we are involved in aims to develop and demonstrate a secure cloud infrastructure. With funding from the European Union and others, the three-year, €10.5 million (US $14.9 million) Trustworthy Clouds, or TClouds, project is a collaboration that includes our group at IBM Research–Zurich, the security company Sirrix, the Portuguese power companies Energias de Portugal and Efacec, and San Raffaele Hospital in Milan, as well as a number of universities and other companies.

TClouds is developing two secure cloud applications. The first will be a home health-care service that will remotely monitor, diagnose, and assist patients. Each patient's medical file can be stored securely in the cloud and be accessible to the patient, doctors, and pharmacy staff. Because of the sensitive nature of this information—not to mention the regulations that apply to patient privacy—TClouds will encrypt the data. The goal is to show how in-home health care can be improved cost-efficiently without sacrificing privacy.

The second application will be a smart street-lighting system for several Portuguese cities. Currently, the streetlights are switched on and off by means of a box that sits in a power station. The TClouds system will allow workers to log into a Web portal and type in when the lights should turn on and off in a given neighborhood; the use of smart meters will help control energy consumption. TClouds will show how such a system can run securely on a cloud provider's computers even in the face of hacker attacks and network outages.

TClouds will also build a "cloud of clouds" framework, to back up data and applications in case one cloud provider suffers a failure or intrusion. Recently TClouds researchers at the University of Lisbon and at IBM Research–Zurich demonstrated one such cloud of clouds architecture. It used a data-replication protocol to store data among four commercial storage clouds—Amazon S3, Rackspace Files, Windows Azure Blob Service, and Nirvanix CDN—in such a way that the data were kept confidential and also stored efficiently.

Although the technology is a major focus of TClouds, it is also addressing the legal, business, and social aspects of cloud computing. Many countries, for instance, have their own data-privacy laws, which will have to be considered carefully in cases where data must cross national boundaries.

Will everyone rush to adopt the kinds of improvements we're working on? Probably not. What we see happening down the road, though, is the diversification of today's one-size-fits-all approach to cloud computing. The demand for basic, low-cost cloud services will remain, but providers will also offer services with quantifiable and guaranteed security levels.

In the future, individual clouds will most likely give way to federations of clouds. That is, businesses will use multiple cloud providers for storage, backup, archiving, computing, and so on, and those separate clouds will link their services. (The social-networking sites Facebook and LinkedIn are already doing this.) So even if one provider suffers an outage, customers will still enjoy continued service.

Ultimately, we believe the cloud can be made at least as secure as any company's own IT system. Once that happens, reaching out to a cloud provider for your computing needs will be as commonplace as getting hooked up to the gas or electric company. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/cloud1211

KEY FIGURES: Ken Thompson [seated] types as Dennis Ritchie looks on in 1972, shortly after they and their Bell Labs colleagues invented Unix. *PHOTO: ALCATEL-LUCENT*

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

main( ) {

The Strange Birth and Long Life of Unix

/* The classic operating system turns 40, and its progeny abound */

by Warren Toomey

}

{

They say that when one door closes on you, another opens. People generally offer this bit of wisdom just to lend some solace after a misfortune. But sometimes it's actually true. It certainly was for Ken Thompson and the late Dennis Ritchie, two of the greats of 20th-century information technology, when they created the Unix operating system, now considered one of the most inspiring and influential pieces of software ever written.

```
11/3/71                                          SYS SEEK (II)

NAME            seek -- move read/write pointer

SYNOPSIS        (file descriptor in r0)
                sys    seek; offset; ptrname    / seek = 19.

DESCRIPTION     The file descriptor refers to a file open for
                reading or writing.  The read (or write) pointer
                for the file is set as follows:

                   if ptrname is 0, the pointer is set to offset.

                   if ptrname is 1, the pointer is set to its
                   current location plus offset.

                   if ptrname is 2, the pointer is set to the
                   size of the file plus offset.

FILES           --

SEE ALSO        tell

DIAGNOSTICS     The error bit (c-bit) is set for an undefined
                file descriptor.

BUGS            A file can conceptually be as large as 2**20
                bytes.  Clearly only 2**16 bytes can be addressed
                by seek.  The problem is most acute on the tape
                files and RK and RF.  Something is going to be
                done about this.

OWNER           ken, dmr
```

MAN MEN: Thompson (ken) and Ritchie (dmr) authored the first Unix manual or "man" pages, one of which is shown here. The first edition of the manual was released in November 1971.

A DOOR HAD SLAMMED shut for Thompson and Ritchie in March of 1969, when their employer, the American Telephone & Telegraph Co., withdrew from a collaborative project with the Massachusetts Institute of Technology and General Electric to create an interactive time-sharing system called Multics, which stood for "Multiplexed Information and Computing Service." Time-sharing, a technique that lets multiple people use a single computer simultaneously, had been invented only a decade earlier. Multics was to combine time-sharing with other technological advances of the era, allowing users to phone a computer from remote terminals and then read e-mail, edit documents, run calculations, and so forth. It was to be a great leap forward from the way computers were mostly being used, with people tediously preparing and submitting batch jobs on punch cards to be run one by one.

Over five years, AT&T invested millions in the Multics project, purchasing a GE-645 mainframe computer and dedicating to the effort many of the top researchers at the company's renowned Bell Telephone Laboratories—including Thompson and Ritchie, Joseph F. Ossanna, Stuart Feldman, M. Douglas McIlroy, and the late Robert Morris. But the new system was too ambitious, and it fell troublingly behind schedule. In the end, AT&T's corporate leaders decided to pull the plug.

After AT&T's departure from the Multics project, managers at Bell Labs, in Murray Hill, N.J., became reluctant to allow any further work on computer operating systems, leaving some researchers there very frustrated. Although Multics hadn't met many of its objectives, it had, as Ritchie later recalled, provided them with a "convenient interactive computing service, a good environment in which to do programming, [and] a system around which a fellowship could form." Suddenly, it was gone.

With heavy hearts, the researchers returned to using their old batch system. At such an inauspicious moment, with management dead set against the idea, it surely would have seemed foolhardy to continue designing computer operating systems. But that's exactly what Thompson, Ritchie, and many of their Bell Labs colleagues did. Now, some 40 years later, we should be thankful that these programmers ignored their bosses and continued their labor of love, which gave the world Unix, one of the greatest computer operating systems of all time.

THE ROGUE PROJECT began in earnest when Thompson, Ritchie, and a third Bell Labs colleague, Rudd Canaday, began to sketch out on paper the design for a file system. Thompson then wrote the basics of a new operating system for the lab's GE-645 mainframe. But with the Multics project ended, so too was the need for the GE-645. Thompson realized that any further programming he did on it was likely to go nowhere, so he dropped the effort.

Thompson had passed some of his time after the demise of Multics writing a computer game called *Space Travel*, which simulated all the major bodies in the solar system along with a spaceship that could fly around them. Written for the GE-645, *Space Travel* was clunky to play—and expensive: roughly US $75 a game for the CPU time. Hunting around, Thompson came across a dusty PDP-7, a minicomputer built by Digital Equipment Corp. that some of his Bell Labs colleagues had purchased earlier for a circuit-analysis project. Thompson rewrote *Space Travel* to run on it.

And with that little programming exercise, a second door cracked ajar. It was to swing wide open during the summer of 1969 when Thompson's wife, Bonnie, spent a month visiting his parents to show off their newborn son. Thompson took advantage of his temporary bachelor existence to write a good chunk of what would become the Unix operating system for the discarded PDP-7. The name Unix stems from a joke one of Thompson's colleagues made: Because the new operating system supported only one user (Thompson), he saw it as an emasculated version of Multics and dubbed it "Un-multiplexed Information and Computing Service," or Unics. The name later morphed into Unix.

Initially, Thompson used the GE-645 to compose and compile the software, which he then downloaded to the PDP-7. But he soon weaned himself from the mainframe, and by the end of 1969 he was able to write operating-system code on the PDP-7 itself. That was a step in the right direction. But Thompson and the others helping him knew that the PDP-7, which was already obsolete, would not be able to sustain their skunkworks for long. They also knew that the lab's management wasn't about to allow any more research on operating systems.

IEEE Spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

/* Unix was small, and you could go through it line by line and understand exactly how it worked */

