THE MAGAZINE OF TECHNOLOGY INSIDERS

ieee spectrum

1.11

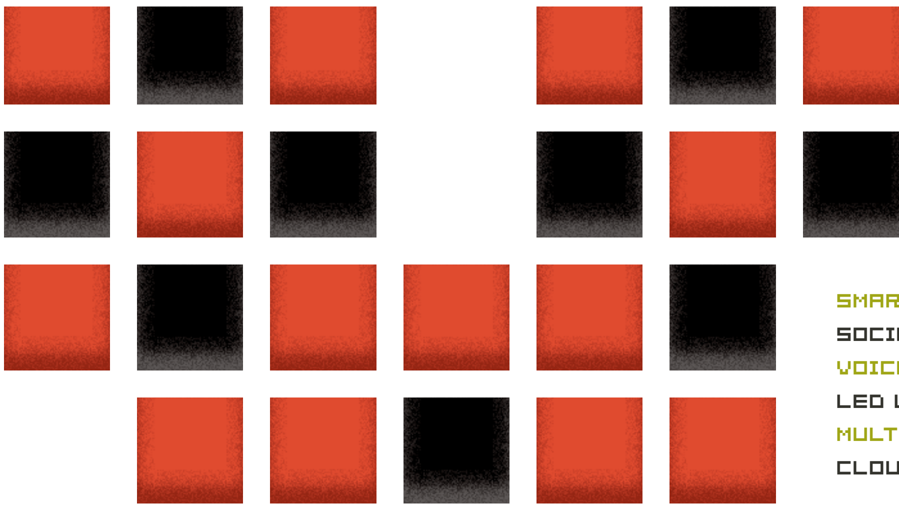

SMARTPHONES
SOCIAL NETWORKING
VOICE OVER IP
LED LIGHTING
MULTICORE CPUs
CLOUD COMPUTING

TOP 11 TECHNOLOGIES OF THE DECADE

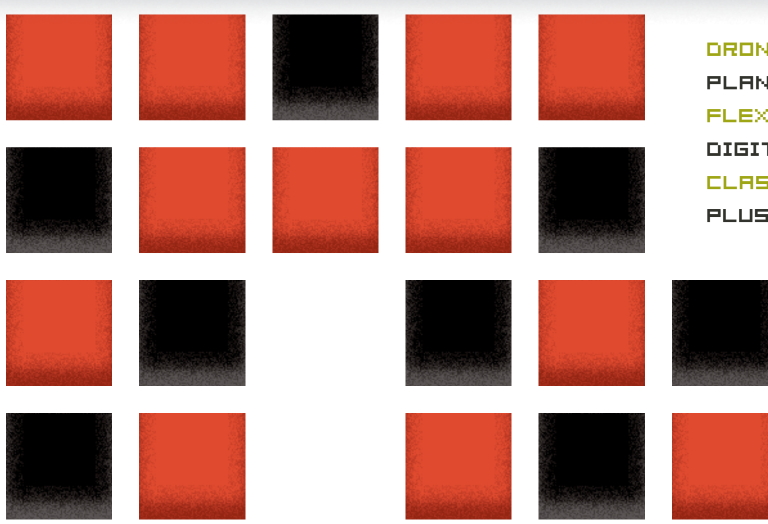

DRONE AIRCRAFT
PLANETARY ROVERS
FLEXIBLE AC TRANSMISSION
DIGITAL PHOTOGRAPHY
CLASS-D AUDIO
PLUS: COMPANIES TO WATCH

U.S.A $3.99
CANADA $4.99
DISPLAY UNTIL
3 FEBRUARY 2011
SPECTRUM.IEEE.ORG

0 1 > 0 74470 64849 9

IEEE

COMSOL MULTIPHYSICS®

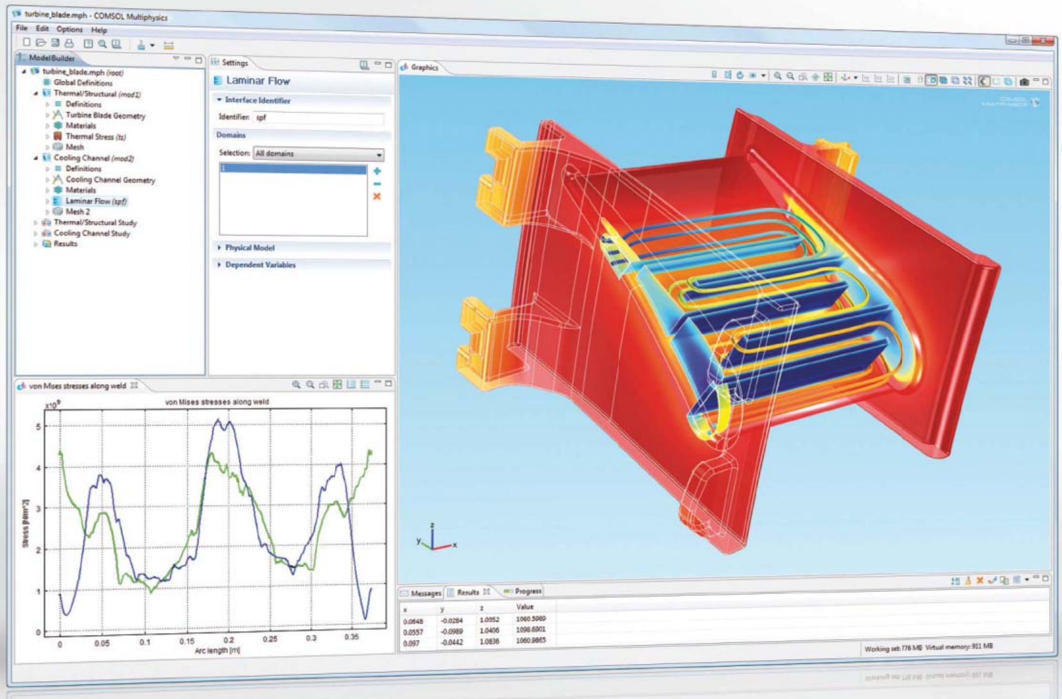

A stator blade in the turbine stage of a jet engine is heated by the combustion gases. To prevent the stator from melting, air is passed through a cooling duct in the blade.

Capture the Concept.

With COMSOL Multiphysics® you are empowered to build the simulations that accurately replicate the important characteristics of your designs. The key is the ability to include all physical effects that exist in the real world. This multiphysics approach delivers results—tangible results that save precious development time and spark innovation.

Watch tutorial

www.comsol.com/showcase

COMSOL

volume 48 number 1 international

ieee spectrum 1.11

COVER: MICHAEL SOLITA

UPDATE

OPINION

DEPARTMENTS

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

volume 48 number 1 international

ieee spectrum 1.11

SPECTRUM.IEEE.ORG

AVAILABLE 5 JANUARY

Live Coverage From CES 2011

It's time once again for that annual extravaganza of gadgets, the Consumer Electronics Show. From 5 to 9 January, *IEEE Spectrum* will be reporting all the tech news live from Las Vegas. Our daily podcasts, videos, and blog posts will keep you up to date with all the latest advances. After the frenzy dies down, check back for an analysis on the best and worst devices at the show. You'll find it all at http://spectrum.ieee.org/CES2011.

LEAVING THE LIGHTS ON

RACKS OF LIGHT SOCKETS fill a 100-square-meter room—just one testing lab at the Rensselaer Polytechnic Institute's Lighting Research Center. See what researchers learn by torturing compact fluorescent and LED lights to the point of failure in this audio slideshow.

CLOCKWISE FROM LEFT: ETHAN MILLER/ GETTY IMAGES; JOSEPH CALAMIA; MICROSOFT

TECH INSIDER WEBINARS

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

- **27 JANUARY:** Highly Efficient Models for Multibody Systems for HIL Simulation of Robotic Systems http://spectrum.ieee.org/webinar/1732611
- Reduce the Costs of Developing DO178B Certifiable Software http://spectrum.ieee.org/webinar/1742631
- Integrated Electronics Cooling Solution for Packages, Boards, and Systems: ANSYS Icepak 13.0 http://spectrum.ieee.org/webinar/1736040
- Capability Versus Complexity: Building Smarter Products http://spectrum.ieee.org/webinar/1730877
- Product & Service Innovation in the Electronics Industry http://spectrum.ieee.org/webinar/1705825
- Four Steps to Success in Prototyping a Medical Device http://spectrum.ieee.org/webinar/1711715
- Smarter Electronics Industry: A Virtual Briefing Center *Sponsored by IBM* http://spectrum.ieee.org/static/ibm-vbc
- Free white papers by experts: http://spectrum.ieee.org/whitepapers
- White papers by Intel: http://spectrum.ieee.org/intel-resource-library
- Check out the design resource library by Texas Instruments: http://spectrum.ieee.org/static/ti-resource-library

IEEE.ORG/THEINSTITUTE

AVAILABLE 6 JANUARY

THE TECH BEHIND KINECT
Read about the real-time motion-tracking technology, developed with the help of an IEEE Fellow, that made Microsoft's controller-free gaming platform Kinect possible.

PICK UP VALUABLE MANAGEMENT SKILLS BY VOLUNTEERING
Serving on an IEEE board can boost your career, according to several members. They say they've gained leadership and organizational skills, honed networking skills, and learned to manage others.

CYCLOTRON HONORED WITH IEEE MILESTONE
In December, the TRIUMF cyclotron was honored with an IEEE Milestone in Electrical Engineering and Computing. Learn the history of the world's largest cyclotron.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2011 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 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 Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

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

back story

Flybot for a Day

When we picked unmanned aerial vehicles as one of the top technologies of the decade, Senior Editor David Schneider leaped at the chance to write about them. Schneider flies radio-controlled model airplanes for fun, and he dearly wanted to see how they stack up against the multimillion-dollar drones flown by the U.S. Army and other militaries. To be sure, some differences are obvious: Hobbyists can't buy a plane capable of annihilating a small building in an instant, for example. But he was surprised to find that recreational model airplanes do have a couple of nice features not found in military UAVs. That became apparent, however, only after he saw the technology up close.

While researching his article, Schneider [above, left] visited the U.S. Army's Unmanned Aircraft Systems Training Battalion at Fort Huachuca in southeast Arizona. There he had the opportunity to inspect a variety of UAVs and the ground stations where their flights are directed by men like Staff Sgt. Frank Peterson, shown above with Schneider in one of the "shelters" used to operate AAI Corp.'s RQ-7B Shadow unmanned aerial vehicle.

Such stations are, of course, far more complex than anything a hobbyist might have access to. Nevertheless, Schneider explains, "They lack some of the features that many modelers enjoy when they fly planes while watching a video downlink—like an easy way to pan and tilt the flight camera."

Are they hard to fly without such niceties? "Not really," Schneider says. "The Army, unlike the Air Force, has gone in for automation in a big way with its UAVs, relieving their operators of having to develop the usual piloting reflexes."

One exception is General Atomics Aeronautical Systems' Sky Warrior (a derivative of the Air Force's well-known Predator UAV), the earlier models of which required the traditional stick-and-rudder skills military pilots have been honing since World War I. Schneider got to inspect such an aircraft up close at the Army's Libby Airfield. His hosts even put him in a ground station and let him take the controls for a short, albeit computer-simulated, flight.

"Takeoff was easy," Schneider reports. "And I even got the thing down on the runway in one piece—although I scraped up the plane's low-hanging tail when I landed it. Good thing Libby has a 12 000-foot runway or I'm sure I would have done a lot more damage." ❑

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 Reflections column is in *IEEE Spectrum,* Vol. 48, no. 1 (INT), January 2011, p. 21, or in *IEEE Spectrum*, Vol. 48, no. 1 (NA), January 2011, p. 25.

ieee spectrum

EDITORIAL

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

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

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

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

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

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

ASSOCIATE EDITORS Joshua J. Romero (Online), j.j.romero@ieee.org; Sandra Upson, s.upson@ieee.org

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

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

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

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

EXECUTIVE PRODUCER, SPECTRUM RADIO Sharon Basco

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

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

IEEE SPECTRUM JOURNALISM INTERNS Ariel Bleicher, a.bleicher@ieee.org; Joseph Calamia, j.calamia@ieee.org

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

ART & PRODUCTION

SENIOR ART DIRECTOR Mark Montgomery

ASSOCIATE ART DIRECTOR Michael Solita

ASSISTANT ART DIRECTOR Brandon Palacio

PHOTO EDITOR Randi Silberman Klett

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy

EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia

SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

WEB PRODUCTION COORDINATOR Jacqueline L. Parker

MULTIMEDIA PRODUCTION SPECIALIST Michael Spector

EDITORIAL CORRESPONDENCE

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

ADVERTISING CORRESPONDENCE

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

DAVID SCHNEIDER

contributors

S. MASSOUD AMIN, considered by many to be the "father of the smart grid," analyzed the 15-year decline in the U.S. grid's reliability for The Data [p. 64]. Before becoming a professor of electrical and computer engineering at the University of Minnesota, he worked at the Electric Power Research Institute, where he led research on self-repairing energy infrastructures. Amin is a senior member of the IEEE.

FRANK CHIMERO gained a renewed appreciation for smartphones after illustrating this month's "Top 11 Technologies of the Decade" [p. 23]. "I'm pretty optimistic about what those little things can do," he says. A teacher and graphic designer in Portland, Ore., Chimero has illustrated articles and covers for *The New York Times, Wired, Time,* and *The Atlantic*. He sometimes sees his work pop up in unexpected places. Once, for instance, an illustration he did for *Newsweek* became the centerpiece of a joke on "The Colbert Report."

PETER FAIRLEY started his career by covering the chemical industry and its environmental baggage. "So when I went freelance a decade ago," he says, "I vowed to focus instead on solutions to problems." While reporting "Flexible AC Transmission" [p. 51], he discovered the social dimension. "By enabling developing countries, such as Brazil and India, to move electricity over long distances in the form of alternating current, FACTS can make transmission more egalitarian. The transmitted power is easy to tap, serving rural communities between the big centers."

RICHARD STEVENSON, a Ph.D. in physics based in Britain, has written many pieces for *IEEE Spectrum* on the incremental improvement of light-emitting diodes. But he says it was only this month's story, "LED Lighting" [p. 34], that allowed him to draw the big picture. He traces how the humble diode graduated from being the "on" button in your stereo to replacing the lightbulb itself through new techniques that squeeze ever more light out of chips.

JAMES MIDDLETON, of London, has covered communications technology for a dozen years, most recently at Telecoms.com. In reporting "Voice Over IP" [p. 30] he finally got to interview the founding fathers of Internet telephony, who recounted the field's unexpected origins. Middleton studied English and drama at university but must have a bit of the engineer inside him: He says he loves learning how things work, whether by hand-coding a Web site or restoring vintage bicycles.

DAVID YELLEN shot portraits of the camera that captured the world's first digital image and its engineer, Steven Sasson, for this month's tribute to digital photography [p. 55]. Before mastering the art and science of digital portraiture, Yellen worked as a musician, fashion designer, and fishing crewman. When he's not posing subjects for *Fortune* or *Rolling Stone*, he still fishes in Sheepshead Bay, near his home in Brooklyn, N.Y.

CLOCKWISE FROM TOP RIGHT: RANDI SILBERMAN KLETT; GIANNA KUNNAS; DAVID YELLEN; RANDI SILBERMAN KLETT; FRANK CHIMERO; UNIVERSITY OF MINNESOTA

Yamagata — Ideal for a Conference

Kaminoyama Castle

ADVERTORIAL

Jane Ambrose looks at the Prefecture of Yamagata and finds that this region of hot springs and natural beauty was perfect for the International Computing & Information Science Conference.

This year's International Computer & Information Science Conference – ICIS 2010 – held under the auspices of ACIS, the International Association for Computer and Information Science, which provides a global forum for research in education and industry. It took place at the Tsukioka Hotel in the beautiful surroundings of Kaminoyama City in the Prefecture of Yamagata, northeast Japan, around 2 hours and 30 minutes from Tokyo station by Shinkansen(bullet train).

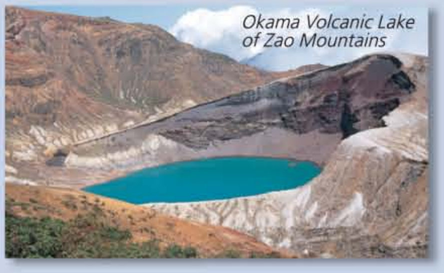

Okama Volcanic Lake of Zao Mountains

Famous for its Onsen (hot springs), the Yamagata region has much to offer to both the tourist and the conference market. Its lush, natural beauty has spectacular scenery throughout its mountains, rivers, and oceanside views.

The Kaminoyama Castle, which rises above the city and gives spectacular views of Mount Zao, is well worth visiting. It is particularly known for its museum, which showcases artifacts from over 400 years of cultural history. Additional fascinating places to visit are: The Harusame-an Tea House (associated with the priest and Zen Master Takuan), the Bukeyashiki (or samurai residences), and Tsurunoyasumiishi (originally the birthplace of the Kaminoyama hot springs).

Zao Onsen Spa

The Conference, which includes such diverse topics as 'Research Issues in Data Intensive Semantics', 'Artificial Intelligence and Emerging Technology' and 'Software Reliability: Past, Present and Future' was particularly well attended this year. Roger Lee, ACIS founder and CEO, Director of Software Engineering & Information Technology Institute, and Professor of Computer Science at the University of Central Michigan believes that it might be because of the Onsen at Kaminoyama. Mr. Lee said: "We normally have 20% of people who don't show up for the Conference, but this time there was only a 4% 'no-show' and I believe that was down to the Onsen." He continued to comment: "Our participants told me that they really enjoyed being here and sooner or later I would like to come back."

Dr. Naohiro Ishii, General Chair of the event along with Roger Lee, Professor in the Department of Information at Aichi Institute of Technology agreed with him and said: "Normal international conferences are held in urban convention centers. This one, held at Kaminoyama, introduced delegates to a bit of Japanese culture which our US delegates told me they had really enjoyed."

One participant from Switzerland, Professor Dr. Robert Winter, Chair of Information Management, Director of the Institute of Information Management and Academic Director, Executive Master of Business Engineering at the University of St. Gallen said: "There is a big difference between a business venue and one that is for leisure but I really believe that it is good to have conferences at the latter. The atmosphere is much more relaxed and therefore interaction becomes easier."

ICIS 2010

Was the Conference a successful one? When asked that question most delegates gave a very positive answer, but one of the more interesting viewpoints came from Mr. Takayuki Ito, Dr. Eng., Associate Professor in the Graduate School of Engineering at Nagoya Institute of Technology. As one of the keynote speakers at the Conference with a lecture entitled 'Multi-Agent Technologies and their Applications' Mr. Ito said: "I wanted to introduce delegates who were coming here to a bit of Japanese culture and I think that was very important. There were many more delegates than I expected wearing the Yukata, which is a casual Japanese summer kimono – and that can only be a good thing. If people are relaxed enough to do that – no matter what their own culture would do – then they are relaxed enough to get the best out of interaction."

