


USING CHAOS TO CRUSH MALWARE ■ HOW TO GAME ONLINE GAMING

IEEE SPECTRUM

THE MAGAZINE OF TECHNOLOGY INSIDERS

DECEMBER 2007

www.spectrum.ieee.org




**THE TASER
SHOCKS
WITHOUT
KILLING.**

OR DOES IT?



Click here to access
the December issue of
**the
institute**



**HIGH-PERFORMANCE
IC DESIGN & VERIFICATION**

**HIGH-PERFORMANCE
SIGNAL & POWER INTEGRITY**

**HIGH-PERFORMANCE
RF & MICROWAVE DESIGN**

**HIGH-PERFORMANCE
EM SYSTEMS DESIGN**

ANSOFT.COM

**SIMULATION SOFTWARE
FOR HIGH-PERFORMANCE ELECTRONIC DESIGN**



IEEE **SPECTRUM**

DECEMBER 2007

VOLUME 44 NUMBER 12 INTERNATIONAL EDITION

Cover Story

18 How Tasers Work

Police officers worldwide are armed with electro-shock weapons, and their use is growing. Here's what happens at each pull of the trigger.

By Mark W. Kroll & Patrick Tchou

Business

26 THE R&D 100

To make money, you have to spend money, and in the engineering world that means spending on R&D. But the companies that spend the most don't always do the best.

By Ron Hira & Philip E. Ross

Hacking

30 PLAYING DIRTY

Automating computer game play takes cheating to a new—and profitable—level.

By David Kushner

Internet Security

36 CONTROLLED CHAOS

Network security researchers study information entropy to fight a new breed of superworms.

By Antonio Nucci & Steve Bannerman

Semiconductors

44 THE SILICON DIOXIDE SOLUTION

Jean Hoerni's planar process revolutionized the manufacture of silicon transistors and microchips.

By Michael Riordan

30 Gamer Richard Thurman hacked *Ultima Online* to generate virtual gold—and real cash.

This month on

SPECTRUM ONLINEwww.spectrum.ieee.org**Gut Check**

Scientists at the Institute for Food Research in Norwich, England, are using an artificial stomach to investigate a new generation of superfoods, as well as ways to fool the human stomach into thinking it's full.


CARS TAKE CHARGE

Check our Tech Talk blog from 2 to 5 December, as *IEEE Spectrum* car editor John Voelcker reports live from Anaheim, Calif., at the 23rd annual International Electric Vehicle Symposium and Exposition.

ALSO ONLINE:

- Webcasts • Radio • News
- Audio Downloads • Podcasts
- Jobs • Career Accelerator
- IEEE Xplore® digital library
- Interviews • Opinions

COVER: MARK LAITA; THIS PAGE:
PHOTO BY JEFF NEWTON; DIGITAL
ILLUSTRATION BY SANDBOX STUDIO

COMSOL
MULTIPHYSICS®

Electromagnetics Simulation goes Multiphysics!

Learn more!



Intro CDs in simulating RF
and AC/DC applications:
www.comsol.com/intro



The COMSOL® environment lets you conduct realistic computer simulations using multiphysics capabilities. COMSOL features a fully flexible model setup, easy-to-use application interfaces, and seamless CAD-import.

USER APPLICATION EXAMPLES

- Microwave and RF heating
- MEMS and RFID tags
- Sea bed logging
- SAR analysis
- Waveguides and photonics
- Antennas
- Semiconductor fabrication
- Plasma modeling
- Induction heating
- Piezoelectric modeling
- Motors and generators
- EMC and EMI
- Import from SPICE netlists
- HVDC

www.comsol.com

COMSOL, COMSOL MULTIPHYSICS, COMSOL REACTION ENGINEERING LAB, AND FEMLAB ARE
REGISTERED TRADEMARKS OF COMSOL AB. COMSOL SCRIPT IS A TRADEMARK OF COMSOL AB.

 **COMSOL**

IEEE SPECTRUM

DECEMBER 2007

VOLUME 44 NUMBER 12 INTERNATIONAL EDITION

NEWS

8 Robot Makers Fight Over \$280 Million Contract

The maker of the Roomba says stolen trade secrets led to its loss of a huge military contract. **By Erico Guizzo**

10 Carbon Nanotubes Cool Chips

11 A Messenger Arrives at Mercury

13 Catastrophic Climate Change

14 Zoo Network Could Halt Spread of Disease

16 THE BIG PICTURE Exploding Watermelon

OPINION

6 FORUM Fond memories of the HP 35 calculator and doubts about Google's green scheme.

7 SPECTRAL LINES Fairchild Semiconductor, an innovator in transistors and ICs, celebrates its 50th anniversary.

64 TECHNICALLY SPEAKING Just as computers turned us into connoisseurs of fonts, so has the Internet made us into cartographers. **By Paul McFedries**

RESOURCES

51 CAREERS An engineer turns her love of games into her livelihood. **By David Kushner**

53 TOOLS & TOYS This gun gets its ammo from a wall socket. **By Paul Wallich**

54 INVENTION Your patent's powers end at the water's edge. **By Kirk Teska**

56 BOOKS Global warming, explained in a nutshell. **Reviewed by M. Granger Morgan**

57 TOOLS & TOYS Pimp your iPhone with audio gadgets. **By Steven Cherry**

5 THE BACK STORY

Available 7 December on

**THE INSTITUTE
ONLINE**

www.ieee.org/theinstitute

**Scanning
The Globe**

Standards and new publications focus on Earth observations.

**EIGHT MENTORING
MYTHS BUSTED**

Get the truth behind some common misconceptions about mentoring.

**GET BUSY, HONOR A
TECH BREAKTHROUGH**

Learn how you can turn a historical achievement into a milestone and find out if one of 10 great breakthroughs awaiting nomination happened in your backyard.

14 This happy hippo will soon be part of a worldwide animal health network.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2007 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: \$19.50 included in dues. Libraries/institutions: \$205. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at W224-N3322 Duplainville Rd., Pewaukee, WI 53072-4195, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of American Business Media, the Magazine Publishers of America, and the Society of National Association Publications.

Streamline Development with NI LabVIEW Graphical Programming


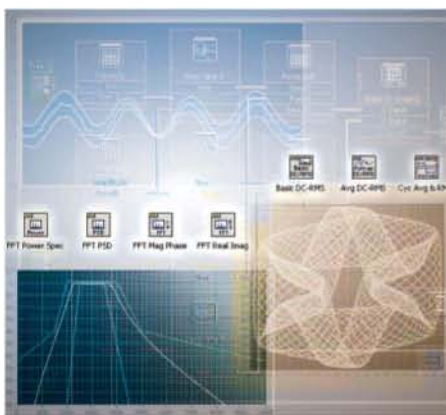
Acquire



Analyze



Present



Instantly Integrate Any I/O Device

- 5,000+ instruments from more than 200 vendors
- 1,000+ smart sensors including cameras
- 1,000+ PAC and PLC devices
- USB, GPIB, serial, Ethernet, PXI, PCI, PCI Express

Turn Raw Data into Critical Information

- 600+ math and signal analysis functions
- 150+ sound and vibration analysis functions
- 120+ communications analysis/design functions
- Integrated text-based math

Share Results Across Your Organization

- Library of user interface controls
- HTML Web-based reports
- Network communication
- Integration with Microsoft Word and Excel

Setting the Standard

Trust your test and measurement projects to National Instruments LabVIEW. With more than 20 years of engineering innovation, NI LabVIEW is the complete solution for acquiring and analyzing data and presenting results. Reduce development times with easy-to-use graphical programming and decrease overall system costs by scaling from simple to complex applications with a single software environment.

Thousands of companies worldwide rely on LabVIEW. See why graphical programming sets the standard.

Take the LabVIEW Guided Tour at ni.com/trylabview.

(800) 453 6202

IEEE SPECTRUM

Editor	Susan Hassler, s.hassler@ieee.org
Executive Editor	Glenn Zorpette, g.zorpette@ieee.org
Managing Editor	Elizabeth A. Bretz, e.bretz@ieee.org
Senior Editors	Harry Goldstein (Web), h.goldstein@ieee.org Jean Kumagai, j.kumagai@ieee.org Tekla S. Perry, t.perry@ieee.org Philip E. Ross (Resources), p.ross@ieee.org William Sweet, w.sweet@ieee.org
Senior Associate Editors	Steven Cherry, s.cherry@ieee.org Samuel K. Moore (News), s.k.moore@ieee.org
Associate Editors	Erico Guizzo, e.guizzo@ieee.org Sandra Upson, s.upsen@ieee.org
Assistant Editor	Willie D. Jones, w.jones@ieee.org
Senior Copy Editor	Joseph N. Levine, j.levine@ieee.org
Editorial Researcher	Alan Gardner, a.gardner@ieee.org
Administrative Assistants	Ramona Gordon, r.gordon@ieee.org Nancy T. Hantman, n.hantman@ieee.org
IEEE Spectrum Journalism Intern	Sarah Adee, s.adee@ieee.org
Interns	Francesco Ferorelli, f.ferorelli@ieee.org Morgen E. Peck, m.peck@ieee.org Joshua J. Romero, j.j.romero@ieee.org
Contributing Editors	John Blau, Robert N. Charette, Peter Fairley, Alexander Hellemans, Jen Lin-Liu, Robert W. Lucky, Paul McFedries, Kieron Murphy, Michael Riordan, Carl Selinger, Seema Singh, John Voelcker

ART & PRODUCTION

Senior Art Director	Mark Montgomery
Assistant Art Directors	Laura H. Azran, Brandon Palacio
Photo Editor	Randi Silberman
Director, Periodicals Production Services	Peter Tuohy
Editorial & Web Production Manager	Roy Carubia
Web Production Coordinator	Jacqueline L. Parker
Web Production Specialist	Michael Spector

IEEE MEDIA

Staff Director; Publisher, <i>IEEE Spectrum</i>	James A. Vick, j.vick@ieee.org
Associate Publisher,	
Sales & Advertising Director	Marion Delaney, m.delaney@ieee.org
Recruitment Sales Development Manager	Michael Buryk, m.buryk@ieee.org
Business Manager	Robert T. Ross
Marketing & Promotion Manager	Blanche McGurr, b.mcgurr@ieee.org
Interactive Marketing Manager	Laura Book, l.book@ieee.org
List/Recruitment Marketing Manager	Ilia Rodriguez, i.rodriguez@ieee.org
Reprint Sales	+1 212 221 9595, ext. 319
Department Administrator	Faith H. Jeanty, f.jeanty@ieee.org
Advertising Sales	+1 212 419 7760
Telephone Advertising	
Sales Representative	John Restchack +1 212 419 7578
Advertising Production Manager	Felicia Spagnoli
Senior Advertising Production Coordinator	Nicole Evans
Advertising Production	+1 732 562 6334
IEEE Staff Executive, Publications	Anthony Durniak

EDITORIAL ADVISORY BOARD Susan Hassler, Chair; Marc T. Apter, Alan E. Bell, C. Gordon Bell, Francine D. Berman, Emmanuel Desurvire, Hiromichi Fujisawa, Kenneth Y. Goldberg, Susan Hackwood, Erik Heijne, David H. Jacobson, Christopher J. James, Ronald G. Jensen, Mary Y. Lanzerotti, Tak Ming Mak, David A. Mindell, Fritz M. Morgan, Leslie D. Owens, Barry L. Shoop, Larry L. Smarr, Harry L. Tredennick III, Sophie V. Vandebroek, Başak Yüksel

IEEE Spectrum Editorial Offices
3 Park Ave., 17th Floor, New York, NY 10016-5997 U.S.A.
Tel: +1 212 419 7555 Fax: +1 212 419 7570
Bureau: Palo Alto, Calif.; Tekla S. Perry +1 650 328 7570
<http://www.spectrum.ieee.org>

IEEE Operations Center 445 Hoes Lane, Box 1331, Piscataway, NJ 08855-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

www.spectrum.ieee.org

THE BACK STORY



Michael Tamburro

Cooking by the Numbers

When we decided to let Spectrum Online visitors interact with six years of data from our annual R&D 100 report [see “The R&D 100,” in this issue], we turned to IEEE member Michael Tamburro and his colleagues at Agile Partners in New York City. We didn’t find Tamburro through his company’s Web site or a request for proposal. We found him in the kitchen.

IEEE Spectrum’s editor, Susan Hassler, interviewed Tamburro for Spectrum Online’s Geek Cooking podcast back in July [<http://spectrum.ieee.org/radio?01.07.07&segStart=2>]. As he recounted, he started cooking in his dorm room at Cornell University, exploring the complex world of Italian cuisine while honing his engineering skills. When he and colleagues Jack Ivers and John Berry founded the software company Agile Partners in 2002, Tamburro was well into perfecting the recipe from his Italian grandmother (his nonna) for a pasta-based dessert called *cruspolà*.

Cooking and software development don’t seem to have much in common at first glance. But as we discovered while working with Agile, whipping up a good Web app requires some of the same techniques used when experimenting with a new recipe. All good chefs taste their food during the course of preparation, adjusting the ingredients on the fly—a dash of salt here, a grind of pepper there, and the flavors, aromas, and colors meld into a feast for the senses. Agile, whose name comes from a programming methodology, emphasizes flexibility and iterations—lots of tasting in other words—before the final product is served. The result is our “R&D 100 Graph-o-Matic,” which you can sample at <http://spectrum.ieee.org/deco7/rndcalc>. As Nonna would say, “Buon appetito!”

CITING ARTICLES IN IEEE SPECTRUM

IEEE Spectrum publishes two editions. In the international edition, the abbreviation INT appears at the foot of each page. The North American edition is identified with the letters NA. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, the first News page is in *IEEE Spectrum*, Vol. 44, no. 12 (INT), December 2007, p. 8, or in *IEEE Spectrum*, Vol. 44, no. 12 (NA), December 2007, p. 10.

FORUM



“Data centers can require upward of 25 000 kilowatts, enough power to serve between 10 000 and 15 000 homes”
—Tom Schaeffer

GOOGLE GOES GREENISH

The article “The Greening of Google” [October] outlined the company’s plan to locate a data center in Oregon near “cheap and abundant” hydropower. This is presented as if that were somehow a positive thing consistent with its conservation efforts. While admittedly renewable, hydro is a limited resource that the company has elected to use up, most likely based on the “cheap” part of the quote.

Data centers can require upward of 25 000 kilowatts, enough power to serve between 10 000 and 15 000 homes or between 50 and 80 typical commercial customers. Most of that load is just for cooling the waste heat from the electronics. The hydro resource Google will consume must be replaced with some type of base-load generating unit that probably won’t have such a positive public relations value. Unless Google pays a “tap” fee equivalent to the tremendous costs associated with building all the facilities required to generate and transmit this additional energy, these costs will be paid for by others in the community. In return, the local community might benefit from as few as 20 new jobs.

The focus should continue to be on energy efficiency and conservation rather than “cheap

and abundant.” While growth and economic development are crucial to communities, electric loads such as these come with system and environmental costs that you never seem to find in a Web search.

Tom Schaeffer
IEEE Member
Longmont, Colo.

In the October issue’s Spectral Lines there’s a picture showing “Google’s Sergey Brin and Larry Page plugging in a RechargeIT hybrid electric car.”

You failed to mention that once the car was charged, they could’ve driven it to the airport to fly on their private Boeing 767 jetliner, which as we all know, is a very fuel-efficient mode of personal transportation.

Joseph Katz
IEEE Fellow
Stony Brook, N.Y.

NOT AS HARD AS WE THOUGHT

In Feng-Hsiung Hsu’s article “Cracking Go” [October], he lists the number of possible game positions for chess and Go at 10^{120} and 10^{170} . I’m perplexed by his number for chess. A simple analysis of Go indicates that, as an upper limit, not accounting for certain “impossible” configurations as restricted by the rules, there are three possible states (black stone, white stone, no stone) in

each of the 19-by-19 interstices. So that comes out to $3^{(19 \times 19)}$ possible game positions, or about 1.74×10^{172} —close enough to the 10^{170} number given by Hsu. If I apply this same analysis to chess, however, I fall quite a bit short. There are six unique white pieces (pawn, knight, rook, bishop, king, and queen) and six unique black pieces, as well as the possibility of no piece in a given spot, on a grid of 64 squares, for a possible 13^{64} game positions, or about 2×10^{71} . Even this number greatly overestimates things, because there cannot, for example, be 64 white pawns on a board! The gap between chess and Go is much wider than Hsu indicates.

Ben Thompson
State College, Pa.

Senior Editor Philip E. Ross responds: As Thompson and several other readers noted, there was an error in our table comparing chess with Go. It turns out that the best estimate of the number of possible game positions in chess is 10^{44} , 76 orders of magnitude lower than the number we listed. We regret the error.

BASIC CALCULATIONS

Thanks to IEEE Spectrum and Kenneth R. Foster for rekindling some old memories with the review of the new HP 35s calculator [Resources, October]. I was a Ph.D. student back in 1975 and somehow managed to find the money to buy a shiny, brand-new HP 35. It made me the envy of most of the other graduate students in the school of engineering at the University of Auckland. This was in the days when programming a computer involved a stack of punch cards, and you got the answer the next day—if you hadn’t made a mistake in your Fortran. I seem to remember the calculator had a

distinctive, almost sweet smell to the plastic.

The HP 35 worked flawlessly for several years and made it with me to Ottawa, where I had taken a job with Miller Communications in 1979. Unfortunately I dropped it one day, and a big ceramic IC cracked. HP wanted \$200 to fix the calculator. A new unit of similar performance (but without the nice smell) was about half that price. Still, I have regretted the decision to buy a second-rate all-plastic IC calculator ever since. I have not managed to kill it, though, even with numerous drops onto concrete.

Philip E.D. Wakeman
IEEE Member
Auckland, New Zealand

Your article reminded me of two calculator-related incidents in my past. First: some years ago, all the engineers in my company were given free TI calculators. I can’t remember the model number. Mine was stolen, and I had to buy my own replacement. Second: I won an HP sweepstakes for subscribers to an HP engineering magazine. The prize was the HP 19B, a business calculator!

You can’t win for losing.

Bob Schuchman
IEEE Life Member
San Diego

CORRECTION

In the October issue, the voltages were reversed in the “Quantum Tunneling Creates Fast Lane for Wireless” diagrams.

Letters do not represent opinions of the IEEE. Short, concise letters are preferred. They may be edited for space and clarity. Additional letters are available online in “And more Forum” at <http://www.spectrum.ieee.org>. Write to: Forum, IEEE Spectrum, 3 Park Ave., 17th floor, New York, NY 10016-5997, U.S.A.; fax, +1 212 419 7570; e-mail, n.hantman@ieee.org.

SPECTRAL LINES

Happy Birthday, Fairchild



The folks at Fairchild Semiconductor Corp. are tooting their horns in celebration of their recent 50th anniversary. It's a big deal for them, but why should we care? Because the computer and electronic devices we use depend in large part on the fundamental breakthroughs the founders of Fairchild made a half century ago. The creation of this company is a remarkable demonstration of how progress advances, in fits and starts, when the right set of individuals have the right conditions in which to work their magic.

For many tech insiders, especially those interested in the roots of modern computing, the Fairchild story is the stuff of legend. It begins with the invention of the transistor at AT&T Bell Telephone Laboratories.

In the years prior to World War II, AT&T's William Shockley probed the possibility of creating a solid-state alternative to the vacuum tube triode. After the war, he was put in charge of the group that developed the first transistor. For their work on this breakthrough, Shockley and his colleagues John Bardeen and Walter Brattain would receive the 1956 Nobel Prize in Physics.

Shockley was a brilliant but difficult man, however, and in 1955 he and Bell Labs parted ways. The following year Shockley convinced businessman Arnold Beckman to back his plans to create an advanced solid-state circuit design that would be as revolutionary to the transistor as the transistor was to the vacuum tube. But when he moved to Palo Alto, Calif., to start his company, not a single researcher in his old Bell Labs group accepted an offer to join him. So Shockley set about hiring some of the brightest young minds in America to form his fledgling company.

Among the 20 prodigies he recruited, Jean Hoerni, Gordon Moore, Robert Noyce, and a handful of others quickly realized that Shockley was not going to be able to advance the field further. And so the dissident staffers decided to start their own company.

In October 1957, the eight members of the newly formed Fairchild Semiconductor Co. opened shop. A furious Shockley labeled them the Traitorous Eight.

The Fairchild team's first effort was to commercialize a new

solid-state device called a mesa transistor. Encountering trouble with its performance, they experimented with novel ways of enhancing its reliability. Jean Hoerni [see Michael Riordan's account of Hoerni's contribution to the semiconductor industry, "The Silicon Dioxide Solution," in this issue] came up with an ingenious alternative. Here's an excerpt from one of Gordon Moore's accounts of the invention of the planar transistor:

Jean was a theoretician, and so was not very useful at the time we were setting up the original facility at Fairchild, building furnaces and all that kind of stuff. He just sat in his office, scribbling things on a piece of paper, and he came up with this idea for building a transistor with the silicon oxide layer left on top over the junctions. Where the silicon junctions come to the surface of the silicon is a very sensitive area, which we used to expose and had to work awfully hard to keep clean. Hoerni said, "Why not leave the oxide on there?" The conventional wisdom from Bell Laboratories had been that by the time you got done, the oxide was so dirty that you wanted to get rid of it. Nobody had ever tried leaving the oxide on. We couldn't try it either, because it required making four mask steps, each indexed with respect to the next with very high precision—a technology that didn't exist.

So we couldn't even try Jean's idea until a year and a half or so after we had gone into business. When we finally got around to trying it, it turned out to be a great idea; it solved all the previous surface problems.

Noyce quickly saw the potential of the new manufacturing method. He realized that by using the planar process a designer could re-create the components found on a typical circuit board of the time and etch them onto the silicon wafer itself. The aluminum layer used to make contact with the base and the emitter of the transistor could also be used to interconnect different electrical components such as resistors and capacitors. It was the second breathtaking advance at Fairchild in a year. Noyce had conceived an integrated circuit that could be commercially developed—thus laying the foundation of modern computing.

The Fairchild Eight went on to pioneer other improvements to microchip technology. Eventually, most of the principals moved on to create other companies. They blazed an entrepreneurial trail through the world of electronics that few have matched [for a look at Fairchild's corporate genealogy, see "Fairchild Turns 50," *IEEE Spectrum*, October].

So when we get giddy about the pace of technological breakthroughs today, we should take a moment to remember that a small band of workers rolled up their sleeves 50 years ago and set our digital world in motion. Thanks, Fairchild!

—Kieron Murphy

The editorial content of *IEEE Spectrum* magazine does not reflect official positions of the IEEE or its organizational units. Please address comments to Forum at n.hantman@ieee.org.



\$280 Million Robot Dustup

Roomba maker accuses military contract winner of stealing trade secrets

Some robotic face-offs take place in gladiatorial arenas, others on ping-pong-table-size soccer fields. Those fought in courtrooms can be just as much fun to watch, because sometimes they come complete with dumpster-diving private investigators, accusations of planted evidence, erased computer disks, shredded data CDs, and trade secrets discussed in closed hearings.

That is the case in a recent dispute between iRobot, in Burlington, Mass., known for its Roomba vacuum cleaner, and a smaller rival, Alsip, Ill.-based Robotic FX. iRobot has filed two lawsuits against Robotic FX and its founder and president, Jameel Ahed, a former iRobot employee, alleging patent infringement and theft of trade

secrets. The suits concern iRobot's PackBot, a bomb-disposal robot widely used in Iraq and Afghanistan. iRobot accuses Robotic FX of using proprietary PackBot technology to create a competing robot called the Negotiator. Robotic FX denies the accusations and says that the lawsuits are an attempt to shut down a competitor that iRobot now sees as a threat.

In August, the two companies competed in a U.S. Army program called xBot, whose goal is to procure a smaller, lighter type of bomb-detection robot than those currently used in Iraq. The Army plans to deploy up to 3000 of the new robots in the next five years as part of efforts to counter improvised explosive devices (IEDs), which are

ROBOTIC FX

responsible for nearly half of all coalition troop casualties in Iraq.

Analysts regarded iRobot, a 375-employee publicly traded company with revenues of US \$189 million last year, as the favored bidder. But privately held Robotic FX, which reportedly has eight employees, won the contract, valued at \$279.9 million. The news knocked iRobot's stock down nearly 30 percent in the following days.

Last month iRobot was granted a preliminary injunction prohibiting Robotic FX from selling the Negotiator in its current design. The Army, which initially opposed iRobot's request for an injunction, arguing that it could delay the delivery of robots to the troops, later decided to freeze the contract and reevaluate Robotic FX's ability to carry out the contract.

The stakes are high for both sides. Ahed revealed during the court hearings that an unnamed major defense company is interested in acquiring Robotic FX, a deal that could reward its founder handsomely. For iRobot, the possibility of its allegedly stolen designs falling into the hands of a large, deep-pocketed competitor is a worrisome development.

THE ARMY CREATED the xBot program to address a pressing need of U.S. troops in Iraq. The robots currently used by specialized bomb-disposal squads are too big and too heavy for regular soldiers on patrol and convoy missions to take with them. A smaller, lighter robot would allow troops to inspect suspicious objects before calling the bomb squad.

To procure the robots, the Army prepared a set of requirements and organized a two-stage competition: a technical test, to see which robots met the requirements, and a reverse auction, in which the participant making the lowest bid would get the contract.

The technical test took place at the Redstone Arsenal in Huntsville, Ala., in August. The robots—required to weigh no more than 22.6 kilograms and feature a manipulator arm, among other things—had to traverse sand, gravel, and water pits, maneuver their arms to lift objects, and position their cameras, in scenarios that simulated IED investigations.

iRobot brought two robots: a lighter version of PackBot and a new 14-kg robot called the Small Unmanned Ground Vehicle, developed under the Army's Future Combat Systems program. Robotic FX brought its Negotiator. All three robots passed the test.

The reverse auction occurred the following month. The starting bid was \$305 million, and the companies kept lowering their bids until iRobot gave up. Robotic FX won the contract with a bid of \$279.9 million (iRobot's final bid was \$286 million).

iRobot, not surprisingly, disapproved of the procurement process. A reverse auction is normally used to buy commoditized products such as spare parts—not advanced systems like robots, says Joseph Dyer, president of iRobot's government and industrial division.

The procurement was also unusual for its brevity and detailed requirements, according to several robotics executives interviewed by *IEEE Spectrum* on condition of anonymity. It appeared, these sources say, that the Army knew exactly which robots it wanted. "It wasn't normal; it was very quick," says a senior executive familiar with the xBot program whose company has contracts with the military.

A spokeswoman for the Program Executive Office for Simulation, TRaining, and Instrumentation (PEO STRI), the Army organization in charge of the xBot contract, says its procurement process "was the best way to satisfy the urgent requirement for robots." PEO STRI officers, based on a "preaward survey," had determined that Robotic FX was capable of fulfilling the 3000-robot contract, but late in October they decided to put the award aside and conduct another assessment.

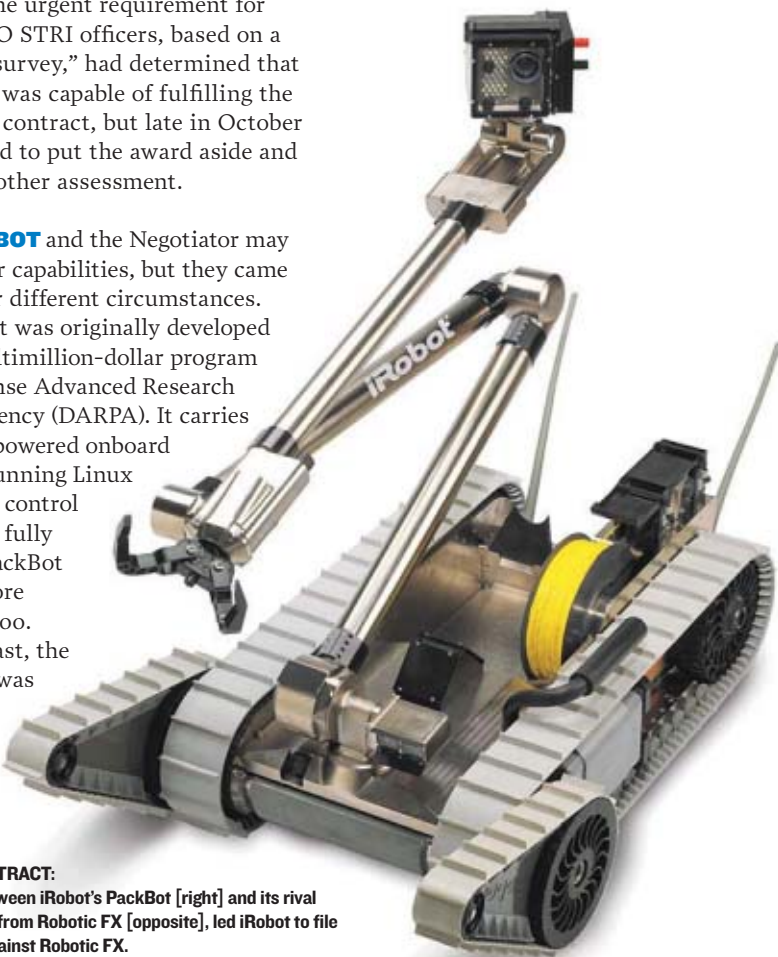
THE PACKBOT and the Negotiator may have similar capabilities, but they came about under different circumstances. The PackBot was originally developed under a multimillion-dollar program of the Defense Advanced Research Projects Agency (DARPA). It carries a Pentium-powered onboard computer running Linux and custom control software. A fully equipped PackBot can cost more than \$100 000.