—Ken Thompson, *Computer,* May 1999

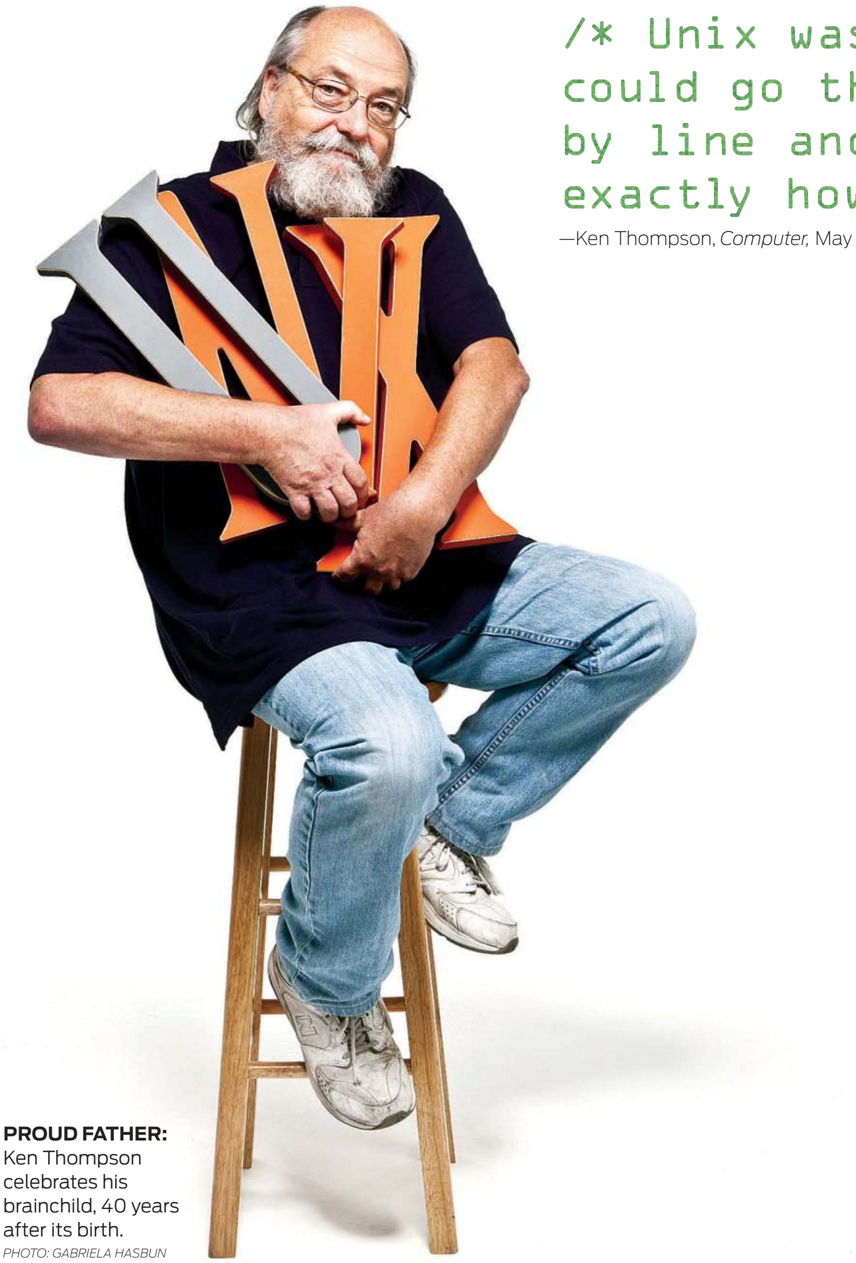

PROUD FATHER: Ken Thompson celebrates his brainchild, 40 years after its birth.

PHOTO: GABRIELA HASBUN

So Thompson and Ritchie got creative. They formulated a proposal to their bosses to buy one of DEC's newer minicomputers, a PDP-11, but couched the request in especially palatable terms. They said they were aiming to create tools for editing and formatting text, what you might call a word-processing system today. The fact that they would also have to write an *operating* system for the new machine to support the editor and text formatter was almost a footnote.

Management took the bait, and an order for a PDP-11 was placed in May 1970. The machine itself arrived soon after, although the disk drives for it took more than six months to appear. During the interim, Thompson, Ritchie, and others continued to develop Unix on the PDP-7. After the PDP-11's disks were installed, the researchers moved their increasingly complex operating system over to the new machine. Next they brought over the roff text formatter written by Ossanna and derived from the runoff program, which had been used in an earlier time-sharing system.

Unix was put to its first real-world test within Bell Labs when three typists from AT&T's patents department began using it to write, edit, and format patent applications. It was a hit. The patent department adopted the system wholeheartedly, which gave the researchers enough credibility to convince management to purchase another machine—a newer and more powerful PDP-11 model—allowing their stealth work on Unix to continue.

}

{

DURING ITS EARLIEST days, Unix evolved constantly, so the idea of issuing named versions or releases seemed inappropriate. But the researchers did issue new editions of the programmer's manual periodically, and the early Unix systems were named after each such edition. The first edition of the manual was completed in November 1971.

So what did the first edition of Unix offer that made it so great? For one thing, the system provided a hierarchical file system, which allowed something we all now take for granted: Files could be placed in directories—or equivalently, folders—that in turn could be put within other directories. Each file could contain no more than 64 kilobytes, and its name could be no more than six characters long. These restrictions seem awkwardly limiting now, but at the time they appeared perfectly adequate.

Although Unix was ostensibly created for word processing, the only editor available in 1971 was the line-oriented ed. Today, ed is still the only editor guaranteed to be present on all Unix systems. Apart from the text-processing and general system applications, the first edition of Unix included games such as blackjack, chess, and tic-tac-toe. For the system administrator, there were tools to dump and restore disk images to magnetic tape, to read and write paper tapes, and to create, check, mount, and unmount removable disk packs. *Continued on page 50*

IEEE spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

WHEN WILL SOFTWARE HAVE THE RIGHT STUFF?

UNMANNED PLANES DOMINATE THE BATTLEFIELD, YET AIRLINERS STILL HAVE PILOTS— AND COPILOTS

BY PHILIP E. ROSS

TIME WAS WHEN A UNIFORMED MAN would close a metal gate, throw a switch, and intone, "Second floor—men's clothing, linens, power tools..." and the carload of people would glide upward. Now each passenger handles the job with a punch of a button and not a hint of white-knuckled hesitation. The first automatic elevator was installed by Otis Elevator Co. in 1924; the things became common in the 1950s.

And back in the day, every train had an "engineer" in the cab of the locomotive. Then robo-trains took over intra-airport service, and in the past decade they have appeared on subway lines in Copenhagen, Detroit, Tokyo, and other cities.

Quietly, automation has taken charge of many other life-and-death functions. It manages white-hot ribbons of steel that shoot through rolling mills. It guides lasers that sculpt the eye and scalpels that excise the prostate gland. It runs oceangoing freighters, the crews of which have shrunk by an order of magnitude in living memory. And, most obviously, it is mastering aerial warfare. Today, the U.S. military trains twice as many ground operators for its unmanned aerial vehicles (UAVs) as pilots for its military jets. Its UAVs started off by flying surveillance missions, then took on ground attack; now they are being readied to move cargo and evacuate wounded soldiers.

In the sphere of commercial flight, too, automation has thinned the cockpit crew from five to just the pilot and copilot, whose jobs it has greatly simplified. Do we even need those two? Many aviation experts think not. "A pilotless airliner is going to come; it's just a question of when," said James Albaugh, the president and CEO of Boeing Commercial Airlines, in a talk he gave in August at the AIAA Modeling and Simulation Technologies Conference, in Portland, Ore. "You'll see it in freighters first, over water probably, landing very close to the shore."

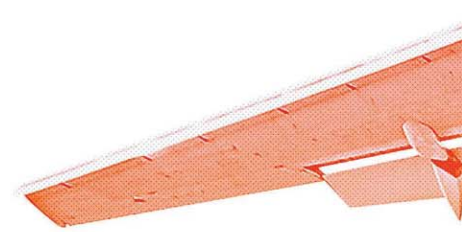

SEAN McCABE

Later, when air-traffic control systems rise to the challenge, pilotless planes will carry stuff to your very doorstep. In the fullness of time, they'll carry you.

Still, UAVs have yet to find a place in even the humblest parts of the aviation business—surveying traffic jams, say, or snooping on celebrity weddings. Such work has not yet been approved for routine purposes, even when the aircraft is small and controlled by a human on the ground—a man-machine meld that keeps a pilot in the loop. Why is it taking so long for the pilot to go the way of the elevator man?

TECHNICAL OBSTACLES can be cited: Fully automated planes can't yet visually identify nearby planes, and as for the remotely piloted ones, the civilian variety can't communicate with ground stations, because they haven't got enough bandwidth. "The name of your magazine is the underbelly of all the problems," says Doug Davis, director of the unmanned aircraft program at New Mexico State University, in Las Cruces. "There's not enough frequency spectrum. We have to compete for it with all these mobile phones and other demands."

But technical problems are hardly the entire explanation. The military has proved this point time and again, by trusting automation to do things that few human pilots could do on a fair day and with a following wind. In a stealth aircraft, for instance, software fiddles from moment to moment with the flight-control surfaces; otherwise, the radar-eluding design would make the plane fall like a brick.

And for nearly two decades, automatic landing systems have been able to drop and stop a jet on the fog-shrouded deck of an aircraft carrier that's barely twice as wide and three times as long as the jet's wingspan—*and* the ship is moving. Meanwhile, the pilot sits in the cockpit, hands folded.

"Look, there's no harder job for a pilot than landing on an aircraft carrier," says Missy Cummings, a former jet jockey for the U.S. Navy and now an associate professor of aeronautics and astronautics at MIT. "It's what Navy pilots have over those in the Air Force. And when I saw an F-18 land itself on an aircraft carrier, I knew my job was soon going to be over." That was in 1994. Automation has gotten rather better since then.

STEALTH DRONE: The U.S. Navy's X-47B, designed to fly from carriers, had its maiden cruise flight in 2011.
PHOTO: NORTHROP GRUMMAN

In fact, the U.S. Army has now decided that if you can trust the life of a pilot to software, you can trust the life of a wounded soldier, too. It's funding R&D into robotic medical evacuation vehicles. One such system, being developed by Piasecki Aircraft Corp. of Essington, Pa., and Carnegie Mellon University, is being designed around a system of sensors and software meant to launch and land a full-size helicopter on cluttered, unmapped ground and also fly the vehicle at low altitudes. Similar systems are under development in other countries. Israel's Urban Aeronautics is experimenting with a pilotless medevac that uses four rotors shielded by cowling so that they won't get tangled in shrubbery.

Given such advances, pilotless commercial flight is overdue, argues Cummings. Civilian UAVs could easily and profitably be deployed to survey infrastructure and carry cargo, she points out. And there's no reason why software, alone or perhaps in conjunction with a quickly mobilized ground controller, couldn't take over a piloted plane should something happen to both the pilot and the copilot. Already, she notes, an airliner's software typically takes over flight seconds after takeoff, handles the landing—and most of what happens in between. The pilot just "babysits," she says.

Of course, software that can meet only "most" of aviation's challenges would hardly satisfy the afraid-of-flying landlubber. That's why the pilot is still there, babysitting, until all the remaining kinks have been worked out. None of the problems are so bad as to prevent civilian pilotless planes from ever happening, but they are real, and they will have to be solved.