Dr. Jonathon Reynolds, Professor of History at Kentucky University had this to add: "Big cities are like big cities everywhere – it's very nice to be able to get out of Tokyo and see the variety of the environment."

JAPAN MICE
JAPAN. A GLOBAL CROSSROADS

CONTACT INFORMATION:

Yamagata Convention Bureau
100 Hirakubo, Yamagata 990-0076, Japan
Fax: +81-23-635-3030
E-mail: sales@convention.or.jp

Japan National Tourism Organization
Japan Convention Bureau, New York Office
Gil Cardon, Convention Manager
11 West 42nd Street, 19th Floor New York, NY 10036, USA
Phone: +1-212-757-5640 Fax: +1-212-307-6754
E-mail: jcb@jntonyc.org
http://www.japantravelinfo.com/

Japan Tourism Agency

spectral lines

MCKIBILLO

The Jacked-In Decade

Connectedness is now a given. Smartphones offer us advice, the Internet carries our voices and video, the cloud archives our data, and e-readers pull it to Earth again. We share all this stuff and more via our social networks. Behind the scenes, brainy grids juggle power, and vast military networks stretch out to robotic planes equipped with devastating firepower above distant battlefields. Nothing seems to happen anymore without global resources coming into play.

These technologies—and a few off-the-grid ones, like dazzling LED light sources—made the cut of our survey of the most consequential innovations to pervade our lives in the decade that just ended. Though many stemmed from discoveries made years or even decades earlier, only in the 2000s did their time truly come.

Our idea was pretty ambitious: to give you a compact and yet comprehensive and compelling guide to the state of technology a little over a decade into the 21st century. Technology moves fast, and although dozens of journals chronicle its esoteric little breakthroughs, and several newsstand magazines detail its human-interest stories and gee-whiz gadgetry, not many magazines ever try to give you an authoritative accounting of what the most important trends have been, where they came from, and where they are headed in the near term.

That's what we did here. You're welcome.

In selecting these 11 technologies for our list and ranking our picks, we strongly considered the technology penetration rate, as the MBAs call it—the speed with which the technology went from a promising possibility to dominance of some domain of human activity. How quickly and intensively the technology insinuated itself into our lives, in other words. And because we didn't want an issue that would have all the excitement of a B-school seminar, in finalizing our list we also included a fudge factor for sheer tech exuberance. That explains the inclusion of, for example, planetary rovers, which have not really changed life for us earthlings but have revolutionized space science, inspired intelligent people everywhere, and probably helped to entice at least a few bright youngsters into becoming engineers rather than hedge fund managers.

> Our amazement at how these 11 technologies have transformed our lives was diminished only by our consideration of what this issue might have looked like had we done it a century ago

Our more attentive readers will note that we've broken with the practice of our past January surveys by ignoring the losers and dealing with only big and broad technology winners. The reason is simple: Here we're giving a backward glance rather than a prediction, so to finger losers would be too easy, even unsporting.

To be sure, any list will annoy the champions of the items that didn't make the cut. How could we possibly snub hybrid electric cars? Because as successful as they have been, hybrid electric cars still account for a very small percentage of auto sales worldwide. Why did we relegate tablet computers to a mere "honorable mention" in our survey? Well, as much as we love the iPad and are awed by its early success, it hasn't quite knocked laptops and netbooks out of the picture yet.

Our amazement at how these 11 technologies have transformed our lives was diminished only by our consideration of what this issue might have looked like had we done it a century ago. In that hypothetical January 1911 issue we'd have covered radio, the automobile, the airplane, the diode, the triode, movies, the mercury vapor lamp, the electrocardiograph, and—drumroll, please—cornflakes. All those things still flourish, and all but the last have markedly improved.

Finally, why 11 technologies? To paraphrase the motion picture *This Is Spinal Tap*, any top 10 list can have 10 entries. Ours has 11. —Philip E. Ross

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

A Less Mighty Wind

Three reasons wind power could wane

Wind turbines wring energy out of a free-flowing fuel supply that may be losing some of its punch. Surface winds appear to be weakening across the Northern Hemisphere, including in the United States, Western Europe, and China—the world's top three markets for wind power. And climate change threatens to weaken them further during this century as faster warming over northern latitudes trims the temperature gradients that energize airflows.

China could be the hardest hit, according to modeling by University of Texas–Austin research scientist Diandong Ren in the November issue of the *Journal of Renewable and Sustainable Energy*. He projects a 4 to 12 percent decrease in wind speeds in China for the last three decades of the 21st century (compared to the corresponding decades of the 20th). Since the energy in wind increases with the cube of the wind speed, Ren estimates that the slower winds would trim power from Chinese turbines by at least 14 percent.

There is now little doubt that China's surface winds are already slowing. Independent analyses published in 2009 and 2010 found that recent readings from weather station anemometers were lower than those taken in the 1960s and 1950s. In both cases, the majority of Chinese stations reported slowing near-surface winds, and the largest declines occurred in the windiest regions—in the north, on the Tibetan Plateau, and along China's coastline.

Comparable stilling is occurring across the Northern Hemisphere, according to an October report by a team centered at France's Laboratoire

CHINA'S LOSS: China is installing wind capacity fast, but the winds over the country seem to be slowing.
PHOTO: DOUG KANTER/BLOOMBERG/GETTY IMAGES

update

des Sciences du Climat et l'Environnement (LSCE). Their report in the journal *Nature Geosciences* found that winds slowed by 5 to 15 percent over almost all continental areas in the northern midlatitudes between 1979 and 2008.

Experts in the wind-power industry pooh-pooh such warnings. Peter Thomas, a senior engineer with the wind energy consultancy GL Garrad Hassan, based in Bristol, England, concedes that the projections are of a scale that could "impact the economics of the wind-power industry." But he questions their veracity.

Thomas argues that data sets from anemometers are not robust enough to support such interdecadal comparisons, because measurement practices were poorly standardized as recently as the 1980s and may be corrupted by construction around weather stations, many of which are at airports or near cities. "It is important to separate these potential influences from the measured data before conclusions are drawn," says Thomas.

That data-quality critique is wearing thin, however, according to Jean-Noël Thépaut, who runs the data division for the European Center for Medium-Range Weather Forecasts in Reading, England, and is a coauthor of the French report. Thépaut says the team applied a stringent screen to remove questionable anemometer data, narrowing its analysis to reports from

NO BREEZE: Anemometers indicate worsening wind-farm prospects.
PHOTO: DAVID PARSONS/ISTOCKPHOTO

just 822 out of roughly 10 000 possible anemometers worldwide. "My colleagues from LSCE have been very careful with the quality control," says Thépaut.

Still, even Thépaut sees unanswered questions, starting with why winds are slowing and whether the stilling will continue. The French study identified climate change as the most likely cause of stilling over central Asia, which means Ren's modeling could well foretell a less productive future for wind power in China.

But the French modelers pegged forests as the primary culprit behind the stilling in other regions. Their modeling showed a correlation between wind reductions and forest regrowth. As they grow, trees increase the roughness of Earth's surface and could be responsible for up to 60 percent of the stilling observed over North America and Western Europe, the modelers estimate. However, the effect on the wind industry might be minimal, because industrial wind turbines tower above most trees, their hubs supported on structures that commonly stand 60 to 100 meters tall. The average anemometer tower is just 10 meters tall. Thépaut says his colleagues plan to figure out if turbines will really be above the fray in the months ahead by analyzing wind-speed data from air-balloon-based weather stations called radiosondes.

There is one source of waning wind that turbines cannot rise above: neighboring wind farms. Here, too, modeling reveals previously unforeseen impacts on wind speed. For example, Arno Brand, a wind modeler at the Energy Research Centre for the Netherlands, in Petten, projects that wind "shadows" behind installed wind farms will sap the productivity of some planned offshore wind projects in the Dutch zone of the North Sea.

Brand's modeling suggests that wind farms must be spaced at least 10 to 30 kilometers apart to keep speed reductions from such shadows below 0.5 meters per second—and even that reduction translates to a 14 percent power loss for a turbine seeing 9.5-m/s wind instead of 10 m/s. Brand says this Dutch problem could become a diplomatic dispute, because the United Kingdom has plans of its own to build what would be three of the world's largest offshore wind farms just upwind of Dutch waters. "These farms are going to produce considerable wind shadows that will affect the most important Dutch zones," says Brand.

While slowing winds could shave value off wind farms or complicate their planning, none of the modelers estimates that the impacts will eliminate the advantage that has made wind power the world's fastest-growing energy source: its supply of virtually carbon-free power at a cost that's comparable to that of fossil fuels. As Thomas points out, it is fossil fuels that are the truly unpredictable fuel source. Even before factoring in their likely contribution to global climate change, the economic cost of fossil fuels is already far harder to predict than the wind will ever be, he says. "In the last 10 years, the cost of a barrel of oil has varied between $20 and $150. The wind is free."

—Peter Fairley

"It's just mind-blowingly awesome"

CEO Elon Musk, following the successful launch, orbit, and splashdown of Space Exploration Technologies Corp.'s Dragon space capsule. It was the first time a commercial company had performed such a feat, paving the way for commercial orbital flights.

Bionic Pancreas

Artificial organ could improve control over diabetes

When Pantelis Georgiou and his fellow biomedical engineers at Imperial College London decided to design an intelligent insulin pump for diabetes patients, they started at the source. "We asked ourselves, what does a pancreas do to control blood glucose?" Georgiou recalls.

The answer is pretty well known: The organ relies primarily on two populations of cells—beta cells, to secrete insulin when blood glucose is high, and alpha cells, which release a hormone called glucagon when glucose levels are low. "We simulated them both in microchip form," Georgiou says. This biomimetic approach diverges from today's dominant method of delivering only insulin using a relatively simple control system.

In a small clinical trial of 10 patients beginning this year, Georgiou and his colleagues Nick Oliver and Pau Herrero Viñas at the Centre for Bio-Inspired Technology will begin testing a device controlled by the microchip's cell models at two London hospitals. The device will infuse both insulin and glucagon by following a pattern that mimics the unique electrical characteristics of alpha and beta cells.

In a patient with type 1 diabetes, the body's immune system attacks and kills the insulin-secreting beta cells, causing an increase in blood glucose; as years pass, the glucagon-secreting alpha cells also tend to fail. So people with type 1 diabetes become prone to occasional—but dangerous—episodes of extremely low blood sugar, a condition that in a healthy body would be prevented by the alpha cells. Periods of extremely low blood sugar can severely affect patients' quality of life and, in the long term, even lead to damage to the heart, kidneys, and eyes. To minimize those complications, the Imperial College device attempts to model the electrophysiology of both types of cells to produce an artificial pancreas with greater fidelity to the real organ.

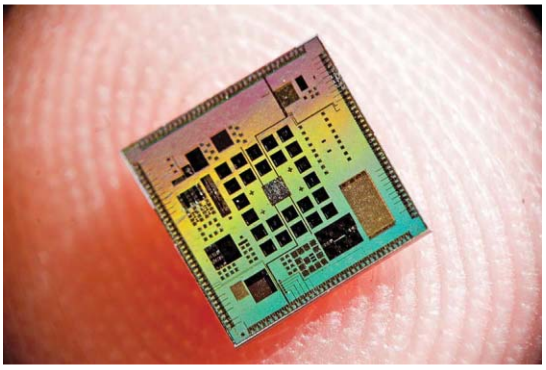

DIABETES DEVICE: A chip mimics the behavior of two types of pancreas cells to fight diabetes.
PHOTO: PANTELIS GEORGIOU

The device consists of an electrochemical glucose sensor that penetrates the skin, the microchip, and two small pumps worn on the body, one for each hormone. Every 5 minutes, the sensor detects the person's glucose level. If the sensor reports a high level of glucose, the silicon beta cell generates a signal that drives a motor. The motor pushes a syringe, dispensing insulin into the tissue beneath the skin until the glucose reading at the sensor drops and the beta cell goes silent. If the sensor reports a low glucose value, the microchip's simulated alpha cell activates the glucagon pump instead.

The microchip's control algorithms were designed to mimic the very different behaviors of the two cell populations. An alpha cell tends to react in spikes: When the concentration of glucose dips below a certain threshold, the electrical potential across the cell's membrane rapidly rises and then falls, releasing a discrete amount of glucagon. The liver then detects the glucagon and unlocks stored glucose in response. As the amount of glucose in the blood falls lower, the spikes become more rapid, releasing more glucagon.

A beta cell, by contrast, tends to react in bursts of voltage spikes punctuated by low-voltage silent periods that last for seconds or even minutes. When glucose concentrations rise, the beta cells remain in the high-voltage burst state longer, secreting more insulin as a result.

Another dual-pump design was reported last April by researchers from Boston University, Massachusetts General Hospital, and Harvard Medical School. In that trial, a similar device infused patients with both insulin and glucagon to show that the two-hormone method could prevent blood glucose concentrations from dropping too low, lending credence to the idea that glucagon could soon play a major role in diabetes management. This device, however, uses a predictive control scheme to determine the release of hormones, rather than copying the behavior of real cells.

If Imperial's preliminary trial is successful, the researchers will then perform more difficult testing on a greater number of subjects. —Sandra Upson

news brief

Squishy Memristors
Researchers at North Carolina State University have demonstrated "soft" versions of electronic components, including memristors. The squishy devices' electrodes are made of a liquid metal alloy, and the material that makes up the memristors is a chemically altered version of a common hydrogel. The scientists hope that such components will work better than traditional electronics to interface with wet squishy things, such as the human brain.
PHOTO: HYUNG-JUN KOO/JU-HEE SO

update

GIGATEPS AHEAD: The Intrepid supercomputer is at the top of the Graph 500 list.
PHOTO: ARGONNE NATIONAL LABORATORY

Better Benchmarking for Supercomputers

The usual yardstick is not a good metric

IMAGINE A WORLD in which a car's performance is judged solely by the time it takes to go from 0 to 100 kilometers per hour, ignoring fuel efficiency and other metrics. This, in essence, is the state of supercomputing today, says a group of U.S. computer scientists. People today typically judge supercomputers in terms of their raw number-crunching power, for example by asking how many linear algebra problems they can solve in a second. But, the scientists argue, the lion's share of challenging supercomputing problems in the 2010s requires quick and efficient processing of petabyte and exabyte-size data sets. And good number crunchers are sometimes bad exascale sifters.

It's time, the researchers say, for high-performance computers to be rated not just in petaflops (quadrillions of floating-point operations per second) but also in "gigateps" (billions of traversed edges per second).

An "edge" here is a connection between two data points. For instance, when you buy Michael Belfiore's *Department of Mad Scientists* from Amazon.com, one edge is the link in Amazon's computer system between your user record and the *Department of Mad Scientists* database entry. One necessary but CPU-intensive job Amazon continually does is to draw connections between edges that enable it to say that 4 percent of customers who bought Belfiore's book also bought Alex Abella's *Soldiers of Reason* and 3 percent bought John Edwards's *The Geeks of War*.

"What we're most interested in is being able to traverse the whole memory of the machine," says Richard Murphy, a senior researcher at Sandia National Laboratory, in Albuquerque, N.M. "There's no equivalent measure for these problems that's accepted industry-wide."

So Murphy and his colleagues from other U.S. national laboratories, academia, and industry have put together a benchmark they're calling the Graph 500. The name comes from the field of mathematics (graph theory) that the benchmark draws most heavily from. And the 500 is, Murphy says, an "aspirational" figure representing what they hope someday will be a "top 500" ratings list of the highest-performing supercomputers around the world, measured in gigateps instead of gigaflops.

The current biannual Top 500 supercomputers list recently made headlines when China's Tianhe-1A took the top position, coming in at 2.57 petaflops. The supercomputers on the list are ranked using a benchmark package of calculation speed tests called the High-Performance Linpack.

Crucially, Murphy says, the point of the Graph 500 is not to run a horse race on a new racetrack. Rather, he says, they've designed the benchmark to spur both researchers and industry toward mastering architectural problems of next-generation supercomputers. And the only way to know if you've solved those problems is for the industry to include those problems in its metrics.

€649 million Fines to be paid by five LCD panel makers in Taiwan and South Korea for conspiring to fix prices from 2001 to 2006. Samsung escaped the fines by ratting out its coconspirators.

In fact, by a Graph 500–type standard, supercomputers have actually been getting slower, says computer science and electrical engineering professor Peter Kogge of Notre Dame University. For the past 15 years, he says, every thousandfold increase in flops has brought with it a tenfold decrease in the memory accessible to each processor in each clock cycle. (For more on this problem, see Kogge's feature article in next month's issue of *IEEE Spectrum*.)

This means bigger and bigger supercomputers actually take longer and longer to access their memory. And for a problem like sifting through whole genomes or simulating the cerebral cortex, that means newer computers aren't always better.

"Big machines get embarrassingly bad gigateps results for their size," Kogge says.

Today only nine supercomputers have been rated in gigateps. The top machine, Argonne National Laboratory's IBM Blue Gene–based Intrepid, clocked in at 6.6 gigateps. But to score this high, Intrepid had to be scaled back to 20 percent of its normal size. (At full size, Intrepid ranks No. 13 on the conventional Top 500 list, at 0.46 petaflops.)

"I think Graph 500 is a far better measure for machines of the future than what we have now," Kogge says. Supercomputing, he says, needs benchmarks that measure performance across both memory and processing.

However, Jack Dongarra, professor of electrical engineering and computer science at the University of Tennessee and one of the developers of the Top 500 list, notes that the Graph 500 isn't the first new benchmark to challenge the High-Performance Linpack. The Defense Advanced Research Projects Agency, the U.S. Department of Energy, and the U.S. National Science Foundation have put forward a different group of benchmarks called the HPC Challenge, aimed at testing both computing power and widespread memory accessibility. Moreover, a coalition of industry partners—the Standard Performance Evaluation Corp.—has also assembled the SPEC set of computing benchmarks, aimed at better measuring the performance of more everyday components like Web servers.

Dongarra says that the Graph 500 may add to the list of metrics that rate a supercomputer's performance. But a Graph 500 score shouldn't be seen as some definitive number any more than the Linpack score used today. "If Graph 500 was the only benchmark we had, we'd criticize that too," he says.

—Mark Anderson

TOILET PAPER: MARK WRAGG/ISTOCKPHOTO; CIRCUIT: ISTOCKPHOTO

Electronics on Anything

Chemical trick puts solar cells and other electronics on rice paper, Saran wrap, and more practical things, too

There's probably not much call for printing solar cells on toilet paper, but a method developed at MIT can do just that, if it's ever needed.

More to the point, oxidative chemical vapor deposition (oCVD) could allow low-cost production of solar cells and other electronic devices on thin, flexible materials that other processes can't easily handle. Miles Barr, a graduate student in the lab of MIT chemical engineering professor Karen Gleason, described the process at the fall meeting of the Materials Research Society, in Boston.