In contrast, the Negotiator was designed around a basic 8-bit

microcontroller, Robotic FX's Ahed told *Spectrum*. It has a modular plug-and-play architecture that accommodates various accessories and sensors. That approach resulted in what Ahed calls "a simple, low-cost system." Robotic FX will not disclose the price of the Negotiator, but in 2005 the Illinois State Police bought six units for a total of \$122 940. The company says it has sold 80 Negotiators to federal, state, and local agencies in the United States since 2004.

iRobot's patent-infringement suit, filed in the U.S. District Court in Birmingham, Ala., focuses not on the robot's brains but on its mobility capabilities. iRobot claims Robotic FX violated its patents No. 6,263,989 and No. 6,431,296, which describe how the PackBot uses a pair of main tracks to move around (like a miniature tank) and a pair of auxiliary tracks, mounted on the sides, to go over obstacles and climb stairs, a capability that made the robot stand out among competitors. The Negotiator has auxiliary tracks similar to those on the PackBot.

In a separate lawsuit over pilfered trade secrets, filed in the U.S. District Court in Boston, iRobot accuses Ahed of



LOOKING FOR

AN ARMY CONTRACT:

Similarities between iRobot's PackBot [right] and its rival the Negotiator, from Robotic FX [opposite], led iRobot to file two lawsuits against Robotic FX.

NEWS stealing proprietary PackBot data and violating the confidentiality agreement he signed while an employee at the company.

Ahed began working at iRobot just after getting his bachelor's degree in biomedical engineering from the University of Illinois, Urbana-Champaign, but resigned two years later, in 2002, to found Robotic FX. iRobot claims that Ahed, who worked on PackBot-related projects, used confidential data to design the Negotiator.

Early this year, iRobot became aware of the Negotiator and obtained a unit to study. Concluding it was a "knockoff version" of the PackBot, iRobot sent a warning letter to Robotic FX and communicated its concerns to the Army PEO STRI officers overseeing the xBot contract. It filed the two lawsuits in August.

According to court records, iRobot hired private investigators to follow Ahed. The day after the lawsuits were filed, iRobot says its investigators saw Ahed load objects into a car and later put them in a dumpster. The investigators retrieved the objects, which included a box marked "iRobot," a robot's wheels and treads, and a welding tool and aluminum molding fixture that iRobot claims are used to make the PackBot tracks.

Ahed said in the hearings that the material dumped was just iRobot "memorabilia" he didn't want to keep anymore and that the aluminum fixture wasn't his, suggesting it had been planted. He acknowledged shredding about 100 data CDs and erasing a laptop's hard drive, but he said they contained only Robotic FX design and financial data and that he was "afraid that someone would come in and steal my work."

In the injunction order against Robotic FX, Judge Nancy Gertner says Ahed's actions undermine his credibility.

At press time, the trade secrets trial was scheduled to begin by 7 April 2008. If similar high-tech legal battles are any indication, the iRobot v. Robotic FX cases will be long and complicated—not to mention hugely expensive for the parties, making settlement out of court likely.

An unlikely but more exciting resolution: the companies could place their robots in a battle arena and let them settle the dispute.

—ERICO GUIZZO

Carbon Nanotubes Take the Heat Off Chips

Purdue scientists find flexible filaments best

Computer chips use only 1 percent of their electrical power to process information. They convert the rest to heat. As chips get smaller and faster, they also get hotter, which has engineers looking to carbon nanotubes and other new technologies to keep them cool.

The promise of carbon nanotubes lies in their high thermal conductivity, the ease with which heat flows through them from one end to the other. Researchers at Purdue University, in West Lafayette, Ind., managed to grow forests of nanotubes directly on a chip, and they found that the key to making nanotubes work as heat conductors is to make them flexible.

Most computer cooling systems work by blowing air over a heat sink—a metal plate usually ribbed like a radiator to dissipate heat into the air—"but the bottleneck is occurring between the heat sink and the chip," says Baratunde A. Cola, the Purdue doctoral student who coauthored a paper about the work in the 26 September issue of *Nanotechnology*. You can't just stick a heat sink directly on a chip, because the sink's microscopic roughness creates air pockets that resist heat flow [see "Beat the Heat," *IEEE Spectrum*, May 2004]. Current systems rely on thermal interfaces such as grease or solder to fill the gaps, but they are far from ideal.

Figuring nanotubes might do a better job, the Purdue team grew between

100 million and 1 billion tubes per square millimeter on test chips. The researchers wanted to see how they could maximize the thermal conductivity of the carbon nanotubes by varying their diameter and defect density. They controlled the tube properties by using a dendrimer template—essentially a chemical structure with uniformly sized cavities, according to Placidus B. Amama, one of the Purdue researchers. They used the dendrimers to place metal seed nanoparticles atop the chip from which the nanotubes grew. The size of the seed nanoparticles, in turn, determined the diameter of the tubes.

To the team's surprise, however, controlling the diameter of the individual tubes was less important than controlling how the tubes made contact with the heat sink. "You need a certain level of conductivity," Cola says, "but once you get past that threshold, it's all about contact."

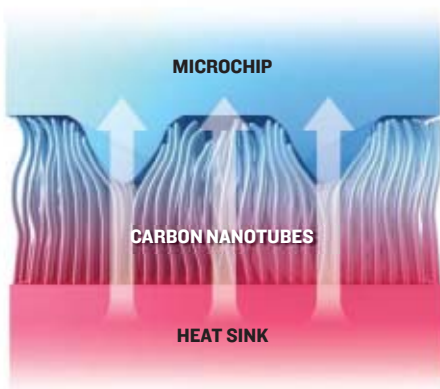
The interface between the nanotube and the heat sink is like that between the bristles on a toothbrush and your teeth, Cola says. The more the nanotubes can bend, the more they find their way into the nooks and crannies of the heat sink surface. To increase the tubes' flexibility, the researchers found that they had to make less conductive, "lower quality" nanotubes with more defects.

Such a carbon nanotube interface is several times as conductive as the thermal

WATCHING NANOGRASS GROW: Baratunde A. Cola [left] and Placidus B. Amama grow nanotubes on chips.



DAVID UMBERGER/PURDUE NEWS SERVICE



FLEXIBLE FIT: Carbon nanotubes can channel heat from a chip into a heat sink but do it best if they can bend enough to fit into the rough spots on the heat sink.

According to Victor Chiriac, a principal scientist at Freescale Semiconductor in Tempe, Ariz., and an expert on thermal management, the Purdue team is among those leading the efforts to make carbon nanotube interfaces practical. Among the other researchers exploring the issue is a team at Stanford University that is experimenting with the concept of growing tubes from both sides of the interface and joining them in the middle.

The nanotube research is still far from seeing use in real products, though, Chiriac says. "It's one thing to build in a lab, and another thing altogether to commercially fabricate a device, as current costs could be prohibitive," he says. He calls the Purdue group's ability to control the diameter, length, and flexibilities of the tubes an important step but just one of many that need to be taken.

—JOSHUA J. ROMERO

greases commonly used now, according to IEEE Fellow Avram Bar-Cohen, chairman of mechanical engineering at the University of Maryland, College Park. Bar-Cohen says carbon nanotubes show promise, especially for passively cooled devices such as cellphones and personal digital assistants, which lack space for a fan.

"Obviously, people want more and more capability in these personal systems," Bar-Cohen says. "You'd like to run the chips at higher power and yet cool them passively."

Messenger Arrives At Mercury

Satellite to zip by sunburned planet next month

Mercury is getting its first man-made visitor in more than 30 years. NASA's *Messenger* space probe is heading for a rendezvous with the planet, where sunlight is 11 times as bright as here on Earth and temperatures can swing from a metal-melting 450 °C in the sunlight to lows of -180 °C in the shade.

In the first of several encounters, *Messenger* (an acronym for *ME*rcury *S*urface, *S*pace *EN*vironment, *GE*ochemistry, and *R*anging) is scheduled to fly by Mercury on 14 January 2008 at a little more than 25 000 kilometers per hour, coming within 200 km of the planet's surface. Because it must perform scientific observations and relay them to Earth while in the scorching glare of Mercury's tight solar orbit, the craft boasts a multi-layer sunshade and the most advanced communications systems ever deployed in an interplanetary mission.

Although Mercury is relatively close to Earth, *Messenger* is just the second craft to visit it. Only about 45 percent of



the planet's surface has been mapped. "We anticipate a flood of data that will provide new insights on the origins and evolution of the [inner] planets—including Earth," says Ralph L. McNutt Jr., the project scientist for *Messenger*, which was designed and built by

the Johns Hopkins University Applied Physics Laboratory, in Laurel, Md.

The January flyby is only a tease for what is to come. *Messenger* is on a 7.9-billion-km trek that is scheduled to take it around the sun 15 times and past Earth once, Venus twice, and Mercury three times, before it finally settles into an orbit around the sun-blasted inner-most planet on 18 March 2011.

To protect *Messenger*'s wiring, electronics, and scientific instruments from the heat of being within 46 million km of the sun, it has a highly reflective and heat-resistant 5-square-meter micrometeorite-proof sunshade. The shade is made from alternating layers of Nextel ceramic cloth (to protect against

NEWS BRIEFS

MORE ROBOTS THAN COSTA RICANS

The International Federation of Robotics gave the latest estimate of the world's currently operating industrial robot population at a record 951 000. By 2010 that should grow to 1.173 million. Add to that the total number of robotic lawn mowers, vacuum cleaners, and other household and professional automats sold before 2006 (no one knows how many of these are still in service) and you get a robot population of 4.5 million. If they were a nation unto themselves, robots would be just ahead of Costa Rica and behind Croatia on the population scale.



TECHNOLOGY INCREASES INEQUALITY

That's one conclusion of the International Monetary Fund's October *World Economic Outlook*. The bank found that, contrary to expectation, technological advances—not the globalization of trade—were most responsible for the recent widening of the income gap. In fact, technology has been the main force driving inequality since the early 1980s, says the bank. Rather than trying to narrow the gap by stopping technological progress, the IMF suggests improving education and job training and making it easier for workers to move from one economic sector to another.

ELECTRIC SPIN Researchers at the Delft University of Technology have managed to control the spin state of a single electron using only electric fields. The feat clears the way for much simpler microchip-based quantum computers. It is easier to generate an electric field in experimental quantum computers than to generate a magnetic field, which scientists have used to control spin up to this point.

NEWS BRIEFS

It will also be easier to control the spins of multiple electrons independently from one another using electric fields [see "Dot-to-Dot Design" *IEEE Spectrum*, September 2007].



JOLTED OUT OF DEPRESSION Northstar Neuroscience, in Seattle, says that its brain implant has succeeded in alleviating the severe, untreatable depression of a small group of patients. Although depression represents a big market, Northstar's real target is stroke patients. It claims that the device, which delivers pulses of electricity to small patches of the outermost region of the brain, can help those paralyzed by stroke to recover movement. The company plans to unveil the results of a trial of more than 150 patients early this month.

INNOVATION NATION With more than 16 R&D personnel per 1000 employed people, Finland beats out its Scandinavian neighbors, as well as the United States and Japan, in the proportion of its residents involved in research, according to the *OECD Science, Technology, and Industry Scoreboard* released in October. Finland, Denmark, Japan, Sweden, and the United States are the only countries in the world with more than six corporate R&D jobs per 1000 employed people.

SUPERFLASH Samsung unveiled a 64-gigabit nonvolatile memory chip, which the firm plans to produce in 2009. Sixteen such chips ganged together would make for a memory card capable of holding 80 DVD movies. ■

any micrometeorite damage) and Kapton plastic insulation (to guard against direct sunlight and radiation). It's the same combination that protects the space shuttle's main engines during reentry into Earth's atmosphere. *Messenger* also has a series of radiators and pipes to divert heat from the spacecraft body. The result is that while the outer layers of the sunshade reach temperatures of 370 °C, the instruments behind it stay a cool room temperature (20 °C).

But not every *Messenger* component can be cooled to that extent. The orbiter's two main communications antennas, situated on each side of the sunshade, will have to withstand temperatures that range from -150 °C to almost 300 °C. Engineers planning for the mission knew that such temperature swings would endanger internal components and the steering mechanism on a conventional gimbaled dish antenna such as the one on NASA's *Mariner 10*—the first craft to reach Mercury, flying by the planet twice in 1974 and once in 1975.


To function in such extreme conditions, *Messenger* carries the first phased-array antenna ever flown in deep space. Although it has no moving parts, the antenna can be electronically steered through a full 90 degrees. Through variations in the phase of signals on different parts of the array, the antenna's radiation is enhanced in one direction and suppressed in all the

others. The array is expected to return around 100 gigabits of data per year.

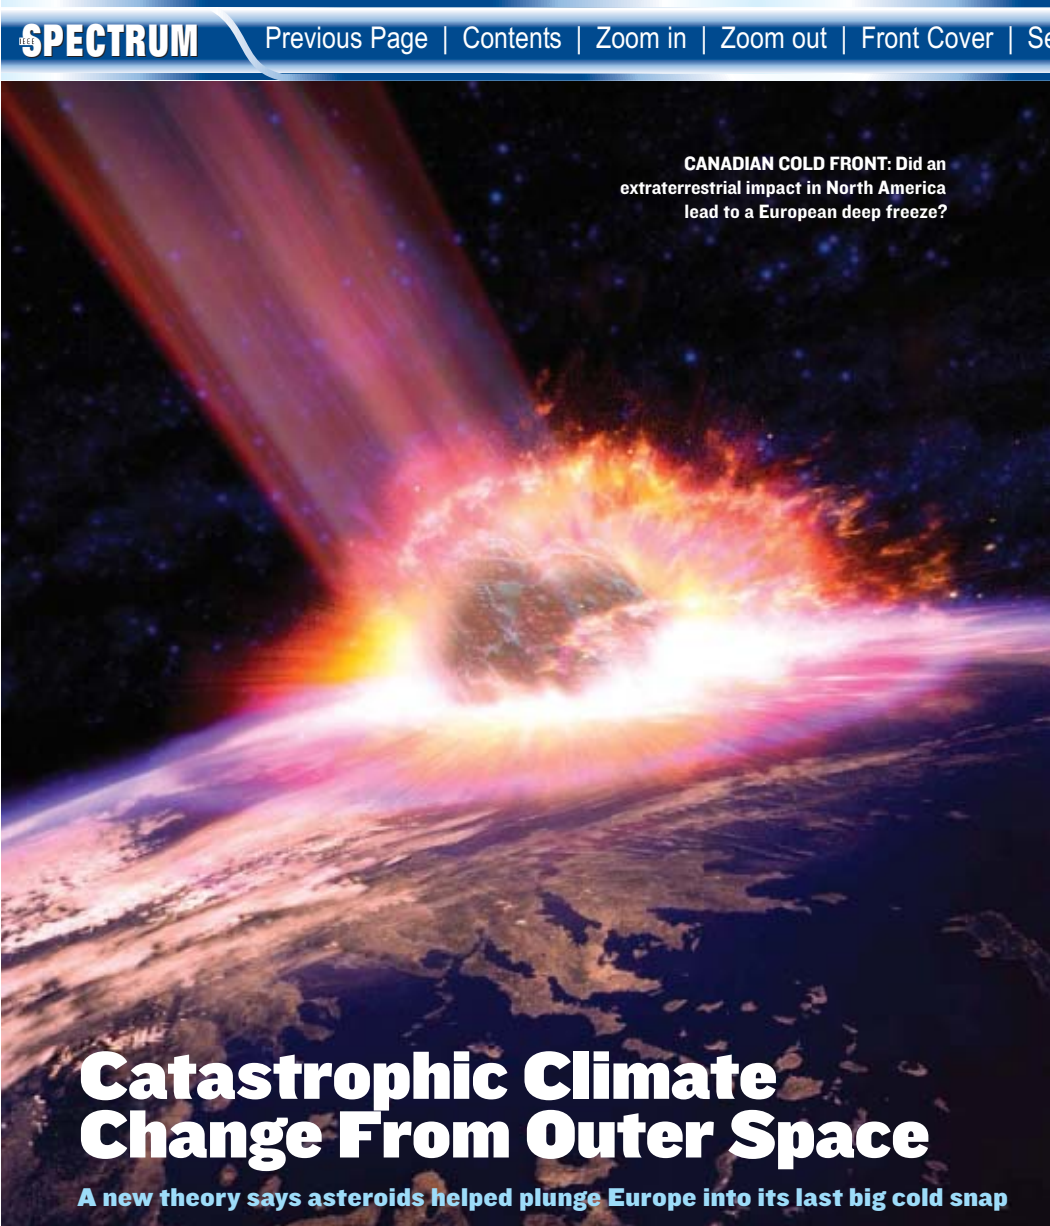
Mariner 10 gathered data and images from less than half of the planet's surface. It left behind many questions for *Messenger* to answer about Mercury's density and geologic history, the nature of its core and magnetic field, and how the solar wind interacts with the planet. Unlike *Messenger*, *Mariner 10* was not equipped to achieve orbit around Mercury. It was the first spacecraft to use the gravitational pull of one planet (Venus) to reach another (Mercury), and engineers of the day didn't feel confident enough in such a maneuver to use it to put the spacecraft into orbit. Orbital mechanics experts have much more experience now with such slingshot maneuvers. *Messenger* will use a gravity assist from Venus plus 16 thrusters to finally drop into a slow, one-revolution-per-year orbit around Mercury. More than half the craft's 1100-kilogram mass at launch was thruster fuel. In order to achieve orbit it will burn nearly 30 percent of this precious resource.

The data the craft gathers could be put to work solving problems here on Earth. *Messenger* will not observe the sun directly, but by examining the effects of the solar wind on Mercury's magnetic field, scientists can make inferences about how solar storms cause radio interference, electricity grid disruptions, and other problems on Earth, McNutt says.

—BARRY E. DIGREGORIO



LEFT: NORTHSTAR NEUROSCIENCE; RIGHT: BRYAN CHRISTIE DESIGN



CANADIAN COLD FRONT: Did an extraterrestrial impact in North America lead to a European deep freeze?

Catastrophic Climate Change From Outer Space

A new theory says asteroids helped plunge Europe into its last big cold snap

During the past two decades—as concern about climate change, and especially abrupt climate change, has mounted—Exhibit A has been a compelling scenario that explains a sharp cold snap called the Younger Dryas period, which occurred in the Northern Hemisphere starting about 12 800 years ago and lasted about 1200 years. In 1987 Wallace S. Broecker postulated that fresh waters from the southern rim of the North American ice sheet spontaneously spilled into the North Atlantic through what's now the St. Lawrence River. Such a deluge would have shut down the salt-and-temperature driven currents that draw warm waters into the ocean there and keep Europe temperate. Broecker, an eminent geochemist and climatologist at Columbia University's Lamont-Doherty Earth Observatory, north of New York City, postulated that the shutdown of the Gulf Stream led to the observed sharp cooling—an 8 °C drop on average. He dubbed the scenario “the biggest chill.”

Since 1987, whenever scientists have produced major new findings about what melting Arctic ice will mean, questions inevitably arise as to whether global warming could produce another big chill, plunging Western Europe into another miniature ice age. (Although the answers are almost always reassuring, the accelerated melting of Arctic ice this year and concerns about the long-term fate of Greenland's ice sheet have kept the issue alive.) When a blue-ribbon panel of scientists produced a report about abrupt climate change for the U.S. National Academies' National Research Council in 2002, naturally the very first reference was to Broecker's work.

Now research reported in the 9 October issue of *The Proceedings of the National Academy of Sciences* suggests that the abrupt onset of the Younger Dryas cooling and the freshwater infusion into the North Atlantic did not arise entirely from the inner workings of earth, ocean, and climate systems—but instead got a big nudge from something extraterrestrial.

The article, signed by 26 scientists from well-known institutions in the United States and Europe, identified a carbon-rich black layer of sediment at about 50 North American sites, dating from roughly 12 900 years ago—near the start of the chill—and strongly argues that the layer was left by something from space. The scientists propose that one or more large, low-density objects—such as a comet or asteroid—exploded over what is now Canada, destabilizing the glaciers there, spilling fresh water into the Atlantic, and triggering the Younger Dryas event.

The new scenario is reminiscent of the theory that a huge impact caused the extinction of the dinosaurs—which itself has become enormously controversial in science circles. The authors of the new theory acknowledge that their idea may never be provable, because virtually all the direct physical evidence of the explosion—such as a crater, if it occurred over the kilometers-thick ice sheets of the time—might have disappeared when the sheets melted. But the sophisticated radiochemical techniques they used, together with the variety of circumstantial physical evidence they discovered, give credence to their claims.

If the new theory gains ground, how much of an amendment will it be to Broecker's scenario, and will it have much influence on public opinion about the possibility of catastrophic climate change? Richard Alley, professor of geosciences at Pennsylvania State University, in University Park, and chairman of the National Academies' abrupt climate change report, argues that some proximate cause had to set Broecker's postulated mechanisms in motion, and that it doesn't really matter if it turns out to have been an extraterrestrial object.

Broecker himself says that if telltale signatures of an impact—such as buckyballs, nanodiamonds, or iridium—are found, then the authors “have something.” But for now, reserving judgment, he declares himself suspicious of “uncritical catastrophists.”

However scientists come down on that issue, might the new scenario have some influence on northern Europeans who have been worrying that runaway greenhouse gas emissions could melt the Arctic ice and plunge them into an ice age? It might.

—WILLIAM SWEET

NEWS



Zoos Network To Contain Disease

Online health records could keep outbreaks from spreading to humans

Keeping track of 100 head of cattle is a job for a cowboy; keeping track of 2 million elephants, emus, octopuses, orangutans, and other animals in the world's more than 700 zoos and aquariums is a job for a coder. Zookeepers in 21 institutions have just begun testing the new Zoological Information Management System (ZIMS), a US \$20 million real-time global network of zoo and aquarium medical files and animal husbandry records. One of the main goals of ZIMS is to monitor the spread of animal diseases that can potentially cross over to humans.

"ZIMS tracks the health and transport of zoo animals through the career of the animal," says Tracee Treadwell, a veterinarian and public health expert at the Centers for Disease Control and Prevention (CDC), in Atlanta. "We have been involved in the initial design of the system and see its potential to identify new and emerging diseases among these animals."

Signs of human disease have shown up in zoo animals in the past. In August 1999, for example, veterinarians at the Bronx Zoo, in New York City, reported the deaths of some flamingos and pheasants. Later that month, two people were diagnosed with a strange neuro-

logical illness at a community hospital in another part of the city. In addition, scores of crows were found dead within the metropolitan area. By 27 September, four human deaths and 37 illnesses had been reported due to what the CDC identified as West Nile virus, the first outbreak ever documented in the Western Hemisphere. It had taken more than a month after the first reported bird deaths for the CDC to make its diagnosis. If the initial flamingo and pheasant deaths at the zoo had been reported to an operational ZIMS network, the response would have been quicker, according to Jaime Meyer, a spokesman for the International Species Information System (ISIS), a nonprofit organization in Eagan, Minn. ISIS maintains computer-



THE DEEP: The Monterey Bay Aquarium, in California, will use ZIMS.

VETTED BY VETS: A new computer network might help keep bird flu and other diseases from spreading among animals in zoos and aquariums.

based information systems used by zoos and aquariums in 72 countries. "An alert would be issued in a matter of minutes. And with a mouse click or two by zoo staff around the world, they could check to see if their zoo had any recent contact or whether an animal was transferred to or from the infected zoo," Meyer says. The information could then be passed on to government health agencies.

Zoos routinely perform disease surveillance and diagnosis, but the process has not been automated. The only way for a zoo or aquarium to obtain information from another zoo has been for it to make a formal request through the mail, by phone, or by e-mail.

ZIMS is designed to replace two DOS-based programs that ISIS has used for the last 20 years: the Animal Records Keeping System (ARKS) and the Medical Animal Records Keeping System (MedARKS), Bronx Zoo veterinarian Paul Calle says. ARKS is used for recording where animals are and where they are transferred. Veterinarians use MedARKS to maintain medical records. ZIMS is 10 times as sophisticated as the two outdated software programs, says Calle. Neither ARKS nor MedARKS contains any mechanism to disseminate information widely. Disease outbreaks in other countries might not be known for days or weeks, he says. Calle is developing some of the clinical computer screens ZIMS users will access to retrieve laboratory results, disease serology, prescriptions, and information on procedures such as anesthesia.

Eighty staff members at 21 zoos and aquariums around the globe have begun testing the software in real-world situations. So far, 143 institutions have contributed to the development of ZIMS, and by fall 2008, ZIMS is expected to begin rolling out to ISIS member institutions at a rate of 20 per month. That's assuming ISIS comes up with the money. "We have essentially found the funding to create the application, but we still must find the money—about \$3 million—to deploy it worldwide and to provide the higher level of technical support that ZIMS will require," Meyer says.

—BARRY E. DIGREGORIO

TOP: NIRELIAS/REUTERS; BOTTOM: JOHN WARDEN/GETTY IMAGES

VISION BECOMES REALITY

CST MICROWAVE STUDIO® 2008

Leading edge 3D electromagnetic simulation

From the first bright spark to the final product, CST MICROWAVE STUDIO® accompanies you from idea to realization.

With a user friendly interface, easy data exchange to and from other software tools, a choice of first class solvers and excellent post-processing tools, you can leverage the latest developments in 3D electromagnetics to bring designs to market faster and with lower risk.

Choose CST MICROWAVE STUDIO® 2008 – complete technology for 3D electromagnetic simulation.



CHANGING THE STANDARDS

THE BIG PICTURE






Melon Blast

High school students at this University of Missouri–Rolla summer camp would be foolish to grow too fond of anything around them, because chances are they will eventually strap it with explosives and send it skyward. The watermelon shown here was only one of the victims in the department of mining and nuclear engineering's 2007 program. Participants blew up concrete columns, underground caverns, quarry rock faces, a Shrek doll, and a Barbie while learning from experts about recreational pyrotechnics and military explosives. The festivities began with a fireworks display by the teachers and ended with one by the students.

It's all meant to whet the appetites of potential engineering students. More than half of the teenagers attending the summer camp end up at the university in one program or another, according to Barbara Robertson, the camp coordinator. Interest in the summer course has grown so much in four years—one mom cried on Robertson's shoulder this year until Robertson let her kid in—that the university plans to increase enrollment from 40 to 60 next year.


For more on the camp, see http://dce.umr.edu/NonCredit/PreCollege/2007_Explosives_Camp.html.

Photo by Peter Newcomb/
The New York Times/Redux



HOW A TASER WORKS THE STUN GUN SHOCKS WITHOUT KILLING—BUT HOW SAFE IS IT? TWO EXPERTS TAKE A LOOK

WEAPONS



MARK LAITA

THE TASER GUN, AN ELECTROSHOCK WEAPON USED BY POLICE DEPART- MENTS WORLDWIDE, IS NO STRANGER TO BAD PRESS.



Last September, campus police officers at the University of Florida scuffled with Andrew Meyer, a student who had just posed a long and angry series of questions to Senator John Kerry (D-Mass.) during a forum at the school. As Meyer finished speaking, officers surrounded him and directed him out of the auditorium. Meyer yelled, resisted them, and demanded to know what he had done wrong. “You’re

going to get Tased if you don’t put your arms behind your back,” an officer said. Meyer continued to struggle and yelled, “Don’t Tase me, bro!” One of the officers fired his Taser Electronic Control Device, and Meyer screamed, his voice breaking.

Within hours, the video record of the event in Gainesville appeared on the Web and became an instant YouTube sensation. The American Civil Liberties Union and Amnesty International chimed in with support for Meyer, whose memorable “Don’t Tase me, bro!” cry leaped into American popular culture on T-shirts and baby bibs. Newspapers across the United States questioned whether the campus police were right to use the Taser, whether it was cruel, and whether Meyer had deliberately provoked the officers into stunning him.