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

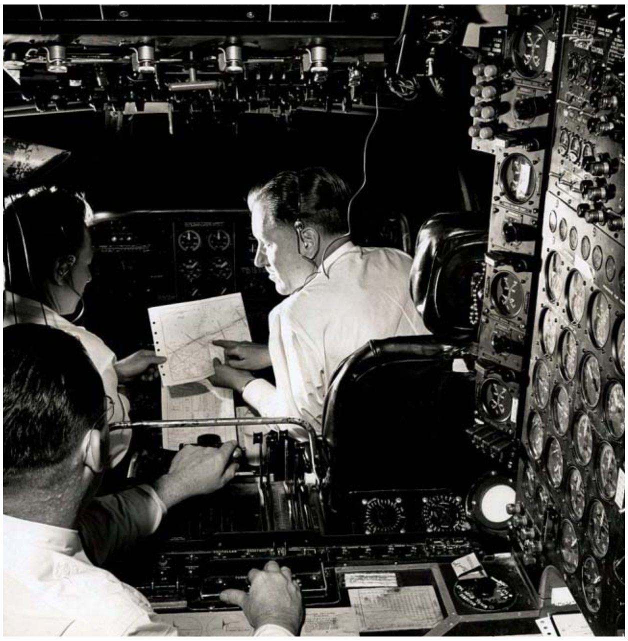

THREE'S A CROWD: A 1950s cockpit held three to five people; automation has pared the number to two—so far.
PHOTO: TOM HOLLYMAN/GETTY IMAGES

FIRST IS THE PROBLEM of "sense and avoid"—the requirement that every airplane be able to spot other aircraft and obstacles and get out of their way. It's a hard job for a machine. "Regulators always require that the pilot be able to look out the window and see if he's about to have a collision," says Rick Prosek, manager of the U.S. Federal Aviation Administration's (FAA) unmanned aircraft program office, which handles the regulations governing where in the civilian airspace UAVs may fly. "There are times when the technology doesn't give you the right answer," he says, and a human should be there "to say, 'Hey, I need to intervene.' "

Even UAVs that are remotely piloted need some sense-and-avoid capability, because there isn't enough bandwidth for sending lag-free, high-definition video to ground-based operators. Compare that with military UAVs, which have a network latency of less than half a second, because they link to their ground controllers half a world away via an owned-and-operated satellite system supplemented by other channels, such as undersea cables.

To go beyond remote control to truly autonomous flight, you'd need onboard software to interpret the data from the aircraft's cameras, radars, and other sensors and then to make good decisions. Autonomous planes would also need to play nicely with piloted planes, keeping their distance in the air, hewing to air-traffic controllers' directives, and avoiding fender-benders on the ground.

Here again the military has lessons for the civilian sector. Northrop Grumman has built some sense-and-avoid savvy into the unmanned helicopters and other UAVs it's developing for the U.S. Navy. The automation does, in three dimensions, pretty much what radar-enabled crash-avoidance systems now do in a few models of luxury cars (see *IEEE Spectrum*, "A Driver's Sixth Sense," October 2011), and like those systems, it keeps a human being in the loop. The company's upcoming X-47B, a fighter-size unmanned combat air system (UCAS), will take autonomous flight one step further. The test will involve taking off from an aircraft carrier, using GPS to fly a predetermined route, and landing on the carrier under light supervision—a minder, somewhere onboard the ship, who stands ready to take control if necessary.

But that's not quite the same thing as flying in crowded airspace, where even a 1 percent error rate would be intolerably high. Even a tiny civilian UAV will need some sort of sense-and-avoid capability before regulators will let it share the air with passenger planes. In this case, at least, the danger is small enough and the potential payoff great enough that agencies like the FAA seem ready to allow such small UAVs to fly within carefully delineated spaces. "In 2013 we anticipate we'll publish a final rule that would allow small UAVs to fly in certain parts of the national air space," says the FAA's Prosek. "Then we could benefit from the data they'd collect for us."

"WHEN I SAW AN F-18 LAND ITSELF ON AN AIRCRAFT CARRIER, I KNEW MY JOB WAS SOON GOING TO BE OVER"

—MISSY CUMMINGS

Having negotiated a sense-and-avoid situation, a pilotless plane would also need enough brains to maneuver away from an uncomfortably close vehicle, notify air-traffic control of what was happening, and then land safely. These challenges, however, are more like flying by the book than by the seat of your pants, and flying by the book is what computers do best.

Take the Global Hawk, a UAV developed by Northrop Grumman for maritime surveillance operations. It's able to fly itself home and land on its own if it loses its satellite link with its ground station. "The plane 'pings' the ground station to maintain communications, called a heartbeat," Cummings explains. "If the heartbeat's lost, the aircraft 'knows' and goes to self-repair mode, trying a second radio, checking circuit breakers, and so on. Then, if nothing works, the aircraft goes to its known profile, follows waypoints, and lands itself with GPS and radar."

Such planes weren't designed to have the parts-per-million defect rate that you'd expect in a passenger plane. (Indeed, the Global Hawk still hasn't fully met the Pentagon's specs, which require that a tag team of several planes provide "near-continuous" surveillance over a particular place for up to 30 days.) But even the Global Hawk's level of autonomy would surely be welcome in an airliner as a last-ditch way of saving it if the pilot and copilot were killed or incapacitated.

AS SIGNIFICANT as the technical hurdles are, however, by far the biggest impediment to pilotless flight lies in the mind. People who otherwise retain a friendly outlook toward futuristic technologies are quick to declare that they'd never board a plane run by software, which they know as the kludgy mess that makes their laptops freeze. But minds can be changed.

Each semester, Cummings asks her students at MIT whether they'd fly in a pilotless plane from Boston to Los Angeles. Two or three hands go up. Then she asks them how

they'd feel if the fare was just 50 bucks. More than half the hands go up. Cummings's little experiment suggests that people's reservations about robo-flight aren't set in stone.

One factor that's often cited for keeping a pilot in charge is what's known as "shared fate." That's the reassurance passengers get from knowing that the human in the cockpit wants to live just as much as they do. But shared fate is not the only way, or even the normal way, to ensure safe service. After all, restaurants don't employ food tasters to reassure diners, nor do losing defense lawyers join their clients in jail. It's usually enough for a professional to demonstrate sheer competence—the "right stuff" of aviator lore. And it's clear that automatic pilots—like those that land F-18s—now have a goodly amount of it.

Even Captain Chesley "Sully" Sullenberger, the pilot who deftly ditched his Airbus 320 airliner in New York City's Hudson River after its jet engines swallowed some geese, owed much to his onboard software—among other things, it managed the plane's angle of descent so as to avoid a stall.

But trusting software to safely shepherd hundreds of passengers across thousands of kilometers? A suspicious public isn't likely to buy into that vision, because safety is one of those things you can't have enough of. And how could a machine ever be as safe as a man-machine combo?

Yet the counterargument is also worth considering: Could the very collaboration of man and machine be causing human skills to wither? The FAA seems to think so. In a draft report cited by the Associated Press in July, the agency stated that pilots sometimes "abdicate too much responsibility to automated systems." Automation encumbers pilots with too much help, and at some point the babysitter becomes the baby, hindering the software rather than helping it. This is the problem of "de-skilling," and it is an argument for either using humans alone, or machines alone, but not putting them together (see *IEEE Spectrum*, "Automated to Death," December 2009).

Consider the 2009 crash of Air France Flight 447 off the coast of Brazil. A report from the French investigators, released this summer, argued that the pilots hadn't been properly trained to handle a malfunction of the airspeed sensors system, the apparent proximate cause of the crash. Thanks to automation, which had taken over many of their routine tasks, some of their skills had atrophied. The report suggested that had the copilot (apparently the one in charge at the time) kept the plane level while trying to figure out what was going on with the airspeed, he could have mastered the situation. Instead, the copilot seems to have held the plane's nose up, causing the very stall he was trying to prevent. The French report proposed that pilots get more experience in flying by the seat of their pants.

But it isn't easy for an airliner pilot to practice flying by the seat of his pants. Simulators hardly satisfy—isn't it the SimFlight world we're trying to get away from?—and even the occasional intervention in the plane's workings can't instill an intuitive feel for flying. And such interventions really are becoming occasional. Airliners—the Boeing 747, for instance—allow even landing to proceed hands free, says Davis, of New Mexico State. We are clearly in a particularly uneasy stage of the man-machine meld.

TO WIN OVER THE PUBLIC, the autopilots of tomorrow will have to start today by exploiting niches where civilian pilots can't or won't work—just as was the case in the military. With time, the systems will improve and eventually fan out to conquer additional segments of the broader market.

Look for the following order of events. Today, small government-run UAVs are plying ocean routes, looking for pirates and lost sailors; next, companies will send the craft into the back country, along routes cleared for their passage by civilian regulators, to check on the state of pipelines and power lines. After that, UAVs will ferry valuable medical samples and packages. A doctor might, for example, put a vial of blood into a UAV and send it to the nearest teaching hospital for analysis; a courier service such as FedEx might fly important packages from Japan to California, using dedicated airfields on each country's shoreline, thus avoiding civilian air traffic altogether. Once they're demonstrated on the battlefield, robotic medevacs will graduate to civilian duty, rescuing people stranded by flood or fire. Maybe then, after seeing such rescues on television, the flying public will finally start to warm to the cold machine.

"We often talk about the 'save little Johnny' scenario, where no human-operated aircraft wants to go out in a terrible storm, so you send out robotic aircraft to save little Johnny, alone in the ocean," said Rodney Walker, in an interview conducted last summer. Regrettably, Walker, a professor of electrical and electronic engineering and of aerospace engineering at the Queensland University of Technology, in Brisbane, Australia, died as this article was being prepared for publication.

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

LOOK MA, NO PILOT: The U.S. Navy's Fire Scout [left] can land on a carrier in a stiff wind; Piasecki Aircraft is testing this autonomous helicopter [center]; the Global Hawk [right] is a kind of pilotless version of the old U-2 spy plane.