The technique deposits conjugated polymers, plastics with good conductivity and semiconductor properties that are also flexible, stretchable, and even foldable. "We're particularly interested in polymers because of their good mechanical properties," Barr says.

update

The process sprays a vapor of a monomer and an oxidizing agent onto a substrate. When they meet on the surface, they polymerize, joining into long chains to form a plastic popularly known as PEDOT. Varying the surface temperature of the substrate between 20 °C and 100 °C dictates how the surface of the film forms; it can range from smooth to studded with nanopores. The polymer is conductive on its own, but lacing the nanopores with silver particles can increase conductivity up to a thousandfold. Barr says the process allows users to synthesize, deposit, and pattern the conjugated polymer all in one step.

To show off oCVD's abilities, Barr and his colleagues used the process on a number of extremely delicate materials. Rice paper, used to make spring rolls in restaurants, would dissolve in most processes, but because this one is free of solvents, it remained intact. A plastic film, such as Saran wrap—hard to coat because it repels water—could be coated with this dry process. The researchers even constructed a solar cell printed on toilet paper.

"This is kind of just to illustrate the versatility, not that these are substrates we necessarily want to process with electronics," Barr says. "You don't typically think of paper as a good substrate for photovoltaics, because it's not very transparent."

There may, however, be applications where the ability to build electronics, such as flexible displays, on fabrics or paper will come in handy. And engineers are increasingly looking to roll-to-roll printing—in which inks are printed onto plastic or another flexible material as it unspools from one machine and is wound up on another—as a faster, less costly method for producing some electronics, including photovoltaics.

The team built solar cells on a commonly used plastic and bent them to a radius of less than 5 millimeters more than 1000 times, then tested them to see if they still worked. Their efficiency was still greater than 99 percent of what it had been before bending, Barr said. Electrodes were bent to a radius of less than 1 mm, creased more than 100 times, and stretched to approximately 200 percent and still maintained high conductivity. A solar cell built on Saran wrap performed well even while it was stretched to about 180 percent, at which point the wrap pulled apart, destroying the cell.

To illustrate the point, Barr printed a solar cell on a piece of paper. In a video he showed at the conference, a student folded the paper into the shape of an airplane, attached leads, and shone a light on the folded device. It still generated current.

"I don't know if paper airplanes are the future of solar cells," Barr concedes. But in case they are, he's got it covered.

—Neil Savage

Chip Champs

THE IEEE's International Solid-State Circuits Conference, which will convene next month in San Francisco, gives the world's top circuit designers their annual opportunity to one-up each other. Who'll walk away with bragging rights? Here's a sample: —WILLIE D. JONES

FASTEST!

IBM is presenting its zEnterprise 196 chip, which it claims is the fastest commercial microprocessor on the planet. The chip, built using IBM's 45-nanometer, silicon-on-insulator process, is the first to break the 5-gigahertz barrier. To make the best use of such high frequencies, IBM engineers had to come up with ways to improve the chip's consumption and distribution of power. The processor also includes a rarity in CPUs: a 30-megabyte high-speed dynamic RAM cache.

MOST CORES!

Researchers at Intel's Bangalore branch are obviously firm believers in the division of labor. They are set to introduce a Xeon processor with ten *x*86 cores, the most ever integrated onto a single server chip. The new Xeon, which is manufactured with a 32-nm complementary metal-oxide-semiconductor process, incorporates several design elements aimed at squeezing more calculations out of every watt of power spent. These include the ability to individually control the flow of power to each core, and a clock scheme for the cache and other shared components that lowers power consumption during idle periods.

MOST ENERGY EFFICIENT!

The Chinese Academy of Sciences in Beijing is showcasing a chip that gets it all done on an austere power budget. The Godson3B processor's 582 million transistors are capable of carrying out 128 billion floating-point operations per second while consuming just 40 watts. That translates to the best energy efficiency for a high-performance computing chip: 3.2 billion floating-point operations per watt.

MOST TRANSISTORS!

How many transistors can you fit on a single die? Well, now we know that it's at least 3.1 billion. That's how many are on Intel's 18- by 30-millimeter Poulson

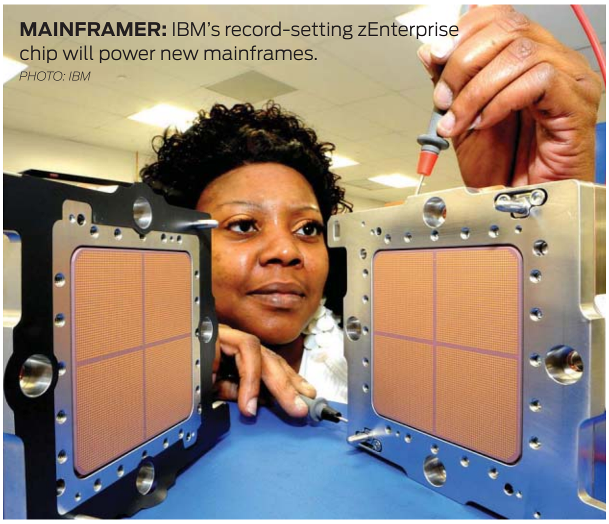

MAINFRAMER: IBM's record-setting zEnterprise chip will power new mainframes.
PHOTO: IBM

Itanium processor. They're arranged into eight processor cores linked by nine layers of copper interconnects. The chip also features four shared caches and high-speed links that raise peak processor-to-processor bandwidth to 128 gigabytes per second.

MOST COMPLICATED!

What would a computer-chip Olympics be without a head-to-head competition between Intel and Advanced Micro Devices? Each chipmaker is presenting a graphics processor making claims to the title of most disparate elements integrated on a single die. Intel's Sandy Bridge processor boasts four *x*86 cores, an optimized graphics processing unit, a dual-channel memory controller, and a 20-lane bus interface, as well as a shared 8-MB cache. AMD's entrant, called Zacate, has 450 million transistors arranged into two *x*86 cores (each with a 512-kilobyte cache), a graphics and multimedia engine, and a high-bandwidth interface.

profiles

DIGITAL DISCO: Dick Burwen's massive audio system sounds best with hundreds of tweeters tweeting—and three disco balls slowly spinning and reflecting their glittery beams.
PHOTO: MARK ANDERSON

DIGITAL MUSIC RENAISSANCE

Dick Burwen's algorithms may bring back analog's rich, complex sounds

NEIL YOUNG has called the digital recording era the "dark ages" of music. Bob Dylan, in a 2006 *Rolling Stone* interview, called music from CDs "static" and "small." Just a couple of grumpy old rockers who can't get with the program—or visionaries who see that the times, they are a-changin'? A retired engineer's basement in Lexington, Mass., is the unlikely wellspring for some technology that could hold the answer.

Richard Burwen, designer of everything from stereo sound chips to the Pioneer spacecraft's magnetometers, has spent nearly 50 years building and tweaking his own 20 000-watt ultrahigh-end hi-fi system. And some of the tricks and algorithms he's developed could someday make your CDs and digital music files sound better than you ever thought they could.

In 1962, Burwen began designing his house around what has become the US $500 000 Burwen Sound Studio—"the ultimate man-cave," as one scribe called it. The rear third of it hosts racks of computers, control panels, and audio components of varying vintage, from 1960s dial-and-needle boxes to modern-day laptops. The rest is an enormous resonant chamber with an Alice-in-Wonderland feel: Because no surface is parallel to any other, there are few troublesome echoes to limit the frequency response of the room.

Three recessed chambers at the far front provide cavities inside which Burwen's homemade left, middle, and right speaker systems reside—150 tweeters, five midrange horns, ten 16-inch woofers, and four 24-inch subwoofers. In front of each speaker Burwen has hung a mirrored disco ball and set a snare drum on the floor to add to the reverberations, which he says make the sound "more musical."

The layout of the studio was largely completed in the mid-1970s. Then came CDs. Burwen became increasingly frustrated at the decreased sound quality of this supposedly superior format. CDs are, he says, "rather screechy to me."

Rob Fraboni, who has produced records for Bob Dylan, The Band, Eric Clapton, The Rolling Stones, The Beach Boys, Joe Cocker, and Bonnie Raitt, says the CD format went to market before it was really ready. The problem, Fraboni says, is the sampling rate: 44.1 kilohertz for CDs. But the sound of the human voice is at its richest and most complex around 1200 Hz. In that range, the CD format leaves just 36 samples to describe a whole waveform. Imagine using just 36 connect-the-dot points to outline the *Mona Lisa*—how beautiful could she possibly be? Worse, MP3s and other compressed formats are just approximations of the CD standard, further distorting and simplifying the image.

But Fraboni says Burwen's CD-remastering software—initially just for his home studio—has played a key role in resuscitating digital music. (For the record, Fraboni was also a paid consultant to Burwen from 2005 to 2008; he's now working on his own separate music-remastering suites that he says leave less of a sonic footprint by not altering the original music quite as much.)

Burwen's CD-remastering algorithm introduces tiny bits of reverb at higher frequencies, where the sample rate starts to get patchy. The ear and brain read these tiny echoes as more organic connections between sample points and, as a result, Burwen's reconstructions sound more natural.

Burwen's entire software suite, a Microsoft Excel workbook containing 1.4 million formulas, is called Audio Splendor and costs $14 000. A home version, Burwen Bobcat, which works only with Windows Media Player, is tentatively priced at $3300.

All in all, says Fraboni, "Dick has made a tremendous contribution toward making music listenable again." If Burwen's algorithms can climb out of his basement cave, perhaps music can emerge from its dark ages.

—MARK ANDERSON

GAME DESIGN: SOMETIMES IT IS ROCKET SCIENCE

Sony executive John Blakely's prior career as a space engineer is more useful than you'd imagine

TWO HUNDRED gaming fans crowded into the Sony Online Entertainment panel at the San Diego Comic-Con last summer to see Luke Skywalker himself—Mark Hamill. What they didn't know was that the panel included someone who had truly harnessed the force of outer space—Sony executive and former space engineer John Blakely.

Blakely was there to unveil *DC Universe Online*, which goes on sale this month. The game—the PlayStation 3's first massively multiplayer online role-playing game (MMORPG in gaming parlance)—allows players to create their own avatars to match wits with Superman, Batman, and other characters from DC Comics. (Hamill was on the panel as the voice of the Joker.)

DC Universe required a design that could imbue its digital images with the laws of physics for hundreds of thousands of players. That, in turn, required servers that could handle 10 times as many computations per second as Sony's most technically advanced MMORPG, *Everquest II*. Such jumps are crucial in a US $8 billion global industry that's facing increasing competition from Asian game developers. The design work took five years.

"We had to design a server that would support the number of transactions needed to keep track of moving objects and how they interact with each player, and communicate that over the network," says Blakely. "That was one of the big pieces of core technology that we had to develop. If I'm in L.A., I could pick up a car in the game world, throw it, and hit a player in New York with it, and he would see the car coming. [That kind of] physics system that's shared by hundreds of thousands of people simultaneously had not been done before."

Blakely credits his gaming acumen to his space career, which included a five-year stint building payloads and embedded systems as a computer programmer and research scientist, first at the University of Alabama's Consortium for Materials Development in Space, and then at Teledyne Brown Engineering, both in Huntsville, Ala. Blakely helped to build the 3-D microgravity accelerometer, a data acquisition system used on the cargo-carrying Spacehab, Spacelab, and rocket missions. He designed software and motherboards and helped flight-qualify the first Macintosh computer in space.

Blakely first combined his interests in space and gaming during a college engineering co-op year at the NASA Johnson Space Center in Houston, a homecoming of sorts. As a child he would visit his engineer father at work there and crawl around the centrifugal force machine the astronauts used for G-force training. He won Co-Op of the Year for designing a space station solar-array simulator. "I used my experience writing gaming code for fun on the side, to dress up the graphics, which didn't exist back then," he says.

Blakely says that both a payload and a gaming console involve embedded systems with limited resources, such as memory. "You have to get everything working perfectly," he says. "You don't want the console game to crash. And in space, you can't reach up there to pull the plug or press the reset button."

Both space simulators and gaming projects require interactive teams of specialized engineers and programmers, and both require physics systems—though not in quite the same way. "Real physics isn't as much fun as entertainment physics," says Blakely. "If you model a game using real physics, it looks as exciting as watching grass grow. Gaming deals with extreme, accentuated physics. You want a superhero hitting the ground with a sense of weight, when real physics would have him shatter. You understand the principles, then throw them away to make them fun." —SUSAN KARLIN

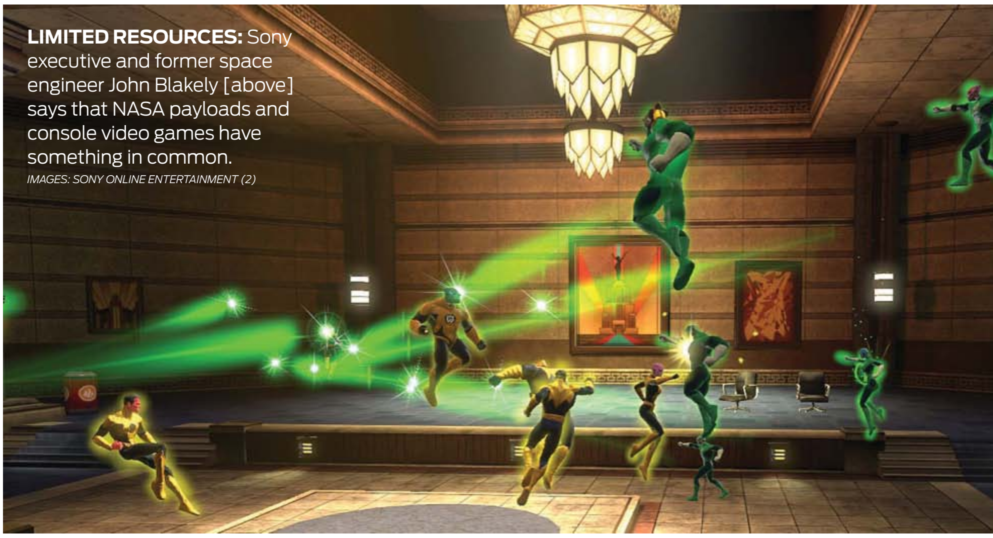

LIMITED RESOURCES: Sony executive and former space engineer John Blakely [above] says that NASA payloads and console video games have something in common.
IMAGES: SONY ONLINE ENTERTAINMENT (2)

tools & toys

THE LIGHTBULB THAT REALLY IS A BETTER IDEA

LED bulbs change the lighting equation

FIVE YEARS AGO they were in the lab; now you can buy LED lightbulbs at a hardware store. Should you? They produce as much light as incandescent bulbs for less than a fifth the electricity and heat, they last up to 20 years, and they fit in standard sockets.

Even more important, today's models—unlike previous generations of superbright light-emitting diodes—produce a light that is natural enough to satisfy most incandescent buyers. Compact fluorescents, even in their warmer incarnations, produce spectra with a handful of sharp peaks. The spectrum of a warm-white LED, by contrast, is relatively smooth, much more like that of a glowing filament. LEDs also turn on instantly, with constant brightness, unlike CFLs. What's more, they beat CFLs where CFLs beat incandescents, by lasting even longer and saving you even more on your electricity bill. And, of course, an LED bulb looks much more like a regular lightbulb than does the CFL corkscrew.

Ironically, although LED bulbs produce far less heat than incandescents, dissipating what heat they do create is the primary factor limiting their brightness. LED bulbs, like other electronic equipment, must remain below 90 °C to get certification from Underwriters Laboratories. The drivers, the LED junction, and the phosphors might be able to withstand higher temperatures, but unless UL changes its rules, that's irrelevant.

That means the maximum heat budget for an LED, once it's stuffed into an A19 envelope (the iconic incandescent bulb shape that's been produced in the billions over the decades), is between 8 and 14 watts, says Ray Chock of Philips Lumileds Lighting Co. The high end of that range, however, requires more heat-sinking than is practical in consumer lighting. Current warm-white LED bulbs top out somewhere under 100 lumens per watt, compared with about 15 for incandescents, meaning the brightest current LED bulb is roughly equivalent to a 60-watt incandescent bulb.

By the end of 2012, lighting engineers expect LEDs to get 150 lumens per watt. But this number is complicated by the strong trade-off between color temperature and lighting efficiency. Color temperature is a measure of the spectral distribution of a light source; for incandescent bulbs it corresponds roughly to the surface temperature of the heated filament. Ordinary incandescent bulbs typically operate at about 2700 kelvin and produce the familiar yellow-white glow, while halogens, whose complex chemical process lets filaments run hotter and bluer, typically operate at about 3200 K.

For LEDs, things work a bit differently. The semiconductor itself produces a much bluer light than most consumers—especially those in the United States—want. So manufacturers lower the color temperature by painting on phosphors that absorb much of the blue and re-emit it as yellow and red light. But as more long-wavelength phosphors get added, more energy gets lost in the conversion process. Consider for example, LED manufacturer Cree's XLamp MX-3 bulb. It comes in a 6500-K version that produces roughly 94 lumens per watt, while a 3500-K version (only slightly bluer than halogen) emits a mere 77. So exactly how much efficiency you get depends on whether you want to feel like you're next to a cozy fireplace or under a cloudless arctic sky.

GETTY IMAGES

A Worthy Investment?

	SYLVANIA ULTRA LED	PHILIPS AMBIENT LED	SYLVANIA SOFT WHITE INCANDESCENT	SYLVANIA MICRO CFL
Wattage or wattage equivalent	40	25	40	60
Cost per bulb at 25 000 hours of use	US $20	$28	$0.75	$2
Number of bulbs	1	1	17	2–3
Cost of bulbs	$20	$28	$12.75	$4–$6
Kilowatt-hours needed	175	125	1000	325
Energy cost (at average U.S. cost of $0.10/kWh)	$17.50	$12.50	$100	$32.50
TOTAL COST	$37.50	$40.50	$112.75	$36.50-$38.50

Note: Prices are approximate and may have changed since retail visits were made in mid-2010. Lamps are not precisely equivalent but are representative of available stock at that time. The incandescent and CFL bulb prices are per bulb, bought in packs of four.

But enough of aesthetics. The real question about these new lights is economic: At US $20 to $70 a bulb, are LEDs worth the investment? Compared with incandescents, they certainly are. A $30 LED that's equivalent to a 40-watt incandescent will save 800 kilowatt-hours or about $80, over its rated 25 000-hour life, plus another $10 or so in replacement bulbs, plus whatever your time is worth to do the replacing. Compared with CFLs, though, LEDs aren't as economical: You'd save maybe 50 kWh and no more than $10 for replacements [see table, "A Worthy Investment?"]. The upshot is a quicker adoption of CFLs and a high-tech catch-22—if more people would acquiesce to the steep initial cost of LED bulbs, economies of scale would kick in, and prices would come down.