The explosion of attention surrounding the incident reflects a deep public ambivalence toward the electroshock weapon and its use. Meyer’s experience is but one of many high-profile cases in which the use of a Taser to subdue a recalcitrant troublemaker may not have been warranted. Last year, a student at the University of California, Los Angeles, was shocked in Powell Library, an event that generated a similar public outcry. The student, Mostafa Tabatabaiejad, was using the computer lab after hours and didn’t show officers his student ID card when asked to do so. His continued refusal to comply or leave the library led the campus police officers to apply Taser shocks to him repeatedly. Reports after the fact acknowledged police error—the officers had overreacted and were too ready to deploy their high-tech gadget in a situation that didn’t call for violence.

The screams of people being shocked by a stun gun sound eerily similar to the blood-curdling cries of torture victims, so incidents that involve unarmed students raise hackles. But there’s another factor underlying the public distrust of Tasers: the possibility that they can kill people.

In the period between 2001 and 2005, Amnesty International reported, 150 people died in the aftermath of receiving shocks from a Taser. In only a handful of the cases did medical examiners cite the shocks received as a cause of death. Even so, the considerable uncertainty surrounding the physiological effects

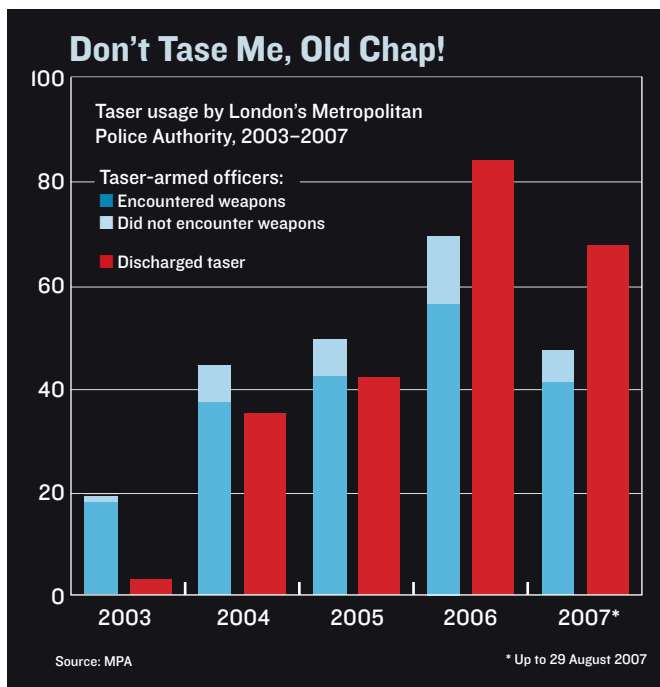
of a Taser shock, as well as ambiguity regarding when it should be used, have bred an atmosphere of distrust and fear.

The electroshock gun used by police—the Taser X26, made by Taser International of Scottsdale, Ariz.—fires barbed electrodes. A shot releases two probes, and those probes must either both make contact with their target, or one must strike the target and the other the ground, to complete the electrical circuit. The electrodes are attached by long, thin wires to a waveform generator that sends muscle-locking electric pulses into the target.

Situations where police have been able to successfully disarm suspects without causing permanent injury are the reason these weapons have gained widespread use. In an October case in the Czech Republic, for example, a kidnapped child was rescued by police who used Taser guns to immobilize her captors. According to a 2006 report by the Police Executive Research Forum, a law-enforcement policy organization in Washington, D.C., more than 8000 police and sheriffs’ offices across the United States have adopted the devices, which are widely used in Canada and the United Kingdom as well [see graph, “Don’t Tase Me, Old Chap!"]. Police departments in Australia, New Zealand, and France started using the devices after Taser International introduced an attachable video camera. The guns also now release bits of identifying confetti with every shot, and the time and duration of each trigger pull is recorded in the gun’s memory. According to Taser, its guns are now fired more than 620 times a day and have been used a total of more than 680 000 times worldwide.

Any new technology that is designed for violent encounters should be carefully assessed. Unlike medical devices, Tasers don’t have to undergo testing and receive approval by agencies such as the U.S. Food and Drug Administration, at least not in the United States. Partly in response, several local and state legislatures have considered introducing laws restricting the stun guns’ adoption, and most police departments, if not all, have instituted guidelines on the proper use of Tasers.

Analyses conducted by British and Canadian police research centers and by the U.S. Air Force concluded that Tasers are generally effective and do not pose a significant health risk to



the recipients of a shock. In Portland, Ore., meanwhile, police found that 25 to 30 percent of the situations in which a Taser was employed met the criteria for the use of deadly force. Other police departments have released statistics showing a decline in the number of deaths of suspects and officers in the months following the introduction of Tasers. But research by the Police Executive Research Forum has raised the concern that multiple activations of Tasers may increase the risk of death.

Even if Tasers are proven to be entirely safe, there's the bigger question of whether the stun guns encourage police brutality. A Taser shock leaves almost no visible scarring or bruising, as a clubbing or a beating typically would. Could the absence of physical scars lift a psychological restraint on officer behavior? Should every Taser gun have a built-in video camera?

Equipping law-enforcement services with Tasers is likely to reduce the number of bullets officers fire from their handguns and therefore the number of serious injuries and deaths. At the same time, it may lead police to inflict an unwarranted amount of pain on individuals who commit only minor crimes.

The broader questions regarding the social effects of stun guns are, however, beyond the scope of this discussion. The two articles that follow investigate the physiological effects of electric shock. The first is by Mark W. Kroll, an electrical engineer who has helped invent numerous electrical medical devices and who sits on the board of Taser International. The second is by Patrick Tchou, a cardiac electrophysiologist at the Cleveland Clinic, who has tested Tasers experimentally on pigs.

—Sandra Upson

CRAFTING THE PERFECT SHOCK BY MARK W. KROLL

YOU KNOW AN ENGINEERING problem is difficult when the prevailing technology dates back to the Stone Age. Let's face it, the police officer's baton is barely more sophisticated than a cave dweller's club, and with

it comes all the same crudeness.

One reason that finding a good replacement has been such a confounding problem is the nature of the task. Police officers often need to take into custody a violent criminal who has overdosed on a stimulant. Most people probably would be surprised to learn that, at present, the main methods police use in such situations all rely on inflicting pain. The old standbys are wrist twists and other forms of joint distortion, pepper spray, and clubbing.


The problem is complicated by the fact that many illegal drugs are painkillers, and as a result standard subduing techniques are frequently ineffective at bringing troublemaking drug users to heel. Even worse, many of the dangerously drug-addled perpetrators exhibit superhuman stamina and strength. There are numerous accounts of a person on a drug overdose manhandling half a dozen law-enforcement officers at once. Many officers are injured along with those they are trying to take into custody.


The ideal arrest tool, then, must meet a number of requirements. First, it must be able to temporarily disable even the largest, most determined drug-anesthetized individual. Second, it must do so without causing serious injury to anyone involved. Third, its effectiveness cannot be dependent on causing pain. Fourth, it must work reliably. And finally, it must be able to be used from a safe distance—let's say 5 meters—so that an arresting officer need not come within range of a suspect's blows.

Some approaches to meeting those criteria have come close, but not close enough. These include powerfully launched nets, which still require an officer to come into contact with a thrashing suspect, and body-immobilizing glues, which don't perform well in cold weather.

A solution that satisfies all the requirements is a device that was once playfully dubbed the "Thomas A. Swift electric rifle" (after the exploits of the fictional Tom Swift, a teenage inventor made famous in a series of juvenile adventure novels published from 1910 to 1941) and is now known as the Taser Electronic Control Device. Under microprocessor control, the device temporarily, and relatively harmlessly, immobilizes a suspect with a carefully engineered electric signal that is specifically designed with human physiology in mind.

WHEN YOU PULL THE TRIGGER of a Taser gun, a blast of compressed nitrogen launches its two barbed darts at 55 meters per





Freeze!

When the trigger on a Taser gun is pulled, the compressed-nitrogen cartridge breaks open. Enough pressure builds up inside the device to launch the two darts. The darts are tipped with barbs that grab hold of a target's clothing, and current travels down the wires to the person. The gun generates a brief arcing pulse to close the circuit, at which point the voltage drops. Shots from a Taser gun land, on average, about 15 centimeters apart on the torso.

The differences between the cells that make up heart muscle and skeletal muscle are key components of the Taser's safety. For example, the cells in the heart generate a longer electric impulse than those in skeletal muscle do, and it takes much more current to trigger cardiac muscle cells.

second, less than a fifth the speed of a bullet from a typical pistol. Each projectile, which weighs 1.6 grams, has a 9-millimeter-long tip to penetrate clothing and the insulating outer layer of skin. Two whisper-thin wires trail behind for up to 9 meters, forming an electrical connection to the gun.

Because the barbs get stuck in clothing and fail to reach the skin about 30 percent of the time, the gun is designed to generate a brief arcing pulse, which ionizes the intervening air to establish a conductive path for the electricity. The arcing phase has an open-circuit peak voltage of 50 000 volts; that is, the voltage is 50 kilovolts only until the arc appears or until the barbs make contact with conductive flesh, which in the worst conditions offers around 400 ohms of resistance [see illustration, “Freeze!”].

The target’s body is never exposed to the 50 kV. The X26—the model commonly used by police departments—delivers a peak voltage of 1200 V to the body. Once the barbs establish a circuit, the gun generates a series of 100-microsecond pulses at a rate of 19 per second. Each pulse carries 100 microcoulombs of charge, so the average current is 1.9 milliamperes. To force the muscles to contract without risking electrocution, the signal was designed to exploit the difference between heart muscle and skeletal muscle.


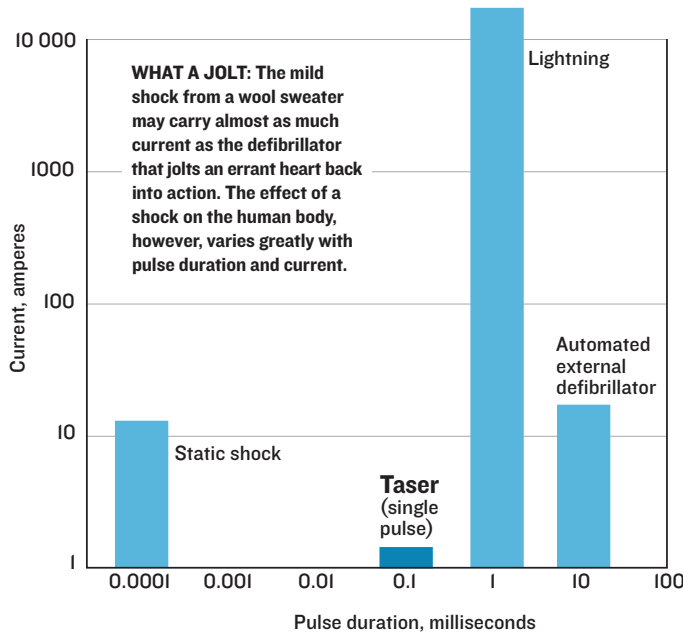
Skeletal muscle constitutes 40 percent of a typical person’s mass and is responsible for making your biceps flex, your fingers type, and your eyelids wink. It’s organized into bundles of single-cell fibers that stretch from tendons attached to your skeleton. When your brain orders a muscle to flex, an electrical impulse shoots down a motor nerve to its termination at the midpoint of a muscle fiber. There the electrical signal changes into a chemical one, and the nerve ending sprays a molecular transmitter, acetylcholine, onto the muscle. In the milliseconds before enzymes have a chance to chew it up, some of the acetylcholine binds with receptors, called gated-ion channels, on the surface of the muscle cell. When acetylcholine sticks to them, they open, allowing the sodium ions in the surrounding salty fluid to rush in.

The movement of those ions raises the cell’s internal voltage, opening nearby ion channels that are triggered by voltage instead of by acetylcholine. As a result, a wave of voltage rolls outward along the fiber toward both ends of the muscle, moving as fast as 5 meters per second. As the voltage pulse spreads, it kick-starts the molecular machinery that contracts the muscle fiber.

By directly jolting the motor nerves with electricity, a Taser can stimulate the muscle and get the same effect.

The force with which a skeletal muscle contracts depends on the frequency at which its nerve fires. The amount of contraction elicited is proportional to the stimulation rate, up to about 70 pulses per second. At that point, called tetanus, contractions can be dangerously strong. (The same thing happens in the disease tetanus, whose primary symptom, caused by the presence of a neurotoxin, is prolonged contraction of skeletal fibers.) The Taser, with its 19 pulses per second, operates far enough from the tetanus region so that the muscles contract continuously but without causing any major damage.

Heart muscle has a somewhat different physical and electrical structure. Instead of one long cell forming a fiber that stretches from tendon to tendon, heart muscle is composed of interconnected fibers made up of many cells. The cell-to-cell connections have a low resistance, so if an electrical impulse causes one heart cell to contract, its neighbors will quickly follow suit. With the help of some specialized conduction tissue, this arrangement makes the four chambers of the heart beat in harmony and pump blood efficiently. A big jolt of current at the right frequency can turn the coordinated pump into a quivering mass of muscle. That’s



LEVELS OF SHOCK: The Taser X26 puts out 2 milliamperes at 19 hertz. The gun packs its current into 100-microsecond pulses, so it can capture muscle with lower current than if it had been delivered as a sine wave, as the rest of the chart shows.

just what electrocution does: the burst of electricity causes the heart’s electrical activity to become chaotic, and it stops pumping adequately—a situation known as ventricular fibrillation.

The Taser takes advantage of two natural protections against electrocution that arise from the difference between skeletal and cardiac muscle. The first—anatomy—is so obvious that it is typically overlooked. The skeletal muscles are on the outer shell of the body; the heart is nestled farther inside. In your upper body, the skeletal muscles are arranged in bands surrounding your rib cage. Because of skeletal muscle fibers’ natural inclination to conduct low-frequency electricity along their length, a larger current injected into such a muscle tends to follow the grain around the chest rather than the smaller current that penetrates toward the heart.

The second protection results from the different timing requirements of the nerves that trigger muscle contractions and the heart’s intrinsic electronics. To lock up skeletal muscle with-

out causing ventricular fibrillation, an electronic waveform has to have a specific configuration of pulse length and current.

The key metric that electrophysiologists use to describe the relationship between the effect of pulse length and current is chronaxie, a concept similar to what we engineers call the system time constant. Electrophysiologists figure out a nerve's chronaxie by first finding the minimal amount of current that triggers a nerve cell using a long pulse. In successive tests, the pulse is shortened. A briefer pulse of the same current is less likely to trigger the nerve, so to get the attached muscle to contract, you have to up the amperage. The chronaxie is defined as the minimum stimulus length to trigger a cell at twice the current determined from that first very long pulse. Shorten the pulse below the chronaxie and it will take more current to have any effect. So the Taser should be designed to deliver pulses of a length just short of the chronaxie of skeletal muscle nerves but far shorter than the chronaxie of heart muscle nerves.

And that's the case. To see just how different skeletal and heart muscles are, let's look at what it takes to seriously upset a heart's rhythm. Basically, there are two ways: by using a relatively high average current, or by zapping it with a small number of extremely high-current pulses.

In terms of average current, the 1.9 mA mentioned earlier is about 1 percent of what's needed to cause the heart of the typical male to fibrillate. So the Taser's average current is far from the danger zone for healthy human hearts.

As far as single-pulse current goes, the Taser is again in the clear. The heart's chronaxie is about 3 milliseconds—that's 30 times as long as the chronaxie of skeletal muscle nerves and the pulse lengths of a Taser. The single-pulse current required to electrocute someone by directly pulsing the most sensitive part of the heart-beat using 3-ms pulses is about 3 A. Because a Taser's 100- s pulses are such a small fraction of the heart's chronaxie, it would take significantly higher current—on the order of 90 A—to electrocute someone using a Taser.

When you factor in that the Taser barbs are likely to land in current-shunting skeletal muscle not near the heart, you wind up with a pretty large margin of safety. For barbs deeply inserted directly over the heart, the margin is slimmer, though, and the key question is whether that margin is adequate. To answer that definitively, one needs to consider what has been learned from the devices' use in everyday life.

In the United States, about 670 people die each year under police restraint, according to the U.S. Department of Justice's Bureau of Justice Statistics. These incidents include arrests and attempts to control an uncooperative person who needs medical assistance, as well as suicides after arrest. Studies have shown that stun guns were used during about 30 percent of in-custody deaths in the United States. Although Tasers were involved in a sizable fraction of these deaths, one should not leap to the conclusion that Tasers caused them. One study found that 100 percent of in-custody deaths involved the use of handcuffs, and one might apply the same faulty logic to argue against "killer cuffs," but that would, of course, be absurd. Medical examiners have cited Tasers as the primary cause of death in only four cases to date, and three of those were later thrown out of court.

There will always be some degree of violence in many police arrests, and a reliance on handguns and hand-to-hand combat can lead to terrible use-of-force dilemmas for police officers. For example, when a suspect brandishing a knife is within striking distance, law-enforcement officers in the United States are trained to shoot that person. Having a Taser gun in their holsters allows those officers an opportunity to disarm suspects in a manner that's likely to be safer for all involved. It's the prevalence of such scenarios that has persuaded so many police departments to pay twice as much for a Taser—on the order of US \$1000 per device—as they do for a traditional handgun. Tasers are expensive and controversial, but in the end it's safety that's on everyone's mind. ■

FINDING THE EDGE OF HEART SAFETY BY PATRICK TCHOU

WITH THE USE OF TASER

Electronic Control Devices by law-enforcement officers on the rise, it's no wonder that questions about the guns' safety come up again and again. As Mark Kroll describes [see "Crafting the Perfect Shock"], Tasers produce uncontrollable muscular contractions, which temporarily immobilize a sub-

ject. Those questions of safety can be answered in two ways: from a medical standpoint—that is, in terms of the bodily harm that can result from a Taser shock—and from the point of view of someone working in law enforcement.

The second perspective is much broader. How would one minimize injury to both the police officer and the person being taken into custody, not to mention bystanders, while restraining a violent and uncooperative subject? To probe further, one must ask how alternative means of restraint compare with the use of a Taser.


As a physician, I contribute to the former perspective by investigating whether Taser shocks can cause serious damage to a heart's normal function.

Let's begin with some basics about how the heart works. Each heartbeat is activated by an electrical impulse that propagates

through the four chambers of the heart [see illustration, "Heart Electronics"]. A number of troubles can throw off the internal rhythm of the impulse as it travels along, and the most dangerous kind of these arrhythmias is ventricular fibrillation, which is typically the cause of death in someone who is electrocuted. What brings on death is the uncoordinated electrical activation of the heart's main pumping chambers. The heart tissue still carries electrical impulses, but they propagate at chaotic and rapid rates, and the heart ceases to function as a pump, so blood pressure quickly plummets. It takes 10 to 20 seconds for a person to lose consciousness, less if he or she is standing.

So the most important question regarding the safety of Tasers is how likely it is that the use of one will induce ventricular fibrillation. Statistics alone suggest that, so far, the incidence of Taser-induced ventricular fibrillation is low. To investigate this question further in a more rigorous experimental setting, my Cleveland Clinic colleagues and I designed experiments to assess the threshold for bringing about ventricular fibrillation using pigs, taking into account the distance between the heart and the Taser darts at the body surface. Taser International covered the costs of the testing equipment and the costs of laboratory use, but none of Taser's funding covered my time or that of any other physicians involved in the studies.

The pigs were under general anesthesia when we performed the experiments. We selected five points on each animal's torso



HEART ELECTRONICS: An electric impulse starts in the SA node and travels to the two atria, which contract, pushing blood into the lower chambers. Current then passes through the AV node to cause the ventricles to contract.

corresponding to sites where Taser darts commonly make contact with human subjects. We used a custom-built circuit that matched the waveform and typical 5-second shock duration of an X26 Taser gun, but our device could deliver a much larger shock. To boost the output current, we increased the capacitor sizes in the device. After inducing ventricular fibrillation, we immediately rescued the animal using an ordinary defibrillator. We then stepped down the current to determine the highest amount that could be delivered without inducing ventricular fibrillation.

We calculated that quantity, cast in terms of multiples of the capacitances, for each of the body sites we'd chosen to test. Of the various positions we examined, some were a mere centimeter or two away from the heart, which sits just under the chest wall, touching it on the inside. Not surprisingly, we found that darts near the heart had the lowest thresholds for inducing ventricular fibrillation. At the closest spots—with one dart hitting at the lower end of the chest wall, and the other at the top of the breastbone—such a cardiac crisis would ensue with about four times the standard Taser capacitance.

Our experiments were the first to document that Taser-like impulses, albeit more energetic ones, applied close to the heart on the chest wall in pigs could have serious cardiac consequences. Even at the standard output of a Taser, we found that current applied to the most vulnerable part of the chest was able to drive the heart to beat up to 250 beats per minute, which is about twice the normal rate for pigs. These experiments also showed us that the onset of ventricular fibrillation is related to how fast the heart is driven by the impulses—which scales with the amount of current used.

Because the standard Taser output proved on average to be one-fourth what was needed to cause fibrillation, one is tempted to conclude that the device is fundamentally safe. But there's another factor to keep in mind: a large portion of the violent individuals with whom the police have to deal are under the influence of cocaine, methamphetamine, or other stimulants. So the Taser has to be safe even for those whose physiology is distorted by the presence of such powerful drugs. Cocaine in particular is

a concern with respect to cardiac complications because it raises heart rate and blood pressure and significantly increases the risk of a heart attack even without any kind of shock.

My colleagues and I supposed that the presence of such drugs would increase the potential for cardiac arrhythmias, and we later tested this hypothesis in a separate study, published in the *Journal of the American College of Cardiology*. To our surprise, the amount of current needed to bring on ventricular fibrillation didn't go down; indeed, it *increased* significantly when the pigs were administered cocaine. After some thought, we realized that our initially puzzling findings were not entirely out of line, because cocaine has certain anesthetic properties that can affect the electrical behavior of the heart in ways that protect it against shocks and decrease its vulnerability to fibrillation. Applying enough voltage to a heart cell will open its sodium-ion channels and start the contraction machinery, but cocaine stops up the voltage-activated sodium channels, making it more difficult for electricity to trigger a muscle contraction.

Another study carried out at our clinic more recently showed that implantable defibrillators and pacemakers function normally after a typical 5-second electric shock from a Taser. It remains to be seen, however, how well such medical devices stand up to repeated or longer shocks.

It is a challenge to relate experiments conducted under controlled laboratory conditions to the vagaries of real life. For one thing, we obtained our results from anaesthetized pigs with ostensibly normal hearts. It's possible that an abnormal or diseased heart—or even a heart under stress or one affected by amphetamines—might be more vulnerable. No one has yet studied the effects of Taser shocks on such hearts, information that is sorely needed to understand what might prove to be the greatest danger from Tasers.

Even so, we were comforted to learn that stun guns do not normally pose any cardiac risk. The full length of the Taser dart tip would have to embed itself into the skin and chest-wall muscle of a relatively small, thin person to get within the range of distances where we found the heart to be most vulnerable. Furthermore, the most sensitive region for the induction of fibrillation covers just a small area. And it is unlikely that two darts would land there.

Much remains unknown about the physiological effects of a Taser shot, but the absence of conclusive medical knowledge doesn't necessarily mean that the devices shouldn't be used—as long as evidence continues to support their safety. Rarely is any biological phenomenon or medical device fully understood and tested, and the Taser is no exception. As more information becomes available, law-enforcement agencies and their officers will better understand the consequences of each pull of the trigger. ■

ABOUT THE AUTHORS

MARK W. KROLL is an IEEE senior member who holds more than 250 U.S. patents as an inventor of electrical medical devices. He sits on the board of Taser International. **PATRICK TCHOU** is a cardiologist who specializes in treating cardiac rhythm disturbances at the Cleveland Clinic, a leading research hospital in Ohio.

TO PROBE FURTHER

The Police Executive Research Forum's report on standards for "conducted energy devices" is on its Web site at <http://www.policeforum.org/library.asp?MENU=356>.

Recent U.S. Department of Justice findings on arrest-related deaths can be found at <http://www.ojp.usdoj.gov/bjs/abstract/ardus05.htm>.

The Institute for the Prevention of In-Custody Deaths has related research available at <http://www.incustodydeath.com>.

BUSINESS

The R&D 100

A company's research budget tells you very little about its prospects **By Ron Hira & Philip E. Ross**

Toyota jumped three places to top the R&D leaderboard this year, just as it was passing General Motors to become the biggest automaker in the world. The coincidence raises a question: Are Toyota's sales so high because of its lavish R&D spending, or is its R&D spending lavish because its sales are so high?

Of course, R&D spending today will affect sales only some years from now, so we must look to the past to understand the present—and history in fact shows that R&D cannot have been the key determinant of success in this case. Five years ago, Toyota ranked fourth among the 12 leading carmakers in R&D spending. It laid out just two-thirds as much as top-ranked Ford—a company that had been the industry's leading R&D spender for five years running while at the same time struggling mightily with declining sales and near bankruptcy.

But let's not jump to conclusions. Could it be that Ford outspends Toyota in absolute terms, but not relative to sales? Nope. Even this metric, which we call R&D Intensity, shows that though Ford has consistently outspent Toyota, with an R&D Intensity that has long been on the high end for automakers, its performance has been subpar. Meanwhile, Toyota's 9.6 percent bump in R&D spending was dwarfed by a 13.8 percent increase in sales, which meant that its R&D Intensity actually fell a bit. No matter how you look at it, Toyota has been consistently on the low end of R&D Intensity for years. You can't make the case that R&D spending accounts for Toyota's edge.

IS THIS FINDING just an anomaly? Apple, the one company perhaps most closely associated with innovation, doesn't even show up on the R&D leaderboard this year. In fact, it hasn't



appeared there since Standard & Poor's and *IEEE Spectrum* began this series five years ago. (To run the numbers yourself, try our interactive calculator at <http://spectrum.ieee.org/deco7/rndcalc>.) Its absence can't be attributed to size, because Apple's sales of US \$19.3 billion surpassed those of 30 of the list's 100 firms. Comparing those sales to the relatively meager \$712 million Apple spent on R&D in 2006 yields an R&D Intensity of just 3.7 percent, a fraction of Nokia's 9.5 percent.

Then there's Google, another firm most people would call innovative. Yet it cracked the top 100 for the first time only this year, coming in 79th. Although it more than doubled its R&D spending, its R&D Intensity still came to just 11.5 percent, lower than that of most software firms.

These examples are part of a larger pattern, identified in a recent study by Booz Allen Hamilton, a consulting firm. The study found that firms whose R&D spending put them in the top 10 percent of their peer group did not outperform those peers in any financial metric. On the other hand, the Booz Allen study found that being a scrooge with R&D is also a bad idea: companies in the bottom 10 percent underperformed their peers.

It isn't clear what the findings mean. It could be that R&D spending is necessary but only up to some ill-defined point. It could also be that unusually low R&D spending is a symptom of larger problems. After all, a company may spend little simply because it is strapped for cash. Finally, a company may so dominate its market that it feels little pressure to come up with new ideas.