PHOTOS, FROM LEFT: KELLY SCHINDLER/U.S. NAVY; PIASECKI AIRCRAFT CORP.; SENIOR AIRMAN NICHELLE ANDERSON/U.S. AIR FORCE

Walker's team, however, continues to work on his more immediately useful project—to use pilotless planes to survey power lines and find early signs of encroaching brush that might take the lines down in a storm. (It was apparently a wayward branch in Ohio that triggered the 2003 blackout in North America.) The idea is perfect for the empty parts of Queensland, a province with nearly three times the space and just a fifth of the population of Texas.

"Our energy distributors are desperately looking for new ways to surveil power lines," Walker told *IEEE Spectrum* in June. "They're taking the automated system we developed, but as a flight-assist system, and putting it into piloted aircraft for two to three years to prove the guidance and the automation. The pilot is largely hands off, but his presence allows us to meet regulations." He added that it would probably take no more than a decade to dispense with the pilot and fly the plane, perhaps autonomously and "certainly by remote control." Even then, such UAVs wouldn't be able to save little Johnny, but you could send out scores of them to sweep the seas for the tyke and then call in a piloted plane to pull him from the water.

Such baby steps in automation will eventually collide with another trend in aviation: the decades-long "decrewing" of airliners. Back in the day, the standard cockpit contained a full cast of characters, including the flight engineer, the navigator, and the radio operator. Technology has replaced those functions, and those people, one after the other.

The copilot is next in line. "Instead of a two-person cockpit, we'll see one person only, and software would serve as a backup," says New Mexico State's Davis. "And it'll come only after a certain period of performance, when the insurance industry is willing to accept the risk."

OF COURSE, money is the necessary fuel for any pilotless air business, and so far, nobody in the industry has spied a huge pile of it waiting to be made. The field still needs a killer application to motivate moneyed interests to lobby the government and win over the flying public.

"The FAA isn't jumping ahead of things, because there hasn't been an extreme business case for doing so," says Davis. "No killer biz application has stood up and said, 'Here I am.' "

Cutting pilots' and copilots' salaries won't save much money: They matter little when spread over a jumbo jet's worth of passengers. Perhaps more significant, though hardly overwhelming, are the indirect costs of having to schedule flights so that fresh, rested crew members are available. Here the savings would be greater because you'd be able to minimize an airliner's downtime.

However, any true killer app would probably dispense with the existing business model altogether and replace it with an entirely new model, one that might start in a small niche and grow to take over commercial aviation. That's how the PC made its way from a glorified word processor to a general-purpose tool for small businesses and, later, for medium-size corporations. At some point it began to eat at the bottom of the mainframe market, in classic "disruptive" fashion.

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/pilotless1211

For robo-flight, the killer app could start off with local air service. In many places just getting to the airport or to the main airport hub is the hardest leg of the trip. If small robo-planes could get you there, air travel would become vastly more attractive.

"Look at five-seat airplanes," says David Vos, senior director of unmanned aircraft systems and control technologies for Rockwell Collins, in Cedar Rapids, Iowa. "You could make a business out of a [pilotless] taxi service from one small airport to another, and it'd be a whole new world."

JUST WHEN WE MIGHT see that world is still the great question. Back in 2002, two illustrious technogeeks bet US $2000 on whether full-fledged pilotless airliners would fly routinely by 2030. And they placed the bet (on the website Longbets.org) before much of the huge progress in military UAVs had even been demonstrated.

Craig Mundie, chief research and strategy officer at Microsoft, thought such an outcome likely. He imagined "arriving at a methodology for system design that yields as much dependability, on an everyday basis, as the triple-redundant computer that flew guys to the moon." Eric Schmidt, now the executive chairman of Google, argued that "the FAA changes so slowly that if this were even all possible, the adoption and certification would all take at least 50 years."

That would be true, perhaps, if the United States were the only force in civilian aviation. But some other country, with lots of space and fewer people, may very well decide that commercial pilotless flight makes sense now. Australia, are you listening? ❑

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

*Nano*DYNAMITE

Fuel-coated nanotubes could provide bursts of power to the smallest systems

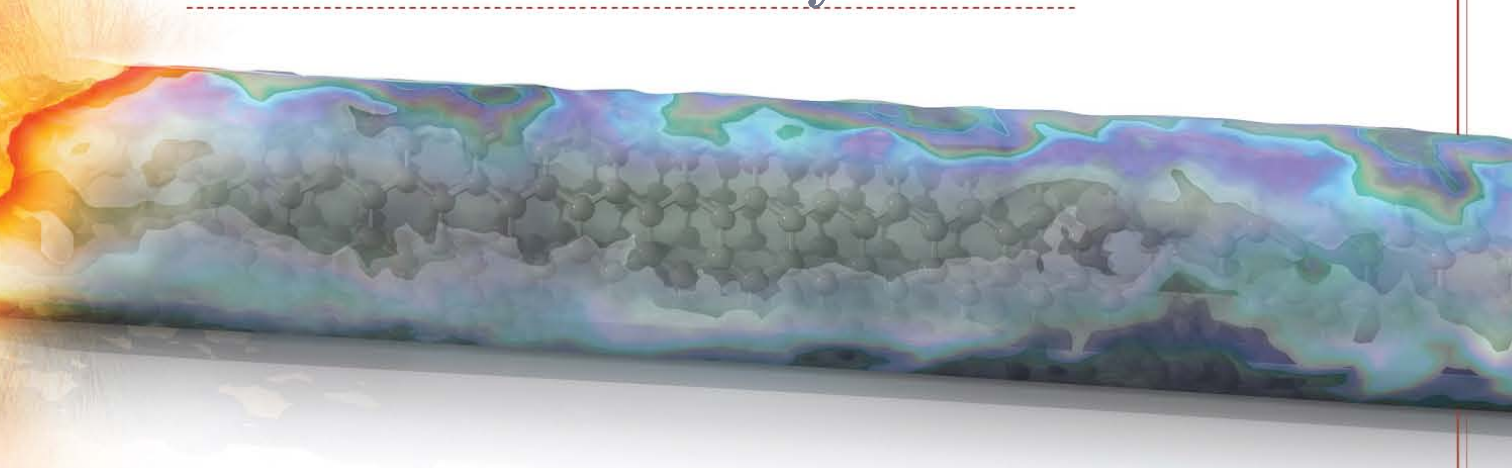

by Michael S. Strano & Kourosh Kalantar-zadeh
Illustrations by Emily Cooper

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

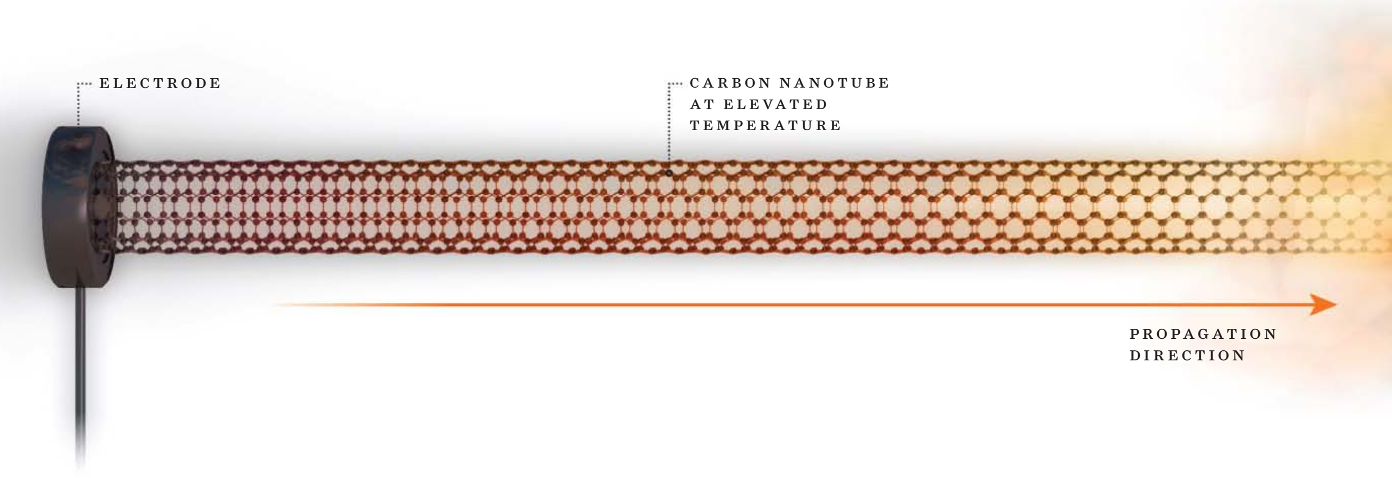

One of the greatest challenges in all of technology right now is improving energy storage.

It's an enormous challenge on many fronts and on many scales, with examples ranging from utility-scale batteries as big as a trailer down to the button-size cells that keep our quartz wristwatches ticking. And while we know how to make batteries bigger—add more cells—we are up against fundamental limits as we try to scale down rechargeable energy sources into the micro realm and beyond.

It's probably not a challenge you've thought much about, unless you're designing the next generation of nanoscale memory-storage devices, circuits, magnets, or biosensors, which all need extremely small and rechargeable power sources. Most of the many researchers working on rechargeable power sources these days are trying to wring a few more kilometers out of an electric-car battery pack or a few more minutes out of a cellphone charge. But the challenges of the micro realm are just as fascinating, even if they're not of commercial consequence yet.

They may soon be. Researchers already envision tiny systems whose realization is blocked by the lack of sufficiently small power sources. These future systems include cardioverter-defibrillators the size of apple seeds, implanted relatively unobtrusively in a heart patient to automatically control heart arrhythmias. Also, micro-size systems could one day attack cancerous cells with localized pulses of energy inside the body. Another possible application for these puny powerhouses is in "smart dust"—sensors that float in the air, collecting information about temperature, airborne pollutants, and other characteristics.