In truth, the economic analysis is a bit more complicated than that, because one of the premises of the payback calculation is that 25 000-hour life span. That's 10 or 20 human years—and not that many people or even businesses stay in the same location that long. In effect, LEDs are a capital investment, while incandescents and CFLs are consumables.

And these economics feed back into the technology in an interesting way. You might think that A19 envelope—and being able to fit your bulb into the billions of standard light sockets around the globe—would be a holy grail for LED developers. Not so. Much of their work involves ways to bypass standard sockets completely.

By making fixtures specifically designed to dissipate LED heat, engineers can pack more lumens into their bulbs. Consider the standard recessed ceiling fixture. For a retrofit LED bulb, it's the worst possible application because there's hardly any air circulation to carry away the heat. On the other hand, an LED-specific fixture can have heat-dissipating fins that extend invisibly for as much as several feet in each direction from the light source, something that Chock says Philips is working on. Custom designs for track lights and floor, table, and desk lamps can bring similar advantages.

So even as you try to figure out which of your current fixtures might be suitable for LED bulbs, expect them to be far more commonly seen at new stores, restaurants, and offices. —Paul Wallich

books

RECENTLY REVIEWED ONLINE
Go to http://spectrum.ieee.org/reviews

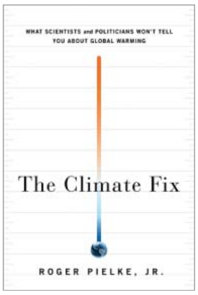

The Climate Fix: What Scientists and Politicians Won't Tell You About Global Warming
By Roger Pielke Jr.;
Basic Books, 2010;
288 pp.; US $26.00;
ISBN: 978-0-4650-2052-2
Reviewed by David Levitan

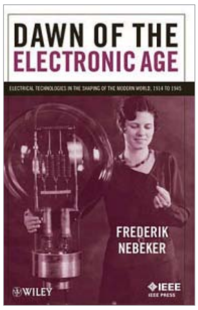

Dawn of the Electronic Age: Electrical Technologies in the Shaping of the Modern World, 1914 to 1945
By Frederik Nebeker;
Wiley–IEEE Press, 2009;
536 pp. (paperback); $57.95;
ISBN: 978-0-470-26065-4
Reviewed by Kieron Murphy

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

tools & toys

Fujifilm FinePix Real 3D W3
US $500; http://www.fujifilm.com

Your Next Camera Will Shoot 3-D

...or at least the one after that

I'VE ALWAYS THOUGHT that 3-D gadgets were like Twitter—impressive and charming, but rarely worth the fuss. In film, Hollywood's garish three-dimensional visuals fail to disguise its typically two-dimensional characters and one-dimensional plots, and I can't be the only person who balks at the idea of donning clunky glasses to watch TV at home. So I approached Fujifilm's new FinePix Real 3D W3 digital camera full of doubt.

I was wrong—this is one 3-D gadget that makes sense. It helps that Fujifilm has ditched the glasses.

Like other cameras with twin lenses—which have been around, in one form or another, for over a century—the W3 combines two separate photographs into a single 3-D view. The W3's "autostereo" screen, however, shows 3-D images instantly, without the need for special spectacles, as lenticular images: A series of tiny lens elements above the LCD directs light to each eye. The W3 is also the first consumer camera to shoot high-definition (720p) movies in 3-D.

The effect was immediately impressive, if disconcerting. The moment I looked into the generous 9-centimeter (3.5-inch) LCD, subjects in the foreground were clearly distinct from the background, which receded into the camera as though I were peering into a tiny puppet theater. However, objects close to the camera had ghostly double outlines that refused to line up, and both the 3-D effect and the screen's overall brightness fluctuated as I moved my head, or the camera, from side to side.

I could adjust the display's level of "3-D–ness" (parallax) with a control lever on top of the camera, but the other problems are inherent to the W3's lenticular technology. Moreover, images are 3-D only horizontally. Turn the camera on its side to capture a portrait view and the effect disappears.

The camera's twin lenses and twin image sensors allow for some other interesting options. In Advanced 2-D mode, you can set the right and left lenses to work as separate cameras. You can vary the sensitivity, the white balance, or even the focal length for each one, allowing ambivalent shutterbugs to shoot simultaneously in wide angle with the right lens and in telephoto with the left.

Its twin lenses aside, the W3 handles like a run-of-the-mill compact from a few years ago. It's heavy (250 grams) and bulky (124 by 66 by 28 millimeters), focusing and playback are sluggish, and the 3x zoom is, by today's standards, slow and short. Images and video are bright and colorful but low on detail, and they suffer from smeary digital noise in low light. More annoyingly, with the right lens situated just below the shutter release, one of my fingers often crept into the shot.

Three-dimensional stills and movies looked fine on the camera, but options for viewing them elsewhere remain limited. A mini HDMI cable (about US $10) allows easy connection to 3-D–ready televisions—but then you're back to wearing glasses. Fujifilm's own Real 3D V1 viewer ($500), has a sharp 8-inch autostereo screen but can't show movies. You can also upload 3-D images to Fujifilm's awkward SeeHere service and receive 5- by 7-inch lenticular 3-D prints through the mail two weeks later, at a pricey $7 each. Fortunately, the W3 saves a standard 10-megapixel JPEG image alongside each 3-D file.

Despite the difficulties of sharing images and the limitations of both screen and camera, the W3 does have a whiff of revolution about it. Even the dreariest holiday snaps and the most pedestrian home movies acquire a naturalistic luster in 3-D, and I defy anyone to use it in public without gathering a crowd of excited passersby. After spending a week in its company, I will find it difficult to return to shooting in a paltry two dimensions. Just don't ask me to tweet about it.

—MARK HARRIS

reflections

BY ROBERT W. LUCKY

Clickphobia

BECAUSE OF a recent experience, I have a raging case of clickphobia—the fear of clicking on Internet links. When I encounter a promising-looking link, my trigger finger freezes.

This started a few weeks ago when I got an alarming e-mail from the security department of my e-mail service provider. It said that they were closing inactive accounts, and unless I verified my account by return e-mail, it would be terminated. I was scared to lose my account and typed my information in, but I was suspicious, and I hesitated before clicking the send button. There was a misspelled word, and the graphic of the company logo was a bit fuzzy. On the other hand, losing my e-mail would be disastrous, and the return address appeared to be legitimate. I asked myself: What harm could result from someone having access to my e-mail, and how quickly could a hacker act before I changed my password?

While I was debating the issue, my index finger twitched and the button was clicked. I visualized the packets speeding around the world, and just for an instant I imagined an "unclick" key that would put out an all-points bulletin to stop them at bordering routers. But it was a fait accompli. Even with all the king's horses and all the king's men, the thoughtless click could not be undone.

Well, I commiserated with myself, what could go wrong? Alas, it was only a matter of hours before I discovered the answers to my two questions about harm and speed—a lot, and fast. I was away on a trip, and friends that I encountered gleefully gave me the bad news; meanwhile, my home phone was ringing off the hook. Everyone—and I mean absolutely everyone I knew—was getting e-mail from me pleading for 2000 euros to be wired to London, where I was apparently stranded and broke.

Frantically, I logged into my e-mail, only to be denied access. As I later discovered, the hacker was busy corresponding with my friends in my name to further plead his case (in my imagination, the hacker was definitely a man). Fortunately, it didn't sound right to ask Americans, who use dollars, to wire euros to London, where they use pounds. And even though the hacker presumably knew his way around computers, his English wasn't up to snuff. One friend showed me the e-mail the hacker sent her when she replied to the original message with, "Bob, is this really you?" The reply was, "Yes, it really are me."

Another friend, having gotten the scam message, did some tricks with it and determined that the hacker was in Lagos, Nigeria. We've all received Nigerian scam e-mails for years, so this

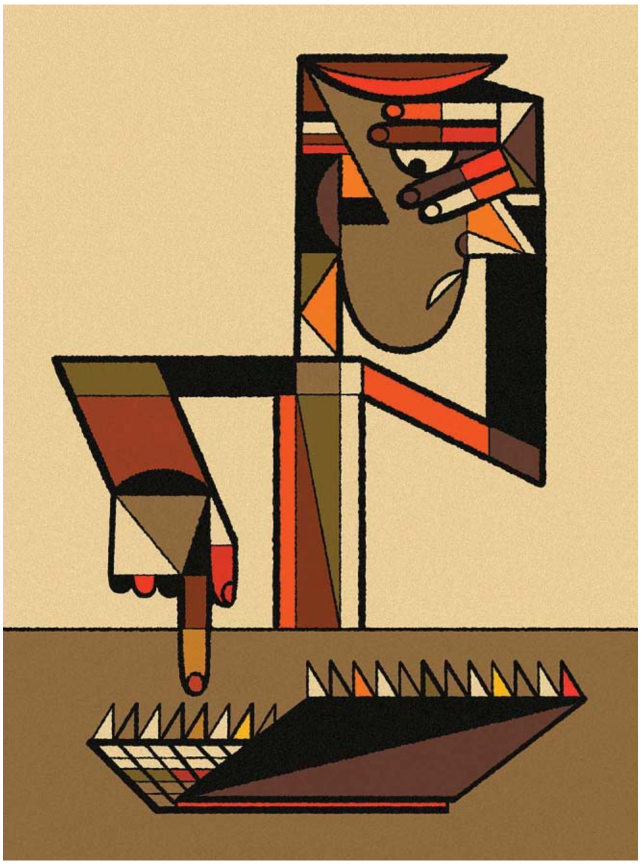

was unsurprising, and in fact the wording and syntax of this message were very similar. It's amazing that after all these years, they—whoever they are—continue to persevere with unlikely, poorly worded scams. Someone must fall for them, but please—just not one of my friends!

Of course, I took a lot of ribbing, especially from friends with whom I have worked on cybersecurity issues. I felt, and still feel, like a complete dunce. Nevertheless, these phishing attacks are insidious. You get a sales confirmation for something you didn't buy, so you click on the link on the bottom that says "report a problem." Big mistake. You get an e-mail from your credit card company saying that there has been suspicious activity on your account, and they need to verify that it is you. Or you hear from your bank about an overdraft.

This all leaves me very mad and feeling helpless. We engineers brought the world closer together with a beautiful network, and a few people are tarnishing it for everyone else. I'm a believer in the wisdom of the crowd, but the corollary is the tyranny of the crowd.

I'm not sure what to do other than accept the risk and live with it. To that end, I've been practicing with the computer turned off. Soon I will turn the computer on, meander over to my e-mail, and see if I've conquered my clickphobia. ❑

GREG MABLY

ORIGIN® 8.5

Data Analysis and Graphing Software

Try out the Origin 8.5 demo today. Call or chat online with us to discuss features, license options and pricing.

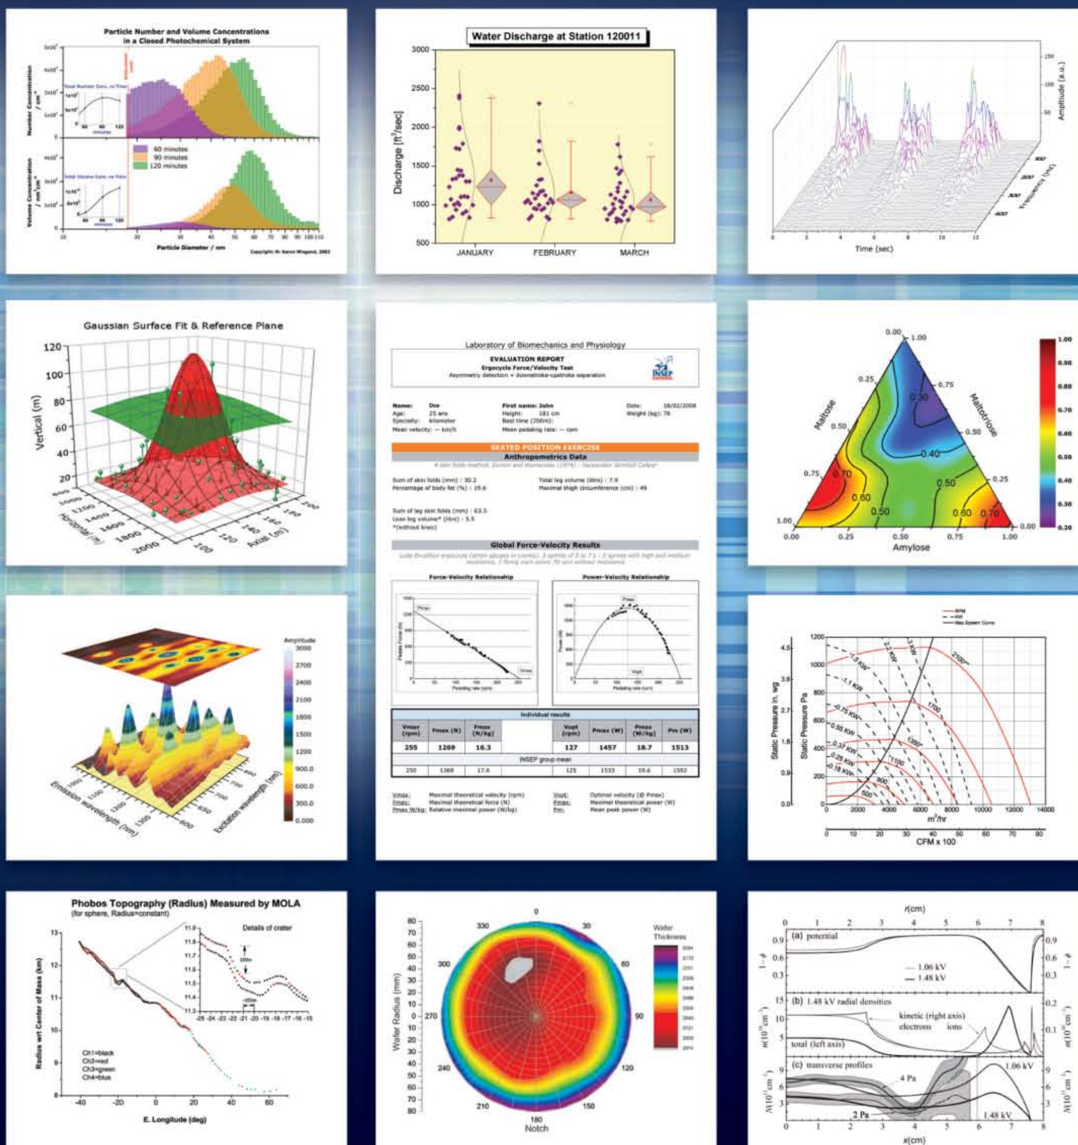

OriginLab

OriginLab Corporation
One Roundhouse Plaza
Northampton, MA 01060 USA

USA: 800-969-7720
FAX: 413-585-0126
EMAIL: sales@originlab.com
WEB: www.originlab.com

New features in Origin 8.5 include...

- Transparency and Gradient Fill
- Embed Equations, OLE Objects
- 3D Intersecting Surfaces
- 3D Vectors and Error Bars
- Waterfall Y or Z Color Mapping
- Improved Label Control
- New Gadgets
- Fitting Function Wizard
- Virtual Matrix from Worksheet
- Thumbnail View in Matrices
- Improved Image Profiling
- Zoom and Pan Support

SMARTPHONES

SOCIAL NETWORKING

VOICE OVER IP

LED LIGHTING

MULTICORE CPUs

CLOUD COMPUTING

DRONE AIRCRAFT

PLANETARY ROVERS

FLEXIBLE AC TRANSMISSION

DIGITAL PHOTOGRAPHY

CLASS-D AUDIO

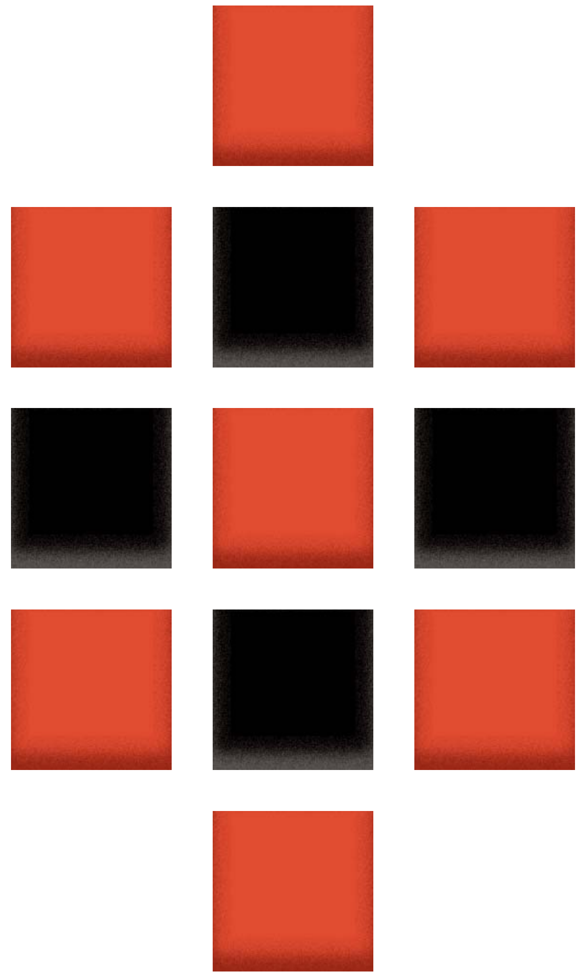

TOP 11 TECHNOLOGIES OF THE DECADE

The most powerful technologies take a while to mature. But when they do, they can rapidly retire mainstays that are decades old

LIFE WAS DIFFERENT A DECADE AGO. Your phone couldn't contain your entire music collection, for example, or guide you to a restaurant in a foreign city. Bomber-reconnaissance planes invariably had pilots on board. And how's this for quaint: Your corner drugstore still stocked photographic film! The technology waves that washed away those realities spread from tremors that occurred years before: The first smartphone was unveiled by IBM in 1993, the first digital photo was taken in 1975, and the first drone aircraft flew during World War II. Clearly, the seeds of the next crop of technology staples have already been planted. Perhaps the first tender shoots can already be discerned among the pages of this issue. *—Philip E. Ross*

NO. 1

Smartphones

IS YOUR PHONE SMARTER THAN A FIFTH GRADER?

Yes

DOUGLAS ADAMS'S *Hitchhiker's Guide to the Galaxy* series is named after a pocketable device that contains everything worth knowing. But that seems almost quaint today, when you can carry the full contents of the Web in your pocket, as well as a telephone, a camera, a radio, a television, and a navigation system. Today's smartphones are marvels of engineering, crammed with more features than the average PC. They've become the prime driver of innovation for both software and hardware.