INVESTORS' ONLY SOURCE on R&D spending is the quarterly and annual reports that publicly traded companies must submit to the agencies that regulate the various stock exchanges, as the Securities

STUART BRADFORD

SPECTRUM'S TOP R&D SPENDERS			R&D EXPENDITURES (US \$ MILLIONS, LESS IN-PROCESS*)			SALES (\$ MILLIONS)		R&D INTENSITY†		R&D EXPENDITURES PER EMPLOYEE (\$ THOUSANDS)	
2006	2005	COMPANY, COUNTRY	2006	2005	% CHANGE	2006	2005	2006	2005	2006	2005
1	4	Toyota Motor Corp., JAPAN	7486	6829	9.6%	201 254	176 789	3.7%	3.9%	25	24
2	2	Pfizer Inc., U.S.	7423	7442	-0.3%	48 201	51 298	15.4%	14.5%	76	70
3	1	Ford Motor Co., U.S.	7200	8000	-10.0%	160 123	176 896	4.5%	4.5%	25	27
4	8	Johnson & Johnson, U.S.	7125	6312	12.9%	53 194	50 434	13.4%	12.5%	58	55
5	7	Microsoft Corp., U.S.	7121	6584	8.2%	51 122	44 282	13.9%	14.9%	90	93
6	3	DaimlerChrysler AG, GERMANY	7007	7425	-5.6%	199 246	196 863	3.5%	3.8%	19	19
7	9	GlaxoSmithKline PLC, UNITED KINGDOM	6611	6108	8.2%	45 263	42 213	14.6%	14.5%	64	61
8	5	Siemens AG, GERMANY	6604	6776	-2.5%	114 779	99 164	5.8%	6.8%	14	15
9	6	General Motors Corp., U.S.	6600	6700	-1.5%	207 349	190 215	3.2%	3.5%	24	20
10	12	Volkswagen AG, GERMANY	6030	5364	12.4%	137 846	125 219	4.4%	4.3%	18	16
11	10	Samsung Electronics Co., SOUTH KOREA	5943	5765	3.1%	91 038	85 927	6.5%	6.7%	69	72
12	14	Intel Corp., U.S.	5873	5145	14.1%	35 382	38 826	16.6%	13.3%	62	52
13	13	Sanofi-Aventis, FRANCE	5823	5315	9.5%	37 293	35 897	15.6%	14.8%	58	55
14	11	International Business Machines Corp., U.S.	5682	5378	5.7%	91 424	91 134	6.2%	5.9%	16	16
15	17	Roche Holding AG, SWITZERLAND	5359	4640	15.5%	34 192	28 882	15.7%	16.1%	72	68
16	18	Novartis AG, SWITZERLAND	5349	4514	18.5%	36 031	32 212	14.8%	14.0%	53	50
17	15	Nokia Corp., FINLAND	5122	5008	2.3%	54 049	44 940	9.5%	11.1%	75	85
18	16	Matsushita Electric Industrial Co., JAPAN	4858	4746	2.4%	76 543	74 746	6.3%	6.3%	15	14
19	20	Honda Motor Co., JAPAN	4638	4289	8.1%	93 174	83 264	5.0%	5.2%	32	30
20	19	Sony Corp., JAPAN	4571	4469	2.3%	69 715	62 822	6.6%	7.1%	29	28
21	21	Robert Bosch GmbH, GERMANY	4401	4039	8.9%	57 418	54 496	7.7%	7.4%	17	16
22	24	Motorola Inc., U.S.	4106	3680	11.6%	42 879	36 843	9.6%	10.0%	62	53
23	30	Cisco Systems Inc., U.S.	4067	3322	22.4%	28 484	24 801	14.3%	13.4%	81	86
24	22	Merck & Co., U.S.	4020	3848	4.5%	22 636	22 012	17.8%	17.5%	67	63
25	25	Telefonaktiebolaget LM Ericsson, SWEDEN	3990	3494	14.2%	25 403	21 693	15.7%	16.1%	63	62
26	23	Nissan Motor Co., JAPAN	3906	3761	3.9%	87 975	79 233	4.4%	4.7%	24	23
27	28	AstraZeneca PLC, UNITED KINGDOM	3885	3379	15.0%	26 475	23 950	14.7%	14.1%	59	52
28	26	Hewlett-Packard Co., U.S.	3591	3490	2.9%	91 658	86 696	3.9%	4.0%	23	23
29	27	Hitachi Ltd., JAPAN	3467	3404	1.8%	86 121	79 540	4.0%	4.3%	10	10
30	45	Amgen Inc., U.S.	3366	2314	45.5%	14 268	12 430	23.6%	18.6%	167	140
31	31	Bayerische Motoren Werke AG (BMW), GERMANY	3344	3239	3.2%	64 404	61 324	5.2%	5.3%	31	31
32	32	Toshiba Corp., JAPAN	3311	3130	5.8%	59 804	53 309	5.5%	5.9%	17	18
33	46	Boeing Co., U.S.	3257	2205	47.7%	61 530	54 845	5.3%	4.0%	21	14
34	39	EADS NV, NETHERLANDS	3231	2727	18.5%	51 832	44 960	6.2%	6.1%	28	24
35	33	Eli Lilly & Co., U.S.	3129	3026	3.4%	15 691	14 645	19.9%	20.7%	75	71
36	36	Wyeth, U.S.	3109	2749	13.1%	20 351	18 756	15.3%	14.7%	62	55
37	37	Bristol-Myers Squibb Co., U.S.	3067	2746	11.7%	17 914	19 207	17.1%	14.3%	71	64
38	43	Bayer AG, GERMANY	3019	2479	21.8%	38 059	35 992	7.9%	6.9%	28	26
39	38	General Electric Co., U.S.	2969	2741	8.3%	160 854	148 559	1.8%	1.8%	9	9
40	34	NEC Corp., JAPAN	2812	2868	-1.9%	39 100	40 548	7.2%	7.1%	18	19
41	42	PSA Peugeot Citroën SA, FRANCE	2625	2483	5.7%	74 386	73 957	3.5%	3.4%	12	12
42	44	Canon Inc., JAPAN	2591	2408	7.6%	34 932	31 549	7.4%	7.6%	22	21
43	40	Renault SA, FRANCE	2580	2674	-3.5%	54 584	54 334	4.7%	4.9%	19	21
44	35	BAE Systems PLC, UNITED KINGDOM	2432	2824	-13.9%	24 035	21 475	10.1%	13.2%	31	35
45	48	Denso Corp., JAPAN	2352	2154	9.2%	30 335	26 794	7.8%	8.0%	21	20
46	41	Nippon Telegraph & Telephone Corp., JAPAN	2286	2592	-11.8%	90 429	90 266	2.5%	2.9%	11	13
47	57	Abbott Laboratories, U.S.	2255	1821	23.8%	22 476	22 288	10.0%	8.2%	34	30
48	55	Oracle Corp., U.S.	2195	1872	17.3%	17 996	14 380	12.2%	13.0%	29	33
49	29	Koninklijke Philips Electronics NV, NETHERLANDS	2192	3356	-34.7%	35 457	39 951	6.2%	8.4%	18	21
50	51	Texas Instruments Inc., U.S.	2190	2015	8.7%	14 195	13 392	15.4%	15.1%	71	57

CONTINUES

and Exchange Commission does in the United States. However, the quality of information in those reports varies greatly.

Toyota's filing provides a lot of detail on the company's R&D strategy. The company goes out of its way to declare that maintaining leadership in R&D is a key to improving overall performance. The argument, however, avoids any acknowledgement of the company's low R&D Intensity relative to that of other carmakers. Toyota's filing asserts that its R&D priorities are to develop such environmentally friendly technologies as hybrid gas-electric drive, fuel cells, and recyclable materials. But it gives little information on the scale of those investments. Of course, Toyota is well known for its Prius hybrid, an iconic symbol for green-conscious drivers. So it has already received large dividends from its R&D in environmental technologies.

In contrast, Ford's description of its impressive R&D program can charitably be described as perfunctory. It simply lists, in two brief paragraphs, how much it spends and where it spends it. Ford does not provide a rationale for last year's 10 percent cut, nor does it try to play up the company's continuing position as one of the most R&D-intense companies in the world. Ford's brevity may reflect other, competing goals in such financial disclosures; the company may, for instance, wish to keep its competitors in the dark.

Fortunately for investors, most other automakers provide a level of detail that is closer to Toyota's than to Ford's. Indeed, many companies cite their R&D ranking as a source of strength. For example, Samsung Electronics Co. says in its filing, "In 2006, the *Financial Times* ranked Samsung Electronics ninth in R&D investment among 1250 companies around the world.... This news-

CONTINUED

SPECTRUM'S TOP R&D SPENDERS			R&D EXPENDITURES (US \$ MILLIONS, LESS IN-PROCESS*)			SALES (\$ MILLIONS)		R&D INTENSITY†		R&D EXPENDITURES PER EMPLOYEE (\$ THOUSANDS)	
2006	2005	COMPANY, COUNTRY	2006	2005	% CHANGE	2006	2005	2006	2005	2006	2005
51	56	Schering-Plough Corp., U.S.	2173	1865	16.5%	10 547	9508	20.6%	19.6%	65	57
52	50	Fujitsu Ltd., JAPAN	2135	2030	5.2%	42 861	40 266	5.0%	5.0%	13	13
53	47	Delphi Corp., U.S.	2100	2200	-4.5%	26 392	26 947	8.0%	8.2%	12	12
54	52	Procter & Gamble Co., U.S.	2075	1940	7.0%	68 222	56 741	3.0%	3.4%	15	18
55	53	Nortel Networks Corp., CANADA	2032	1924	5.6%	11 829	10 906	17.2%	17.6%	60	54
56	49	Sun Microsystems Inc., U.S.	2008	2046	-1.9%	13 873	13 068	14.5%	15.7%	53	54
57	54	Alcatel-Lucent, FRANCE	1906	1897	0.5%	16 143	17 265	11.8%	11.0%	21	33
58	61	SAP AG, GERMANY	1755	1431	22.6%	12 358	11 189	14.2%	12.8%	46	40
59	64	BASF AG, GERMANY	1679	1398	20.1%	69 149	56 183	2.4%	2.5%	18	17
60	59	STMicroelectronics NV, SWITZERLAND	1659	1627	2.0%	9838	8876	16.9%	18.3%	32	33
61	58	Infineon Technologies AG, GERMANY	1630	1700	-4.1%	10 422	8884	15.6%	19.1%	39	47
62	62	Takeda Pharmaceutical Co., JAPAN	1625	1426	13.9%	10 968	10 187	14.8%	14.0%	108	95
63	69	Sharp Corp., JAPAN	1596	1297	23.0%	26 285	23 506	6.1%	5.5%	33	28
64	76	Genentech Inc., U.S.	1588	1151	38.0%	9284	6633	17.1%	17.4%	151	121
65	65	United Technologies Corp., U.S.	1529	1367	11.9%	47 715	42 584	3.2%	3.2%	7	6
66	89	Qualcomm Inc., U.S.	1516	1011	50.0%	7526	5673	20.1%	17.8%	135	109
67	60	Fujifilm Holdings Corp., JAPAN	1488	1531	-2.8%	23 384	22 417	6.4%	6.8%	20	20
68	68	Daiichi Sankyo Co., JAPAN	1434	1334	7.5%	7811	7781	18.4%	17.1%	93	72
69	72	3M Co., U.S.	1427	1242	14.9%	22 923	21 167	6.2%	5.9%	19	18
70	74	Astellas Pharma Inc., JAPAN	1411	1194	18.2%	7737	7390	18.2%	16.2%	102	80
71	84	Honeywell International Inc., U.S.	1411	1072	31.6%	31 367	27 653	4.5%	3.9%	12	9
72	73	Nestlé SA, SWITZERLAND	1410	1219	15.7%	80 077	74 072	1.8%	1.6%	5	5
73	92	BT Group PLC, UNITED KINGDOM	1349	947	42.4%	39 412	38 030	3.4%	2.5%	13	9
74	70	Bayer Schering Pharma AG, GERMANY	1349	1291	4.5%	7449	6977	18.1%	18.5%	58	53
75	81	Caterpillar Inc., U.S.	1347	1084	24.3%	41 517	36 339	3.2%	3.0%	14	13
76	67	E.I. du Pont de Nemours & Co., U.S.	1302	1336	-2.5%	28 356	27 516	4.6%	4.9%	22	22
77	90	EMC Corp., U.S.	1254	1005	24.8%	11 155	9664	11.2%	10.4%	40	38
78	78	Medtronic Inc., U.S.	1239	1113	11.3%	12 299	11 292	10.1%	9.9%	33	31
79	119	Google Inc., U.S.	1218	578	110.9%	10 605	6139	11.5%	9.4%	114	102
80	77	Advanced Micro Devices Inc., U.S.	1205	1144	5.3%	5649	5848	21.3%	19.6%	73	116
81	82	AB Volvo, SWEDEN	1194	1080	10.5%	36 984	34 372	3.2%	3.1%	14	13
82	71	Unilever NV, NETHERLANDS	1191	1253	-4.9%	52 105	52 144	2.3%	2.4%	6	6
83	83	Dow Chemical Co., U.S.	1164	1073	8.5%	49 124	46 307	2.4%	2.3%	27	25
84	80	Akzo Nobel NV, NETHERLANDS	1163	1096	6.1%	18 056	17 087	6.4%	6.4%	19	18
85	87	Lockheed Martin Corp., U.S.	1139	1042	9.3%	39 620	37 213	2.9%	2.8%	8	8
86	95	Applied Materials Inc., U.S.	1138	941	21.0%	9167	6992	12.4%	13.5%	81	73
87	94	France Telecom SA, FRANCE	1125	941	19.6%	67 956	64 455	1.7%	1.5%	6	5
88	99	Novo Nordisk A/S, DENMARK	1122	897	25.1%	6833	5954	16.4%	15.1%	48	41
89	115	Broadcom Corp., U.S.	1117	681	64.0%	3668	2671	30.5%	25.5%	213	159
90	85	Hyundai Motor Co., SOUTH KOREA	1116	1068	4.5%	67 830	62 696	1.6%	1.7%	20	20
91	79	Mitsubishi Electric Corp., JAPAN	1115	1098	1.6%	32 403	30 289	3.4%	3.6%	11	11
92	88	NEC Electronics Corp., JAPAN	1107	1016	9.0%	5818	5429	19.0%	18.7%	46	43
93	86	Sanyo Electric Co., JAPAN	1070	1065	0.4%	19 401	20 878	5.5%	5.1%	11	10
94	66	LG Electronics Inc., SOUTH KOREA	1046	1357	-22.9%	49 384	47 365	2.1%	2.9%	34	43
95	108	Electronic Arts Inc., U.S.	1041	758	37.3%	3091	2951	33.7%	25.7%	132	105
96	75	Fiat SpA, ITALY	1032	1186	-13.0%	68 127	61 177	1.5%	1.9%	6	7
97	116	Boston Scientific Corp., U.S.	1008	680	48.2%	7821	6283	12.9%	10.8%	35	34
98	91	QinetiQ Group PLC, UNITED KINGDOM	1005	997	0.8%	2240	2050	44.9%	48.6%	79	87
99	93	Altria Group Inc., U.S.	1005	943	6.6%	70 324	68 920	1.4%	1.4%	6	5
100	96	Merck KGaA, GERMANY	988	937	5.4%	8226	7716	12.0%	12.1%	33	32

*Less In-Process—R&D expenditures less those accrued as a result of a merger or acquisition. †R&D as a percentage of total sales.

Data comes from fiscal 2005 and 2006 for all the companies except Microsoft and Sun Microsystems, for which it comes from fiscal 2006 and 2007.

Because companies have restated figures for 2005, rankings may differ from those published last year. Source: Standard & Poor's

paper reported that over the past four years, Samsung's massive investment in R&D has had a great impact on the electronics industry, prompting competitors to spend more on R&D." As *Spectrum* reported last year, Samsung surpassed Intel as the leading spender on R&D in the semiconductor industry, a position the South Korean company maintained this year in spite of double-digit growth in Intel's spending [see "IBM Takes the Guesswork Out of Services Consulting," *Spectrum Online*, December 2006].

Another metric of R&D is the number of patents that come out of it. Of course here, too, there must necessarily be a lag between the investment and the result. Samsung boasts of regis-

tering 2474 new patents in the United States during 2006, raising it three notches to place second, behind IBM, the world leader for the 14th year in a row. The problem is that the business value of patents varies greatly, so the sheer number of patents correlates loosely at best with a firm's actual performance [see "Keeping Score in the IP Game," *Spectrum*, November].

The filings also provide a window into R&D collaborations between companies. For example, Motorola highlighted its creation of a new joint research facility with Huawei Technologies, in Shanghai, to bring the Universal Mobile Telecommunications System and High-Speed Downlink Packet Access cellphone

technologies to market. Intel touts its collaboration with Micron Technology to develop NAND flash-memory technologies.

One more point: the year-by-year filings provide a kind of slide show of the process of globalization. The top R&D spenders are all multinational corporations, and their R&D operations are themselves increasingly dispersed around the world. Microsoft has R&D facilities in Canada, China, Denmark, India, Ireland, Israel, and the United Kingdom. Even upstart Google has R&D centers in China, India, Israel, Japan, and Russia.

HOW WALL STREET'S PROS EVALUATE R&D depends on the industry, the company, and the individual analyst. R&D Intensity is always the analyst's main benchmark, but the key question is what, exactly, it is supposed to be measuring. R&D is an investment in the future and also an expense against current earnings. An analyst may choose to favor either side of the equation.

Stephen R. Biggar, director of U.S. equity research for Standard & Poor's (which is separate from the data-generating arm of the company that supplied the R&D statistics for this article) says that Wall Street's desire for "instant gratification" is too high, and it's getting higher. He blames a short-term outlook that puts pressure on companies to shoot for ever-quicker payoffs, which in turn tends to make them shortchange R&D. The reason is that lag time again: it takes a long time to yield profits—up to 15 years in the pharmaceutical industry. That's forever to most analysts, who generally forecast revenues and earnings just two or three years out. He says those pressures are stronger in the United States than in other countries (perhaps because boards of directors in those countries are less in thrall to shareholders).

Biggar says pressure for quick payoffs isn't all bad. It induces companies to try to squeeze what they can out of the plant and equipment they already have, which is good for efficiency so long as no technological revolution intervenes to render that equipment obsolete. He also notes that R&D is viewed differently in each sector. He says a good rule of thumb is that the higher a sector's average R&D Intensity tends to be, the more important R&D will be to analysts.

Consider these three industries: pharmaceuticals, semiconductors, and software. They all have very high R&D intensities, but each invests in R&D in its own way, shaped by its own risks, time to market, industrial organization, regulatory regimes, and business models.

Pharmaceuticals companies live and die on R&D: their R&D Intensity averages 16.4 percent. Because the vast majority of apparently promising compounds end up as failures, a firm must sink billions of dollars over many years just to get one or two successes. That's why a drug company's fortunes can turn on the result of a single patent trial; it's also why pharmaceutical analysts work ferociously to track R&D projects as they snake their way through the many stages of the pipeline.

Herman Saftlas, who covers some major pharmaceutical companies for Standard & Poor's, says most firms highlight their pipelines much as a manufacturer might account for back orders. Pfizer, for example, documents its drug pipeline in detail in a special report (<http://www.pfizer.com/pipeline>). The transparency of pipelines allows analysts an opportunity to evaluate the productivity of a firm's R&D, and Saftlas has concluded that some firms are simply better at getting more for their R&D buck than others. (He notes, for instance, that Pfizer's R&D performance has been below average for its sector, in part because its 2003 acquisition of Pharmacia Corp. hasn't worked out as well as it expected; Merck, by contrast, has bettered the average.) Firms have adjusted their R&D strategies to fit the changing marketplace for drugs. The

attrition rate for candidate drugs has risen so high that firms are beginning to give up hope of bringing in billions with blockbuster drugs, and are instead settling for mere tens of millions in niche markets. It's the difference between, say, a Viagra and a longer-acting antihistamine pill. The development timeline for such niches is shorter and less risky.

Semiconductor manufacturers are putting a high and rising share of their resources into R&D, for a research intensity averaging 17.1 percent. Still, Clyde Montevirgen, who covers the sector for Standard & Poor's, says R&D is not his single most important metric—sales are. He reasons that R&D still constitutes a small part of overall costs, and one that is hard to gauge because most firms jealously guard their data, rarely talking at all about developments that are more than a year away from the market. Montevirgen therefore measures R&D effectiveness by "design wins," in which a firm announces a partnership with a hardware device maker. For example, in August, STMicroelectronics, No. 60 on the leaderboard, announced that Garmin had selected STM's chips for its new range of portable and handheld GPS and navigation devices.

Market changes since 2000 have pushed semiconductor firms to ratchet up their R&D budgets, Montevirgen says. Because an ever-broadening array of products incorporate semiconductors—think PDAs, smart phones, GPS, iPods—many in the industry hope that the market will grow even faster than it already has. However, Montevirgen cautions that the semiconductor industry is a highly cyclic market and R&D spending will surely follow the cycle. If the market turns downward, expect R&D spending to drop back.

Jim Yin, who covers major software firms for Standard & Poor's, puts less emphasis on R&D spending than his two colleagues do, even though the sector's average R&D Intensity of 17.9 percent puts it higher than those of semiconductors and pharmaceuticals. He looks, instead, mainly at sales.

Because R&D spending varies widely, he says, it's hard to compare numbers across the industry. Investments increase rapidly before a product launch and then fall quickly afterward. A young firm spends all it has to make its first launch, while an established firm can get sales out of existing lines with far less R&D investment. Software projects are notoriously prone to late delivery, and that makes it even harder to forecast the payoffs of R&D.

To limit the confusion, Yin focuses not on total R&D but just on the portion budgeted for a specific product launch he is following. He also divides software into categories requiring differing R&D investment. For example, a game company may need to spend more per product because each game has features that cannot be shared easily with others. Yin adds that Microsoft's immense profitability allows it to place bigger R&D bets than other firms do and to wait longer for those bets to pay off.

SO, KEEP THE FOLLOWING IN MIND as you review our R&D leaderboard. R&D spending and intensity do matter, but different industries require different R&D investments, as do companies in different niches within an industry and at different stages of development within a niche. Whether a firm emphasizes quick or long-term payoffs depends on its country, its industry, its maturity, and its strategy. Finally, what matters isn't how much money you spend but how wisely you spend it. Unfortunately, our leaderboard doesn't have a column for wisdom. ■

ABOUT THE AUTHOR

RON HIRA is an assistant professor of public policy at the Rochester Institute of Technology, in New York (rhira@mail.rit.edu). He is past chairman of the Research & Development Policy Committee of IEEE-USA.




```
<html><head><title>=
"Playing Dirty"
//Automating computer
game play takes
cheating to a new-
and profitable-level//
<By David Kushner>
```

Richard Thurman is like a lot of 35-year-old guys. He's married. Has a couple kids. When he wants to blow off steam, he flops into his chair in front of his PC, and he fires up a computer game.

But Thurman is no ordinary player. In the weird and burgeoning virtual universe, he's a former outlaw. While earnest gaming geeks spend hours slaying dragons to earn booty playing Sony's *EverQuest*, Blizzard Entertainment's *World of Warcraft*, and other multiplayer online games, Thurman spent years using his coding chops to cut to the chase: rigging his computers to play games automatically and rake in gold. It took three months and 50 000 lines of code to pull off the feat. And it was all perfectly legal, at least in the real world.

In December 2003, however, when a person or persons unknown of the gaming underworld began threatening Thurman's real-world family, he

unplugged his operation and took a programming job with a major corporation, which he'd prefer not to name. Now that he's out of his gaming business, he agreed to give *IEEE Spectrum* an inside look at his pioneering automated gold-farming system. The games today have changed, but the way a person profits from them remains very much the same.

Players and game makers despise the kind of hacking that was Thurman's specialty, because it makes their lives more difficult. That doesn't bother Thurman. "I'm a metagamer," he says. "Game companies lay down their rules. Some play by them, and some don't."

Thurman wasn't hacking for fun. In the new online economy, virtual cash, earned in games by killing a monster or performing a service, has real-

PHOTO: JEFF NEWTON; DIGITAL ILLUSTRATION: SANDBOX STUDIO



GAMEBOTS: When Thurman went into automated gold farming, he purchased stacks of computer gear [bottom] and set up a bank of computers [top] to play *Ultima Online* for virtual currency, which he later exchanged for real money.



world value, thanks to sites specializing in what are called real-money transactions (RMTs). People covet the jewel-encrusted super-sword in a game but can't spare the time to log the kind of hours they'd need to actually earn the virtual gold to buy it. So they obtain it the newfangled way: with their credit cards. In other words, they pay real money to buy virtual things.

Edward Castronova, an associate professor of telecommunications at Indiana University, in Bloomington, and author of *Synthetic Worlds: The Business and Culture of Online Games* (University of Chicago Press, 2005), puts the annual total market value for virtual assets between US \$200 million and \$1 billion. Although that may sound like small potatoes—the cellphone ringtone market is roughly \$5 billion per year—the cheating is

already wreaking havoc in the virtual worlds. In one episode a few years ago, cheaters unleashed fake currency into the world of *EverQuest*, one of the most popular online games, inflating its economy by 20 percent.

Gamers and game makers are feeling swindled. "It's criminal, in the context of a virtual world," says Scott Hartsman, senior producer and creative director of *EverQuest II* at Sony Computer Entertainment America, in Foster City, Calif. "The entire reason societies have laws and mores is to protect people from getting hurt. By definition, people are getting hurt."

There are odd and controversial real-world repercussions to the cheating. News accounts during the past year have described the rise of sweatshops in Asia, especially China, where low-paid workers play online games for 12 hours a day to amass virtual goods to be sold on the black market.

"This is evidence that there really isn't anything special about virtual worlds," Castronova says. "We've been reading about globalization of labor markets, about software engineers in India taking jobs, and this is just another example of that phenomenon. Americans will spend money for online goods; wage rates are lower in Shanghai. The Internet allows [these transactions] to happen. It's the globalization of the labor market."

Isn't this unauthorized activity illegal? Aside from possible violation of local labor laws, the answer is no. No real-world laws cover online gaming, so the players and makers instead rely on their own terms of agreement, which users accept when they install games on their home computers. The agreements basically state that everyone will play by the rules—and allow the delicate balances of make-believe worlds to survive. But none of it is legally binding anywhere in the world.

Thurman was one of the first geeks to take breaking the rules of virtual worlds to a new level by engineering the automation of gold farming. Many others followed his lead. Although no one knows for sure how many gold farmers there are, Thurman guesses as many as a million worldwide. Their shadowy world has become big enough to have its own published manifesto: Gary McGraw and Greg Hoglund's *Exploiting Online Games* (Addison-Wesley, 2007).

Thurman has been part of it from the start. You might even say he helped establish it. A software specialist with a bachelor's degree in business information systems and a master's in computer science, both from the University of Phoenix, he spent three years applying himself to milking *Ultima Online*, then one of the most popular multiplayer games, for all he could. At his peak, he had a fleet of 30 computers automatically raking in game gold, earning him more than \$25 000 per month.

Subverting video games isn't new. Geeks have been figuring out how to exploit game technology to their advantage for decades, giving themselves extra "lives" in *Pac-Man* or switching into invincible "God mode" in *Doom*. When massively multiplayer games such as *Ultima Online*, from Electronic Arts of Redwood City, Calif., and *EverQuest* came onto the scene during the last decade, the emergence of virtual economies raised the stakes. You weren't just competing for ego anymore; you were gaming for dollars.

Other factors helped attract hackers. For example, economies of scale. Online games are not just for nerds. The action is mainstream. Hordes of engineers, accountants, lawyers, and

other wannabe knights and knaves do battle in *EverQuest* (dubbed “EverCrack” for its addictiveness), *World of Warcraft*, and other games. Schoolchildren, college students, and GenXers are playing such online games as *Halo 3* on the Xbox 360 or *Madden NFL 2008* on the Playstation 3. Many graying gamers take to casual online games, such as bridge and chess. It doesn’t take much more than a computer and an IP address to access your passion.

Thurman started playing *Ultima Online* as an undergraduate in 1997. He couldn’t help but wonder if, through a few hacks, there was a way to make his game-playing experience better. After surfing around, he came upon software such as UOAssist and EasyUO. When run in conjunction with a game, those programs gave players advanced macros, which are keyboard shortcuts to speed up mundane tasks such as healing yourself after battle. He realized he was on to something.

Thurman left Phoenix in 1998, moved to Dallas, and began working full time as a support engineer for a large software company, which he also prefers not to name. He continued thinking about hacking *Ultima Online*, and he became aware of the growing real-world market for virtual gold. The problem was that he couldn’t amass it fast enough to make a decent buck. But, he thought, if he could create an auto-playing robot, something that could basically play the game for him—then maybe he could cash in.

Drawing on his programming knowledge and with the help of DIY hacker sources online, such as Fravia.com, Thurman got to work. He started by shelling out \$800 for a reverse-engineering software tool called IDA Pro from DataRescue of Liège, Belgium. IDA Pro lets users see the structure of a program’s logic. Point it at a program, and it creates a flowchart of how the software works. Thurman directed the tool to the “client” software he’d downloaded to his PC to let him to play *Ultima Online*. (The client software is what every player downloads in order to play.)

Basically, IDA Pro reverse-engineered *Ultima Online*’s inner workings. Not only did it let Thurman see the basic functions of the client software, it also let him see the specific memory addresses where the software stored key variables such as the player’s location in the game world, an inventory of the player’s possessions, and the status of the player’s health.

That information led Thurman to write a chunk of C++ code that he inserted into the client software to allow it to communicate with Microsoft.Net, a development environment for Windows computers. In effect, the C++ code functioned as a kind of outlet to the servers running the game. With that done, he needed, essentially, to write a plug to stick into the outlet. He wrote that plug in Visual Basic. Once complete and installed in his machine, it could exchange information with the *Ultima Online* client in his computer and, through that client software, the *Ultima Online* servers at the Redwood City headquarters of Electronic Arts. In other words, he got access to the brains running the game.

Next, Thurman set up his bank of computers. He chose the cheapest off-the-shelf PCs available that had enough power to run *Ultima Online*, and he bought 30 of them. Each was equipped with an Intel Pentium 4 or a Celeron processor, a gigabyte of RAM, and a 20-GB hard drive. He connected the bank of PCs to three monitors and a network of six cable modems, four routers, and a Toshiba tablet PC that he used to manage the whole operation.