We think that a solution to this long-standing problem is finally within reach. Working at MIT's Strano Laboratory, along with colleagues from the Royal Melbourne Institute of Technology, in Australia, we have been pursuing a new approach to energy storage and power generation. Our experimental system, based on one of the new materials that have come from nanotechnology—carbon nanotubes—generates power in a way that has no macroscopic analogy. By coating a nanotube in fuel and igniting one end, we set off a combustion wave along it and learned that a nanotube is an excellent conductor of the heat from the burning fuel. Even better, the combustion wave creates a strong electric current.

BEFORE WE EXPLAIN how our nanotube-based power source works, it is worth considering the alternatives. For instance, why not just keep scaling down chemical batteries? The simple reason is that battery performance begins to degrade when engineers reduce a battery's size to several tens of micrometers. The flow of ions is disrupted, sharply reducing the battery's power density—the amount of power it produces per unit of mass.

What about other energy-storage technologies, such as rocket engines, which convert chemical energy to mechanical work? They don't typically scale to the sizes needed for the next generation of technologies, either. To date, scaled-down rocket engines can sustain power densities of only around 0.1 watt per kilogram—a minute fraction of the 200 W/kg that average lithium-ion batteries achieve.

A few newer technologies have the potential to power future micro- and nanoscale systems. Very small fuel cells, about the size of a peppercorn, might yet make promising power sources. But despite many years of research and innovation, fuel cells of this size have yet to produce sufficiently high levels of power per unit of mass. Fuel cells designed to work with microscale systems produce only 100 W/kg.

To understand how our tiny power sources work, first consider the carbon

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

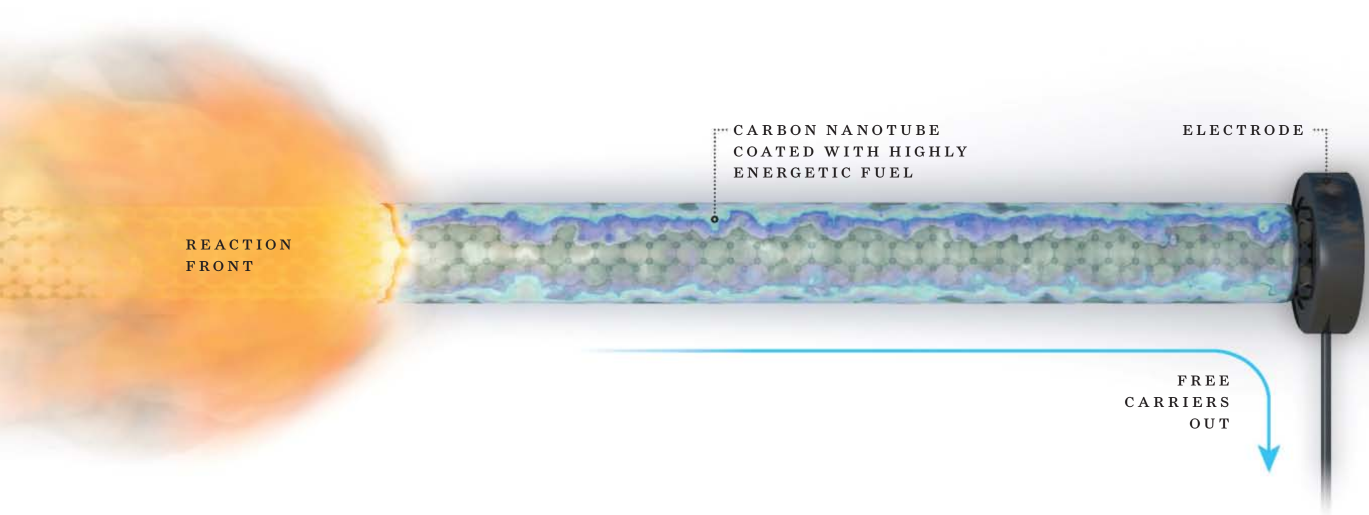

nanotube, a hollow, submicroscopic tube made of a chicken-wire-like lattice of carbon atoms. A single nanotube has an average diameter of 5 or 6 nanometers, but some have diameters as small as 1 nm. Their lengths vary: A nanotube can be as short as tens of nanometers, but researchers have grown some that are many millions of times as long—up to tens of centimeters. Nanotubes are extremely strong and have a high density of electrical carriers—electrons and electron deficiencies, or holes—which enable the nanotubes to conduct both electricity and heat very well.

Researchers have been studying carbon nanotubes for decades. But new advances over the past few years have made them easier to produce and use experimentally. These factors prompted us to look for ways to use carbon nanotubes to harness thermal energy in a previously undiscovered way. Basically, we wanted to take advantage of their shape and strength to propagate an explosive chemical reaction along the outside of the tube. If we could use the tube to harness the energy of such strong reactions, we reasoned, we would have a small system with exceptionally high power density.

In 2009, we began coating carbon nanotubes with a highly energetic fuel. When we ignited one end using either a laser beam or a hot wire, we found that the chemical reaction set off a wave of heat that propagated along the nanotube's walls like a flame along the length of a lit fuse. The resulting combustion wave traveled down the nanotube at speeds 10 000 times as fast as the fuel would burn in open air—from 0.01 up to 2 meters per second. That's because the core of the system—the nanotube—conducts heat so well. The heat entering the carbon nanotube propagates much faster than within the fuel itself, allowing the heat to ignite more fuel as it travels. The nanotube guides the decomposition of the fuel on the surface while keeping the reaction moving in one direction, thus serving as a guide for the wave and allowing a large current to flow unimpeded.

We call this combustion wave a thermopower wave because as it transmits energy from one place to another, it couples with the nanotube's electrical carriers, setting them in motion along the conducting tube. The moving heat source sweeps the coupled electrical carriers along, an effect we call electron entrainment. The carriers all move together in a single direction, creating an electric current that is extremely large relative to the mass of the system. In some of our latest experiments, the power generated exceeded 7 kilowatts per kilogram, or about three to four times what is possible with the best lithium-ion batteries currently available.

AS FAR AS WE KNOW, we are the first to harness these thermopower waves to convert chemical energy to electric power, and we're still trying to understand some of the fine points of how these waves propagate heat and electrical carriers to produce electricity. What we do know is that the propagation of the charge carriers depends on the thermoelectric effect—the voltage that results from a steady temperature difference across a wire. Depending on the nanomaterial, either electrons or holes flow from the hot side to the cold side; their density determines the current.

The main challenge with any power generator is, of course, to maximize its output. One way to accomplish this is by choosing materials with good thermoelectric properties. To understand what makes a good thermoelectric material, recall that power equals voltage times current. For a thermoelectric power generator, this translates to the voltage differential across the wire multiplied by its induced current. That means there are only two ways to increase power output—increase the voltage by increasing the temperature differential between the two ends of the wire or increase the current by decreasing the wire's resistance. So in choosing

How to Catch a THERMOPOWER Wave

A SINGLE CARBON NANOTUBE becomes a thermopower wave generator when it's coated in fuel and ignited on one end, which sets off a chemical reaction. The reaction starts moving along the nanotube's walls even faster than the fuel itself is consumed. This results in a high-speed thermopower wave, in which heat from the reaction feeds back to the fuel. As this wave travels, it pushes electrical carriers to the end of the nanotube, thus creating the electric current.

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

COMBUSTION on Camera

CANADIAN PHOTONICS LABS' high-speed CCD digital camera, rigged to begin taking photos at the start of a thermopower wave reaction, shows the waves in a generator in real time. The camera captured 4000 frames per second.

the material for the wire, researchers strive to maximize both the temperature differential—you want the greatest possible difference between the hot end and the cold end—and minimize the resistance.

It's no small feat. That's because materials that are highly conductive of electricity—metals, for example—are usually also good conductors of heat. So in almost all materials with low electrical resistance, it is impossible to sustain a large temperature differential between two points on the material—the high thermal conductivity precludes such a differential. That's exactly what French physicist Jean-Charles Peltier and Estonian-German physicist Thomas Johann Seebeck quickly learned when they independently discovered the thermoelectric effect in the first half of the 19th century—many metallic materials are weak when it comes to maximizing thermoelectric properties.

The first real breakthrough in designing materials for thermoelectric power generation came in 1958, when Tasmanian physicist H. Julian Goldsmid, now emeritus professor at the University of New South Wales, noticed that certain semiconducting materials—bismuth telluride and antimony telluride—have something unusual in common: high electric conductance and low thermal conductivity. This meant it might be possible to maintain a large temperature gradient in these materials. Early experimenters even dreamed of producing unlimited amounts of power; they imagined very long wires and rods, with one end in a hot environment, such as a geyser or volcanic crater, and the other end in a cool environment, perhaps several meters underground. Free and everlasting electric power!

Well, not quite. Although semiconducting materials like bismuth telluride and antimony telluride have lower thermal conductivity than any metal except mercury, even a small amount of conductivity is enough to prevent a material from sustaining a temperature difference between its two sides forever. Eventually, heat moves from one side to another, and a temperature balance is reached.

Researchers have tried for years to exploit the thermoelectric effect by manipulating materials with maximal electrical conductivity and minimal thermal conductivity. But despite extreme efforts, they have been stymied by the difficulty of that process. Ten years ago, scientists at the Research Triangle Institute, in North Carolina, were able to combine bismuth telluride and antimony telluride to make a material with almost no thermal conductivity at all. However, the fabrication process faced practical difficulties, and it's still far from being cost-effective for power generation. Researchers at Boston College; Wayne State University, in Michigan; Nanjing University, in China; and the Korea Institute of Science and Technology are also studying similar synthesized structures, but none are viable for everyday use yet.

That's where carbon nanotubes come in. Carbon nanotubes have extremely high thermal conductivity because of their crystal-like molecular structure. And like bismuth telluride, they also have high electrical conductivity. Our discovery that a thermopower wave works best across these tubes because of their dual conductivity turns conventional thermoelectricity on its head: It's the first nanoscale approach to power generation that exploits the thermoelectric effect but sidesteps the feasibility issues associated with minimizing thermal conductivity.

Here's why: The thermoelectric effect says that the more efficiently a material can sustain a temperature differential across it, the more electric potential it has. So materials with very little thermal conductivity but high electrical conductivity can sustain high electric potential, which determines a material's voltage. The thermoelectric effect suggests that the faster a material conducts heat, the more quickly it loses its electric potential.