It took half a century to shrink the mainframe from the size of a living room to that of a suitcase. It took another decade to make it smaller than a wallet. The smartphone has swallowed and assimilated functionality from music players, remote controls, gaming consoles, even printed maps and news publications. And now that smartphones are serving as Wi-Fi hot spots, they can replace wireless routers and modems, too. Smartphones are becoming as essential as keys or a wallet, and they'll soon replace those as well [see "Company to Watch"].

This has some real consequences. Unlike its predecessors, the smartphone is an inherently personal device: Not only is it always on, it's always somewhere on us. Without realizing it, we've let smartphones usher us into an age of ubiquitous, pervasive computing that technologists, as well as science-fiction authors, have been dreaming about for years.

"Smartphones help users stay connected to information at any given time, any given location," says Dilip Krishnaswamy, a Qualcomm engineer and associate editor in chief of *IEEE Wireless Communications*. "The information is just there when you need it."

We've come to rely on such connectivity. There's no need to pack a map or directions when an app can guide you in real time, nor to consult a restaurant guide before leaving the house. In these and a thousand other ways, the smartphone, more than any other technology to have emerged in the past decade, is the one that has most changed our lives.

To be sure, back in 1973, Motorola's Martin Cooper didn't set out to build an always-connected, portable computing device. He was simply trying to shrink the car phone down to the size and weight of a luggable brick. But once the cellphone had earned a permanent place in our pockets, it became an unavoidable platform for innovation, upstaging the PC. If Starbucks wants to make it quicker and easier to pay for a cup of coffee, why not do it through the phone? If *The New York Times* wants to get away from paper, well, everyone's already carrying around a perfectly readable screen.

Smartphones are more than just bells and whistles—they actually change behavior. With a traditional mobile phone, users spend most of their time making calls and sending text messages. On a smartphone, basic communication takes a back seat to Internet browsing, e-mail, entertainment, and games. This difference is due to three key ingredients, each of which has seen tremendous advances in the last decade: hardware, software, and network infrastructure.

The hardware is the most obvious. Thanks to high-resolution displays with touch screens or QWERTY keyboards and tiny camera lenses on the outside and gigahertz processors, radio antennas, and image sensors on the inside, the phones hardly resemble their modest predecessors.

But at least as important is the software. "The operat-

NOT ALL SMART YET:

Projected smartphone penetration

PERCENTAGE OF ALL MOBILE PHONES

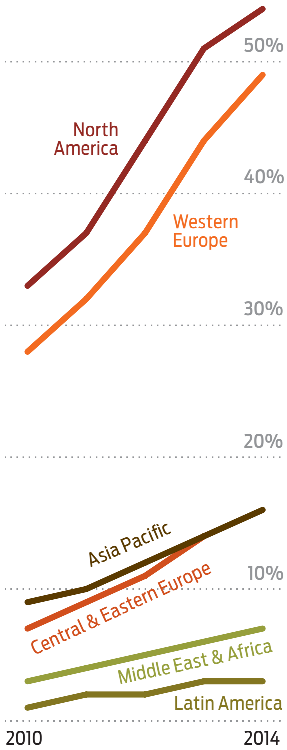

Smartphones are proliferating rapidly, but they still make up a minority of all mobile phones. Customers in North America have been especially quick to embrace them, but Italy still has the highest concentration in the world. The low percentage in Asia is due partially to the popularity of advanced feature phones (which have many capabilities but lack a true general-purpose operating system) in Japan and South Korea.

Source: Cisco Visual Networking Index, February 2010

FRANK CHIMERO

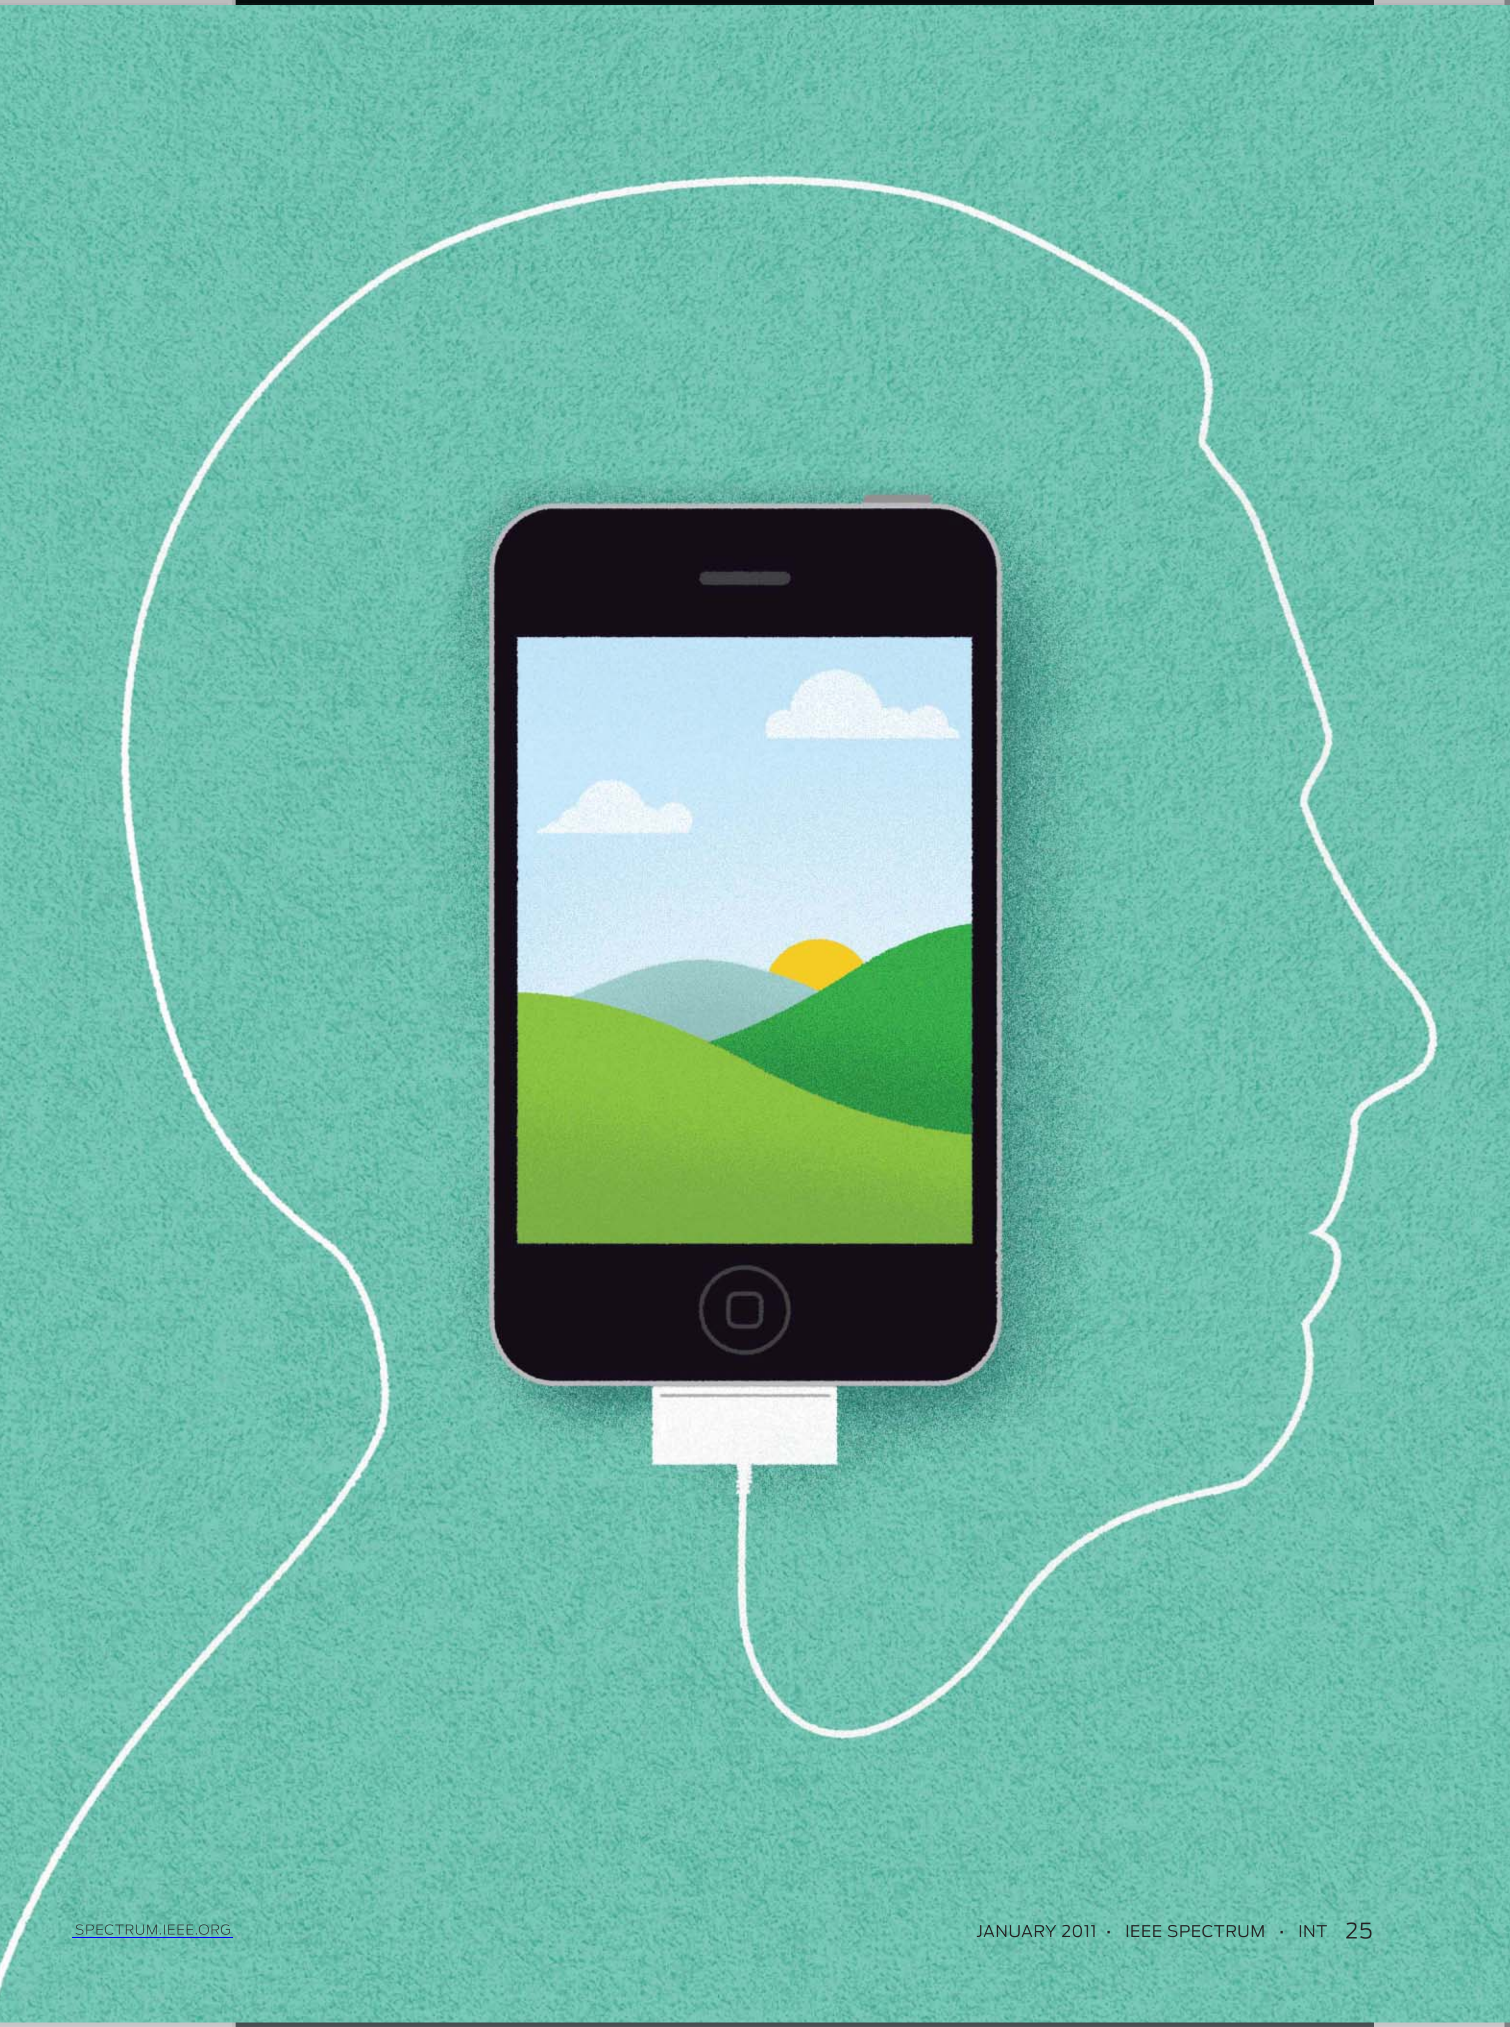

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

COMPANY TO WATCH: BROADCOM CORP., IRVINE, CALIF.
If you want your phone to replace your wallet and keys, it will need a near-field communications (NFC) chip. Broadcom already supplies companies like Apple with integrated Wi-Fi and Bluetooth chips, and thanks to its August 2010 purchase of UK-based Innovision, the company should soon be able to add NFC for less than US $1 per unit.

ing system is the foundation for everything else in a smartphone," says Donna Dubinsky, a cofounder of Palm, the company that first succeeded in cramming computer functionality into a pocketable device. Every major smartphone operating system now supports third-party applications that extend the phone's capabilities far beyond what any one manufacturer can do.

And in addition to using the resources of the phone itself, these apps can off-load data storage and processing to the cloud [see "It's Always Sunny in the Cloud," in this issue], in the form of server farms around the world, thanks to ever increasing wireless bandwidth—the third key development behind the smartphone. "People naturally want to focus on the device itself," says Dubinsky, "but what's important is the complete system, including hardware, software, and application development environment."

Current smartphones quietly shift between Wi-Fi and 3G so that users are always connected to the best available network. Remember that the iPhone, still less than four years old, didn't even access 3G networks when it launched. By October of last year, you could get a 3G signal at the top of Mount Everest, and now the first 4G networks are emerging.

Today we're seeing only hints of how a smartphone world will be different. With their numerous sensors, they will form nodes in a vast and unprecedented data collection network. Researchers have already used phones' accelerometers to follow basic health indicators (such as a patient's gait), their GPS to monitor crowd and pedestrian traffic patterns, and their microphones to track bird migrations. Several app developers have created the first useful examples of augmented reality—letting you point your phone at a restaurant and see a bunch of customer reviews, for example.

These capabilities come with strings attached, notably the addictive effects of always-on connectivity. BlackBerries are rightly nicknamed "CrackBerries" for the way they feed a workaholic's addiction. Krishnaswamy notes that we're training ourselves to always be ready for the next e-mail or status update, and we're disappointed when one doesn't arrive. And not everyone likes it when people interrupt dinner to surf the Web to fact-check the conversation.

Some experts even worry about a new digital

BET YOU DIDN'T KNOW

South Korea, which has long boasted the world's fastest data connections, saw average Internet connection speeds slow by 24 percent in 2009. Blame the iPhone: It made its debut on the country's slow wireless networks and was then widely adopted.

HOW PHONES GOT SMART: EIGHT MILESTONES

IBM SIMON PERSONAL COMMUNICATOR [1993]

IBM was ahead of the curve back in 1993. The Simon was a touch-screen phone with a calendar, address book, calculator, and even the capability to send and receive e-mails and faxes. It was impressive, but also bulky, and came with a hefty price tag: US $899.

NOKIA 9000 COMMUNICATOR [1996]

If you wanted a smartphone in the late 1990s, your choice was pretty much limited to the pricey Communicator. The device was like a personal digital assistant and a mobile phone sandwiched together. (It's rumored that an early prototype was actually a Hewlett-Packard PDA connected to a Nokia phone by a hinge.)

ERICSSON R380 [2000]

The R380 was one of the first phones actually marketed as a "smartphone." It was the first commercial phone to run the Symbian OS. Symbian has long been the most common smartphone operating system, but its market share has rapidly declined since it began facing competition from Apple's iOS and Google's Android.

SONY ERICSSON P800 [2002]

The P800 came with several new features that we consider standard for smartphones today: It could play MP3s, came with a camera, and featured a color touch screen (although it supported only 4 096 colors).

RIM BLACKBERRY 6210 [2003]

Research in Motion's early communicators offered two-way text paging, but the company soon realized that push e-mail service was its killer app. RIM developed a QWERTY keyboard for quick thumb typing and BlackBerry Enterprise Servers to tap into existing e-mail infrastructure. In addition to e-mail and basic Web browsing, the 6200 series were the first BlackBerries that were also fully integrated phones (earlier models required a headset).

PALM TREO 600 [2003]

The Treo combined some of the best features of the Palm Pilot PDAs with a mobile phone shape, creating a smartphone that had appeal beyond business users. With its Palm OS, the Treo supported lots of third-party apps. Palm OS also gave the Treo integrated functions, like the ability to dial directly from the contact list and check the calendar while on the phone.

APPLE IPHONE [2007]

Apple's much anticipated iPhone immediately obliterated the competition. Its multitouch, capacitive touch screen encouraged users to interact using their fingertips, and the mobile Safari browser made mobile Web browsing fun, not just possible. Apple also used its clout to upend the mobile ecosystem, wresting unprecedented control from carriers. The iPhone 3G upgraded the data connection and launched the wildly successful app store.

HTC EVO 4G [2010]

At launch, the EVO 4G was arguably the best Android smartphone around: It had an 8-megapixel camera that could shoot HD video, HDMI output, and a 4.3-inch touch screen. But its best trick was the ability to act as a mobile hot spot connected to Sprint's WiMax network, the fastest data network in the United States at the time.

SIMON: BCOS47/WIKIPEDIA; ALL OTHER PHONES: PROVIDED BY THE MANUFACTURERS divide between those who can afford smartphones and data plans and those who can afford only basic mobile phones. In fact, many smartphones cost more than low-end computers, once you take away the subsidized prices that wireless carriers offer for them. Yet in rural and impoverished areas, they represent a much better investment because they're self-contained, needing neither additional network infrastructure nor even reliable power.

In any case, many high-end features will inexorably filter down to low-end phones, as they have in the camera market, and what begins as a luxury will quickly become a necessity. In 2007, sales of smartphones surpassed sales of laptops, and some predict that by 2014 more people will browse the Internet by phone than from traditional computers.

The drive to communicate on ever-wider scales has shaped many of our technological advances, and these in turn have shaped how we communicate. Moving from text messages to Twitter updates, from voice to video chat seems to be part of our evolution. "The interesting thing is how it's changing human behavior itself," says Krishnaswamy. "Smartphones will become a sixth sense for the user, gathering information from wireless sensors in the user's environment and from the network, interpreting the information, and providing valuable feedback to the user."