Then he got down to business. The plan was that each

of the 30 PCs would play the game individually, creating a character and then using that character to perform tasks that would earn gold. Thurman wrote software to randomly generate details about the characters—names, classes (fisherman, say, or fighter), and skills (such as magic or cooking), saving him the trouble of creating each character manually. He cloaked his identity by purchasing anonymous gift cards to set up accounts rather than paying for them with a personal credit card (the gift cards are no longer being sold).

Once his computers logged into a game, communication between them and the game server was fairly straightforward. For every action happening in the game that involved one of Thurman’s 30 characters, the game server sent the details back to the relevant client computer, and vice versa. The details included the skills of a character, the status of its health, and the size of its bank account. Thurman eliminated the human element—cut out the middleman, you might say—by programming his computers to automatically respond to the incoming data from the game server.

The application performed the functions that a normal player would have to do with many repetitive keystrokes (*Ultima Online* players use keyboards, not joysticks). One thing the program couldn’t do was sniff out moneymaking opportunities, so Thurman did that himself. But once he identified an opportunity, he would quickly write code that told his characters what to do to capitalize.

For example, in *Ultima Online*, gamers can make money by cooking and selling chickens to tavern keepers. Thurman programmed his characters to buy raw birds from the butcher and then prepare the food. Ordinarily, a gamer can cook only one bird at a time, but Thurman automated the process so that his 30 PCs could cook as many as 500 birds at a time; he sold them in huge quantities to the taverns. In minutes, his bank of computers could rack up an amount of virtual money that it would take an individual player weeks to earn.


But wouldn’t it be easy to spot a user who was cooking and selling, in minutes, enough chicken to feed an army? Absolutely. And that’s where the real finesse of being a game hacker comes in. A big part of the tradecraft is simply managing to avoid getting busted by the company game masters, whose job it is to prowl for hackers. If they even suspect illicit activity, they look up the associated Internet Protocol address and can take action. “They would mass-ban your accounts,” Thurman notes.

So he installed countermeasures. First, he got a separate account for each of the 30 computers. He had six cable modems, with five accounts tied to each one. He also paid his Internet service provider an extra \$16 per month to get four IP addresses to use (most households have just one), and wrote software to instruct the modems to release one of those IP addresses every six hours and grab a new one to replace it. In a network with dynamically assigned IP addresses, any modem outage and reboot results in a new address assigned; Thurman effectively generated his own outages so that he could get new IP addresses. His constantly shifting array of IP addresses made it hard for the sleuths at Electronic Arts to notice the fantastic quantities of chicken he was selling, to say nothing of the ore he

was mining, melting into ingots, and exchanging for game currency.

But churning the IP addresses wasn’t a fool-





MAP HACK: Thurman's game-playing system earned gold by making maps of the virtual world and selling them at virtual cartography shops. He simplified the mapmaking process with a hack that displayed buildings without their roofs.

proof countermeasure, he realized. Just in case his activity aroused suspicion, he rigged his bank of computers to alert him via text message or instant message to odd bursts of activity—for example, when a person from Electronic Arts was confronting one of his automated systems to see if it was, in fact, a real player or just a proxy.

That happened a few times, Thurman says, and they were close calls. One time he was traveling in Arizona when an instant message came through on his phone. The game server had sent a message to his client indicating that a game master, an employee or volunteer who, in the form of a game character, roves the game enforcing rules, was on screen. Game masters are identifiable by a special flag their avatars carry. "GM Alert!" the message read. Thurman had set up the machines to automatically log out his other characters when that happened, just in case. But he left his one character online with the GM because it'd be too suspicious if he suddenly vanished.

Game masters try to verify that players are in front of their monitors, often by challenging them with questions that they presumably could answer only if they were sitting in front of the screen. But Thurman had anticipated such a challenge, and he had rigged his instant messaging system so that it could send crude but useful screen shots to his laptop computer. "Are you there?" the GM asked. "Yes," Thurman replied. "Prove it," the GM replied. "What color is my shirt?" No problem. "Red," Thurman typed after glancing at the screen shot. And the GM went on his way.

It took Thurman nearly two years, from February 2002 to December 2003, to perfect his system. The "labor of love," as he describes it, paid off. Soon he was making 45 000 units

of gold per hour and, eventually, as much as 2 million units of gold every 15 minutes. All told, that translated into as much as \$2400 per hour of *real money*: \$80 per hour per character, and Thurman had up to 30 characters at his disposal. It was around then that he quit his day job as a software consultant.

With "game gold" in hand, the next step is converting the virtual cash to real-world money. Dozens of companies are happy to help gamers do that. The biggest is Hong Kong-based IGE, which Thurman compares to Wal-Mart. The company employs more than 800 people in Seoul, Hong Kong, and Shanghai. Founder Brock Pierce said in a phone interview last year that the site brokers real-money transactions, taking a piece off the top as it connects sellers of virtual gold, earned legitimately, with buyers. He put the estimated annual earnings at \$700 million. (IGE did not respond to requests for an update.)

But the secondary market is, to put it mildly, shadowy. It revolves to some extent around hackers who scoff at efforts by online game companies to fight against automated software. It also depends on hundreds of loosely organized gold farmers in China, who game for money around the clock and then cash out their winnings to online brokers. They may not be breaking any rules, technically, but they are sure violating the spirit of the games. In a sense, such people constitute a manual version of the automated software written by the likes of Thurman.


Documentary filmmaker Ge Jin has been chronicling the gold farms in China for a movie to be released next year. He says that while gold farming may be an oddity—if not anathema—in the West, it's more widely accepted abroad. "The unemployment rate is soaring in China," he says, "so [hired gold farmers] are happy to have a job, which pays no less than other jobs available to them. The majority of them are game fans anyway; they are happy that they can be paid for playing games and can enjoy games that are expensive to subscribe to or even those not imported into China."

According to a June report in *The New York Times Magazine* by Julian Dibbell, a typical gold farmer in China works 12-hour days for weeks on end, with only a few days of rest per month. The farmers work at long tables strewn with computer monitors and keyboards in small rooms crowded with dozens of people and thick with cigarette smoke. Dibbell estimated that 100 000 such workers are employed in what are called *youxi gongzuoshi*, or gaming workshops.

Unlike Thurman, the Chinese workers actually do go out into the "worlds" and game. But they do so in teams—which gives them a distinct advantage in certain situations. For example, they can gang up on giant monsters whose slaughter will be rewarded with big piles of gold. "Gold farmers attack high-level mothers," Thurman says, a little enviously. "They're not cooking birds."

Patrick Bernard, 31, joined a worldwide gold-farming team after working as a product manager for a Silicon Valley dot-com (he declines to say which one). The gold-farming work quickly became tedious, he says. "We just pooled monsters and killed them for dozens of hours," he says. "I could generate \$1000 in gold per hour; my

pay, at the time, was \$15 per hour." Bernard now works on the other side of the business, running *Gamer's Loot*, an online



RMT service—one of the companies that convert game gold into real money and vice versa.

While there's no law against real-money transactions, game companies are understandably uncomfortable with the whole idea. "We all admit that sort of thing is out there," says David Swofford, spokesman for NCsoft of Seoul, maker of the online game *Lineage*. "But it's not anything we endorse.

"It's a hazard of the business. What we're trying to do is have games create the best possible experience. If people are doing things in violation of rules...we don't want them in our game."

Game companies and hardware manufacturers such as Intel are going after hackers with varying degrees of aggressiveness. Among the most intense is Sony. The company says that during the past few years it has booted out more than 20 000 players suspected of farming gold in *EverQuest*; *Star Wars: Galaxies*; *Vanguard: Saga of Heroes*; and other Sony online games. And the game companies don't take kindly to operations like IGE. "They claim they don't have any employees doing farming," Sony's Hartsman says, "but they have thousands of contractors doing it. You push a button that says, 'I would like to sell a coin'; within 5 minutes you have people respond. And they're not asking where that coin came from."

But here's a hint at how alluring, and maybe insidious, gold conversion is: for all its prosecutorial zeal, Sony itself has succumbed to the temptations of gold conversion. It now has its own service, called the Sony Exchange, which allows players to buy and sell virtual items online. Sony gets a cut, of course.

Real gamers are fed up. "It's disconcerting to find out that the warrior decked out in purple epic bling bought all his kit on eBay," complains Drew Shiel, webmaster of a *World of Warcraft* fan site called the Wizard of Duke Street (<http://www.dukestreet.org>). "Having someone come in and buy a high-level character makes a mockery of the effort that other people have put into their own characters."

But as online gaming worlds become more realistic, there's little chance of getting rid of the perceived criminal element. If anything, the most gamers and companies can hope for is that the metagamers eventually do what Thurman did: grow up and log off.

After a couple of years of gaming for dollars, Thurman got tired of living on the edge. The clincher came when a competing gold farmer began sending him messages threatening him and his family. "We thought he'd show up at our house and kill us," Thurman says. Even the biggest sword in *Ultima Online* would not be able to protect him from that. ■

ABOUT THE AUTHOR

Contributing Editor **DAVID KUSHNER** blogs for Spectrum Online at <http://blogs.spectrum.ieee.org/gizmos/>. His latest book is *Jonny Magic and the Card Shark Kids* (Random House, 2005).

TO PROBE FURTHER

To check out the largest online retailer of virtual items for games, go to <http://www.ige.com>.

Julian Dibbell's book on his adventures inside the virtual game economy, *Play Money: How I Quit My Day Job and Made Millions Trading Virtual Loot*, was just released as a paperback by Basic Books. Dibbell's *New York Times Magazine* article, "The Life of the Chinese Gold Farmer," appeared in the 17 June 2007 issue.

www.spectrum.ieee.org

The Oldest Profession

Of all the weird new businesses supported by virtual gold, one of the strangest is related to prostitution. Online, real-world women (or men posing as women) give gamers text-based and animated cybersex in exchange for game gold. The sex workers then trade that virtual dough for real cash through real-money transaction sites.

Under the nom de jeu of Stacey Sugar, a 25-year-old from England runs a popular strip joint in Second Life, an online virtual community where players customize their characters and environments. Industrious players create elaborate strip

clubs and brothels.

If your character doesn't have genitals, no problem, you can buy yourself whatever you need in a Second Life store that sells virtual doodads in exchange for virtual dollars.



Sugar runs the Club XTC Elite, a virtual brothel complete with "pole dancing, lap dancing, table dancing, booth dancing, cage dancing, shower dancing, a champagne room, and four private fully fitted and animated sex suites," she says. Visitors pay in game cash, called Linden Dollars, so that their game characters, or avatars, can have "sex"—expressed through titillating chat and character animations—with virtual prostitutes controlled by real people. Sugar gets a 20 percent cut. Translated into real money, she's earning roughly \$7 per hour for work that's less tedious than flipping burgers at McDonald's.


The people who control the virtual prostitutes get all the training and accessories they need. "All girls who come to us complete an employment application form," Sugar says, "and if accepted go into our mentoring program [called 'Bimbos-R-Us'], which provides the avatar with a revised body, sexy skin and makeup, shoes, a large range of clothes, bling, and hair—everything they need to look gorgeous fast."

One of the most notorious video-game tarts is Kannea Suntzu. In real life, she's a 30-something divorcée living on disability checks in the Netherlands. In the game world, she readily parts players from their in-game cash. It's not just a matter of slinging salacious type or wiggling her cartoon character on yours. For the right price, Suntzu takes the action out of the game and into real life via the online telephony software Skype.

But that's where the virtual story ends.

—D.K.





CONTROLLED CHAOS

We need to exploit the science of order and disorder to protect networks against coming generations of **SUPERWORMS**

BY ANTONIO NUCCI & STEVE BANNERMAN

Internet security professionals are, by occupational temperament, a pretty nervous bunch. But lately they've had more reason than ever to be jumpy. Early this year, a new kind of worm, known as Storm, began to sweep through the Internet. It hasn't received much attention in the mainstream press, but it has given security professionals more than a few sleepless nights. Storm is far more sophisticated than previous worms, because it uses peer-to-peer technologies and other novel techniques to evade detection and to spread. The popular press hasn't paid much attention to Storm, because it has yet to wreak devastating havoc on businesses, as some previous worms have. But we shouldn't be fooled by that relative quiet: Storm's designers appear to be biding their time, building an attack network far more disruptive than any before seen.

HARRY CAMPBELL

Storm methodically infiltrates computers with dormant code that could be used to take down the entire network of a corporation, creating opportunities for blackmail or for profiting by selling the company's stock short. And Storm's creators, whoever they are, continue to modify and refine their malevolent progeny even as it already stands as a dark cloud poised over the Internet.

Network security software products on the market today offer only limited defense. They use firewalls, which simply block access to unauthorized users, and software patches, which can be created only after a worm or virus's unique bit pattern is discerned. By the time this laborious process of hand coding is complete, the infestation has had hours and hours to spread, mutate, or be modified by its creators.

A new kind of answer is needed. Network security researchers—including ones at our company, Narus, in Mountain View, Calif.—are developing software that can rapidly detect a wide variety of intrusions from worms, viruses, and other attacks without the high rate of false alarms that plagues many conventional Internet security products. These new programs can detect anomalous network behavior in seconds, as opposed to hours or days—even on so-called backbone networks running at 10 billion bits per second. That means the software is fast enough to block threats that can span the globe in minutes, a rate that far outpaces what a firewall can monitor.

This new generation of algorithms is based on concepts related to the thermodynamic concept of entropy. Often defined briefly as a measure of the disorder of a system, entropy as a cornerstone of thermodynamic theory goes back more than a century and a half. But as a construct of information theory it is only 60 years old, and its application to data communications began only in the last decade or so.

In essence, an entropy-based defense works because a worm's malicious activity changes, in subtle but unavoidable ways, the character of the flow of data on a network. Those data flow changes alter, in clearly measurable ways, the entropy of the network—a measure of the endlessly shifting ebb and flow between the predictability and randomness of the movement of data on the network.

Researchers at Intel, Microsoft, Boston University, and the University of Massachusetts are among those plumbing the mysteries of randomness and order in data flows to get a leg up on network attackers. Although ours is the only company we know of whose commercial products apply entropy to network security, we are confident that the approach will find much wider favor in the next few years.

We'll have lots more to say about entropy and how algorithms that measure changes to the order and disorder of a network can detect a worm outbreak long before traditional methods can. But to get a grip on those algorithms, first consider how viruses and worms attack.

VIRUS OR WORM? Security experts distinguish between them, but their differences are less important than their similarities. Either can render computers on a network unstable, and in many

cases unusable. A virus is a program that can copy itself and infect a computer without the knowledge of the user. It can, and often does, damage a computer's files or the hardware itself. A worm is, similarly, a self-replicating computer program that uses a network to send copies of itself from one computer, which we will call a "host" of the infection, to other computers on the network. Worms usually harm the network, if only by consuming bandwidth.

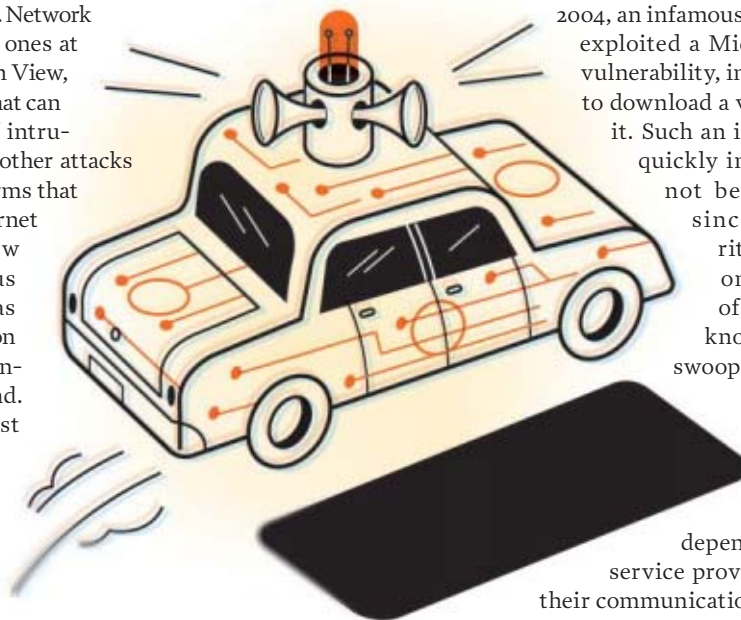
An e-mail worm, the most common kind, spreads slowly, because users have to click on an attachment to become infected or to propagate the worm. Storm is one example; it uses a variety of means to get installed on a host, but the most common one is the e-mail attachment. Not all worms spread by e-mail; in

2004, an infamous worm called Sasser instead exploited a Microsoft Windows network vulnerability, instructing infected systems to download a viral code and then execute it. Such an infestation can spread very quickly indeed. Although there has not been a catastrophic worm since Sasser, network security systems still have to be on guard against this sort of attack, because we never know when the next one will swoop down on us.

Faced with attacks that could occur too quickly for their firewalls to cope with, companies and governments are now depending on Internet and other service providers from which they buy their communications bandwidth to "clean the traffic" before it ever reaches their front doors. The world's largest carriers, such as AT&T, BT, Korea Telecom, NTT, and Verizon, strive mightily to do that. They are the backbone of the Internet; they carry most of the world's traffic every day. Yet their unique position, that of owning the largest, most complex networks in the world, also makes screening this traffic no easy feat—for two reasons.

First, these global networks have hundreds of entrances and exits. BT Global Services, for example, operates in 170 different countries around the world, connecting to hundreds or thousands of large corporations and service providers in each one. Yet firewalls and other security technologies are designed to protect a single "link" or connection to the Internet—the point at which an organization's wide area network exchanges its data with the carrier. Second, firewall devices are designed to operate at the speeds of corporate networks, not backbone networks of the sort operated by AT&T, NTT, and so on. Corporate networks generally operate at speeds below 1 gigabit per second. Commercial firewall products designed for them simply cannot protect networks containing thousands of links that operate at core speeds 10 to hundreds of times that fast.

Using principles of entropy to protect a network begins with knowing a great deal about how traffic moves around that network, from hour to hour and minute to minute. Network security systems, including ours, operate inside the data center of a large Internet service provider or carrier. They run on standard off-the-shelf servers from, say, Dell or IBM, and collect data about traffic from a variety of key locations, called nodes, on the network. To collect these data, the carrier has to properly



configure its network routers, the servers that direct data traffic throughout the network. The routers must be configured to send “streams” of traffic statistics, a capability that is built into them. These data provide detail about traffic features such as the source and destination Internet Protocol (IP) addresses of packets in the traffic, the source and destination port numbers, the type of protocol, the number of bytes per packet, and the time elapsed between packets.

It is around traffic features such as these that our entropy algorithms first build a profile of the network’s normal behavior. This profile serves as a baseline in detecting anomalies. Our system also collects other data from the network’s routers that provide detail about how hard the routers themselves are working—data such as CPU and memory usage—and some additional detail about the volume of traffic on each of the router’s interfaces to the network. We then correlate all of these router statistics to verify anomalies detected by our algorithms, identify their root cause, and even suggest mitigating actions to cleanse the traffic.

Today, the big carriers collect some of the same data, but by and large they rely on “behavior-based” systems to protect their backbone networks. These systems are based on algorithms that focus primarily on changes in the volume of traffic at specific points on the network, ones where large companies and Internet service providers connect to the carrier. For example, during a denial-of-service attack, traffic between a single source (the attacker’s computer) and a single destination (the victim’s Internet servers) surges precipitously, reflecting an attempt to flood that destination and cut it off from users. In order to maximize the mayhem, attackers spread out their attacks by hijacking unprotected machines on the Internet and planting code that recruits them as “zombies” (or “bots,” short for robots). These computers in effect form armies (“botnets”)

that number in the tens of thousands and can be orchestrated to launch attacks that emanate from multiple sources. Known as distributed denial-of-service attacks, these actions concentrate the damage into a period lasting minutes or even seconds rather than hours.

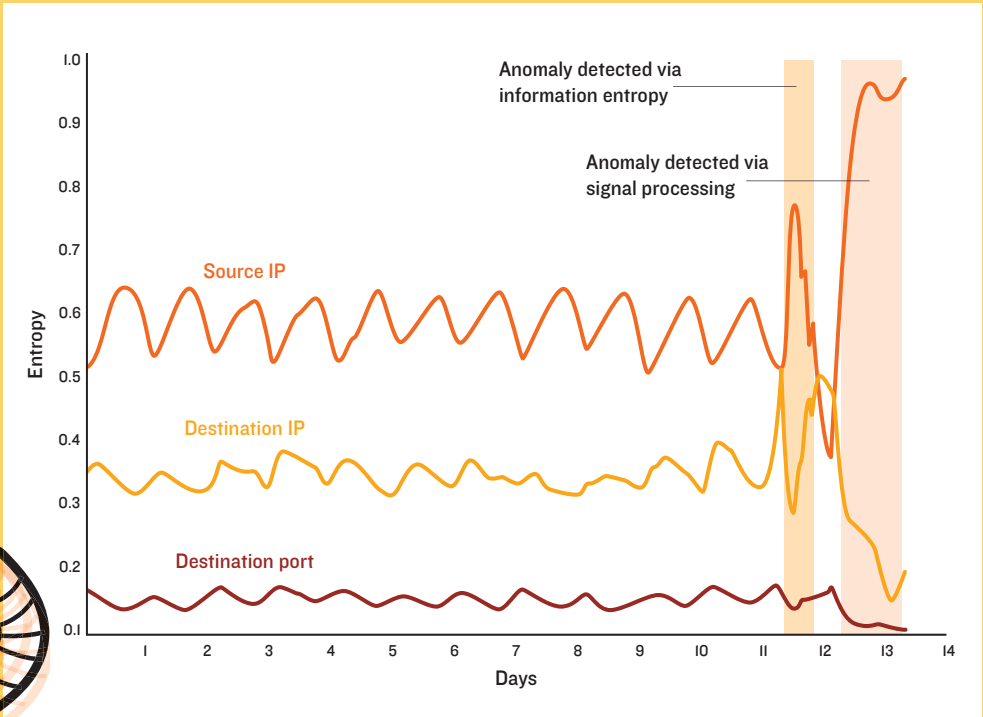

Traditional behavior-based systems detect such a sudden increase in traffic volume at the customer link and promptly alert the operator. Of course, this is just the start of an extended cat-and-mouse game. Attackers then devise new and clever methods to fly under the radar and avoid detection. Some intruders, for example, strive to consume the resources of a Web server located within the victim’s network. They don’t need to flood the server with traffic. Instead, they simply identify the Web pages that are the biggest drain on memory and CPU time—ones containing video clips, for example—and coordinate their armies to request frequent access to those pages only.

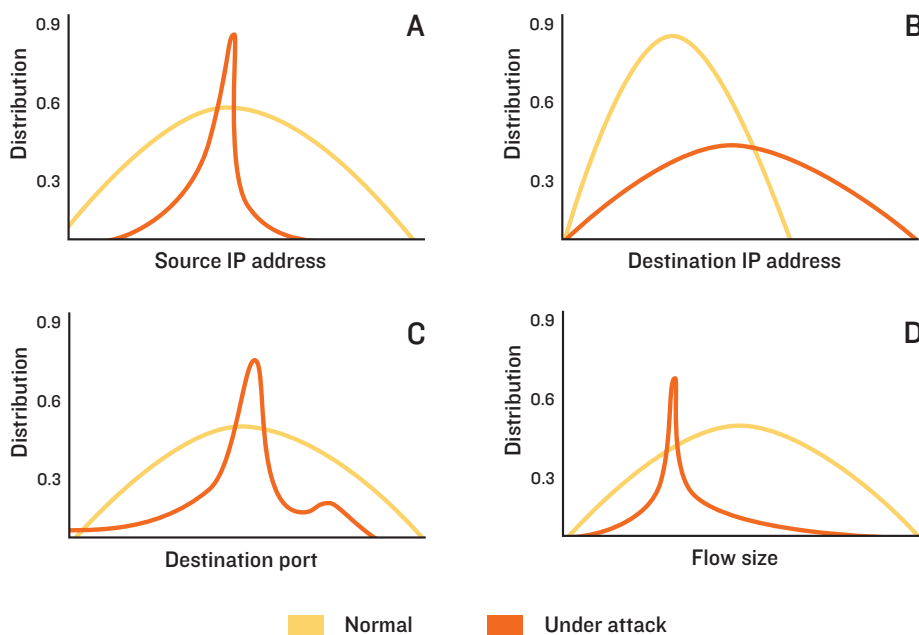
In such a case, the overall volume of traffic into the network looks normal, yet the attack is effective because it degrades or even chokes off service. Similarly, a flood of spam messages, for example, can overwhelm a mail server. It might seem like such a flood would unavoidably trigger a detection system. But that’s not always the case. The load on the server depends on the number of messages, not the quantity of data, which is what the detection system is measuring. If a spam attack is written so that each spam message consists of only a few data packets, then the overall traffic never rises to the threshold level. Internet telephony spam can similarly clog a network.

Many network operators have responded to these sorts of attacks by lowering the threshold of their behavior-based systems in an attempt to detect more subtle changes in traffic volume. This threshold change, however, tends to create false positives, in which the system often mistakenly takes nonmalicious fluctuations in the volume of traffic to be an

SASSER'S ENTROPY

An entropy-based security system would have quickly detected the 2004 Sasser worm outbreak, shown here as it targeted hosts in the network of a North American wireless service provider. As soon as a few infected machines try to spread the infection, their information entropy deviates from their norm, signaling a problem. The information entropy associated with destination IP addresses rises suddenly, indicating an increase in randomness in traffic destinations due to the scanning initiated by the infected machines, as it looks for new victims. Similarly, the entropy associated with the source IP addresses suddenly drops, indicating a decrease in randomness as the already infected computers initiate a higher than normal number of connections. The destination port exhibits a drop in its entropy, but much later than the other two features. Overall, the network is forced into a new internal state never observed before.





FINGERING THE CULPRIT: Examples of network fingerprints during well-behaved traffic (distinct traffic-feature distributions in yellow) and during a worm attack (distinct traffic-feature distributions in orange). Note the changes in shape of the distributions during malicious activity. The spikes in graphs A, C, and D show a change in network entropy, as does the flattening of the expected high curve in B.

attack. Such fluctuations are common; think of the flood of traffic that ensues when a Web page on a site with modest traffic is cited on a popular bulletin-board site such as Slashdot or Digg. The problem is that such false positives prevent operators from trusting the system, forcing slow and expensive case-by-case human intervention.

To avoid false positives, security software needs to monitor Internet traffic across the entire network, as opposed to a single link at a single time, and then correlate all the events it detects. Only then can a model of the traffic behavior on the entire network be created, allowing security algorithms to focus on the structure and composition of the traffic and not just its volume.

In other words, a security system must monitor the actual entropy of the network itself.

IN THERMODYNAMICS, entropy refers to changes in the status quo of a physical system—a cup of ice water, the gas in a balloon, a solar system. It is a measure of “molecular disorder.” In 1877, Ludwig Boltzmann visualized a probabilistic way to measure the entropy of an ensemble of gas molecules. Boltzmann showed that the ensemble’s entropy was proportional to the number of microscopic states such a gas could occupy. More precisely, entropy is a function of $k \log p$, where k is a constant and p is the probability of a given configuration of molecules.

What exactly is a configuration of molecules? Consider the temperature of air, which is determined by the average speed at which its molecules are moving. The temperature of a room might be 20 °C, but some molecules will be moving very quickly, for example in the sudden draft when a door opens or in the vicinity of a hot burner on a stove. Entropy reflects the amount of uncertainty about which exact molecules are moving at what speed.

For a given set of macroscopic quantities, such as temperature and volume, entropy measures the degree to which

the probability of the system—in this case, the air in the room—is spread out over different possible states. To take a much simpler example, if you roll a pair of dice, there are 11 different outcomes, some more likely than others. The complete array of possibilities and probabilities—only one way to get a 2, for example, but five chances of a 6 and six for a 7—is a probability distribution. Similarly, each gas molecule in that room has a number of different possible locations and speeds, just as the two dice each have six possible values.

For the entropy of the distribution of possible outcomes of a single die, each possible outcome has the same probability ($1/6$), so the distribution is flat. In this case there is nothing we can predict about the outcomes in the distribution. They are completely random, and the entropy of the distribution is very high—at its maximum, in fact. In the case of two dice, on the other hand, there are several possible combinations or outcomes that have a higher probability

than others. The probability of a 7 is much higher than that of an 11, for example. So if you roll two dice 25 times, the results will be less random than if you rolled one die 25 times. Another way of putting this is that the two-dice system has less entropy than the one-die system. We can guess more reliably about specific outcomes.

That is the principle behind our entropy algorithms. Malicious network anomalies are created by humans, so they must affect the natural “randomness” or entropy that normal traffic has when left to its own devices. Detecting these shifts in entropy in turn detects anomalous traffic.