But if, as with our system, voltage is related to how fast electrical carriers are moving—and the carriers are moving because of a high-temperature reaction—then thermal conductivity isn't a drawback at all. In fact, it's essential. It turns out that sustaining a temperature difference across a material isn't really important: It's all about how fast you can move electric potential down a wire. Our trick of using an explosive reaction exploits a carbon nanotube's thermal conductivity and its high concentration of electrons to speed up the chemical reaction. That reaction—rather than the ends of the tube—is what provides the needed temperature difference.

We're not the only ones trying to better understand how carbon nanotubes can generate power by means of

JOEL ABRAHAMSON/STRANO RESEARCH GROUP

the thermoelectric effect. Researchers at institutions such as Sungkyunkwan University, in South Korea; Texas A&M; Wake Forest University and NanoTechLabs, in North Carolina; and Victoria University, in New Zealand, are now looking at why carbon nanotubes can produce thermoelectric power and how to harness it better. What we know so far is that with thermal waves, the movement of a temperature gradient from one end of a material to the other is what creates voltage. And the faster it moves, the better: Higher thermal conductivity and higher temperatures mean a stronger electric current.

A thermopower-wave generator produces up to 0.2 to 0.3 volts and 0.1 to 0.2 amperes of electric current. The current is generated as a pulse, typically several milliseconds long. We can also increase the capacity and current by arranging nanotubes in parallel or increase the voltage by putting a number of nanotubes in series.

The current appears to scale up or down with wave velocity. We initially assumed that the true mechanism here wasn't wave velocity but rather wave-front temperature, itself the effect of a stronger thermal reaction that swept more electrical carriers along the tube. But after testing several fuels that reacted at different speeds on a nanotube, we confirmed that wave velocity is in fact the more important factor. We are still trying to explain why this is the case in order to understand more about how thermal waves couple with electrons and holes. One of our future goals is to find equations that help us calculate the relationship between the temperature, the voltage, and the wave-front behavior. So far, we've seen that the fuel reaction needs to generate localized temperatures around a thousand degrees Celsius in order to start a reaction that's fast enough to kick off the thermal wave.

While that might seem like an absurdly high temperature for use in any practical application, the high heat is contained within an area smaller than a cell. It's so highly localized and insulated by the nanotubes that we think the high temperature would be safe in almost any device, even ones inside the human body.

EARLY IN OUR RESEARCH, we began to see ways that we could modify and control carbon nanotubes, hoping to demonstrate their usefulness in future systems. Many electronics applications require only the sorts of power pulses that thermopower wave generators provide.

Because the length of the nanotube determines the duration of a reaction and how many charge carriers are entrained by the system, it therefore determines the energy and duration of the power pulse. By changing the length and choosing a fuel that supplies the right energy, we can effectively set the system's propagation velocity.

During our experiments, we made thermopower wave generators that had many different dimensions and properties and tested them in different conditions. Interestingly, the smallest generators produced the largest power densities.

Thermopower wave generators are also remarkable in other ways. For instance, before the nanotubes are ignited, the chemical energy can be stored in the fuel coating indefinitely. Batteries can leak or erode over time, but carbon nanotubes stay completely intact until lit, as well as after a reaction. Thermopower generators also have simple designs, which can take advantage of standard industrial micro- and nanofabrication technologies, and they have potential as nanogenerators. And they're rechargeable—reapplying the fuel is all it takes to launch another thermopower wave. Carbon nanotubes are sturdy enough to remain intact after the reaction, even though reaction temperatures exceed 1000 °C around the nanotube, and the devices made from them operate in open air.

Encouraging as some of these results have been, there is much we need to learn before we can turn carbon nanotube generators into a commercially viable power source. For instance, the 200 to 300 millivolts our systems have been able to put out so far isn't enough for most applications. We hope to find out whether different materials or mixtures of materials could produce more voltage. We'd also like to try liquid and gaseous fuels, which could work together with microfluidic systems.

On another front, we recently discovered that changing the conductivity of carbon nanotubes, by doping or other means, alters the propagation velocity of the thermal waves along the tube. And by changing certain properties of the fuel it should, in theory, be possible to produce both alternating and direct current with a single reaction; we've already seen that the rate of a reaction can decrease and increase as it travels.

Right now, we're trying to better understand the physics of thermopower waves. Already there are multiple angles to explore when it comes to taming these exotic waves and, ultimately, finding out if they're the wave of the future. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/nanodynamite1211

The Strange Birth and Long Life of Unix

Continued from page 37

UNIX ENGINE: Early versions of Unix ran on Digital Equipment Corp.'s PDP-11 minicomputers. *PHOTO: MARK RICHARDS/COMPUTER HISTORY MUSEUM*

Most important, the system offered an interactive environment that by this time allowed time-sharing, so several people could use a single machine at once. Various programming languages were available to them, including BASIC, Fortran, the scripting of Unix commands, assembly language, and B. The last of these, a descendant of a BCPL (Basic Combined Programming Language), ultimately evolved into the immensely popular C language, which Ritchie created while also working on Unix.

The first edition of Unix let programmers call 34 different low-level routines built into the operating system. It's a testament to the system's enduring nature that nearly all of these system calls are still available—and still heavily used—on modern Unix and Linux systems four decades on. For its time, first-edition Unix provided a remarkably powerful environment for software development. Yet it contained just 4200 lines of code at its heart and occupied a measly 16 KB of main memory when it ran.

UNIX'S GREAT influence can be traced in part to its elegant design, simplicity, portability, and serendipitous timing. But perhaps even more important was the devoted user community that soon grew up around it. And that came about only by an accident of its unique history.

The story goes like this: For years Unix remained nothing more than a Bell Labs research project, but by 1973 its authors felt the system was mature enough for them to present a paper on its design and implementation at a symposium of the Association for Computing Machinery. That paper was published in 1974 in the *Communications of the ACM*. Its appearance brought a flurry of requests for copies of the software.

This put AT&T in a bind. In 1956, AT&T had agreed to a

U.S government consent decree that prevented the company from selling products not directly related to telephones and telecommunications, in return for its legal monopoly status in running the country's long-distance phone service. So Unix could not be sold as a product. Instead, AT&T released the Unix source code under license to anyone who asked, charging only a nominal fee. The critical wrinkle here was that the consent decree prevented AT&T from supporting Unix. Indeed, for many years Bell Labs researchers proudly displayed their Unix policy at conferences with a slide that read, "No advertising, no support, no bug fixes, payment in advance."

With no other channels of support available to them, early Unix adopters banded together for mutual assistance, forming a loose network of user groups all over the world. They had the source code, which helped. And they didn't view Unix as a standard software product, because nobody seemed to be looking after it. So these early Unix users themselves set about fixing bugs, writing new tools, and generally improving the system as they saw fit.

The Usenix user group acted as a clearinghouse for the exchange of Unix software in the United States. People could send in magnetic tapes with new software or fixes to the system and get back tapes with the software and fixes that Usenix had received from others. In Australia, the University of Sydney produced a more robust version of Unix, the Australian Unix Share Accounting Method, which could cope with larger numbers of concurrent users and offered better performance.

By the mid-1970s, the environment of sharing that had sprung up around Unix resembled the open-source movement so prevalent today. Users far and wide were enthusiastically enhancing the system, and many of their improvements were being fed back to Bell Labs for incorporation in future releases. But as Unix became more popular, AT&T's lawyers began looking harder at what various licensees were doing with their systems.

One person who caught their eye was John Lions, a computer scientist then teaching at the University of New South Wales, in Australia. In 1977, he published what was probably the most famous computing book of the time, *A Commentary on the Unix Operating System*, which contained an annotated listing of the central source code for Unix.

Unix's licensing conditions allowed for the exchange of source code, and initially, Lions's book was sold to licensees. But by 1979, AT&T's lawyers had clamped down on the book's distribution and use in academic classes. The antiauthoritarian Unix community reacted as you might expect, and samizdat copies of the book spread like wildfire. Many of us have nearly unreadable *n*th-generation photocopies of the original book.

End runs around AT&T's lawyers indeed became the norm—even at Bell Labs. For example, between the release of the sixth edition of Unix in 1975 and the seventh edition in 1979, Thompson collected dozens of important bug fixes to the system, coming both from within and outside of Bell Labs. He wanted these to filter out to the existing Unix user base, but the company's lawyers felt that this would constitute a form of support and balked at their release. Nevertheless, those bug fixes soon became widely distributed through unofficial channels. For instance, Lou Katz, the founding president of Usenix, received

```
/* Unix is indeed one
of the most influential
operating systems ever
invented           */
///////////
```

a phone call one day telling him that if he went down to a certain spot on Mountain Avenue (where Bell Labs was located) at 2 p.m., he would find something of interest. Sure enough, Katz found a magnetic tape with the bug fixes, which were rapidly in the hands of countless users.

By the end of the 1970s, Unix, which had started a decade earlier as a reaction against the loss of a comfortable programming environment, was growing like a weed throughout academia and the IT industry. Unix would flower in the early 1980s before reaching the height of its popularity in the early 1990s.

For many reasons, Unix has since given way to other commercial and noncommercial systems. But its legacy, that of an elegant, well-designed, comfortable environment for software development, lives on. In recognition of their accomplishment, Thompson and Ritchie were given the Japan Prize earlier this year, adding to a collection of honors that includes the United States' National Medal of Technology and Innovation and the Association of Computing Machinery's Turing Award. Many other, often very personal, tributes to Ritchie and his enormous influence on computing were widely shared after his death this past October.

}

{

UNIX IS INDEED one of the most influential operating systems ever invented. Its direct descendants now number in the hundreds. On one side of the family tree are various versions of Unix proper, which began to be commercialized in the 1980s after the Bell System monopoly was broken up, freeing AT&T from the stipulations of the 1956 consent decree. On the other side are various Unix-like operating systems derived from the version of Unix developed at the University of California, Berkeley, including the one Apple uses today on its computers, OS X. I say "Unix-like" because the developers of the Berkeley Software Distribution (BSD) Unix on which these systems were based worked hard to remove all the original AT&T code so that their software and its descendants would be freely distributable.