—Joshua J. Romero

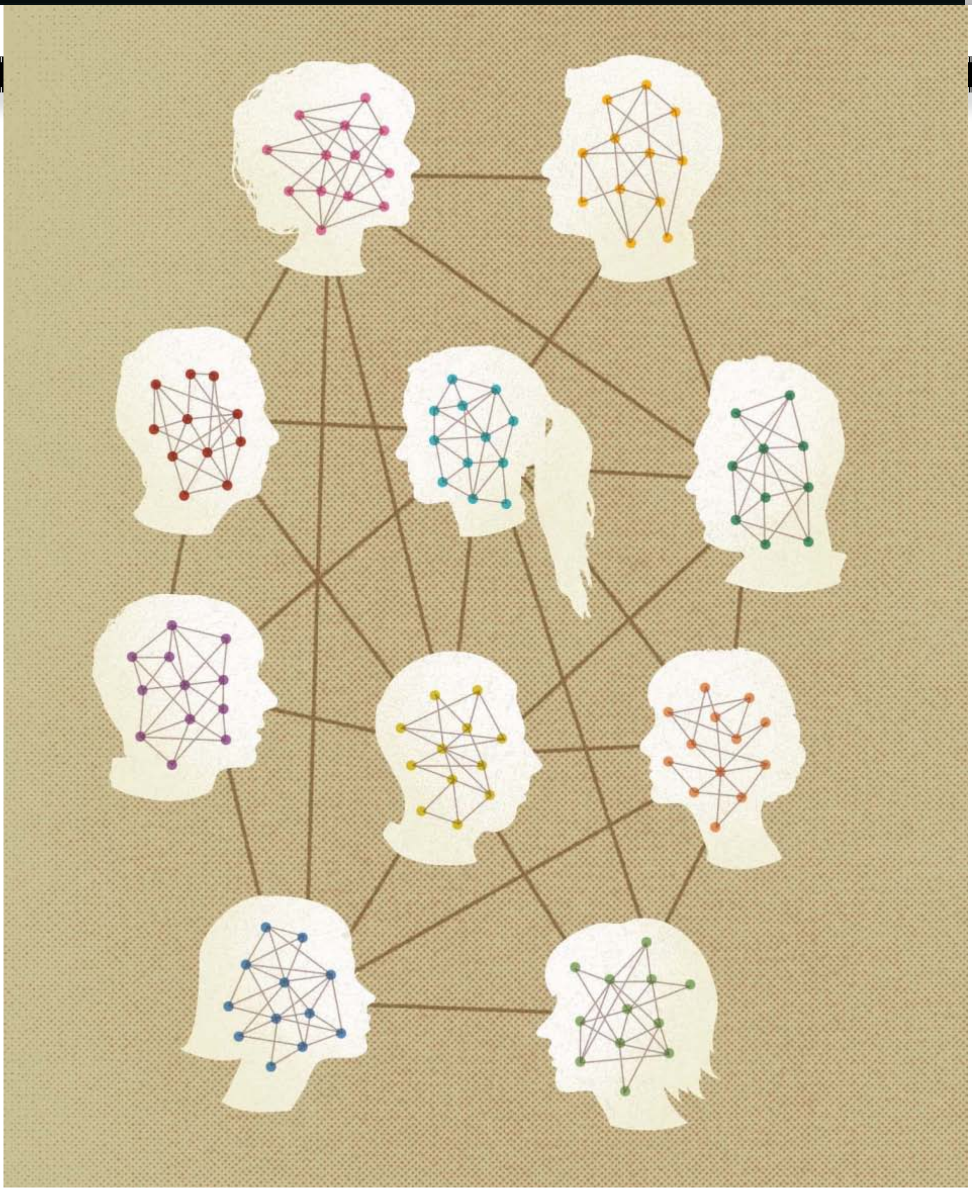

FRANK CHIMERO

NO. 2 Social Networking

FRIENDED

Bandwidth, digital cameras, and a hunger for connectedness have created a virtual dinner party

A decade ago, it might have taken a new person in town months to make contacts, find places to hang out, and meet like-minded people. Now, with a few clicks of the mouse, you can get the job done through social networking—a communications revolution that began in fits and starts in the late 1990s and reached recognizable form in March of 2003, with the public launch of Friendster.

"The idea was to have the Internet do the work of a dinner party," says Kent Lindstrom, a former Friendster CEO. A user could set up a profile, with personal facts and a picture, and invite friends to join. Friendster's servers

COMPANY TO WATCH: NING, PALO ALTO, CALIF.
Launched in 2005 by Gina Bianchini and Netscape cofounder Marc Andreessen, Ning is betting on small, exclusive networks. Already the company hosts 2 million customized social networks, including Shred or Die, for extreme skateboarders, and GovLoop, for anyone working in government.

would then generate a continually updated list of her friends as well as her friends' friends, mapping relationships out to four degrees of separation. Within nine months, Friendster had a million members.

Around that time, *Fortune* magazine prophesied, "There may be a new kind of Internet emerging—one more about connecting people to people than people to websites." Indeed, such an Internet has emerged, though, as so often is the case, the first movers have been left behind. The big winner has been Facebook, founded less than a year after Friendster went public. Facebook has 540 million users who spend about 700 billion minutes on the site every month; if it were a country, it'd be the third most populous in the world.

YOU CAN LIVE much of your life in that strange, virtual country. Matthias Galica, 26, founded ShareSquare in May 2010 to enable users to print bar codes to "geotag" objects in the real world so they could be followed in the virtual one. The day he moved the business into a loft in downtown Los Angeles, he got an e-mail message through his Facebook account from one of his more than 1400 "Facebook friends" inviting him to join the Facebook group DTLA.

He clicked the "Accept Invitation" icon, which connected him instantly to hundreds of like-minded Los Angelenos. Later that night, at a concert, he used his iPhone to log into Foursquare, a friend-locator service, which automatically posted the message "Matthias just checked-in @Shrine Auditorium" to his Facebook page, prompting an impromptu get-together with another member of DTLA.

"Connections that would've taken months now happen in a space of hours," Galica remarks.

This is exactly the kind of life promoted by Friendster, which was left in the dust by MySpace, which was lapped by Facebook. Yet even Friendster climbed up on the shoulders of still-earlier pioneers. The most notable was Sixdegrees, arguably the first true social network. Twenty-eight-year-old New York businessman Andrew Weinreich launched Sixdegrees at a party in 1997, announcing that "with the click of a button," the site would revolutionize human networking. Like Friendster, Sixdegrees let users identify their friends, their friends' friends, and so forth. At its peak in 1999, it had attracted 3.5 million members.

But each of these early actors came too early or learned a critical lesson too late to dominate. For Sixdegrees, the major flaw was the absence of photographs. "We had letters coming in all the time asking, 'If I mail in a picture, can you scan it for me and put my picture on your site?' " Weinreich recalls.

FRIENDSTER AND ITS rivals came around at just the right time. Point-and-shoot cameras were ubiquitous. Broadband was cheap and available. Users could be sure that the friends they invited to see their profiles and chat online had high-speed Internet connections and plenty of photos to share.

But Friendster made a single, master mistake: It took on customers faster than it could manage them. By the end of 2003, Friendster was acquiring some 9500 users a day, and the company was struggling to buy enough servers to keep up with the growth. Users complained that their home pages regularly took more than a minute to load.

Their frustration paved the way for MySpace, a Los Angeles start-up founded by hackers Chris DeWolfe and Tom Anderson. They did away with pesky friend

HONORABLE MENTION

E-paper

Paper or Plastic?

In 2020, newspapers will still be with us, but they won't be paper

Despite the increasing sophistication of electronic displays, with their staggering color palettes and expanding contrast levels, Gutenberg's 15th-century technology didn't begin to give way until the 2000s. That was when electronic paper made its debut in digital book readers like Amazon's Kindle and Barnes & Noble's Nook.

E-paper exploits a phenomenon called electrophoresis, discovered in 1807 and revived at Xerox PARC in the 1970s but put to practical use at MIT's Media Lab only in 1997. It uses a jolt of current to make the black dye inside thousands of microcapsules sandwiched between flexible polymer sheets rise to the top of the capsules so that the "ink" becomes visible through one of the sheets. A great advantage of this method is that it draws power only when updating the image on the screen. Recent models offer a contrast ratio similar to that of a newspaper.

Speaking of the dailies, they just may supplant books as the killer app for e-readers. Uploading a digital version of your hometown broadsheet on a plastic sheet that you can roll up and tuck in your bag would eliminate the cost of printing and distribution, and it'd also save a lot of trees.

—Willie D. Jones

MCKIBILLO

BET YOU DIDN'T KNOW

Social networks are proving to be gold mines for predicting human behavior. In a proof-of-concept study, scientists at HP Labs used movie chitchat on Twitter to accurately predict box office hits. And the police department of Richmond, Va., is now using network-analysis software to monitor Facebook, MySpace, and Twitter for evidence of crime incubators, such as rowdy parties.

calculations and instead let users cruise the entire public site, eavesdrop on strangers, create bizarre "fakester" profiles—notably those of Jesus and of the Burger King—and make their own connections. MySpace's meteoric success (72 million members in 70 countries by 2007) led to a boom of businesses that New York University social-software analyst Clay Shirky labeled "YASNS," as in "Yet Another Social Networking Service."

For instance, there's the messaging service Tencent QQ, which is popular in China; Google's Orkut, which is popular in Brazil; and Twitter, a blogging network that limits posts to 140 characters, which is popular just about everywhere. There's Flickr for photo sharing and YouTube for video sharing. There's LinkedIn for job networking and Classmates for finding long-lost school friends. And there's a whole host of niche networks: Coastr for beer aficionados, Goodreads for bookworms, ResearchGATE for scientists, and Dogster for dog lovers.

The reign of MySpace, however, didn't last long. By September 2009, the company could claim just 30 percent of the American social-networking market, down from 67 percent the year before, according to the research firm Hitwise. Its biggest mistake was losing control of the site's usability. Rather than develop tools that would help users organize the vast amounts of shared information, MySpace dumped all its resources into new features it thought would drive traffic: bulletin boards, job listings, horoscopes, even a YouTube-esque video-sharing service called MySpaceTV. The site got so cluttered that many users left to look for something simpler.

FACEBOOK CONQUERED in part because it took to heart the lessons of its predecessors' mistakes. For instance, its founder, Mark Zuckerberg, expanded from the company's base at Harvard by adding one university at a time, ensuring that no new customers would come online until the servers could handle the additional traffic.

Facebook also lured Internet users with its sleek, easy-to-use interface and engineering wizardry. One of its most innovative features was Multi-Feed, which searches your friends' databases for new updates and streams them to your home page as a continuous news feed. Facebook now contends with some 30 billion shared updates a month—a monumental processing feat that requires tens of thousands of servers.

When the feed was introduced, some veteran users derided it as an intrusive assault on their privacy. But most of them stayed to gawk, and many new users poured in because of it. Each new attempt to expand Facebook's reach has prompted renewed criticism that the company allows users too little control over their own data.

Rival sites have rushed to fill the possible privacy gap. Diaspora, a New York University student-built network, lets users program their own servers rather than rely on a centralized system. Ning, based in Palo Alto, Calif., offers a do-it-yourself social-networking platform for tight-knit communities. Next up could be Google Me, the search giant's most recent foray into the social-networking business. However, it seems that Facebook is here to stay.

"Their scale makes them formidable," says Weinreich, who believes that Zuckerberg has successfully constructed the ultimate "social graph"—a term for the web of relationships the 26-year-old billionaire popularized in May 2007, when he announced to a warehouse full of programmers in San Francisco that he was opening up Facebook's infrastructure to anyone wanting to run programs on it. Little more than three years later, Facebook's social graph has become a powerful business platform. Independent software developers have coded a cumulative 550 000 applications for Facebook users. Some are absurd. SuperPoke, for example, lets users virtually "kiss" or "spank" their friends. Others, particularly multiplayer games, such as Zynga's *Mafia Wars* and *FarmVille*, have cultivated multimillion-dollar businesses.

To spread Facebook's influence even further, its engineers created Facebook Connect, a software-to-software communication tool (an application programming interface, or API) that lets Facebook users log in to their iPhones or other Web sites—the news aggregator Digg, for example—with their Facebook identities.

Joe Stump, the chief technology officer of SimpleGeo, which helps location-based services collect and manage data connected to places, predicts that soon your Facebook identity will follow you wherever you go. You'll use it to decide, based on friends' recommendations, which book to buy or which doctor to see. It's likely, then, that geo-networking apps such as Foursquare and its closest competitor, Gowalla, will be the next big thing online.

Again, it will be the mass adoption of new consumer technologies—this time, smartphones and cloud computing—that will allow people to collect and compute vast amounts of data about places and things. The trend has already begun. "We're heading into a third phase of the Web," Stump says.

The webs between humans and machines, and between humans and the people they trust, have already been spun. The Web of the future will instead connect our physical world with our virtual one, enabling our online social interactions to give us valuable insights into our off-line lives.

—Ariel Bleicher

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

NO. 3

Voice Over IP

SETTING PHONE SERVICE FREE

How Ma Bell's cash cow became a free software app

FOR GENERATIONS it was possible to grow up and grow old without outgrowing the telephone of your youth. Handsets stayed tethered to the wall, "long distance" remained a concept, costly service was a given. Then in the 1990s, when competition and cellphones began to free us from wires and monopolistic pricing, we marveled at what seemed a revolution. We were wrong: The real revolution began only in the 2000s, when it became clear that the Internet would be the telephone network of the world.

Look at this one fact: Skype, which launched only in 2003, now has more than half a billion registered users, making it the largest provider of telephony services in history. It's good for Skype, but better still for its users, who when they talk Skype-to-Skype pay exactly nothing, no matter where they may be on the planet. The company charges only when you call a phone listed on a "real" telephone network.

It is at that junction, in the increasingly hazy no-man's-land between the telecom and Internet worlds, that the revolutionary insurrection is playing out.

EXACTLY WHEN Internet telephony began is unclear. By all accounts the guys behind ARPANET, the predecessor to the Internet, experimented with digital signal processing in the 1970s. But the main story begins, at the earliest, in 1989, when Alon Cohen and Lior Haramaty, two graduates from the Israeli Defense Forces telecommunications wing, founded VocalTec Communications, in Herzliya, Israel. Following on from what they'd seen in the military, where packetized voice was already being used to deliver sensitive orders and intelligence, the entrepreneurs experimented with commercial sound cards for desktop PCs. They became experts on digital sound at a time when most technologists were still squinting at green monochrome monitors; in fact, people used to ask Cohen and Haramaty why they weren't trying to deliver color rather than sound to the PC terminals of the day.

Meanwhile, they faced down potential disaster when a better-funded competitor, Creative Technology, based in Singapore, arrived on the scene with the Sound Blaster, a sound card targeted to gamers. Sound Blaster delivered synthesized music to desktop PCs and bundled it with a musical instrument digital interface and a joystick port. It took the market by storm.

"We were wondering what to do to survive," Cohen recalls. "And we realized that telephony was a huge market. We envisioned a day when all computers would have integrated telephony—you wouldn't need to install new hardware."

Cohen and Haramaty first aimed at the level of the local area network, which in those early days meant the in-house telephone systems of companies. They relied on the User Datagram Protocol (UDP), a quick but not particularly reliable way to break a voice signal into packets of data and reassemble them at the other end of the call. They anticipated that packets could get lost, might arrive late and out of sequence, or be corrupted by jitter—the loss of transmitted data between networked devices.

Therefore, the technologies to solve these problems were all in VocalTec's products from day one, although it turned out they weren't needed. "When we tried the service on the local area network, the network was so good that we didn't need any of these mechanisms at all," Cohen says. "So we started selling this product, and people started using it between their office networks in the UK and the U.S. to save money on long-distance telephone calls, because calls were very expensive back then. And that was the trigger for delivering this same service over the public Internet."

TIMELINE

- 1970s: ARPANET experiments with digital signal processing (DSP)
- 1974: The IEEE publishes a paper titled "A Protocol for Packet Network Interconnection"
- 1995: VocalTec launches iPhone, the first commercial VoIP app
- 1999: Mark Spencer creates Asterisk
- 2002: Vonage launches
- 2003: Skype launches
- 2004: Microsoft Xbox Live launches
- 2007: VoIP over Wi-Fi apps appear for smartphones
- 2009: VoIP over 3G apps appear
- 2010: First LTE deployments take place

FRANK CHIMERO

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

COMPANY TO WATCH: VIVOX, NATICK, MASS.
Founded in 2005, Vivox provides voice chat services for online games and virtual worlds. Supporting over 25 million users in more than 180 countries and delivering over 3 billion minutes of voice chat a month, the Vivox Network claims to be the world's largest voice network for gamers.

BACK THEN, in the early 1990s, modems crawled along with just 14.4-kilobit-per-second bandwidth, and the 386 and 486 processors in desktop PCs weren't powerful enough to compress the traffic efficiently. So VocalTec got to work again, developing a technology to squeeze the data into a 10-kb/s channel, leaving CPU power available for other applications. "That's when we thought, now we can put this on the Internet," Cohen says.

In 1995, VocalTec launched the world's first commercial voice over Internet Protocol (VoIP) application. "You could just download it; you didn't need to buy any hardware," Cohen says. "Then suddenly it all became clear. All the technologies we created initially, to deal with packet loss and jitter, all that was needed to make this work. Without these, it would not work at all."

The service was still rather rudimentary: You could use your PC to call someone only if that person subscribed to the same VoIP service as you—in this case, VocalTec's Internet Phone, or iPhone, service. There was no interoperability between the various VoIP providers. Other VoIP providers were out of bounds, as was anybody using a standard phone service. If a VoIP provider wanted to connect calls between, say, the United States and Brazil, it had to set up hardware in both markets.

To get around the problem, in 1997 VocalTec and Daniel Berninger, an American entrepreneur who'd worked with AT&T, established a global VoIP exchange called ITXC. This scheme effectively stuck the Internet in the middle of a long-distance call: Because any VoIP service provider could connect to ITXC, a subscriber to one VoIP service could call someone in any other service. Also, and more significantly, standard, public switched telephone networks could also connect to ITXC, so a subscriber to a VoIP service could at last call somebody in the regular phone network with relative ease. ITXC brokered all the interconnect deals between VoIP providers worldwide, giving them greater reach, and made similar deals with standard telecom players so the VoIP start-ups didn't have to. In one stroke, ITXC pushed VoIP on its way to becoming a bona fide alternative to the public switched network.

ITXC was the model for Vonage Holdings Corp., in Holmdel, N.J., a company Berninger later helped found as the first consumer VoIP service provider. It signed up its first residential subscriber in March 2002, a full year ahead of Skype's beta debut.