Getting back to the gas example, the array of all possible locations and speeds creates a probability distribution for the gas. Because entropy theory is really designed to describe the configuration of a system based on a series of outcome probabilities, we can relate high or low entropy to the high or low probability of an outcome. So there’s a rough equivalence between thermodynamic entropy, understood as the probability that the molecules in a gas are in a predicted state, and the amount of information we have about a system.

Information entropy was originally conceived by Claude Shannon in 1948 to study the amount of information in a transmitted message. If the two states of a digital signal, 0 and 1, have exactly the same probability of appearing in the signal, then our uncertainty about which bit we will receive next is maximized—like throwing a single die that has only two sides. On the other hand, if the 1 has a higher probability of appearing, then there is slightly less uncertainty about what the next bit will be. That is, if the next bit has a greater chance of being a 1, entropy is reduced. When the information entropy is low, we are less ignorant of the details of the digital communication signal being transmitted.

Much the same can be said about traffic patterns on the Internet. More specifically, an enormous amount of information can be gleaned by observing traffic flows on a data network.

If we observe enough of them, we can come up with historical averages for inbound and outbound data packets, noting such key features as which Internet addresses the network receives packets from and which ones it sends packets to. We can also note how many packets are sent in accord with which Internet protocols at various times of the day and the overall traffic volume. At any given time, the probability distributions of the flow of traffic through the network will be characterized by distinct curves [see graphs, “Fingering the Culprit”]. In fact, the shape of the curve shows the entropy of the system. If the shape of the curve is uniform, then entropy is high. If there’s a spike, then a low-probability event has occurred, and the entropy is correspondingly low.

Internet traffic is dynamic and constantly evolving. Nevertheless, over the course of, say, a year, some consistent patterns emerge. These patterns are driven mainly by the mixture of applications generating the traffic, such as Web surfing, e-mail, music downloading, or Internet telephony, though seasonal and geographical factors also affect them. The first step in using these patterns to spot anomalous activity is to develop a probability distribution for each of the characteristics. When these distributions are taken together, they uniquely profile the traffic and create a “fingerprint” of the network under consideration and what we might call its internal state—the sum total of these network characteristics.

If we have monitored and measured a system long enough, we know which internal states are associated with well-behaved Internet traffic. Any malicious activity introduced into the network alters the nature of the Internet traffic, because it has a designed, premeditated outcome that is different from any of the network’s normal states. Even if an attack came in the form of an activity that fits within network norms—say, downloading a number of music files—the fingerprint of the network would look unusual, because it would differ in some way from the network’s established patterns of usage, if not in terms of volume, then time of day, source, or some combination of those or other characteristics.

Paradoxically, Internet traffic has features of both randomness and structure, and a worm, for example, will alter both, making the traffic appear in some respects more uniform or structured than normal Internet traffic, while appearing more random in others.

Packets flowing into a server seem to come from random locations. For example, requests for Web pages typically come from surfers all over the Internet. More will come from some people and networks than from others, to be sure, but a graph of them will normally be a fairly uniform curve. If a worm is loose on the Internet, however, and the packet flows from infected hosts grow to be a significant part of the set of total traffic flows, then the addresses of those hosts will show up disproportionately in any distribution graph—indicating how many flows have come from a given source.

During a worm infestation, hosts that have been maliciously co-opted connect to many other hosts in a short period. The number of open connections from infected hosts become dominant, and entropy decreases. Similarly, the target IP addresses seen in packet flows will be much more random than in normal

traffic. That is, the distribution of destination IP addresses will be more dispersed, resulting in higher network entropy.

Most malicious attacks tend to seek out and exploit certain vulnerabilities in the implementation of an Internet protocol. Two of the most important of these are the Hypertext Transfer Protocol, HTTP, which downloads Web pages, and the Simple Mail Transport Protocol (SMTP), for sending e-mail. Besides the protocol, specific operating-system ports are used to send and receive traffic. We can think of the protocol as a means of transit, such as an ocean freighter or a yacht, while the port (as the name suggests) terminates the data’s journey at the computer’s equivalent of a berth number at a marina.

Fingerprinting is also possible at the port level. An attacker can scan for a specific vulnerability by sending packets looking to see whether they are received and what response they get; these scans often have to go to a specific target destination port. If the traffic that results from this scanning becomes a significant component of the overall network traffic, then this will create an unusual fingerprint. Lastly, the flow size—the number of packets in the flow—of the malicious worm activity will become more dominant and will alter the distribution of flow size observed during a normal network operation.

The Sasser worm, one of the largest and best-studied infestations in Internet history, is an ideal example of this port-specific approach. It began by scanning the computers on whatever network it had infiltrated. Whenever a connection was made, the worm sent a piece of code. The goal of the code was to cause the infected host computer to accept commands on TCP port 9996. Sasser then created a small program named cmd.ftp on the host computer, which then executed it. The “ftp” in this script’s name stands for the File Transport Protocol. The FTP script instructed the victim machine to download and execute the worm from

the infected host without human intervention. The infected host accepted this FTP traffic on still another port. To spread itself even faster, Sasser spawned multiple threads, finding and capturing as many vulnerable computers within an organization’s network as possible.

Each of Sasser’s activities created a unique network fingerprint. Information entropy can capture the dynamics of such fingerprints by extracting any sudden change in the shape of the distributions constituting that fingerprint. There is little that the attacker can do to control the information entropy associated with the fingerprint and thereby conceal the attack.

The Sasser worm significantly affected the information entropy of a large North American wireless service provider network [see graph, “Sasser’s Entropy,” based on an analysis done after the attack]. Notice that traffic is much heavier during the day, as reflected by the information entropy: high during the day, low at night. When the Sasser worm invaded this wireless carrier’s network, the behavior-based security systems were unable to detect the outbreak until the network became saddled with more than 30 times its normal traffic volume. Behavior-based systems cannot detect the initial attack, because the traffic generated by one infected machine is negligible. Within minutes, however, that one machine has infected 10 others, and

**PARADOXICALLY,
INTERNET TRAFFIC
HAS FEATURES OF
BOTH RANDOMNESS
AND STRUCTURE,
AND A WORM WILL
ALTER BOTH**

those 10 infect 10 more, and so on, each generating its share of data. By the time the behavior-based system can generate an alert, the traffic is overwhelming.

Sasser quickly infected some 20 000 computers. Patches were soon created, but the worm was relaunched in multiple waves, spawning nine more variants in just over 30 days and infecting hundreds of thousands more machines.

The traditional defense against Sasser worked, eventually. But had the worm been detected earlier, for example by a system based on network entropy, Sasser would have been much more limited in its damage and probably would have spawned fewer variants.

IF TRADITIONAL DEFENSES struggle to keep up with traditional viruses, they fall far behind when it comes to new, more sophisticated forms of attack, such as the Storm worm. In broad outline, Storm shares some of the typical characteristics of an e-mail worm—users click on an attachment, which opens a file that places new code on the user's computer. The code then causes the computer to join an existing botnet, hooking itself up as a slave machine to a master computer out on the Net.

But Storm differs from earlier worms in a number of important ways. First, it does an excellent job of getting people to click on the attachment—it employs some clever social engineering by using subject lines and file names related to a hot topic in the news, such as a major storm or hurricane warning (this is where its name came from).

More significantly, Storm cleverly hides its network activities. While Sasser, for example, created a lot of new—and easily detected—traffic on TCP port 9996, Storm first looks to see what ports and protocols a user is already using. If Storm finds a file-sharing program—such as eDonkey, a popular program for trading music and videos—it uses that program's port and protocol to do its network scanning. The resulting minor increase in activity on that port would be missed by a conventional intrusion detector. Storm also looks to see what IP addresses the file-sharing program has already exchanged data with, instead of suddenly communicating with a whole bunch of new IP addresses, which would again be easily detected.

Finally, traditional worms spread as fast as they can, generating a fingerprint that is easily seen by a network-entropy security system. Storm, on the other hand, has a dormant mode and a waking mode. For example, every 10 minutes it will try to gather information. Then it will go quiet, and then start again.

Storm is now hibernating in millions and millions of computers in North America, Europe, and Asia. It is flying under the radar of current detection systems by tailoring its behavior to its victims' existing patterns of network usage. Its methods are changing in subtle ways over time as its creators stay one step ahead of their adversaries. Its botnet is poised to strike at major networks at any time.

How do we deal with this monster? We look, as we do with all worms, for changes in the entropy of the network. After all, Storm still has to alter certain things about a user's behavior in ways that can be detected. For example, during the 10-minute

periods that Storm is active, the victim's computer will send a lot of e-mail, much more than it normally does, typically about 30 messages per minute, or about 300 in one of its 10-minute active phases. Nobody sends out that much e-mail. And the e-mail goes out on a port that doesn't usually get e-mail traffic—in our example, the port that eDonkey uses. Normally, eDonkey traffic is very dense—bulky audio and video files—while e-mail is very low-volume data. These are all ways in which the network's entropy has changed.

A victim's communication to a master computer will also differ from previous usage patterns, maintaining a connection for days or even weeks of very low volume. Consider, too, that when two computers on the same network, say Nucci's and Bannerman's, are both victims of Storm, their behavior will suddenly be quite similar, whereas previously they were very different (perhaps Nucci downloaded lots of TV programs, while Bannerman did not).

In all of these ways, Storm is altering the network entropy of the victim's computer traffic. The increase in e-mail activity, for example, biases traffic in favor of port 25, the usual e-mail port—decreasing entropy, because port-25 activity becomes more predictable. Similarly, it becomes more predictable than previously that e-mail is sent rather than received. Our existing code, at Narus, will detect these changes in entropy, even if the classification of the changes is confusing—the victim

behaves somewhat like an e-mail spam generator, somewhat like a worm victim, and somewhat like a botnet member.

Traditional behavior-based Internet security servers cannot detect these attacks accurately enough and early enough to mitigate them before they achieve their goals. Today's carriers and other large network owners need a new approach to security that can correlate traffic data at extremely high speeds. Systems based on information entropy can do that, and the security of these most critical networks depends on it. ■

ABOUT THE AUTHORS

ANTONIO NUCCI is the chief technology officer, and STEVE BANNERMAN is vice president of product management, of Narus, in Mountain View, Calif. The company produces software for analyzing data traffic and protecting networks. Both authors are IEEE members.

TO PROBE FURTHER

The BBC had a good account of the Sasser worm's destructive assault on the Internet; see <http://news.bbc.co.uk/1/hi/technology/3682537.stm>.

See "Hackers Attack via Chinese Websites" (*Washington Post*, 25 Aug 2005, http://www.washingtonpost.com/wp-dyn/content/article/2005/08/24/AR2005082402318_pf.html) regarding the 2005 attack on the U.S. Department of Defense.

For a general discussion of information entropy, see Charles Seife's 2006 book, *Decoding the Universe*, published by Viking Penguin; in particular, check out pages 46, 47, and 71. As well, the Wikipedia entry for Entropy is helpful, see <http://en.wikipedia.org/wiki/Entropy>.

TRADITIONAL DEFENSES FALL FAR BEHIND WHEN IT COMES TO NEW, MORE SOPHISTICATED FORMS OF ATTACK, SUCH AS STORM

Why We Joined...

Sameet Shriyan

IEEE Graduate Student Member

"IEEE Membership has instilled in me a sense of professionalism and leadership, and provides a platform from which I can network with fellow researchers and industry professionals."



Why We Stay...

Fritz Morgan


IEEE Member

"The half life of cutting-edge technology is only a few years and rapidly getting shorter. You need a way to ensure that your skills remain current and relevant. IEEE has always been that place for me."



IEEE Membership: Connecting Professionals, Advancing Technology
For more information, visit: www.ieee.org/benefits





PRACTICAL THEORIST: On 1 December 1957, physicist Jean Hoerni conceived the planar process, a technique used to manufacture essentially all silicon transistors and microchips today.

HOERNI: WAYNE MILLER/MAGNUM PHOTOS; PHOTO-ILLUSTRATION: BRANDON PALACIO

The Silicon Dioxide Solution

How physicist Jean Hoerni built the bridge from the transistor to the integrated circuit **BY MICHAEL RIORDAN**

Not plastic bags, nor metal screws, nor cigarette butts. No, the commonest human artifact today is the transistor—invented 60 years ago this month by Bell Labs physicists John Bardeen and Walter Brattain. Millions of these subminiature switches populate computers, cellphones, toys, domestic appliances, and anything else that carries a microchip. Exactly how many transistors are around is hard to know, but several years ago Gordon Moore, a founder of Intel Corp. and author of the famed Moore's Law, made an educated guess: more than 10^{18} —that's one *quintillion*—transistors are produced annually. "We make more transistors per year than the number of printed characters in all the newspapers, magazines, books, photocopies, and computer printouts," Moore told me recently. "And we sell these transistors for less than the cost of a character in the Sunday *New York Times*."

Behind the explosive growth that transistor production has seen since 1960 is a major technological achievement. Today, chipmakers essentially *print* transistors on silicon wafers. It's a manufacturing method rooted in the mechanical printing process originated by Johannes Gutenberg more than 500 years ago—though far more complex, of course. Moore himself played a lead role in developing transistor-fabrication technology during the 1960s when he was research director at Fairchild Semiconductor Corp., in Palo Alto, Calif. But much of the credit for that revolutionary advance belongs to a lesser-known semiconductor pioneer and Fairchild cofounder. The unsung hero of this pivotal chapter in the history of electronics—the invention of the planar transistor—is Jean Hoerni.

A Swiss-born theoretical physicist, Hoerni, along with seven other determined, like-minded rebels—Moore, Robert Noyce, Jay Last, Sheldon Roberts, Eugene Kleiner, Julius Blank, and Victor Grinich—founded Fairchild in 1957 [see photo, "The Fairchild Eight"]. They all contributed, directly or indirectly, to the new

technology, but none so much as Hoerni [see photo, "Practical Theorist"]. Fifty years ago, sitting alone in his office, he elaborated a radically new kind of transistor: a more compact, flatter device whose sensitive parts were protected beneath a thin layer of silicon dioxide. Hoerni's brilliant idea, more than any other single factor, allowed the fledgling firm to begin printing transistors on silicon. Planar transistors would prove to be much more reliable and perform far better than other designs, in effect rendering the competition's offerings obsolete.

The planar process also made it easy to interconnect neighboring transistors on a wafer, paving the way to another Fairchild achievement: the first commercial integrated circuits. As other companies realized the great advantages of planar technology and began adopting it on their own production lines, Hoerni's elegant idea helped to establish Silicon Valley as the microelectronics epicenter of the world.

THE FINAL MONTHS OF 1957 were a time of anticipation at Fairchild as the founders organized the new firm's labs and production lines in a group of buildings at 844 Charleston Road in Palo Alto. In September of that year, the eight scientists and engineers had resigned en masse from Shockley Semiconductor Laboratory, in Mountain View, about 2 kilometers away. They were rankled by the heavy-handed management style of its founder, transistor pioneer William Shockley, and his pursuit of difficult R&D projects at the expense of useful, salable products. So they persuaded the Fairchild Camera and Instrument Corp. of Syosset, N.Y., a firm looking to diversify its business, to found Fairchild Semiconductor. The eight founders planned to use the silicon processing techniques they'd learned under Shockley to make and sell advanced, high-speed transistors.



THE FAIRCHILD EIGHT: From left, Gordon Moore, Sheldon Roberts, Eugene Kleiner, Robert Noyce, Victor Grinich, Julius Blank, Jean Hoerni, and Jay Last.

Their timing could not have been better. On 4 October 1957, the Soviet Union launched *Sputnik I* into orbit, igniting a frenzied space race with the United States. Millions worldwide gazed skyward to watch the awesome, undeniable evidence that the Soviets had a big head start. Meanwhile, Senator Lyndon B. Johnson (D-Texas) spearheaded congressional investigations into how the Eisenhower administration could ever have permitted such a “missile gap” to arise. With the USSR holding a major advantage in the greater thrust of its missiles, the U.S. aerospace industry sought every imaginable way to reduce the size and weight of its payloads and satellites. “There was a great deal of talk about the packing density of electronic functions in the late 1950s,” Noyce recalled in a 1975 interview, which is archived in the IEEE History Center. “It was the Missile Age, and transportation costs from here to Russia were very high.” The need for small, ultralight electronic circuits based on reliable silicon transistors made these devices a promising market for Fairchild.

That fall, the Fairchild founders worked feverishly to get everything up and running. Moore set up diffusion furnaces designed to impregnate silicon wafers with micrometers-thin layers of impurities—chemical elements such as boron, phosphorus, or aluminum that alter silicon’s electrical characteristics to form a transistor’s building blocks. Metallurgist Sheldon Roberts took on the task of growing high-purity silicon crystals from which the wafers could be sliced. Noyce and Last developed

methods to do photolithography and oxide masking, by which they could define precise openings in a thin silicon-dioxide layer on the wafer surface; the impurities would diffuse through these openings into the underlying silicon. Other cofounders dug into manufacturing, testing, and selling the high-tech devices to aerospace customers.

And then there was Hoerni. A theorist with not one but two doctorates, from the Universities of Cambridge and Geneva, he had come to the United States to pursue postdoctoral studies at Caltech. In 1956, Shockley lured the 32-year-old physicist away from academia and assigned him to do theoretical calculations of diffusion rates. At first, Hoerni was cloistered in a separate office, but he kept coming around and snooping in the lab in the main building—which gave him valuable insights into solid-state diffusion. Later, at Fairchild, while the others worked on building or installing equipment, he mostly sat in his office and “scribbled in his notebook,” Moore told me.

On 1 December 1957, Hoerni grabbed his crisp new lab notebook and began writing an entry titled “Method of protecting exposed *p-n* junctions at the surface of silicon transistors by oxide masking techniques.” In a loose, fluid scrawl interspersed with three simple drawings, he described a revolutionary new way to fabricate transistors—unlike anything ever before attempted.

The most advanced silicon transistors at that time were called mesa transistors because they resembled the plateaus

of the American Southwest, the impurity layers running laterally like the colorful rock strata [see illustration, “Mesa vs. Planar”]. These transistors basically consisted of three impurity layers piled up vertically, each rich in either electrons (*n*-type) or electron deficiencies, better known as holes (*p*-type). The main drawback of the mesa structure is that its *p-n* junctions, the interfaces between layers where the transistor’s electrical activity occurs, are exposed at the edges. Bits of dust or drops of moisture can contaminate the sensitive interfaces and disrupt their normal electrical behavior.

Hoerni’s idea was to protect the *p-n* junctions by keeping the oxide layer in place upon the silicon after the diffusion process; the standard practice at the time was to etch that layer away, baring the junctions. “The oxide layer so obtained is an integrant [sic] part of the device,” he wrote in his notebook that December day, “and will protect the otherwise exposed junctions from contamination and possible electrical leakage due to subsequent handling, cleaning, and canning of the device.”

It was a brilliant conception but too far ahead of its time. Hoerni’s approach would require additional fabrication steps, and making mesa transistors was already at the limits of the possible. Bell Labs and Western Electric had produced prototypes of mesas, but no company had sold one on the open market.

In early 1958, Fairchild secured its first purchase order for silicon transistors from IBM’s Federal Systems Division, which planned to use them in the onboard computer it was designing for the B-70 bomber. Fairchild, which didn’t even have prototypes, faced the formidable challenge of delivering real working devices. To maximize the chances of success, the cofounders decided to develop two different kinds of mesa transistors. A group under Moore pursued the *n-p-n* transistors, which were thought to be easier to fabricate, while Hoerni formed another group to delve into the *p-n-p* versions.

Crucial to both efforts was the work Last and Noyce were doing on the optical methods needed to transfer the patterns defining a transistor’s features onto the silicon wafer. On a trip to San Francisco, they purchased three 16-millimeter lenses from a camera store and used them to fashion a step-and-repeat camera, a contraption that produced rectangular arrays of tiny, identical images on photographic plates, called masks. Workers shone light through the masks onto a special photosensitive resin that had been deposited on the wafer’s oxide surface layer. When they subsequently rinsed the wafer in a powerful acid, it etched the illuminated areas away, exposing the silicon beneath them. Thin layers of impurities were then diffused into the silicon through the resulting openings. Using such techniques, Fairchild could batch-process hundreds of identical transistors on a single wafer.

Another breakthrough was the use of a single metal to make the electrical connections to both *n*-type and *p*-type silicon, an approach that greatly simplified the manufacturing process. Moore had been struggling with this issue, trying many different metals, when Noyce happened by his lab early one day and suggested aluminum. As a *p*-type impurity, aluminum easily bonds to *p*-type silicon but often sets up a current-blocking *p-n* junction when it is deposited on *n*-type silicon. Moore found a way around this problem by starting with *n*-type silicon that had more impurities than usual. Moore’s group got its *n-p-n* transistors into production in May 1958, well ahead of Hoerni’s team, which had opted to use silver for electrical contacts.

To protect the mesa’s sensitive junctions, each transistor was packaged into a pea-size hermetically sealed metal can and then tested. Fairchild shipped the first hundred of them to IBM on

schedule that July, billed at US \$150 apiece. The next month, at the WESCON electronics trade show, the founders discovered to their delight that they were the only ones with silicon mesa transistors on the market. “We scooped the industry!” Noyce said, exulting at a Fairchild meeting a few days later.

ABOUT THE ONLY PERSON at Fairchild not celebrating was Hoerni. A proud, charming, but irascible and often volatile man, the scion of a Swiss banking family, he was miffed that his *p-n-p* approach had been passed over. But he was also a hardheaded contrarian whose creative fires were stoked by adversity. Hoerni not only didn’t give up, he set out to develop an even better transistor. Later that year, he returned to the ideas written down in the opening pages of his notebook. Could the oxide layer in fact be used to protect the sensitive *p-n* junctions? There were indications it might. That spring, reports had come in from Bell Labs that the oxide layer indeed protected the silicon underneath. Why not the junctions, too?


With a doctorate in crystal physics, Hoerni realized that the impurity atoms coming through the tiny openings in the oxide layer would diffuse sideways nearly as well as downward into silicon’s crystal structure. Which meant that the junction interfaces would curl up *under* the oxide layer surrounding an opening, just micrometers farther out from its edges. If left in place instead of being etched away, he figured, the oxide layer could protect those junctions.

But the device Hoerni envisioned would not only be more difficult to fabricate, its structure flew in the face of conventional wisdom. Especially at Bell Labs and Western Electric, the oxide layer was considered “dirty”—filled with impurities after the diffusion process—and thus had to be removed.

Meanwhile, serious concerns began to emerge in late 1958 and early 1959 about the mesa transistors Fairchild was selling. Some of the devices were experiencing amplification instabilities, and others were malfunctioning. One important customer reported that a transistor had suddenly stopped working altogether. A Fairchild technician eventually traced the failures to tiny dust particles and solder fragments trapped inside the cans. The specks were attracted to the junctions by the strong electric fields there. In a subsequent quality-control procedure that became known as the tap test, workers would tap on the cans with pencil erasers, trying to dislodge any bits that might short out the junctions. If that happened, the transistor was discarded. Those were anxious days for the brash young firm, for such failures in its only product threatened its very existence.

Hoerni’s single-minded pursuit of a more reliable transistor proved timely indeed. In what Moore described to me as a “kludge experiment” intended to assess Hoerni’s ideas, a technician deliberately left the oxide layer on top of one of the *p-n* junctions in a mesa transistor. When tested, it had substantially better amplification stability—suggesting that Hoerni was truly onto something. On 14 January 1959, he had two of his notebook pages typed up as a formal disclosure and sent to John Ralls, Fairchild’s patent attorney. Other than a few minor corrections and better drawings, it was identical to the notebook entry he had written more than a year earlier.

One problem with Hoerni’s approach—and part of the reason nobody attempted it at first—was that his transistor structure was more complex than the mesa’s, requiring a fourth photolithographic mask to fabricate it. Last and Noyce’s step-and-repeat camera could accommodate only three masks. But that February, Last “jury-rigged a fourth mask” for this purpose, he



recalled in a recent telephone interview. On 2 March, Hoerni wrote another entry in his notebook titled “A method of manufacture of PNP transistors with oxide protected junctions.” In two more pages of text and drawings, he indicated specifically how to fabricate such a device, though still stubbornly using silver for the electrical contacts on the top side. By then, his technicians were already transforming his novel ideas into actual fabrication processes.

But all that progress came at a time of upheaval at Fairchild. The same week that Hoerni was jotting down his fabrication ideas, Edward Baldwin, who had been hired from Hughes Electronics Corp. to serve as Fairchild’s general manager, departed abruptly to found Rheem Semiconductor in Mountain View, taking with him five key people from the manufacturing division. After persistent urging by the other Fairchild cofounders, Noyce stepped up to replace him, and Moore took over Noyce’s position as research director.

The following week, Hoerni invited several colleagues to watch a demonstration of his new prototype transistor. Under a microscope it appeared unlike any other Fairchild device. Less than a millimeter across, it was completely flat—no mesa protruded in the middle. All that was visible was a circular metallic dot with a metal ring around it, plus the oxide surface layer between them. It resembled a bull’s-eye target with a portion of it pulled out like a teardrop, making it easier to attach a wire [see photos, “Silicon Flatland”].

What happened next is unclear. Some observers have claimed that Hoerni suddenly spat on his transistor, to demonstrate that such outrageous abuse had no ill effects on the oxide-protected junctions. But Last and Moore don’t recall him actually spitting, and Moore points out that saliva would have shorted out the metal wires on the device. Even so, the demonstration was dramatic and convincing. Last told me, “Gee, it’s too bad Baldwin had to leave last week,” he recalls joking afterward.

Things moved swiftly after that. It was obvious that Hoerni’s creation was far more rugged and reliable than the mesas. And it also proved to have much lower leakage currents—small, wrong-way trickles that can seriously degrade transistor performance. In a Fairchild report released the following year, Hoerni

observed that the leakage currents in his device were usually less than a nanoampere, or as little as 1 percent of those in mesa transistors.



The nagging question on everyone’s mind was, Can we manufacture these transistors in quantity? Initially, the planar process yielded only a few working transistors in every 100—much worse than the mesa process. But as various problems, such as pinholes in the oxide layer, were resolved, yields rose and doubts evaporated. In April 1960, Fairchild sold its first planar transistor, the 2N1613—a metal cylinder about half a centimeter in diameter and almost as high, with three little metal legs sticking out beneath it.

A few months later, Noyce and Moore decreed that henceforth all the company’s transistors would be planar. While other semiconductor firms such as Rheem, Motorola, and Texas Instruments had begun churning out competitively priced mesa transistors, Fairchild boldly struck out in a promising new direction. Soon avionics manufacturers began to demand planar transistors because of their unmatched reliability. The Autonetics division of North American Aviation, for example, insisted on using Fairchild’s planar transistors in the guidance and control systems for the Minuteman missile.

Fairchild eventually licensed the planar process to other transistor makers—even Bell Labs and Western Electric. Either the other firms followed Fairchild’s lead or they exited the industry.

WELL BEFORE FAIRCHILD succeeded in commercializing Hoerni’s device, Noyce had begun thinking about what *else* the company could do with the planar approach. In his 1975 interview, he credited patent attorney Ralls with challenging the Fairchild team to consider other applications that could arise from the new way of making transistors. Noyce realized that by leaving the oxide layer in place, “the surface of the silicon then had one of the best insulators known to man covering it.” Which meant that the electrical connections could be made by depositing strips of metal—such as the aluminum contacts Moore’s group had perfected—on top of the oxide layer. The strips would be automatically insulated from the components underneath.

LEFT: FAIRCHILD SEMICONDUCTOR. RIGHT: STANFORD UNIVERSITY ARCHIVES



On 23 January 1959, not long after Hoerni had his patent disclosure typed up, Noyce penned an entry in his own notebook: “In many applications now it would be desirable to make multiple devices on a single piece of silicon in order to be able to make interconnections between devices as part of the manufacturing process, and thus reduce size, weight, etc., as well as cost per active element.” His entry went on for another four pages and included the crucial idea of using the oxide layer as an insulator underneath the connections. He also described a way to isolate the circuit elements—not just transistors but also resistors, capacitors, and diodes—from one another by inserting between them extra *p-n* junctions, which permit current flow in only one direction.

Did Noyce recognize the significance of these ideas at first? In those days, researchers at Bell Labs, Fairchild, and elsewhere often had a colleague immediately witness and sign important, potentially patentable ideas. Noyce, for instance, had witnessed Hoerni’s entry back in December 1957. Curiously, however, nobody witnessed Noyce’s entry, suggesting that he did not consider it all that important when he wrote it.