The effectiveness of those efforts were, however, called into question when the AT&T subsidiary Unix System Laboratories filed suit against Berkeley Software Design and the Regents of the University of California in 1992 over intellectual property rights to this software. The university in turn filed a counterclaim against AT&T for breaches to the license it provided AT&T for the use of code developed at Berkeley. The ensuing legal quagmire slowed the development of free Unix-like clones, including 386BSD, which was designed for the Intel 386 chip, the CPU then found in many IBM PCs.

Had this operating system been available at the time, Linus Torvalds says he probably wouldn't have created Linux, an open-source Unix-like operating system he developed from scratch for PCs in the early 1990s. Linux has carried the Unix baton forward into the 21st century, powering a wide range of digital gadgets including wireless routers, televisions, desktop PCs, and Android smartphones. It even runs some supercomputers.

Although AT&T quickly settled its legal disputes with Berkeley Software Design and the University of California, legal wrangling over intellectual property claims to various parts of Unix and Linux have continued over the years, often involving byzantine corporate relations. By 2004, no fewer than five major lawsuits had been filed. Just this past August, a software company called the TSG Group (formerly known as the SCO Group), lost a bid in court to claim ownership of Unix copyrights that Novell had acquired when it purchased the Unix System Laboratories from AT&T in 1993.

As a programmer and Unix historian, I can't help but find all this legal sparring a bit sad. From the very start, the authors and users of Unix worked as best they could to build and share, even if that meant defying authority. That outpouring of selflessness stands in sharp contrast to the greed that has driven subsequent legal battles over the ownership of Unix.

}

{

THE WORLD of computer hardware and software moves forward startlingly fast. For IT professionals, the rapid pace of change is typically a wonderful thing. But it makes us susceptible to the loss of our own history, including important lessons from the past. To address this issue in a small way, in 1995 I started a mailing list of old-time Unix aficionados. That effort morphed into the Unix Heritage Society. Our goal is not only to save the history of Unix but also to collect and curate these old systems and, where possible, bring them back to life. With help from many talented members of this society, I was able to restore much of the old Unix software to working order, including Ritchie's first C compiler from 1972 and the first Unix system to be written in C, dating from 1973.

One holy grail that eluded us for a long time was the first edition of Unix in any form, electronic or otherwise. Then, in 2006, Al Kossow from the Computer History Museum, in Mountain View, Calif., unearthed a printed study of Unix dated 1972, which not only covered the internal workings of Unix but also included a complete assembly listing of the kernel, the main component of this operating system. This was an amazing find—like discovering an old Model T Ford collecting dust in a corner of a barn. But we didn't just want to admire the chrome work from afar. We wanted to see the thing run again.

In 2008, Tim Newsham, an independent programmer in Hawaii, and I assembled a team of like-minded Unix enthusiasts and set out to bring this ancient system back from the dead. The work was technically arduous and often frustrating, but in the end, we had a copy of the first edition of Unix running on an emulated PDP-11/20. We sent out messages announcing our success to all those we thought would be interested. Thompson, always succinct, simply replied, "Amazing." Indeed, his brainchild was amazing, and I've been happy to do what I can to make it, and the story behind it, better known. ❑

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/unix1211

ieee spectrum | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

Tufts UNIVERSITY | School of Engineering

John R. Beaver Professorship and Department Chair

The Department of Mechanical Engineering at Tufts University is pleased to announce the establishment of the *John R. Beaver Endowed Professorship in Mechanical Engineering*. The inaugural holder of this Professorship will also serve as Chair of the Department of Mechanical Engineering and lead the faculty in continuing to grow and strengthen a dynamic engineering program. We are seeking a forward-looking individual who is internationally recognized for the impact of his/her scholarship and is committed to interdisciplinary research and education. We expect to fill this position at the academic rank of full professor. Primary evaluation criteria will include a demonstrated ability to build and sustain cross-disciplinary research programs, leadership potential, and a track record of successful and inclusive mentoring of graduate and undergraduate students. Outstanding candidates from any subfield of mechanical engineering are encouraged to apply. It is anticipated that the successful candidate will contribute to one or more of the SOE's three strategic focal areas: Engineering for Sustainability, Engineering for Human Health, and Engineering the Human/Technology Interface.

The past seven years have been a period of extraordinary growth for Tufts' School of Engineering, witnessing recruitment of over half of its current tenured and tenure-track faculty members and close to a three-fold increase in research expenditures. The SOE distinguishes itself by the interdisciplinary and integrative nature of its engineering education and research programs within the environment of both a "Research Class 1" University and a top-ranked undergraduate institution. We offer the best of a liberal arts college atmosphere, coupled with the intellectual and technological resources of a major research university. As home to seven graduate and professional schools across three campuses, Tufts University prides itself on its culture of cross-School partnerships. Located on Tufts' Medford/Somerville campus, only six miles from historic downtown Boston, SOE faculty members have extensive opportunities for academic and industrial collaboration, as well as participation in the rich intellectual life of the region.

The Department of Mechanical Engineering currently consists of 13 full-time faculty, and the new Chair will have the opportunity to recruit a number of new faculty members to its ranks in the next few years. Faculty scholarship spans a broad spectrum of interests including materials engineering, human factors, biomechanical systems, electro-mechanical systems, and sustainable energy. The Department offers an ABET-accredited BSME, and MEng, MS, and PhD degrees in mechanical engineering, as well as undergraduate and graduate degrees in human factors. Recognized for its synthesis of theory and practice, the Department has an enrollment of approximately 200 undergraduate students and 90 graduate students.

See **http://engineering.tufts.edu/me/** for more information.

Candidates should possess an earned doctorate in Mechanical Engineering or a related discipline. Screening of applicants will begin immediately and continue until the position is filled. Applicants should submit a cover letter, curriculum vitae, a vision for research and education within the discipline of mechanical engineering, a statement of leadership style, and names and contact information for at least three references to Professor Jeffrey Hopwood, via email: **Jeffrey.Hopwood@tufts.edu** with copies to **Lorin.Polidora@tufts.edu**. Tufts University is an Affirmative Action/Equal Opportunity employer, committed to excellence in teaching and scholarship, and to building a faculty that reflects the diversity of both its students and the world for which it is preparing them. Members of underrepresented groups and women are strongly encouraged to apply and are invited to identify this status in their cover letters.

ELECTRICAL AND COMPUTER ENGINEERING
UNIVERSITY OF MICHIGAN, ANN ARBOR

The Electrical and Computer Engineering (ECE) Division of the Electrical Engineering and Computer Science Department at the University of Michigan, Ann Arbor invites applications for junior or senior faculty positions.

Successful candidates will have a relevant doctorate or equivalent experience and an outstanding record of achievement and impactful research in academics, industry and/or at national laboratories. They will have a strong record or commitment to teaching at undergraduate and graduate levels, and to providing service to the university and profession. Research areas of particular interest are networks and communications; electromagnetics and remote sensing; and nanotechnology.

The ECE Division (**www.eecs.umich.edu/ece**) prides itself on the mentoring of junior faculty towards successful careers. Ann Arbor is highly ranked as a best-place-to-live and for its family friendly atmosphere.

Please see application instructions at **www.eecs.umich.edu/eecs/jobs**

For full consideration applications must be received by January 6, 2012.

The University of Michigan is an Affirmative Action, Equal Opportunity Employer with an Active Dual-Career Assistance Program. The College of Engineering is especially interested in candidates who contribute, through their research, teaching, and/or service, to the diversity and excellence of the academic community.

Through Inspiration, Discovery
King Abdullah University of Science and Technology

Faculty Positions: Open Rank

KAUST is an international graduate-level, merit-based research university dedicated to advancing science and technology through bold and collaborative research and to addressing challenges of regional and global significance. Located on the Red Sea coast of Saudi Arabia, KAUST offers superb research facilities, generous assured research funding, and internationally competitive salaries. Further information can be obtained by visiting the following link: http://www.kaust.edu.sa.

The EE program currently has 13 full-time faculty and is recognized for its vibrant research programs and collaborative environment. EE research is strongly supported by KAUST's international research collaboration networks and KAUST's advanced research facilities, including the Nanofabrication, the Imaging and Characterization, and the Supercomputing Core Facilities. More information about the EE academic programs and research activities are available at http://ee.kaust.edu.sa.

The disciplines of EE critically support KAUST's interdisciplinary research initiatives in energy, environment, food and water, especially by advancing green technology, and energy efficient computers and communication systems through the use of novel nano, bio, info, and opto device technologies. To accelerate these initiatives, the EE Program seeks to work with other Programs and Centers in KAUST to aggressively strengthen research expertise in these integrative areas. Priority will be given to candidates with research interests in areas the following areas:

- **Circuits and Microsystems**
- **Nano-semiconductor material growth using MBE for photovoltaic**
- **Nanoelectronics and photonics applications**
- **Spectroscopy**
- **Nanoelectronics and nanoelectronics integrated systems**
- **Photonics devices**
- **Optical integrated systems**
- **Communication networks**
- **Wireless sensor networks**
- **Energy harvesting**
- **Robotics/mechatronics**

All candidates should have the ability to pursue a high impact research program and have a commitment to teaching at the graduate level. Applicants should submit, as a single PDF file to the Chair of the EE Search Committee (**ee@kaust.edu.sa**) including, a complete curriculum vitae with a list of publications, a research plan, a statement of teaching interests, and the names and contact information for at least 3 references for an Assistant Professor position or a list with the names and affiliation of potential referees for Associate Professor and Full Professor positions. Applications received by **Jan 30, 2012** will receive full consideration and positions will remain open until filled. Please identify the position in which you are applying to in the subject heading of the email.