MEANWHILE, in the background, more technical challenges were appearing. Throughout the world, broadband was exploding, cable and digital subscriber line (DSL) connections were becoming commonplace, and people were getting comfortable with having wireless routers in their homes. A VoIP application now had to get past a router's firewall as well as any other software firewalls a user might deploy against viruses; it was a difficult task to perform without an IT professional easily at hand.

Next, VoIP had to jump the hurdle posed by network address translation. NAT accounts for the fact that devices on the home network—laptops, PCs, Wi-Fi routers, and VoIP handsets—are invisible to the outside Internet. In essence, NAT substitutes a different IP address for the original addresses assigned when the data packet leaves its origin. So, while Vonage and

HONORABLE MENTION

Tablets

The New Computing Covenant

Apple brings tablets down from the mountain

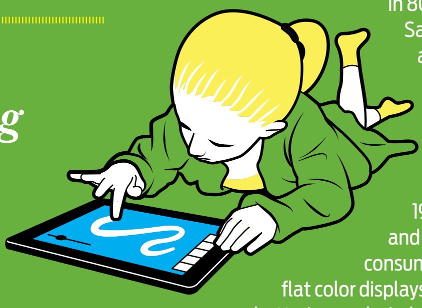

MCKIBILLO

In his 1972 article "A Personal Computer for Children of All Ages," computer scientist Alan Kay imagined the DynaBook, then a computer of "science fiction," around 2 centimeters thick, weighing less than 2 kilograms, and the size and shape of a paper notebook, with a power-charging connection for use at work or in libraries. Arguably, the first machines even to resemble Kay's tablet device emerged in the 1980s, as digital clipboards for insurance adjusters and salespeople.

Then, in April 2010, Apple released the iPad. Only 1.34 cm thick, the Wi-Fi–enabled and 3G-capable model weighs just 0.73 kg. Apple sold some 2 million units in 60 days and 3 million in 80 days. Meanwhile, HP, Dell, Samsung, Notion Ink, Asus, and Lenovo, among others, all have similar tablet computers on the market or are planning them.

According to Jeff Hawkins, developer of the 1989 GridPad, an early tablet, and later of the Palm Pilot, the consumer versions had to wait for flat color displays, low-power CPUs, better batteries, and wireless networks. "The technology didn't exist 20 years ago that was necessary for successful consumer tablets," he says.

Still, Kay says, tablet computers haven't yet reached his vision of the DynaBook. "What end users can make, and what it takes to make something, are both woefully inadequate on today's machines," he says. With a different approach, he adds, tablets such as the iPad could be "one of the greatest educational amplifiers for children ever made." —*Joseph Calamia*

BET YOU DIDN'T KNOW

In 1995, VocalTec launched the world's first commercial VoIP application. The downloadable app was called Internet Phone, or iPhone for short.

most of its competitors got bogged down in disputes over setting standards for traversing the firewall and NAT, Skype secured itself a commanding lead in the burgeoning market by bypassing the standards, relying instead on its own technology.

The founders of Skype, Niklas Zennström and Janus Friis, managed to solve these problems with the peer-to-peer technology they had pioneered at Kazaa Media Desktop. Their technology employed decentralized routing, in which every node—in this case, a subscriber's PC—uses an encrypted channel to keep track of all other users and resources in the network, sort of like opening a tunnel through the firewall and NAT barrier. Wideband codecs—compression systems that capture a wide range of frequencies—enabled Skype to deliver better audio than a fixed-line telephone could manage.

The service caught on fast. In September 2005, Zennström and Friis sold the company to eBay for US $2.5 billion. Zennström, Friis, Berninger, Cohen, and Haramaty were part of a new breed: twentysomethings with the vision and technical brilliance to overturn the established order. Another member of that club was Mark Spencer, who had started up Digium, in Huntsville, Ala., a company that provided Linux tech support, while he was still a computer engineering student at Auburn University.

Spencer balked at paying tens of thousands of dollars for a telephone system—that is, a PBX, or private branch exchange—for his start-up company, so he wrote his own software-based switchboard. He called the software Asterisk, after the Unix symbol for "everything." It was only a few years later, in the early 2000s, that it dawned on him that people were more interested in the phone system than in the tech support service.

For the traditional telephone companies, this was the second blow of the old one-two punch: First they'd had some of their business siphoned off by VocalTec and ITXC; now Digium's open-source software was cutting the cost of equipping a telco in the first place. Now pretty much anybody could set up shop as a VoIP provider.

"Telecom products were really expensive, and there was a real need for customization, especially in other countries [outside the United States]," Spencer recalls. "All these things lined up just right, so that when Asterisk came out, it was able to win a lot of attention."

IN 1932, AT&T's bean counters had the first and last word on telephonic voice quality: They said they wanted the worst quality that paying customers would tolerate. So Bell Telephone Laboratories engineer Harvey Fletcher (who would later invent an electronic hearing aid) truncated the high and low frequencies—below 300 hertz and above 3300 Hz—in the process removing subtlety and emotion from telephonic voices. Fletcher's legacy continues in all the real telephony networks in the world, and it's not the only throwback. In the 15 years since Cohen and his peers decided to take on the telecom giants on their own turf, the user experience has hardly changed. A phone call is much the same as it was in 1994. Only the price has fallen—and not fast enough to prevent other communications media, like e-mail and instant messaging, from squeezing out voice. For voice to get back in the game, Berninger argues, it must offer more.

It could, above all, offer better voice quality. Today's networks can do the job because they're far more advanced than the ones that carried Cohen's first croaking attempts at packetized telephony. According to Skype's chief technology officer, Jonathan Rosenberg, the highest-quality voice calls last around 31 minutes on average, compared with 21 minutes for the low-quality ones. This phenomenon is forcing the traditional telcos to fall in line and adopt VoIP themselves. Operators can no longer get away with the attitude of "if it ain't broke, don't fix it."

Ironically, for a technology that started life with a dependency on wires, nowhere is the evolution of voice services happening at a faster pace than in the wireless world. Mobile phone connections are racing past the 5 billion mark worldwide, and the telco community is consolidating on a common IP-based platform for the fourth generation (4G) of mobile telephony, LTE (Long Term Evolution), in which VoIP is no longer an option—it's a requirement.

According to Eric Ericsson, head of telephony evolution at telephony equipment vendor Ericsson, the mobile industry's Voice over LTE initiative will allow carriers to deliver better service than Skype and its ilk at a lower cost than can be managed today. It also is supposed to make it easier to add new features, such as cheaper roaming, video, and "presence," which routes the call to the device that's most convenient for the user at the moment.

To be sure, a complete switchover to VoIP or LTE won't happen overnight. Operators have invested too much in their legacy voice networks to kill them off in favor of what is still seen as an unproven technology. Though 4G services are being launched right now, they'll coexist with 3G and even 2G for much of the next decade.

But as communications adopt an all-IP architecture, it will get easier to overlay voice on top of other technologies. Digitized voice is finding its way into dozens of applications, including social networking, online gaming, videoconferencing, even advertising. Voice will be reduced to a commodity, like electricity or water. A good thing, too.

—JAMES MIDDLETON

NO. 4 LED Lighting

BLUE + YELLOW = WHITE

Giving LEDs the blues was the key to replacing the incandescent bulb

BACK IN THE 20th century, just about the only LED you normally saw was the one that lit up when your stereo was on. By the noughties, tiny light-emitting diodes were also illuminating the display and keypads of your mobile phone. Now they are backlighting your netbook screen, and soon they'll replace the incandescent and compact fluorescent lightbulbs in your home.

This revolution in lighting comes from the ever-greater bang the LED delivers per buck. With every decade since 1970, when the red LEDs hit their stride, they have gotten 20 times as bright and 90 percent cheaper per watt; the relation is known as Haitz's Law, and it applies also to yellow and blue LEDs, which were commercialized much later.

The forerunners of the white LEDs that are now going into lightbulbs were the chips that backlit handsets starting about a decade ago. Back then, they used tens of milliamps and consumed a watt for every 10 lumens of light they produced. They were also tiny—just 300 micrometers on a side.

Since then, the chips have more than tripled in size, to a millimeter square or more, current has shot up to an ampere or so, and efficiency has rocketed to around 100 lm/W. They now have everything they need to dominate lighting, except for a low enough price. But that, too, will soon come.

Even now, white LEDs are competitive wherever replacing a burned-out lamp is inconvenient, such as in the high ceilings and twisty staircases of Buckingham Palace, because LEDs last 25 times as long as Edison's bulbs. They have a 150 percent edge in longevity over compact fluorescent lights, and unlike CFLs, LEDs contain no toxic mercury. That means it isn't a pain to dispose of them, and you don't have to worry that your house has become a hazard zone if one breaks.

Making these white-emitting chips bigger and driving them harder has been quite easy; it was increasing the efficiency that required a radical redesign of the device's architecture. To produce the first generation of white LEDs, engineers would deposit a stack of carefully chosen gallium nitride and indium gallium nitride layers on a semitransparent substrate to yield blue-emitting devices; then they'd add a yellow-emitting phosphor on top to turn the output white. However, this design traps a lot of light within the chip and sends another fraction in the wrong direction, through the substrate.

To address both weaknesses, engineers coated the nitride film—a combination of GaN and InGaN layers—with a metal that acts as a mirror, then flipped the assembly over, removed the substrate, and roughened the underlying surface. In the resulting chip, because most of the rays impinge on the textured

COMPANY TO WATCH: SORAA, FREMONT, CALIF.

Soraa, a California-based university spin-off cofounded by LED trailblazer Shuji Nakamura, is pioneering a different approach to making gallium nitride devices. By making light-emitting chips on a different cut of the GaN crystal, it can build brighter devices, free from the strong internal electric fields that hamper light emission.

FRANK CHIMERO

top surface at a shallow enough angle to avoid reflecting back, nearly all the light can get through to the world outside.

Europe's leading LED manufacturer, Osram Opto Semiconductors, in Germany, and the two U.S. LED giants Cree and Philips Lumileds are all using variations of this approach. Japan's Nichia, the world's biggest LED manufacturer, has a different way of doing things. Its engineers also roughen the top surface, but they do this by etching a hexagonal pattern into the substrate, which they do not subsequently remove from the gallium nitride film.

These second-generation white LEDs hit the market three or four years ago. Since then interest has rocketed: "If you go to any [lighting] show now, they might as well be called the LED show," says Rick Hamburger, director of segment lighting at Philips Lumileds.

Commercial success followed. White LEDs now illuminate parking lots, streets, and civic buildings. Exactly when they make it into most homes will depend on the price. I just bought a really high-quality, warm-white LED bulb in the United Kingdom from Philips for about US $55; my lamp consumes just 7W while emitting as much light as a 40-W incandescent. (Products that give off a harsher, blue-tinged light go for as little as $10.) I calculate that if I use it for 4 or 5 hours a day, it should pay for replacing the incandescent bulb in about five years.

The manufacturing cost should fall as production yields rise and substrates grow. "At the moment, one-third of the LEDs in the world are made on 2-inch wafers," says Mark McClear, who directs new business development at Cree. Toolmakers are now offering equipment for 3-inch, 4-inch, and even 6-inch substrates, and Cree plans to start using the largest of these platforms in the next 18 months.

Another way to drive down the cost is to increase the light output at a particular current. That's one of the goals of the U.S. Department of Energy's 2010 solid-state lighting road map, which calls for more than doubling the lumens per watt in commercial LED products by 2015. Or engineers could try to build better packages for handling additional heat so they could crank the current up higher and get more light out of each LED.

Of course, LEDs give us much more than just a more efficient, longer-lasting bulb. They're small, cool in operation, and easy to place in walls, automobiles, appliances, even the heels of children's shoes. When designers fully exploit their potential, LEDs will light up places we'd never thought to use them, thus changing the look of our world.

With price the only remaining hurdle and falling all the time, it's clear that this technology will be a winner in the long run. The one potential casualty that no one is talking about: jokes about changing the lightbulb, which may be heading the way of the dodo.

—Richard Stevenson

BET YOU DIDN'T KNOW

LEDs glow brighter if you put them in the fridge. Wal-Mart has cottoned on to this, and since 2006 it has been ripping the fluorescent tubes out of its refrigerators and putting white LEDs in their place.

HONORABLE MENTION

MP3

Compress Me a Song

A German researcher took us from albums to algorithms

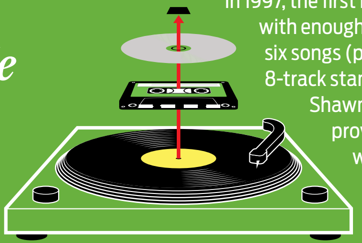

When my children are my age, they will surely look back with bemusement at the crude means by which their elders entertained themselves. They'll laugh at the thought of vinyl platters spinning like carousels, cartridges containing spools of magnetic tape, and even laser-etched metal sheets embedded in plastic discs. They'll take for granted that you can carry every bit of music you own in your pocket.

And it's no sweat now, because NAND flash memory sells for about a dollar per gigabyte. But back in 1989, when a German researcher introduced the idea behind the MP3 compression algorithm, flash memory cost several hundred dollars *per megabyte*. Several things happened at the close of the last decade that together marked the dividing line between B.C. (before compression) and A.D. (Apple domination).

In 1997, the first MP3 player was introduced, with enough storage capacity for about six songs (paltry, it's true, even by 8-track standards). Two years later, Shawn Fanning rolled out Napster, providing an easy way for people with MP3 files to share the songs in their collections. Around that same time, the first USB flash memory key drives arrived, further stoking demand for nonvolatile memory.

In 2001, when Apple introduced the first iPod, consumers' expectations were irrevocably changed. (Pay for music online? Perhaps. Buy a whole LP? Fuhgeddaboutit.) By 2006, digital music had unseated its predecessors: Five billion digital music files were swapped on peer-to-peer networks that year, and the ubiquitous iPod, aided by Apple's iTunes online music store, was fast approaching 100 million units a year in sales.

Think: When was the last time you even saw someone carrying a portable CD or cassette player or found yourself browsing the shelves in a record store? *—Willie D. Jones*

MCKIBILLO

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

NO. 5

Multicore CPUs

PROCESSOR PROLIFERATION

From multicore to many-core to hard-to-describe-in-a-single-word core

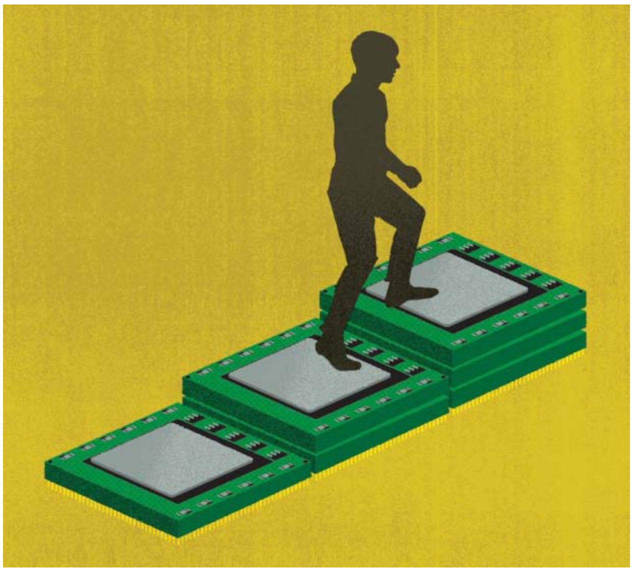

FRANK CHIMERO

BACK IN 1994, programmers figured that whatever code they wrote would run at least 50 percent faster on a 1995 machine and 50 percent faster still on a '96 system. Coding would continue as it always had, with instructions designed to be executed one after the other.

But Kunle Olukotun, then a newly minted professor of electrical engineering at Stanford, saw that the party couldn't go on forever. The microprocessors of the day couldn't scale up as efficiently as you'd expect through the mere addition of ever more and ever faster transistors, the two things that Moore's Law provided.

To solve that problem, Olukotun and his students designed the first general-purpose multicore CPU. This idea, more than any other in the past decade, is what has kept the semiconductor industry climbing the Moore's Law performance curve. Without multicore chips, the computing capability of everything from servers to netbooks would not be much better than it was a decade ago. Everyone's happy—except perhaps for the programmers, who must now write code with threads of instructions that must be executed together—in pairs, quartets, or even larger groupings.

It's not that old, single-core CPUs weren't already doing some parallel processing. When Olukotun began his work, most microprocessors had a "superscalar" architecture. In the superscalar scheme, the CPU contained many replicated components, such as arithmetic units. Individual instructions would be parceled out to the waiting components. Scaling up such "instruction-level parallelism" meant building in more and more such components as the years rolled by.

Olukotun argued that within a few more generations, it wasn't going to be worth the effort. You needed to provide a quadratic increase in resources for a linear increase in performance, he said, because of the complexity of the logic involved in parceling out and keeping track of all the instructions. If you combined that with the delays inherent in the mess of interconnects among all those parts, it seemed a losing proposition. Doug Burger and Stephen Keckler, both computer scientists at the University of Texas, Austin, put a finer point on it later in the decade, calculating that instead of the 50 percent improvements everyone had gotten used to, the computing industry should start thinking 12.5 percent. And 12.5 percent isn't much of a reason to buy a new computer, is it?

Olukotun's answer was Hydra, a processor whose parallelism came not from redundant circuits within a single complex CPU but from building four copies of a simpler CPU core on one chip. That way, you save on interconnects and on the time lost casting instructions out and reeling answers back in. In Hydra, you got parallel processing without all the delay-inducing complexity. In 1998 "we wrapped up the hardware portion of the project and declared victory," says Olukotun.

It was a quiet victory. In the computing environment of the 1990s, Hydra seemed a little crazy, Olukotun says. Superscalar designs were still delivering 50 percent performance improvements every year. "It was by no means clear at the time that our view of the world was going to win," he recalls. And indeed it would be years before processor giants like Intel, Advanced Micro Devices, and IBM got the multicore religion Olukotun preached. And when they did, it would largely be for a reason he had hardly considered: power.

COMPANY TO WATCH: TILERA CORP., SAN JOSE, CALIF.

In 2008, MIT professor Anant Agarwal transformed an academic project to efficiently make use of lots of simple cores connected in a mesh into Tilera, a company whose commercial processor has one of the highest core counts of all. It's selling a 64-core product now, the 100-core Tile-Gx starts sample shipments in mid-2011, and the company plans a 200-core product in 2013.

It turned out that the rising density of transistors on a chip intensified the hot spots in CPUs. This, even more than the resource-to-performance ratio that had bothered Olukotun, was the problem that seemed most likely to stop Moore's Law in its tracks. In presentations in 1999 and later, Intel engineers showed that if trends in microprocessors were to continue, by 2010 they'd burn as hot as the surface of the sun.

The answer was clear: Slow down the CPU's clock and add more cores. That way, you'd gain more from the extra parallelism than you lost from the slower processing. The chip would gobble less power and generate less heat.