Around that time, the “monolithic idea” of fabricating complete, rugged electronic circuits in a single chunk of silicon, germanium, or other semiconductor was becoming fashionable. The U.S. Army, Navy, and Air Force were each promoting their own pet approaches and funding R&D contracts in industry. Monolithic integration was considered a way to overcome the “tyranny of numbers” bemoaned by Bell Labs Vice President Jack Morton. He had warned that as the number of circuit components increased, so did the likelihood of circuit failure [see “How Bell Labs Missed the Microchip,” *IEEE Spectrum*, December 2006]. But what if you fabricated reliable components and interconnected them in a single semiconductor chip? Then your odds of building successful complex circuits might be much higher.

In August 1958, Jack Kilby at Texas Instruments had conceived a way to make such integrated circuits in silicon. He even built a prototype oscillator based on the idea, using germanium mesa transistors, which were then readily available at TI. But while Noyce’s subsequent approach involved metal strips deposited on an oxide layer, Kilby’s device used “flying wires” to make the electrical connections. TI publicly announced this

SILICON FLATLAND: Above from left, an early prototype planar transistor made by Fairchild in the spring of 1959; a cutaway model of the company’s first commercial planar transistor, the 2N1613, initially marketed in April 1960; one of the first integrated circuits made by Jay Last’s development team in the spring of 1960; and a prototype planar flip-flop circuit fabricated in the fall of 1960.



breakthrough on 6 March 1959 at a gathering of the Institute of Radio Engineers (a predecessor of the IEEE) in New York City. TI President Mark Shepherd boasted that it was “the most significant development by Texas Instruments since we divulged the commercial availability of the silicon transistor.”

News of TI’s achievement reached Fairchild just as its management turmoil was winding down and Hoerni was about to demonstrate his new transistor. Later that month, Noyce called a meeting to discuss how to respond to TI and revealed his thoughts about how to interconnect multiple devices in silicon. By then it was becoming obvious that Hoerni’s planar process offered major advantages in fashioning such integrated circuits. Hoerni, Last, Moore, and the other cofounders discussed that possibility extensively, with the emphasis on the pragmatic. “Any one of us could think of ten things we might do, but then we’d rule out nine or even ten of them as impractical,” Last said in a recent phone conversation. “We were focused on making things that worked.”

Out of this creative stew emerged another crucial concept, which historians have so far overlooked. With the planar transistor, it was now easy to put all three electrical contacts—to the emitter, base, and collector—on *one side* of the silicon wafer. At first glance, it might seem just a marginal improvement, but this feature, plus the fact that a single metal such as aluminum could be used to form the connections, meant that Fairchild could now, in effect, print electrical circuits—transistors and all—on silicon. Like the typographic patterns of ink impressed onto paper by a printing press, the patterns of the individual semiconductor devices and metal interconnections could now be imposed photolithographically on a single side of a wafer.

Hoerni was the first to publish the concept of putting all the electrical contacts on one side. In his patent application for a “Method of Manufacturing Semiconductor Devices,” filed on 1 May 1959, he presented the idea almost as an aside, after revealing a structure

MESA VS. PLANAR: Side views of a mesa [left] and a planar transistor, from a report Hoerni prepared in 1960.



closer to that of the mesa, with contacts on both sides of the wafer. In Noyce's much more famous patent, "Semiconductor Device and Lead Structure," filed three months later, the single-side feature is a fundamental aspect of his planar integrated-circuit structure. But neither man's lab notebook mentions the idea—suggesting that it probably emerged from the fertile give-and-take discussions that spring and was later added to the patent applications.

In any event, that special feature of the planar process gave Fairchild a tremendous advantage in realizing the monolithic idea.

TO IMPLEMENT THIS NEW TECHNOLOGY, Last formed a group in the fall of 1959, aiming to manufacture integrated circuits based on Hoerni's planar process. It took another 18 months before the first commercial microchips, Fairchild's Micrologic series, reached the market. But Fairchild still came out with its microchip more than six months ahead of TI, which succeeded only after it began using the planar technology it had licensed from Fairchild.

To achieve this goal, Last's team had to overcome several significant obstacles. Tolerances were a lot tighter on positioning the physical features of these chips—which meant the various masks had to be aligned more precisely. Finding a way to isolate their components electrically was also a thorny problem. Noyce's idea of inserting back-to-back p - n junctions between individual components proved an effective solution, opening the door to commercialization in March 1961.

But Hoerni and Last were not around to share in the celebrations. They had become disenchanted with the increasingly stratified Fairchild hierarchy and the worsening relations with its New York parent. They also felt that Fairchild's marketing department opposed microchips because they'd compete directly with the company's principal products—transistors and diodes. So Hoerni and Last departed to start yet another semiconductor operation, the Amelco division of Teledyne, with the goal of producing integrated circuits.

Close friends since their days at Shockley Lab, the two often spent their weekends together hiking in the deserts and mountains of the Southwest. Last remembers that Hoerni had incredible stamina and could hike for hours on little food or water. To lighten his load, he carried only a skimpy old sleeping bag. When temperatures got too cold, he'd stuff it with newspapers—once claiming that *The Wall Street Journal* provided the most extra warmth.

In two years, however, Hoerni began to have problems with the new company. In the midst of a cash crunch in April 1963, Teledyne executives suggested that he be reassigned from general manager of Amelco to director of research as a cost-cutting measure. The moody Swiss physicist did not warm to the idea. Instead, he decided to leave the firm and began casting around for other business alternatives.

Although their relationship was "rather frosty" after Hoerni's decision, Last says, they still headed out that spring for a 3000-meter climb in the Inyo Mountains east of the Sierra Nevadas. Exhausted, they reached the summit at dusk, just before a cold front pushed through and temperatures plummeted. Despite their differences, the two huddled together the rest of the night to keep from freezing. "We climbed down the next morning, drove back to the Bay Area, and continued our frosty business discussions," Last recalled years later, during a memorial service for his friend.

While Last remained with Teledyne until the late 1970s, Hoerni went to work for Union Carbide, setting up its semiconductor division. In 1967 he ventured out in yet another direction, founding Intersil Corp., with European investors, to make microchips for digital watches; it was the first company to produce such low-voltage, low-power circuits based on CMOS (complementary metal-oxide-semiconductor) technology. The following year Moore and Noyce abandoned Fairchild to launch Intel in Santa Clara, Calif., at the heart of what soon became known as Silicon Valley.

For the next three decades, Hoerni remained active as an investor and consultant in the semiconductor industry. He also became involved in philanthropic initiatives and continued trekking throughout the world. He died in Seattle on 12 January 1997, the year the transistor turned 50. Although often overlooked in semiconductor history, he should be remembered as the person who engineered the all-important bridge from this revolutionary solid-state device to the integrated circuit, which has become so ubiquitous today.

ABOUT THE AUTHOR

Contributing Editor **MICHAEL RIORDAN** teaches the history of physics and technology at Stanford University and the University of California, Santa Cruz.

TO PROBE FURTHER

For excellent accounts of the early Fairchild work and its wider implications, see Gordon Moore's "The Role of Fairchild in Silicon Technology in the Early Days of 'Silicon Valley'" and Jay Last's "Two Communications Revolutions," both in *Proceedings of the IEEE*, Vol. 86, No. 1 (January 1998).

Two recent books that go into great detail about planar technology and the origins of the silicon integrated circuit are Christophe Lécuyer's *Making Silicon Valley: Innovation and the Growth of High Tech, 1930–1970* (MIT Press, 2006) and Leslie Berlin's *The Man Behind the Microchip: Robert Noyce and the Invention of Silicon Valley* (Oxford University Press, 2005). A review of both books appeared in the April 2006 issue of *IEEE Spectrum*.

The Computer History Museum explores semiconductor history at <http://www.computerhistory.org/semiconductor>.

RESOURCES

The Player

An engineer reinvents herself as a video-game developer

BY DAVID KUSHNER



FROM HOBBY TO JOB:
Tobi Saulnier test-markets
her new game.

www.spectrum.ieee.org

December 2007 | IEEE Spectrum | INT 51

RESOURCES

These are good times to be an electronic-game developer. With sales hitting US \$7 billion per year, the industry is becoming a serious part of what is commonly called “the media.” Steven Spielberg, Peter Jackson, and other high-profile filmmakers are pursuing game ventures.

CAREERS

Movies based on game franchises top the charts. And there are more platforms for games than ever before: computers, handhelds, and cellphones. But breaking into the industry is a metagame unto itself, and for every thumb jockey who finds a career making games, there are thousands of cubedwelling Dilberts who dream of a break that never comes.

They have much to learn from Tobi Saulnier, an IEEE member who loved video games from the start but figured she'd better chase a seemingly more practical career. She studied electrical engineering at Rensselaer Polytechnic Institute, in Troy, N.Y., and then worked for a decade at General Electric, where she ended up managing research and development in embedded and distributed systems. But through a bit of social engineering and serendipity, Saulnier was able to leave her safety net behind for the path less traveled—which took her to video-game development. She's the founder and chief executive of 1st Playable Productions, a video-game developer in Troy. As one of the few engineers, let alone female engineers, to make the leap, she is living proof that creating video games needn't be a pipe dream.

“I'm very attracted to anything that has a problem that requires creativity to solve,” Saulnier says. “Sometimes we have problems that are technical. Sometimes we have problems that are aesthetic. Engineering gives you a solid base for this kind of problem solving.”

VIDEO GAMES were hardly on Saulnier's mind in 1984 when she took her first job, as a contract programmer at GE's research center. Of the 1500 people on staff, she discovered, more than 100 had doctorates. It didn't take long for her to realize that her studies had really just begun. “I was inspired to go get graduate degrees because I felt I had learned nothing,” she says. “I was surrounded by extremely smart and educated people.”

The birth of her daughter stretched her Ph.D. program at Rensselaer to 10 years, ending in 1995. She then led a small group at GE that explored both the tools of network simulation and the application of the technique to satellite and utility systems. The company has used this technique in products as varied as medical scanners and locomotives.

She relished the challenge of fixing machines, but the grind of working as a cog in one wore her down. “When you're at a big company,” she says, “you have to implement corporate policies—some of which you may not agree with.”

When Saulnier observed some of her colleagues leaving and joining start-ups, she felt the itch to join in. “They were having all these adventures at small companies,” she recalls. A former colleague suggested she interview at Vicarious Visions, a video-game developer in Troy (which has since moved to Albany, N.Y.), and she quickly found that the video-game business had plenty of room for engineers like her.



CHAINED TO HER DESK?
No, Tobi Saulnier just loves her work.

“There's a near endless demand for engineers and game programmers,” says Jason Della Rocca, executive director of the International Game Developers Association (IGDA), in Mount Royal, N.J. “Nearly every development studio has job openings for engineers.”

Saulnier, an avid violin player, had the kind of creative background that Vicarious Visions preferred, as well as greatly needed patent expertise. In 2000 she signed on as the company's vice president of product development, organizing multidisciplinary teams of artists, coders, and designers to work on a variety of gaming titles, including children's games (*Blue's Clues*) and first-person shooters (*Doom 3*).

She refined her engineer's perspective on the science of video games. “A video game is an embedded system, but with art,” she says, “plus it has the hard specification that it has to be fun.”

The embedded nature of the hardware, she discovered, provided delectable challenges. New video-game consoles and handheld systems come to market only every four years or so, and that means game developers learn to do the best with the

hardware they have. "As a game developer, you can't change the hardware," Saulnier says. "You must find out what it's good at and maximize it, very much like working with a satellite system you shoot into space, or a utility system, where approval cycles are long. You have to work with those constraints."

GETTING PRODUCTS to market in good shape can also be a challenge, due to the complications surrounding software updates and patches. Although networking makes it possible to patch a console game, there's always a potential for disaster. "You don't want to make it that easy to patch a system, because an ill-intentioned person can take the system down," Saulnier says. "Video-game quality standards are high and extremely tightly managed because the stakes are high."

In April 2005, after Activision, a video-game publisher in Santa Monica, Calif., bought Vicarious Visions, Saulnier

launched her own firm, 1st Playable, which focuses on kids' games—a natural for her, as she was by then a mother of two. Titles include a spin-off of the Cabbage Patch franchise and a mind bender called *Puzzle Quest*. She keeps up the energy level by having her staff of 35 share an open space.

She signed on as a member of the IGDA board with a particular interest in encouraging women to join the field. Her choice of position with IGDA reveals her professional pedigree: "I'm treasurer because I made the mistake of asking about the numbers," she jokes. "But for me it's second nature. I'm an engineer. Numbers are interesting." ■

ABOUT THE AUTHOR

Contributing Editor DAVID KUSHNER blogs for Spectrum Online at <http://blogs.spectrum.ieee.org/gizmos>. In this issue, he also wrote "Playing Dirty," about how one man turned virtual dross into gold.

Build Yourself An Electric Gun

Why? Because you just plug it in, aim...and fire

BY PAUL WALLICH

TOOLS & TOYS

A gun that uses moving electrons instead of messy chemicals to throw a slug has been a staple of speculative fiction since the days of Edison and Tesla—and not only of fiction. Electrically operated projectile launchers—variously known as Gauss rifles, railguns, and mass drivers—have both fascinated and frustrated military researchers the world over [see "For Love of a Gun," *IEEE Spectrum*, July].

So of course I jumped at the chance to build one.

After skimming the Web for sources, I settled instead on a design from the optimistically titled book *Mechatronics for the Evil Genius*, by Newton C. Braga (McGraw-Hill, 2006), and nipped out to my local electronics shop for some parts.

These shops aren't what they used to be. There was a drawer for silicon-controlled rectifiers (SCRs)—I'd need one for a fast, high-current switch—but it was empty, and the clerk said there were no plans to restock the item. On top of that, the biggest capacitor on the rack offered a piddling 4700 microfarads. I was lucky to get a transformer, a few packets of resistors within shouting distance of the values I needed, and some wire that just might be suitable for winding a coil.

Even online the pickings were slim: most vendors cater to buyers willing to place bulk orders with plenty of lead time, not writers

on deadline. I finally found an outlet that had 22 000- μ F capacitors and the SCR I needed and promised to deliver them fast. Then I got the soldering iron and heavy-gauge wire out of the basement and went back to wiring the rest of the circuit and winding my solenoid.

When the capacitors and the SCR arrived, I was eager to get everything hooked together. On the incoming side of the circuit, I had a 12-volt transformer (to make sure I didn't kill myself), a fairly hefty diode to transform ac into pulsed dc, and a 10-watt, 50-ohm brick of a resistor to limit the charging current for the capacitor so the wires wouldn't melt. On the output side, I had my coil—160 turns wound around a transparent plastic tube, chosen so I could see the projectile move—and the SCR with a push-button switch controlling voltage to the gate. (You can also substitute a photoresistor for the gate switch, to trigger the SCR automatically.)

I plugged in the transformer, threw the

switch to charge the capacitor, waited with bated breath for it to reach maximum voltage, then touched the firing contact.

Tick.

I closed the circuit again.

Tick.

The scrap of metal inside my magnet coil moved perceptibly each time, but that was about it. I guess my concerns about the danger of this home-built electromagnetic cannon were overblown.

It turns out I should have spent a little more time jotting calculations on the back of an envelope. That 22 000- μ F capacitor stores a little more than 1/50 of a joule for each volt of potential across it. At the 15 to 20 volts my slapdash circuitry was willing to generate, a perfectly efficient transfer of energy would propel a 25-gram projectile at a blistering 3 meters per second. My toddler can throw harder than that. But I wasn't getting 3 meters per second. I might not even



QUICK ON THE DRAW:
...for power, that is.

RESOURCES

have gotten 3 centimeters per second.

Back to the envelope. The ideal coilgun uses the interplay between the current-induced magnetic field pulse inside the coil and the movement of the ferrous projectile to maximize the energy transferred from wire to slug. But that requires the slug to zip through the coil in a fraction of a second. My energy transfer was abysmal, as I could tell by the spark when I closed the contact for a second time. Essentially all the energy of the current pulse was winding up right back in the capacitor.

I needed to put more turns in my coils and to compress my magnetic field to a smaller volume. That way, I could get the projectile moving fast enough to play effectively with the emerging magnetic field.

The right approach would have been to find another tube and wind my wire carefully around it or even to get wire better suited for winding solenoid coils. Instead, I just took a

fresh spool, unwound a few inches from the outside, and soldered a few inches of heavy-gauge wire onto the nib that projected into the hollow core of the spool. The projectile is smaller, but now when I close the connection it hits the other side of the desk with a satisfying tink.

It's not going to shoot down an incoming ballistic missile or even seriously annoy our cat, but it'll do as proof of concept. If I add a second spool and capacitor (or third, or fourth) that can be triggered by a circuit that detects the projectile emerging from the previous one—using, say, a bright light, a photoresistor, and a thin coat of white paint on the slug—I could get some real velocity. I bet I could get the total kinetic energy up to well over a joule.

This little toy also points up many of the reasons that more-powerful Gauss rifles and other electronic projectile throwers still haven't changed the face of battle. The

current through my circuit peaks somewhere around 10 amperes, which is almost 20 times as much as the wire in the coil is rated to carry in continuous duty. A real weapon would be discharging hundreds or even thousands of amperes at hundreds of volts (albeit for only milliseconds at a time) with corresponding stress on capacitors, coils, and switches. That's fine for a small electric power substation but not much fun to carry over your shoulder.

Still, in the back of my mind I have visions of a Mark II home version. Maybe a huge bank of capacitors scavenged from defunct PC power supplies. Or a bicycle wheel, reinforced and wired into a high-current generator. The underlying idea is so attractive that there has to be a way.... ■

ABOUT THE AUTHOR

PAUL WALLICH is a science writer who lives in Montpelier, Vt.

The Foreign Patent Money Trap

You may well need patents in many countries, but that doesn't mean you can afford them

BY KIRK TESKA

INVENTION

A lot of people simply assume that a U.S. patent provides protection outside of the United States or that there is some kind of a "European patent" or even a "world patent." These beliefs are dead wrong. Your U.S. patent gives you no legal recourse should a company based overseas sell your invention overseas; in fact, it even provides that competitor with a free blueprint of your technology.

A recent U.S. Supreme Court decision not only showed how short the arm of U.S. patent law can be, it shortened it a bit more. AT&T had alleged that Microsoft had violated its U.S. patent for speech-processing software. While the question inside the United States was uninteresting—Microsoft's U.S. sales were indeed found to be a violation—the real issue was Microsoft's supplying the code to non-U.S. manufacturers for installation on computers sold abroad. Those non-U.S. activities, the Supreme Court held, were beyond the reach of AT&T's U.S. patent.

The upshot is that your U.S. patent provides you essentially no protection beyond U.S. borders. Other countries' patents also have little force outside their own domains.

So, you say, why not just go out and get a passel of international patents? The catch is the cost. A 2002 report by the U.S. Government Accountability Office, supplemented in 2003, estimated that getting a single patent in the United States and maintaining it for 20 years would cost a small business about US \$10 000, and that extending that same patent to nine other countries would add between \$160 000 and \$360 000. That's



why an engineering manager has to know how patent laws vary between countries.

The first key point has to do with filing deadlines. Most countries don't let you publicly disclose your invention before filing for a patent on it. International treaties, however, allow you to file in the United States, then disclose your invention or sell a product based on it, and then take up to a year to file in other countries.

The second point is the cost of filing and when it will be incurred. You have to know this before beginning the process, or you may find out too late that you can't afford to complete it. Typically, the costs are low to moderate at the beginning, and then they ramp up. An application filed under the Patent Cooperation Treaty (PCT) can designate numerous countries for


MICK WIGGINS

Up-to-date, Relevant Information

- Driving the Bottom Line
- Fueling Imagination

“Saving time in any way you can is critical. Access to IEEE articles and papers is key in this regard.”

– Jon Candelaria, Project Manager, Motorola



From Imagination to Market

Today's top companies rely on IEEE information to drive new research and generate cutting-edge patents.

- Periodicals and conference proceedings that define the future of innovation
- Over 1.5 million documents in the IEEE Xplore® digital library
- Top cited journals in the field

Free Trial!

Experience IEEE – request a trial for your company.

www.ieee.org/innovate

IEEE Information Driving Innovation



RESOURCES

just a few thousand dollars. A year and a half later, however, when the PCT application must be filed in all the countries where a patent is desired, the cost can run \$50 000 or more for only a handful of countries.

Third, be very careful where you file, considering for each country the level of patent protection, the total cost, the size of the market, the number of potential licensees, whether your company will manufacture or sell a product incorporating the patented invention in that country, and whether competitors will likely do the same. You need to know all these things to calculate whether you'll get a return on your patent-investment dollars.

Also, consider the quality of the patent. Only a small fraction of all patents provide any real return on investment. Therefore, you must analyze the strength of the patent you are likely to obtain—which may well not be as broad in its coverage as the patent you have in the United States—and check your analysis with an expert.

Suppose, for example, that the deadline for international filing is fast approaching (remember the one-year rule) and you have two U.S. applications pending. Application A broadly covers the core technology underlying your flagship product, but application B narrowly covers only one functional feature of the product. You might file A in numerous countries, but if competitive products can likely compete using functional features other than yours, then B might be filed in only a few, if any, non-U.S. countries.

In the end, the non-U.S. filing decision carries some risk, because you can't always know what you need to know when you need to know it. For instance, a non-U.S. filing decision must sometimes be made before a market is clearly defined or a product is actually ready for production. Also, the filing decision must sometimes be made before you know whether the effort and cost will bear fruit. Of course, such uncertainty also applies to anyone filing for a patent in his or her home country, but there the timing is easier and the market far more familiar.

This leads us to one final point: things change. As time marches on, all of the above considerations have to be reevaluated. A patent portfolio needs to be weeded out from time to time, freeing up money that would have gone to maintenance fees for better use elsewhere, perhaps even for new non-U.S. filings.

ABOUT THE AUTHOR

KIRK TESKA, adjunct law professor at Suffolk University Law School, Boston, is managing partner of landiorio & Teska, an intellectual-property law firm in Waltham, Mass. His book *Patents for the Too Busy Manager* is scheduled to be published in the spring by Nolo, in Berkeley, Calif.

TO PROBE FURTHER

The Web sites of most intellectual-property law firms contain helpful tutorials regarding non-U.S. patent filings. The GAO's reports mentioned above are free, at <http://www.gao.gov>, appearing under the labels GAO-02-789, July 2002, and GAO-03-910, June, 2003.

Kirk Teska's own *Patent Savvy for Managers* (Nolo, 2007) covers these issues in greater depth; it is available online and at most major bookstores.

Fire and Ice

This book puts the global warming controversy in a very small nutshell
REVIEWED BY M. GRANGER MORGAN

BOOKS On Amazon.com, this book looks like any other: attractive cover, nice promotional blurbs. Imagine my surprise when my copy arrived: 12 by 18 centimeters, 85 pages, wide margins, and only about 140 words to the page. Still, I am a great fan of Kerry Emanuel's previous semitechnical book, the beautifully produced *Divine Wind: The History and Science of Hurricanes* (Oxford, 2005), so I sat down to read the seven brief chapters with interest. I was not disappointed.

Emanuel, a professor of atmospheric science at MIT, explains how and why the climate system has varied enormously over geologic time scales, and he does so in language so clear and concise that any college grad should be able to understand. With echoes of Robert Frost's apocalyptic poem "Fire and Ice," he outlines why, despite these great swings, Earth has never careened into a permanent deep freeze or become overheated, as in the case of Venus, whose surface is hot enough to melt solder.

Next comes a clear explanation of planetary heat balance and the greenhouse effect, with a discussion of the central role of water vapor and clouds, a topic too often glossed over in descriptions for lay readers. Then, using the analogy of the paths traced over time by leaves falling into a turbulent brook, chapter 3 provides a masterly explanation of chaotic systems and why there is an inherent limit to how far ahead we can forecast the weather. This chapter is less successful in moving from the moment-by-moment characterization of weather to the more stable average values that make up climate.

Chapter 4, which discusses



WHAT WE KNOW ABOUT CLIMATE CHANGE
By Kerry Emanuel, MIT Press, Cambridge, Mass., 2007; 85 pp.; US \$14.95; ISBN-10: 0-262-05089-7

climate models, includes the only plot in the book, comparing the output of a climate model based on "natural" radiative forcing with one to which human emissions of greenhouse gases and aerosols have been added. Only the latter successfully replicates the past three decades of climate records. The author conveys this conclusion clearly in a diagram (despite the reference to "colored curves" in what is, in fact, a black-and-white plot).

Having laid out the basics, in chapter 5 the author explains why even a modest increase in average global temperature may—by melting ice sheets, thus raising sea level—intensify droughts, floods, hurricanes, and other disasters. These changes would, of course, have a profound impact on the things people value.

The final two chapters describe in a brief and balanced way the nature and sources of the political controversy that has swirled around the issues of climate science.

There is a good discussion of the self-correcting nature of the scientific method, and of how virtually all serious climate scientists now agree that the average atmospheric temperature is rising and that human activities are contributing to this trend. It also makes clear that sorting out some of the additional details will probably involve uncertainties that are irreducible on the time scale of the geophysical experiment we're running with planet Earth. There is also an appropriate, if brief, discussion of the interest-group politics that have complicated public discourse, notably the role played by the handful of professional deniers who are forever being quoted by scientifically naive journalists seeking "balance."

Unfortunately, the final 15-page afterword by two other scholars was quite a letdown and detracted from the book's overall effectiveness. Judith Layzer, an assistant professor of environmental policy at MIT, and William Moomaw, a professor of international environmental policy at Tufts

University, in Medford, Mass., note that addressing the problem of climate change will require a fundamental restructuring of our energy system. They correctly suggest that if done properly this restructuring "could be relatively painless." However, as at the end of Al Gore's documentary, *An Inconvenient Truth* (2006), when a somewhat random list of remedies scrolls past viewers, Layzer and Moomaw bounce through a potpourri of technologies and policies. This treatment fails to paint a compelling picture of what the United States or the rest of the world should do.

The book would have benefited from a one-page list of suggested further readings and Web addresses for those who do not know their way around this subject and would like to read more. ■

ABOUT THE AUTHOR

M. GRANGER MORGAN, an IEEE Fellow, is head of the department of engineering and public policy at Carnegie Mellon University, in Pittsburgh.

A Word In Your Ear

This iPhone headset is so light you can barely feel it

BY STEVEN CHERRY

TOOLS & TOYS

Like most things Apple designs these days, the iPhone Bluetooth Headset (US \$129) is elegant and simple. Most of all, it's tiny—the smallest Bluetooth headset I've seen. It's so small, in fact, that Apple designed a special travel cable for it. Instead of the usual 30-pin USB cable—good for iPods and the iPhone—the 30-pin end has, at the back, a slot for the headset's 2.3- by 5.0-millimeter slanted tip.

The advantages when traveling are enormous. Now, not only can you charge the iPhone directly from your computer with just this one cable, you can also charge your headset without bringing along a power cord for it.

Even though there's no separate cord for the headset, there are other charging

CALL FOR PARTICIPATION



CTS 2008



The 2008 International Symposium on Collaboration Technologies and Systems

May 19 – 23, 2008

The Hyatt Regency Irvine Hotel,
Irvine, California, USA

<http://cisedu.us/cis/cts/08/>




IEEE



In cooperation with the ACM, IEEE, IFIP

RESOURCES



options. The headset comes with a dock. Also, at the USB end of the 30-pin cable there is the power adapter that came with your iPhone. I was able to charge the headset from zero to full in 90 minutes, just as the diminutive manual had claimed.

So how does it rate as a headset? It's incredibly light, clinging to your ear by only its speaker, an arrangement I found secure enough when walking down the street but rather less so when engaging in more strenuous activity. Salespeople at the Apple store said that the precarious fit explains why a lot of people prefer designs that hook around the ear, such as the Jabra JX10.

Logistically, the headset worked fine. It pairs easily with the phone when the two are attached at the same time to the cable or dock. As with other Apple designs—the one-button mouse and the one-button iPhone—the headset has a single button, at its tip. It also has an LED that goes from red

to yellow to green during charging, or when you turn the headset on and off.

By itself, the one-button design left me uncertain as to whether the headset was paired with the phone or whether it was on. However, the headset also generates rising or falling tones to indicate whether it's on or off, and these tones provided better guidance than the LED. The phone itself can help. Once the devices are paired, the iPhone lets you select, at any time during a call, from its three possible sources of sound: the headset, the regular phone speaker you hold up to your ear, or its built-in speakerphone.

The manual's instructions showed me how to pair the headset with my laptop, so that I could use the headset for a voice-over-IP phone call with Skype.

I would recommend the Apple headset for an iPhone owner who wants to cut down on cables when traveling. However, make sure that it fits you well and won't fall out.

WHAT'S THE OPPOSITE of a Bluetooth headset? External speakers, of course. The Dutch design firm Boynq (<http://boynq.com>)

makes a variety of highly portable speakers that marry attractive form and quality sound. The sleek black-and-silver model of the iCube II (\$69) shown here is indeed a cube, 10 centimeters on a side, that also serves as a recharging station for your iPod, iPhone, or iTouch.