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

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

the data

Tech Is Invented Globally but Adopted One Country at a Time

EVEN IN ANCIENT GREECE it was a truism that "Nothing endures but change," but according to the experts at Rand Corp., the future impact of new inventions and innovations will vary widely, depending on where you look on the globe.

A new analysis of a 2006 Rand report titled *The Global Technology Revolution 2020, In-Depth Analyses* reveals that only a handful of countries will be able to surpass the limitations found in the original report.

The 2006 study identified 16 revolutionary technologies that would be widely available in a decade. However, "not all countries will necessarily be able to acquire [a technology]—much less put it widely to use—within that time frame," the study's executive summary stated.

Rand chose a total of 29 countries representing various geographical regions, levels of social and economic development, and science and technology capacity. A country's technical capacity determines its ability to acquire a technology, either through its domestic R&D efforts, international R&D collaboration, or technology transfer, or simply by buying commercial systems from other countries. Of course, simply acquiring a technology does not mean a country will be able to put it to use. That depends on factors such as financial resources, policies, social values, and demand for the technology.

Only seven of the countries will have the capability to assimilate all 16 technologies. Eleven "scientifically lagging" countries, most from Africa and the Middle East, will be able to acquire only the five simplest applications, and the four most sophisticated technologies might prove too complicated for even the "proficient" countries.

Five years later, the report's lead author, Richard Silberglitt, says three tech applications are moving along more speedily than anticipated: rapid bioassays, pervasive sensors, and ubiquitous information access. What's changing even more drastically, he says, is the prowess of some of the "scientifically developing" countries. Brazil, Chile, South Africa, and Turkey "are really developing fast," he says. "They're not China and India yet but are emerging."

—*Prachi Patel*

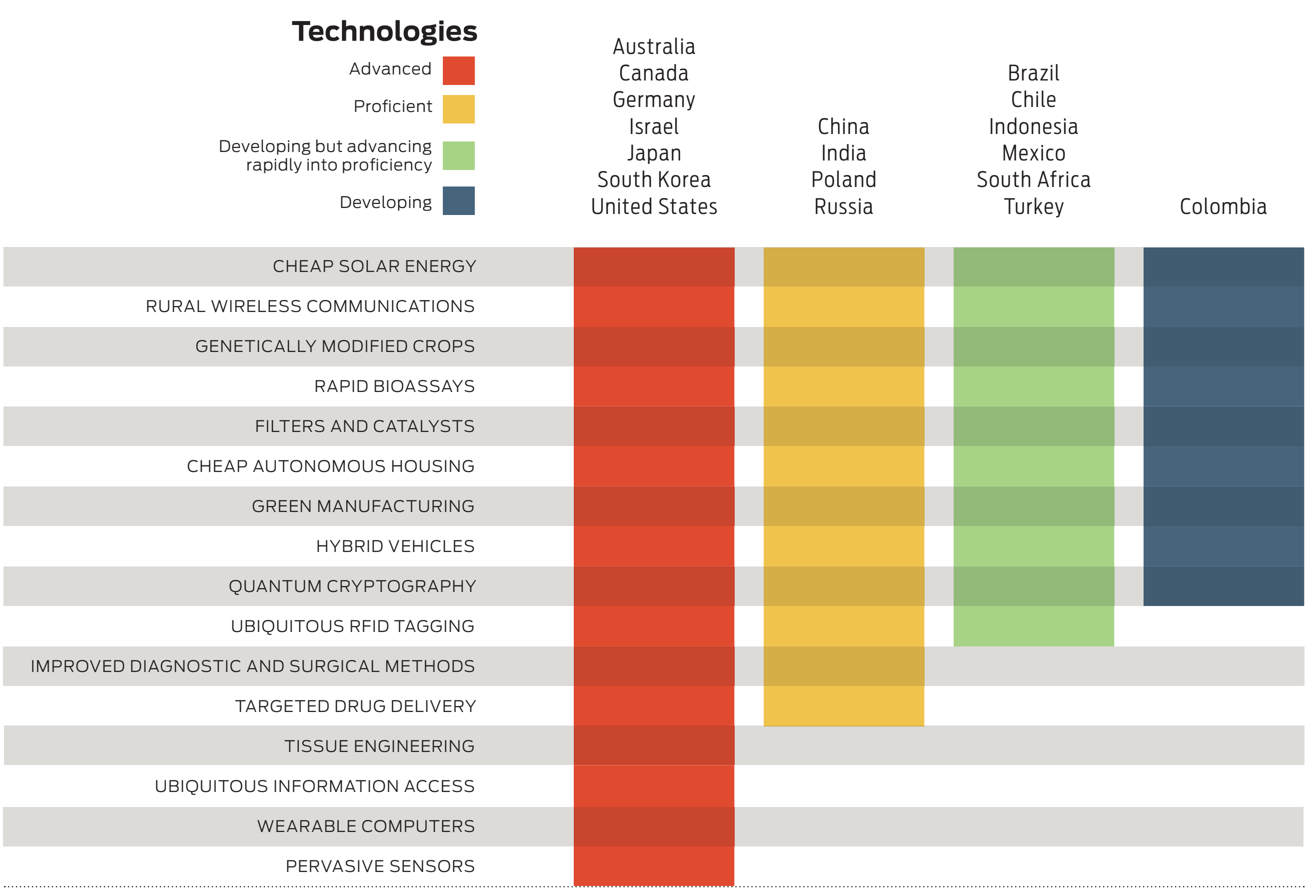

SOURCES: THE GLOBAL TECHNOLOGY REVOLUTION 2020 (RAND CORP. 2006); RICHARD SILBERGLITT

IEEE spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®


How can you get your idea to market first?
On the road to innovation, speed wins.
Accelerate your R&D and beat the competition to market. Instant access to over 3 million top-cited technology research documents can save hours—and keep your ideas in the fast lane.
IEEE Xplore® Digital Library
Discover a smarter research experience
Request a Free Trial
www.ieee.org/tryieeexplore
Follow IEEE Xplore on
IEEE
Advancing Technology for Humanity

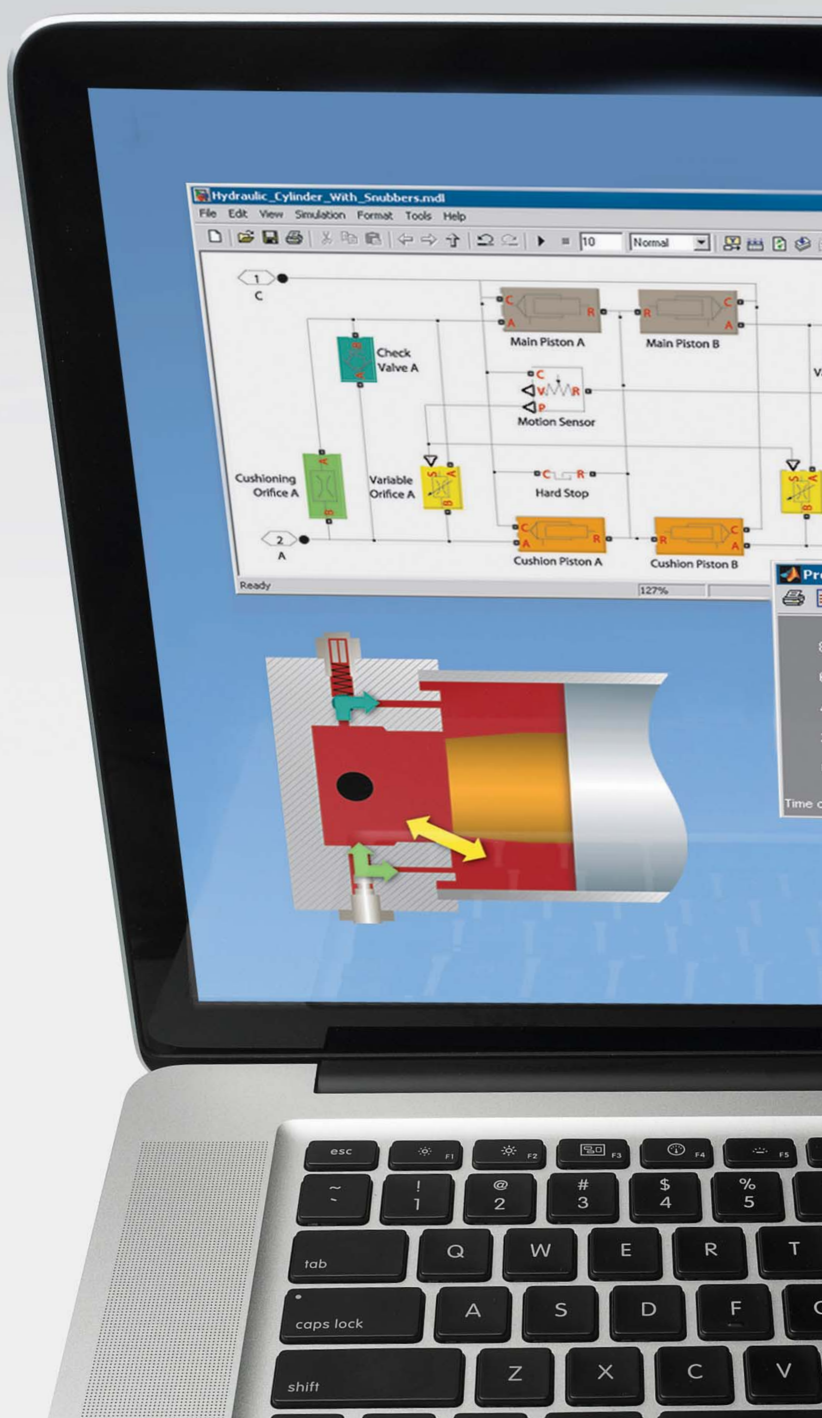
Hydraulic_Cylinder_With_Snubbers.mdl
Main Piston A
Main Piston B
Check Valve A
Motion Sensor
Cushioning Orifice A
Variable Orifice A
Hard Stop
Cushion Piston A
Cushion Piston B
Ready