FOR ONE AND ALL: The iPhone Bluetooth Headset (left) and the iCube II complement one another.



LEFT: APPLE; RIGHT: YOUSB.B.V.

IEEE INFORMATION

IEEE BOARD OF DIRECTORS

PRESIDENT & CEO Leah H. Jamieson
Phone: +1 732 562 3928 Fax: +1 732 465 6444
E-mail: president@ieee.org

PRESIDENT-ELECT Lewis M. Terman
TREASURER David G. Green
SECRETARY Celia L. Desmond
PAST PRESIDENT Michael R. Lightner

VICE PRESIDENTS Moshe Kam, Educational Activities; John Baillieu, Publication Services & Products; Pedro A. Ray, Regional Activities; George W. Arnold, President, Standards Association; Peter W. Staecker, Technical Activities; John W. Meredith, President, IEEE-USA

DIVISION DIRECTORS Steven J. Hillenius (I); Thomas G. Habetler (II); Mark J. Karol (III); Edward Della Torre (IV); Oscar N. Garcia (V); Irving Engleson (VI); William D. Kennedy (VII); Thomas W. Williams (VIII); Richard V. Cox (IX); William A. Gruver (X)

REGION DIRECTORS Barry L. Shoop (I); John C. Dentler (2); George F. McClure (3); Robert J. Dawson (4); Robert A. Scollis (5); Loretta J. Arellano (6); Robert A. Hanna (7); Jean-Gabriel Remy (8); Luiz A. Pilotto (9); Janina Mazierska (10)

DIRECTORS EMERITUS

Eric Herz
Theodore W. Hissey

IEEE STAFF

Jeffrey W. Raynes, CAE; Executive Director & COO
+1 732 562 5400, j.raynes@ieee.org
Betsy Davis, SPHR; Human Resources
+1 732 465 6434, e.davis@ieee.org
Anthony Durniak, Publications
+1 732 562 3998, a.durniak@ieee.org
Sally Ericksen, Chief Information Officer
+1 732 562 5345, s.ericksen@ieee.org
Douglas Gorham, Educational Activities
+1 732 562 5483, d.g.gorham@ieee.org
Judith Gorman, Standards Activities
+1 732 562 3820, j.gorman@ieee.org
Cecelia Jankowski, Regional Activities
+1 732 562 5504, c.jankowski@ieee.org

Matthew Loeb, CAE; Corporate Strategy & Communications
+1 732 562 5320, m.loeb@ieee.org
Richard D. Schwartz, Business Administration
+1 732 562 5311, r.schwartz@ieee.org
Mary Ward-Callan, Technical Activities
+1 732 562 3850, m.ward-callan@ieee.org
Chris Brantley, Managing Director, IEEE-USA
+1 202 785 0017, c.brantley@ieee.org

IEEE PUBLICATION SERVICES

& PRODUCTS BOARD John Baillieu, Chair; Tayfun Akgul, Duncan C. Baker, John T. Barr IV, Mohamed E. El-Hawary, Gerald L. Engle, Gerard H. Gaynor, Roger A. Grice, Marion O. Hagler, Jens Hannemann, Donald N. Heirman, Evelyn H. Hirt, Hirohisa Kawamoto, Phillip A. Laplante, Mary Y. Lanzarotti, Michael R. Lightner, George F. McClure, Adrian V. Pais, Roger D. Pollard, Saifur Rahman, Suzanne M. Rivoire, Jon G. Rokne, W. Ross Stone, James M. Tien, Robert J. Trew, Stephen Yurkovich, Amir I. Zaghloul

EDITORIAL CORRESPONDENCE

IEEE Spectrum, 3 Park Avenue, 17th Floor, New York, NY 10016 Attn: Editorial Dept. Responsibility for the substance of articles rests upon the authors, not the IEEE or its members. Articles published do not represent official positions of the IEEE. Letters to the editor may be excerpted for publication.

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.

ADVERTISING CORRESPONDENCE

IEEE Spectrum, 3 Park Avenue, 17th Floor, New York, NY 10016, Attention: Advertising Department, +1 212 419 7760. The publisher reserves the right to reject any advertising.

COPYRIGHTS AND TRADEMARKS

IEEE Spectrum is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Careers, EE's Tools & Toys, EV Watch, Innovations, Newswlog, Progress, Reflections, Software reviews, Speakout, Spectral Lines, and Technically Speaking are trademarks of the IEEE.



I needed to get half-way across the country for an important presentation, but I didn't have a big budget for travel. IEEE got me where I needed to be, and saved me money in the process.

I Made It. Thanks, IEEE.

In nearly every case, **IEEE Travel Services** beat the prices available through other services. Find out how we can help you get where you need to be.

www.ieeetravelonline.org • travel-team@ieee.org
+1 800 TRY IEEE (879 4333) US & Canada
+1 732 562 5387 elsewhere



Are You Ready to Find Your Next Job?



You don't need to be actively seeking a new job to explore the engineering opportunities that might be right for you. The IEEE Job Site can help you find out what you're worth and increase your chance of finding your next job. You owe it to your career. You owe it to yourself.

Visit the IEEE Job Site Today! www.ieee.org/jobs



Forschungszentrum Jülich

in der Helmholtz-Gemeinschaft



The Research Centre Jülich is a member of the Hermann von Helmholtz Association of National Research Centres and is one of the largest research institutions in Europe. We work in the fields of energy and environment, information, key technologies and health, and cooperate intensively with the universities in the federal state of North Rhine-Westphalia.

For our "Central Institute for Electronics" (ZEL) we are seeking a scientist as the

HEAD OF DEPARTMENT

ZEL is a scientific and technical centre of excellence, which pursues research and development projects in cooperation with the institutes of the Research Centre and with external partners. ZEL's tasks largely concern the research programmes of the institutes at the Research Centre.

ZEL's expertise focuses on the development of scientific instruments and is to be found in the fields of analogue and digital electronics, large integrated logic components (FPGA systems), measuring and automatic control, detector technology, imaging systems, information technology as well as signal and image processing.

Candidates should be scientists with broad-based experience in the fields of scientific instrumentation and electronic systems engineering. The successful applicant is expected to be able and willing to supervise and support all activities at ZEL. The project-oriented work at ZEL requires proven organizational talent and business management skills as well as the willingness to undertake intensive cooperation with all the institutes at the Research Centre, as well as with external partners and universities in the region. In view of the size of ZEL with a staff of more than 70, leadership experience is an important requirement.

The implementation of equal opportunities is a cornerstone of our staff policy at the Research Centre, for which we have received the "TOTAL E-QUALITY" Award. Applications from women are therefore particularly welcome. We also welcome applications from disabled persons.

Salary and social benefits will conform to the provisions of the German civil service.

Applications comprising curriculum vitae, list of publications and a short summary of scientific achievements should be sent by 15 January 2008 to

Vorstand der Forschungszentrum Jülich GmbH
52425 Jülich
Germany

Further information at: www.fz-juelich.de

Faculty Positions in Biomedical Engineering and Advanced Materials



The Electrical Engineering Program of the College of Engineering and Applied Science at the University of Wisconsin-Milwaukee (UWM) invites applications from outstanding candidates to fill several faculty positions at all ranks in the fields of biomedical engineering and advanced materials. An earned doctorate degree in Electrical Engineering or an equivalent degree in a closely related field is required. Qualified candidates should have a strong commitment to research, and to teaching undergraduate & graduate courses. Candidates with an established research program and a strong record of extramural funding are preferred and invited to apply for senior positions. Exceptional junior candidates will also be considered. Competitive salary and a generous start-up package will be provided. Screening begins January 15, 2008 and continues until positions are filled. Visit www.uwm.edu/CEAS/EE for further information.

To apply, send a cover letter with a statement of research and teaching interests, a curriculum vitae, and a list of at least three references to: Chair of the Electrical Engineering Search Committee, Department of Electrical Engineering and Computer Science, College of Engineering & Applied Science, University of Wisconsin - Milwaukee, PO Box 784, Milwaukee, WI 53201-0784. **AA/EOE**



SHANGHAI JIAO TONG UNIVERSITY School of Microelectronics

Applications are invited for full-time faculty positions at the ranks of assistant, associate or full professor. Potential candidates should have a Ph.D. degree in Electrical Engineering, Computer Science or Computer Engineering, or are near the completion of their Ph.D. program. Desirable areas of research include but are not limited to RF/Mixed-Signal/Digital IC design, semiconductor device and fabrication, IC packaging and testing, microprocessor architecture, and embedded systems. The successful candidates should be able to demonstrate strong potential and abilities in both teaching and research. They are expected to teach both undergraduate and graduate courses, to supervise graduate students, and to develop and maintain leading research programs.

For more information on the SJTU School of Microelectronics, please visit <http://ic.sjtu.edu.cn/sub/some/index.asp?id=921>

Applications should be forwarded to Mr. Larry Luan, School of Microelectronics, Shanghai Jiao Tong University, Dong Chuan Road No.800, Shanghai, 200240, China, or by email to some-hr@ic.sjtu.edu.cn

Inquiries are welcome by telephone to 86-21-34204546 or 34204748 ext1020.

The Edward S. Rogers Sr. Department of Electrical and Computer Engineering at the University of Toronto invites applications for tenure-stream Assistant or Associate Professor positions, beginning July 1, 2008, in three areas:

ROGERS DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING UNIVERSITY OF TORONTO

1. Information Security

Research areas of interest include identity, privacy and security information technologies for: computer networks, distributed systems, sensor and networked systems, embedded systems, computer architecture and system survivability. Applications and references for this position should be addressed to Professor Dimtrios Hatzinakos, Chair of Search Committee and sent to: InfoSecSearch@ece.utoronto.ca.

2. Computer or FPGA Architecture

Research areas of interest include, but are not limited to: multi- or single- core processor architecture, FPGA architecture and CAD, embedded processor design, programming and compiler support for multi-core and novel processors, memory systems, programmable architectures for integrated digital circuits, and power-aware and power efficient architectures. Applications and references for this position should be addressed to Professor Tarek Abdelrahman, Chair of Search Committee and sent to: CompFPGAsearch@ece.utoronto.ca.

3. Power Engineering

Research areas of particular interest include: operation and control of future transmission and distribution networks, distributed generation and energy storage systems. Applications and references for this position should be addressed to Professor Peter Lehn, Chair of Search Committee and sent to: EnergySearch@ece.utoronto.ca.

Candidates must have (or are about to receive) a Ph.D. in the relevant area.

The department ranks among the top 10 ECE departments in North America. It attracts outstanding students, has excellent facilities, and is ideally located in the middle of a vibrant, artistic, and diverse cosmopolitan city. The department offers highly competitive salaries and start-up funding, and faculty have access to significant Canadian research operational and infrastructure grants. Additional information can be found at: www.ece.utoronto.ca.

The successful candidates are expected to pursue excellence in research and teaching at both the graduate and undergraduate levels.

Applicants must submit their application by electronic email to one of the three email addresses given above. Please submit only Adobe Acrobat PDF documents. Applicants will receive an email acknowledgement.

All applications should include:

a curriculum vitae, a summary of previous research and proposed new directions, and a statement of teaching philosophy and interests. In addition, applicants must arrange to have three confidential letters of recommendation sent directly (by the referee) by email to the correct address given above.

Applications and referee-sent references should be received by January 15, 2008.

The University of Toronto is strongly committed to diversity within its community and especially welcomes applications from visible minority group members, women, Aboriginal persons, persons with disabilities, members of sexual minority groups, and others who may contribute to the further diversification of ideas.

All qualified candidates are encouraged to apply; however, priority will be given to Canadian Citizens and Permanent Residents.



UNIVERSITY OF TORONTO


The Edward S. Rogers Sr. Department of
Electrical & Computer Engineering
10 King's College Road
Toronto, Ontario, Canada M5S 3G4

COLLEGE OF ENGINEERING ALFAISAL UNIVERSITY Riyadh, Saudi Arabia

Alfaisal University is a private, not-for-profit, research university, comprising the Colleges of Engineering, Science and General Studies, Medicine, and Business, and will commence its programs in Fall 2008. The language of instruction is English and modern learning outcomes, paradigms and technologies are used. The university was founded by King Faisal Foundation along with organizations such as Boeing, British Aerospace, Thales, and King Faisal Specialist Hospital & Research Center, who serve on its Board of Trustees.

The College of Engineering will offer undergraduate and graduate programs in the following disciplines and subdisciplines: **ELECTRICAL** (power, communications, signal processing, electronics, photonics), **COMPUTER** (intelligent systems, language and speech, computer systems, computation), **MECHANICAL** (applied mechanics, product creation), **AEROSPACE** (thermo/fluid systems, aerospace systems, transportation, system dynamics and control), **MATERIALS** (materials processing, materials properties and performance, polymers, nanoscience and technology), **CHEMICAL** (catalysis, reactor design, design-systems, polymers). All programs have been developed by renowned scholars from leading universities in the US and the UK, and are designed to be qualified for accreditation according to US and UK standards and requirements.

Alfaisal Engineering seeks candidates for the following positions, commencing in August 2008: **FOUNDING SENIOR FACULTY** (with instructional, research, and administrative responsibilities), **RESEARCH SCIENTISTS** (academics with research focus), **LECTURERS** (academics with instructional focus), **POST-DOCS** (Doctorate degree holders with research focus), **INSTRUCTORS** (Masters degree holders with instructional focus), and **ENGINEERS** (Bachelors degree holders). Attractive salary and start-up support is provided. Queries and Applications should be sent to engnr_recruiting@alfaisal.edu. The subject line should specify the discipline, subdiscipline, position, and the advertisement reference. The deadline for applications is 31 December 2007. Interviews for leading positions will be conducted in January and February 2008 in Cambridge, MA, USA, and Cambridge, England, UK.



A World Class University in the Heart of the Middle East

The American University in Cairo, founded in 1919, is incorporated in the State of Delaware, operates in Cairo, Egypt and is accredited by the Middle States Association of Colleges and Schools. Its primary mission is to advance the ideals of liberal arts and professional education and of life-long learning. Construction of a new 260 acre campus on the outskirts of Cairo is now nearing completion and will enable the University to enroll more than 6,000 degree-seeking students, with its large school of continuing education serving roughly 40,000 adult, part-time learners from downtown Cairo.

For more information visit our website at www.aucegypt.edu

An independent, non-profit, apolitical, non-sectarian and equal-opportunity institution.



The Department of Electronics Engineering has an opening for a faculty position.

The successful candidate will be expected to teach undergraduate courses in Circuits, Electronics, Digital Logic and Microprocessors in the context of the Computer Engineering Major in the Department of Computer Science and Engineering and/or the Mechatronics Concentration in the Department of Mechanical Engineering. Participation in research is expected.

Position # EENG-1-08

- PhD is required.

One-, two- or three-year appointment, subject to mutual agreement will begin September 2008. Renewal of an appointment depends upon institutional needs and/or the appointee's performance.

Salary and rank are according to scale based on qualifications and professional experience. For expatriates, benefits include housing, annual round-trip air travel for appointee and qualifying dependents, plus schooling for the equivalent of up to two children at Cairo American College. In view of AUC's protocol agreement with the Egyptian government, which requires specific proportions of Egyptian, U.S., and third-country citizen faculty, at this time preference will be given to qualified applicants who are U.S. citizens.

For complete details and application information:

<http://aucegypt.interviewexchange.com>



The fastest growing independent provider of advanced yield ramp consulting services and software for major IC manufacturers. PDF is seeking both recently graduated and experienced candidates with PhD, MS or BS degree. Engineering positions in San Jose, San Diego and Dallas, Tx.

Sr. Feol or Beol Integration Eng
Device Eng w process exp
Test Chip Layout Eng
Data Integration Eng
YRS Apps Eng
Semi conductor Process
Integration Experts
Fab Management
Sr. Engagement Dir.
Yield Ramp Eng
Software Eng
CV Test Chip R&D Eng

For more information visit
<http://www.pdf.com> or
 call 408-938-6432.

Attn: Staffing.

Email: jobs@pdf.com

Mail: PDF Solutions, Inc.
 333 W San Carlos St Suite 700
 San Jose, CA 95110

CREATE-NET



CREATE-NET is an International Research Center based in northern Italy (Trento) with expertise in the areas of Applied ICT, Broadband & Wireless Communication, Smart Environments, Pervasive Computing & Communication and Security. We are part of a multicultural environment under the Bruno Kessler Foundation umbrella consisting of several hundred researchers, engineers and scientists. Our objective is to attract and retain the best professionals in the field from all around the world.

We are now actively looking for qualified senior and junior candidates for funded research positions in the following areas:

- **Wireless** – focusing on UWB communications, cognitive radio/networks and coexistence and interference mitigation
- **Security** – focusing on pervasive and mobile computing security, including wireless sensor networks and their applications
- **Pervasive** – focusing on wireless mesh networking as an access architecture for the next-generation Internet.
- **Multimedia and Smart Environments** – focusing on context-aware technologies for improving quality of life by means of pervasive smart environments

Candidates must have an excellent academic record and affinity to research with an industrial focus (PhD or equivalent and publications). English language fluency (spoken and written) is essential as are good writing and oral communication skills, and strong interpersonal skills (proactive, creative and ambitious). You can visit our website (www.create-net.org) or simply send your resume/CV directly to careers@create-net.org.

Important: due to Italian Privacy Protection Law n.196/03 any CV not mentioning explicitly the following wording: "I authorize the use of my personal data in accordance with Italian Privacy Protection Law (30/06/2003, n.196/03)" will be automatically deleted from our database and not taken into consideration.

CREATE-NET is an inclusive, equal opportunity employer offering attractive conditions and benefits appropriate to an international research organization.

IEEE Job Site



Top Jobs.

Top Prospects.

www.ieee.org/jobs

The Department of Mechanical Engineering at National Taiwan University: is seeking potential faculty candidates at all levels starting in August of 2008. The application deadline is 31st of January, 2008. Applicants with the following backgrounds are preferred: System and Control, Opto-Mechtronics, Nano Technology, Bio-Technology, Signal Processing, Communication, Fluid and Thermal Sciences, and Mechanical Engineering related areas. A Ph.D. Degree in ME, EE, or its equivalent is required. Additional information can be found on <http://www.me.ntu.edu.tw>.

University of Waterloo: The Department of Electrical and Computer Engineering invites applications for faculty positions in most areas of computer engineering, software engineering, and nanotechnology engineering, and in VLSI/circuits, information security, photonics, MEMS, control/mechatronics, signal/image processing, and quantum computing. Please visit <https://eceadmin.uwaterloo.ca/DACA> for more information and to apply online.



UNIVERSITY OF
NOTRE DAME

Bettex Chair in Electrical Engineering

The University of Notre Dame invites applications and nominations for the Leonard C. Bettex Chair in Electrical Engineering.

The successful candidate will have a distinguished and internationally recognized record of research achievement in the general area of communication systems and networking. The Bettex Chair holder will be part of a Notre Dame team working on the most important and exciting problems in 21st century digital communications and networking; the successful candidate will be expected to play a leadership role in that effort – securing external funding, supervising graduate students, and carrying out research projects that impact the evolving information infrastructure. In addition, the Bettex Chair holder will be expected to meet and exceed Notre Dame's traditionally high standards for educating both graduate and undergraduate students.

Applications consisting of a curriculum vitae and a cover letter (preferably in PDF format) should be sent via e-mail to bettexchair@ee.nd.edu. Alternatively, a hard copy may be sent to Prof. Tom Fuja, Chair – Department of Electrical Engineering, 275 Fitzpatrick Hall, University of Notre Dame, Notre Dame, IN 46556. Consideration of applications will begin January 1, 2008 and will continue until the position is filled.

The University of Notre Dame is an equal opportunity employer. We particularly invite applications from women and members of groups that are underrepresented in science and engineering.

THE GEORGE
WASHINGTON
UNIVERSITY
WASHINGTON DC

Department of Electrical & Computer Engineering

BIOMEDICAL ENGINEERING FACULTY POSITION

Tenured Full Professor

The Biomedical Engineering Program in the Department of Electrical and Computer Engineering at The George Washington University invites applications for one faculty position, to begin as early as Spring 2008, at the rank of tenured Full Professor.

Applicants must have a doctoral degree in Bioengineering, Biomedical Engineering, Biophysics, or any related discipline that includes depth in both engineering and biology. Teaching responsibilities include the development of both undergraduate and graduate courses. The successful candidate will be expected to develop a vigorous, externally funded research program that supports masters and doctoral students, and to promote the growth and visibility of the program through high-impact journal publications and presentations.

Additional information and details on position qualifications and application procedure are available on <http://www.ece.gwu.edu>. Electronic applications are encouraged and must be sent to: korman@gwu.edu. Review of applications will continue until the position is filled.

THE GEORGE WASHINGTON UNIVERSITY IS AN EQUAL OPPORTUNITY/AFFIRMATIVE ACTION EMPLOYER.



Texas
A&M
UNIVERSITY AT QATAR

Applications are invited for faculty positions for all ranks:

Professor, Associate Professor and Assistant Professor. The positions are available at Texas A&M University at Qatar (TAMUQ) which prepares graduates that satisfy identical requirements and receive Texas A&M University degrees. The TAMUQ campus is situated within a brand new building and is part of Education City, Doha, Qatar, a consortium of educational and research institutions hosted by the Qatar Foundation for Education, Science and Community Development.

TAMUQ began teaching undergraduate students in chemical, electrical, mechanical, and petroleum engineering in the Fall of 2003. As the first students progress through the curriculum, faculty are being added to teach all necessary courses in these majors. The electrical engineering program at TAMUQ offers for the moment only BS degrees but a graduate program in electrical engineering is scheduled to start in the near future. All formal instruction is given in English. More information about TAMUQ can be found at <http://www.qatar.tamu.edu/>. Specific areas of interest are listed below for guidance but are not intended as a limitation.

Computer Engineering (at any rank): Computer architecture, software engineering, operating systems, distributed computing, compiler techniques, network/processor security, microprocessors and embedded systems, and mobile computing.

Power Systems (at any rank): Generation, distributed generation, transport and distribution of electrical energy, deregulation, forecasting, electrical installations, intelligent buildings (industrial and commercial).

Electronics (at any rank): Analog and Mixed Signal Circuits and Systems, Design, implementation and application of CMOS wireless transceivers; CMOS sensors and circuitry, such as ADC, instrumentation and power amplifiers for a variety of applications, including telemetry.

Applicants must have a Ph.D. or equivalent degree, or completion of all requirements by date of hire. For senior positions, applicants should have a proven record of scholarly contributions and a proven ability to attract research funding. For junior positions, candidates should have demonstrated potential for quality teaching and research.

Starting rank and salary will depend on qualifications and experience. The appointment includes the following benefits according to TAMUQ's policy: air tickets to Doha on appointment; annual home leave allowance for all family members; coverage of the local tuition fees for school-age dependent children; and local transportation allowance. Fringe benefits include health and medical insurance as well as an enrollment in a retirement plan. Initial appointment will normally be on a two-year contract. Re-appointment will be subject to mutual agreement.

Applications, including a full curriculum vitae with list of publications, statement of teaching, statement of research as well as the names, addresses (regular mail and E-mail), fax, and phone numbers of three references to should be sent to:

Dr. Costas N. Georgiades, Department Head
c/o Ms. Debbie Hanson
Department of Electrical and Computer Engineering
Texas A&M University
College Station, TX, 77843-3128.

Texas A&M University at Qatar is an equal opportunity/affirmative action employer and actively seeks the candidacy of women and minorities. The deadline for applications is February 15, 2008 but applicants will be considered until the positions are filled.

TECHNICALLY SPEAKING

By Paul McFedries

The New Geographers

"The territory no longer precedes the map."

—Jean Baudrillard

Mapmaking seems like a quaint art that ought to have died off at the turn of the millennium, if not before. Yes, things change—the Czechs and the Slovaks part company; Burma becomes Myanmar, and Bombay becomes Mumbai; the Aral Sea shrinks to a quarter of the size it was 50 years ago. Major events all, but mere tweaks in the mapping world.

I actually have no idea whether *analog* maps are bombing, but I am certain that **digital maps** are booming, and they're generating tons of new words and phrases as a result. But digital two-dimensional representations of the world, also known as **Web maps**, are only the beginning. Such services as Google Maps, MapQuest, and Yahoo Maps are redefining how we look at—and get around in—the world. They provide so-called **base maps** as starting points for more detailed **map mashups**, which plot the locations of user-generated content, such as apartment rentals, weather forecasts, traffic data, and photos. You can now buy digital cameras that come with built-in Global Positioning System (GPS) receivers that note the exact longitude and latitude of each picture, data that are readable by map services. This is called **neogeography**, and it has become absurdly easy thanks to such annotation services as Platial ("The People's Atlas") and the annotation features built into Google Maps and others. **Neogeographers** annotate maps to create their own **ground truth**, that is to say, the world as they see it—their **autobiogeographies**. These **geo enthusiasts** may also engage in **collaborative annotation**, in which a number of people add **geotags** to a single map. (If they **geotag** a location on a Platial map, the location is said to have been **platialized**.) The combination of all available base maps and **geotagged public maps** is sometimes called the **networked atlas**, the **geo ecosystem** or, more commonly, the **geoweb**.

It's not just **neocartographers** who are making newfangled **map worlds**; companies are also automating the process. For example, there is software available that can analyze the text of, say, a book, extract the place names mentioned in the text, and then plot them on a map, a function known as **geoparsing**. As maps become searchable according to such **geodata** as ZIP codes and latitude and longitude coordinates, users can tailor their searches to specific places, a process called **geosniffing**. Companies also offer interactive programs that display a series of digital maps annotated with local lore, facts, and historical data, creating a new genre called **map-based storytelling**, or **geostorytelling**. We're starting to see **location-aware**



devices such as GPS-enabled mobile phones running services that display annotated maps of the user's current location, a technology combination known as **mobile augmented reality** [see "Is It Live or Is It AR?" *IEEE Spectrum*, August]. People also play **geocaching**, a scavenger hunt in which participants receive the geographical coordinates of a cache of items and then use GPS and other such **geotools** to locate them.

The three-dimensional equivalent of the digital map is the **digital globe**, which incorporates photos and 3-D modeling technologies to produce an immersive environment for exploring nearly any part of the world. That **virtual globe** is most famously found in Google Earth, but Microsoft's Virtual Earth is similar. As with

2-D digital maps, 3-D digital globes can be tagged by users and by automated means, a process known as **geocoding**. Photography plays a big part in these virtual worlds, particularly satellite and aerial imagery, although both Google Earth and Virtual Earth are starting to incorporate ground images as well, a competition sometimes called the **3-D data arms race**. (And not without controversy: the first Google Earth ground images included embarrassing shots of people hanging around outside strip clubs.) The images are augmented with **geospecific** simulations of actual sites, in contrast to generic, or **geotypical**, environments.

In some cases, markers are not to the virtual world but to the real world itself—buildings, bridges, and equipment. Companies attach sensor chips to these and countless other objects to watch over them, but we're starting to see the first signs of technology that blends sensor data with 3-D maps, a technique called **reality mining**. The U.S. military hopes to capitalize on such data to generate what it calls **geoint** (geographical intelligence).

If there's a killer app for **geospatial** data it may be **virtual tourism**, which lets people "travel" to any part of the world without the agony of airline food. Virtual tourism is also called **virtual globetrotting** and **Google sightseeing**.

With all this digital mapmaking activity, you can see that maps and atlases printed on (scoff!) paper are so last century. The new arts of neogeography and neocartography are thriving in their stead, and they will soon be annotating, augmenting, tagging, coding, and parsing your reality. ■

PAUL MCFEDRIES is a technical and language writer with more than 40 books to his credit. He also runs Word Spy, a Web site and mailing list that tracks new words and phrases (<http://www.wordspy.com>).

LOU BEACH

Included
with IEEE
membership

ieee.tv™

Internet television for technology professionals



Program Series Include:

Meet the Authors

- Carl Selinger
"Stuff You Don't Learn
in Engineering School"

Conference Highlights

- Do the right thing:
Social Implications of
Technology

Careers in Technology

- Careers in IT
- Art of the Start:
Entrepreneurship

Additional Programs

- Analog to Digital
- Ethernet in the First Mile
and MORE


Tune in: www.ieee.org/ieeetv





Sprechen Sie MATLAB?

Over one million people around the world speak MATLAB. Engineers and scientists in every field from aerospace and semiconductors to biotech, financial services, and earth and ocean sciences use it to express their ideas. Do you speak MATLAB?



Modeling electric potential in a quantum dot. Contributed by Kim Young-Sang at HYU.

This example available at mathworks.com/ltc

MATLAB®
The language of technical computing.

Image: Kim Young-Sang, Jeong Hee-Jun, Quantum Device Lab, Hanyang Univ. ©2007 The MathWorks, Inc.