It was a daunting engineering job, but the big processor makers were more prepared than you might expect, because they'd already redesigned the way that CPUs communicate with other chips. For Intel, the solution, called the front-side bus, debuted in 1996, in the Pentium Pro. According to Intel senior fellow Steve Pawlowski, the bus was, in large part, originally meant to save on testing and validation costs. It was a very convenient piece of luck, because when the time came to get two cores working together, the front-side bus was there, waiting to link them up. And in 2005 Intel released its first dual-core component, the Pentium D, which was really two single-core chips in the same package, tied together by the front-side bus.

Engineers at AMD—influenced by Olukotun, Burger, and Keckler—were more purposeful. They prepped the initial, single-core version of AMD's breakout server chip, the Opteron, with a redesigned communications component that would make a multicore version easy. That version came out in 2005. The component was the chip's "northbridge," a switchyard that acts as the chip's gateway to other chips in the computer.

IBM was, arguably, even more on top of the multicore revolution. Around the same time that Intel's Pentium Pro was released, the company began work on its Power4 processor. Looking for an advantage, IBM entertained a number of cutting-edge ways to enhance instruction-level parallelism in single cores, according to Jim Kahle, chief architect of that design. But, deciding to play it safe, his team rejected each. "Turned out to be a good idea," he says. The most conservative option was a dual-core processor. And so Power4, released in 2001, became the first mainstream computer processor with more than one core on a single die.

Olukotun himself wasn't absent from the revolution he predicted. In 2000, he took the lessons from Hydra and founded Afara Websystems. That start-up was acquired by Sun Microsystems in 2002, and its technology became Sun's powerful Web server CPU, the eight-core UltraSparc T1 (also known as Niagara), released in 2005.

Once the multicore revolution got going, it had a natural momentum. "As soon as we got to two cores, it became obvious we needed to start thinking about going to four," says AMD corporate fellow Chuck Moore. "And as soon as we got to four, we started thinking about going to six or eight."

So today programmers can again count on a solid 50 percent annual gain in effective processing power, driven not by raw speed but by increasing parallelism. Therein lies the rub. Back when Olukotun worked out Hydra, "it was unclear if you could take advantage of all the parallelism," he says. "It's still unclear today."

So where does it end? Sixty-four cores? Already there. Start-up Tilera Corp.

HONORABLE MENTION

DVDs

Time to Eject

The rise and fall of the optical disc

MCKIBILLO

This year *Star Wars* fans will have the chance to buy yet another version of the films, this time on high-definition Blu-ray. Although the new disc may shoulder aside the DVD, it cannot equal it as a game changer.

Take for example Netflix, the movie rental company. It owes its existence to the DVD, because this disc, unlike the VHS tape, was not too expensive to ship. Now, Netflix typically mails 2 million DVDs a day, says Steve Swasey, vice president of corporate communications. When cofounder Reed Hastings first imagined movies by mail in the 1990s, DVDs were so new that he didn't own any. He tested the system by mailing CDs to himself.

Blu-ray required lasers capable of creating and reading smaller marks—stamped pits, dye discolorations, or quenched metal coatings—on the surface of a disc. A DVD uses the red beam of a 650-nanometer laser; a Blu-ray uses a blue 405-nm beam. Shuji Nakamura developed the laser diode required for Blu-ray while working at Nichia Corp., in Tokushima, Japan. He says he chose to study the "poor" semiconductor gallium nitride to make laser diodes in the 1980s because he wanted opportunities to publish. "The major companies were using zinc selenide. There were too many papers," Nakamura says.

Making more marks in even more layers might lead to future generations of optical discs, says Barry Schechtman, executive director emeritus at the Information Storage Industry Consortium. But he wonders whether the disc is already coming to the end of its momentous life, given external hard drives, flash drives, and streaming. "When we will need another generation beyond Blu-ray—and even if we would need it—is still a big question," he says. —*Joseph Calamia*

"Cores are the new transistors"

KUNLE OLUKOTUN, STANFORD UNIVERSITY

is selling it (see "Company to Watch"). Two hundred? One thousand? "Cores are the new transistors," jokes Olukotun.

Just adding traditional cores isn't going to be enough, says AMD's Moore. The scheme may have saved the power-versus-performance curve for a time, but it won't do so forever. "These days, each core is only getting 8 or 10 watts," he says. "In some sense we're running back into that power wall." With its new Bulldozer architecture, AMD has managed to buy some breathing room by finding a set of components that the cores can share without seriously degrading their speed. But even so, Moore's best guess is that 16 cores might be the practical limit for mainstream chips.

Intel's Pawlowski won't put a number on it, but he will say that memory bandwidth between the cores is likely to be the big constraint on growth.

What will keep computing marching forward, according to Moore, is the integration of CPUs and graphics processing units (GPUs) into what AMD calls an accelerated processing unit, or APU. Say you want to brighten an image: Just add 1 to the number representing the brightness of every pixel. It'd be a waste of time to funnel all those bits single file through a CPU core, or even 16 of them, but GPUs have dedicated hardware that can transform all that data practically at once.

It turns out that many modern workloads have just that kind of data-level parallelism. Basically, you want to do the same thing to a whole lot of data.

That key insight drove AMD to acquire a leading GPU maker, ATI Technologies, and start work on jamming their two products together. So a future processor, from AMD at least, would probably contain multiple CPU cores connected to several GPU elements that would step in whenever the work is of a type that would gum up a CPU core.

With Cell, the processor released in 2006 to power the PlayStation 3, IBM has already gone in that direction. Instead of actual GPU functions, it developed a more flexible core that specializes in executing the same instruction on several pieces of data at once. IBM, with help from Toshiba and Sony, stuck eight of the new cores on the same chip with a more traditional processor core. But that's not quite where Kahle, who led the Cell project, sees things going in the future. Instead he expects to see a mix of general-purpose cores and cores specialized for one task—encryption, decryption, video encoding, decompression, anything with a well-defined standard.

Olukotun agrees that such a heterogeneous mix of cores is the way forward, but it's not going to be easy. "It's going to make the programming problem much worse than it is today," he says. "Just as things were getting bad for software developers, they have the potential to get worse." But don't worry. They're working on it.

—Samuel K. Moore

FRANK CHIMERO

Cloud Computing NO. 6

IT'S ALWAYS SUNNY IN THE CLOUD

Cloud computing puts your desktop wherever you want it

JUST 18 YEARS ago the Internet was in its infancy, a mere playground for tech-savvy frontiersmen who knew how to search a directory and FTP a file. Then in 1993 it hit puberty, when the Web's graphical browsers and clickable hyperlinks began to attract a wider audience. Finally, in the 2000s, it came of age, with blogs, tweets, and social networking dizzying billions of ever more naive users with relentless waves of information, entertainment, and gossip.

This, the adulthood of the Internet, has come about for many reasons, all of them supporting a single conceptual advance: We've cut clean through the barrier between hardware and software. And it's deeply personal. Videos of our most embarrassing moments, e-mails detailing our deepest heartaches, and every digit of our bank accounts, social security numbers, and credit cards are splintered into thousands of servers controlled by dozens—hundreds?—of companies.

Welcome to cloud computing. We've been catapulted into this nebulous state by the powerful convergence of widespread broadband access, the profusion of mobile devices enabling near-constant Internet connectivity, and hundreds of innovations that have made data centers much easier to build and run. For most of us, physical storage may well become obsolete in the next few years. We can now run intensive computing tasks on someone else's servers cheaply, or even for free. If this all sounds a lot like time-sharing on a mainframe, you're right. But this time it's accessible to all, and it's more than a little addictive.

The seduction of the business world began first, in 2000, when Salesforce.com started hosting software for interacting with customers that a client could rebrand as its own. Customers' personal details, of course, went straight into Salesforce's databases. Since then, hundreds of companies have turned their old physical products into virtual services or invented new ones by harnessing the potential of cloud computing.

FUN FACT:

Transmitting a terabyte of data from Boston to San Francisco can take a week. So the impatient are returning to an old idea, "Sneakernet": Put your data on a disc, take it to FedEx, and get it to a data center in a day.

Consumers were tempted four years later, when Google offered them their gateway drug: Gmail, a free online e-mail service with unprecedented amounts of storage space. The bargain had Faustian overtones—store your e-mail with us for free, and in exchange we'll unleash creepy bots to scan your prose—but the illusion of infinite storage proved too thoroughly enthralling. This was Google, after all: big, brawny, able to warp space and time.

COMPANY TO WATCH: F-SECURE, HELSINKI, FINLAND

F-Secure Corp. uses the cloud to protect the cloud. Its global network of servers detects malicious software and distributes protective updates in minutes. To assess a threat, it uses the Internet itself: A widely available application is more likely to be safe than a unique file.

GMAIL'S INFINITE storage was a start. But the program's developers also made use of a handy new feature. Now they could roll out updates whenever they pleased, guaranteeing that Gmail users were all in sync without having to visit a Web site to download and install an update. The same principle applied to the collaborative editing tools of Google Docs, which moved users' documents into the browser with no need for backups to a hard drive. "Six years ago"—before the launch of Docs—"office productivity on the Web wasn't even an idea," recalls Rajen Sheth, a product manager at Google.

Docs thus took a first, tentative bite out of such package software products as Microsoft Office. Soon hundreds of companies were nibbling away.

Adding new features and fixing glitches, it turned out, could be a fluid and invisible process. Indeed, sites like the photo storage service Flickr and the blog platform WordPress continually seep out new products, features, and fixes. Scraping software off individual hard drives and running it in anonymous data centers obliterated the old, plodding cycles of product releases and patches.

In 2008, Google took a step back from software and launched App Engine. For next to nothing, Google now lets its users upload Java or Python code that is then modified to run swiftly on any desired number of machines. Anyone with a zany idea for a Web application could test it out on Google's servers with minimal financial risk. Let's say your Web app explodes in popularity: App Engine will sense the spike and swiftly increase your computing ration.

With App Engine, Google began dabbling in a space already dominated by another massive player, Amazon.com. No longer the placid bookstore most customers may have

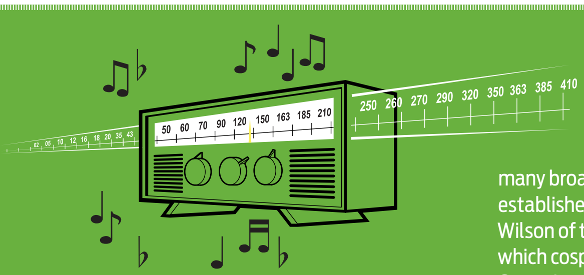

MCKIBILLO

HONORABLE MENTION

HD Radio

The End of Analog

AM and FM go HD

In 2002, HD Radio promised Americans FM-quality sound on AM channels and CD-quality sound on FM channels—with no subscription charge. In 2006, retailers sold 28 000 HD Radio receivers for nightstands and dashboards, reports iBiquity Digital Corp., the Columbia, Md., company that developed the system. In the first three quarters of 2010, sales hit 1.2 million. Radio was finally going digital.

Well, almost digital. Called IBOC for in-band on-channel, today's "hybrid" setup requires broadcasters to continue transmitting the old analog signal. Upgrading to the iBiquity system allows them also to transmit a digital signal, centered on the analog channel and in the frequencies slightly above and below. For FM channels, that's about 101 to 200 kilohertz from the analog center. To avoid interference, broadcasters transmit these digital signals at a lower power than their analog neighbors.

The IBOC system appeals to many broadcasters because it keeps intact their established territories and frequencies, says Dave Wilson of the Consumer Electronics Association, which cosponsors the National Radio Systems Committee, one of the organizations that tested the system. It also takes up no more of the radio spectrum than did the old analog system.

The original point of going digital was to get improved sound quality, but ambitions grew along with the technology. Because today's compression programs can stuff far more information into a given bandwidth than had been dreamed possible, new things are being found to stuff: captioning for the deaf, iTunes song "tagging," album art, and even "multicasting"—different broadcasts layered into one station, found by tuning a bit above and below the center FM channel. The system is also more efficient than streaming Internet or satellite services, says Robert J. Struble, president and CEO of iBiquity. It takes only one tower on the top of the Empire State Building, he notes, to dole out HD Radio to all of New York City.

To upgrade, broadcasters must buy new transmitting equipment and pay a fee to iBiquity, but the company reports that over 2000 U.S. stations have made the switch. "They feel this is a question of survival," says Ann Gallagher at the Federal Communications Commission, audio division, regarding a 2010 broadcasters' panel discussion. "Radio doesn't want to be the last analog medium."

—Joseph Calamia

assumed it to be, in 2000 Amazon had begun to use its sales platform to host the Web sites of other companies, such as the budget retailer Target. In 2006 came rentable data storage, followed by a smorgasbord of "instances," essentially slices of a server available in dozens of shapes and sizes. (Not satisfied? Fine: The CPU of an instance, which Amazon calls a compute unit, is equivalent to that of a 1.0- to 1.2-gigahertz 2007 Opteron or 2007 Xeon processor.)

To get a flavor of the options, for as little as about US $0.03 an hour, you can bid on unused instances in Amazon's cloud. As long as your bid exceeds a price set by Amazon, that spare capacity is yours. At the higher end, around $2.28 per hour can get you a "quadruple extra large" instance with 68 gigabytes of memory, 1690 GB of storage, and a veritable bounty of 26 compute units.

In a sense, the cloud environment makes it easier to just get things done. The price of running 10 servers for 1000 hours is identical to running 1000 machines for 10 hours—a flexibility that doesn't exist in most corporate server rooms. "These are unglamorous, heavy-lifting tasks that are the price of admission for doing what your customers value," says Adam Selipsky, a vice president at Amazon Web Services.

As unglamorous as an electric utility, some might say. Indeed, Amazon's cloud services are as close as we've gotten to the 50-year-old dream of "utility computing," in which processing is treated like power. Users pay for what they use and don't install their own generating capacity. The idea of every company running its own generators seems ludicrous, and some would argue that computing should be viewed the same way.

Selling instances, of course, is nothing like selling paperbacks, toasters, or DVDs. Where Google's business model revolves around collecting the world's digital assets, Amazon has more of a split personality, one that has led to some odd relationships. To help sell movies, for example, Amazon now streams video on demand, much like companies such as Netflix. Netflix, however, also uses Amazon's servers to stream its movies. In other words, Amazon's servers are *so* cheap and useful that even its competitors can't stay away. But to understand what's truly fueling the addiction to the cloud, you'll need to glance a bit farther back in time.

IN THE MID-1990S, a handful of computer science graduate students at Stanford University became interested in technologies that IBM had developed in the 1960s and '70s to let multiple users share a single machine. By the 1980s, when cheap servers and desktop computers began to supplant mainframe computers, those "virtualization" techniques had fallen out of favor.

The students applied some of those dusty old ideas to PCs running Microsoft Windows and Linux. They built what's called a hypervisor, a layer of software that goes between hardware and other higher-level software structures, deciding which of them will get how much access to CPU, storage, and memory. "We called it Disco—another great idea from the '70s ready to make a comeback," recalls Stephen Herrod, who was one of the students.

They realized that virtualization could address many of the problems that had begun to plague the IT industry. For one thing, servers commonly operated at as little as a tenth of their capacity, according to International Data Corp., because key applications each had a dedicated server. It was a way of limiting vulnerabilities because true disaster-proofing was essentially unaffordable.

So the students spawned a start-up, VMware. They started by emulating an

SGI

WIZARDS OF AN ONLINE OZ

The Internet may sometimes seem like a massive astral projection, but it, too, must have its little men behind curtains, operating the gears and levers of the great and powerful cloud.

Not just Google and Amazon, but Microsoft, Rackspace, SGI [above], and many other companies have used striking economies of scale to acquire cheap commodity hardware and build sprawling, multihectare server farms. James Hamilton, a vice president and distinguished engineer at Amazon, once estimated that data centers with tens of thousands of computers could use their clout to buy hardware and network bandwidth for one-seventh the price a piddling 1000-server facility could exact.

"Twitter, Facebook, Google—their performance needs rise astronomically every day," notes Michael Coté, a cloud computing analyst at RedMonk, an IT consulting firm. To keep up, companies now buy servers by the shipping container with all the power and networking connections preconfigured. An IT manager plugs the container into a power supply, a network connection, and a cold water tap, and presto, 1400-odd more servers come online.

Now a data center can materialize overnight. Furthermore, the closed environment of a can turns out to be well suited to maintaining exquisite control over airflow for cooling—a major issue in the blazing hot environment of a data center. Following the typical architecture of alternating "hot aisles" full of server racks and "cold aisles" for ventilation, vents blow air in a loop through each container while radiators release it into "cold" aisles chilled to a brow-mopping 32 °C. *—Sandra Upson*

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

Intel *x*86 microprocessor's behavior in software. But those early attempts didn't always work smoothly. "When you mess up an emulation and then run Windows 95 on top of it, you sometimes get funny results," Herrod, now VMware's chief technology officer, recalls. They'd wait an hour for the operating system to boot up, only to see the Windows graphics rendered upside down or all reds displayed as purple. But slowly they figured out how to emulate first the processor, then the video cards and network cards. Finally they had a software version of a PC—a virtual machine.

Next they set out to load multiple virtual machines on one piece of hardware, allowing them to run several operating systems on a single machine. Armed with these techniques, VMware began helping its customers consolidate their data centers on an almost epic scale—shrinking 500 servers down to 20. "You literally go up to a server, suck the brains out of it, and plop it on a virtual machine, with no disruption to how you run the application or what it looks like," Herrod says.

Also useful was an automated process that could switch out the underlying hardware that supported an up-and-running virtual machine, allowing it to move from, say, a Dell machine to an HP server. This was the essence of load balancing—if one server started failing or got too choked up with virtual machines, they could move off, eliminating a potential bottleneck.

FUN FACT:

Dude, where are my bits? In the growing obfuscation of who's responsible for what data, Amazon recently deployed its storefront platform on privacy-challenged Facebook for the first time. The irresistible business case? Selling Pampers diapers.

You might think that the virtual machines would run far more slowly than the underlying hardware, but the engineers solved the problem with a trick that separates mundane from "privileged" computing tasks. When the virtual machines sharing a single server execute routine commands, those computations all run on the bare metal, mixed together with their neighbors' tasks in a computational salad bowl. Only when the virtual machine needs to perform a more confidential task, such as accessing the network, does the processing retreat back into its walled-off software alcove, where the calculating continues, bento-box style.

Those speedy transitions would not have been possible were it not for another key trend—the consolidation of life into an Intel world. Back in virtualization's early days, a major goal was to implement foreign architectures on whatever hardware was at hand—say, by emulating a Power PC on a Sun Microsystems workstation. Virtualization then had two functions, to silo data and to translate commands for the underlying hardware. With microprocessor architectures standardized around the *x*86, just about any server is now compatible with every other, eliminating the tedious translation step.

VMware no longer has a monopoly on virtualization—a nice open-source option exists as well—but it can take credit for developing much of the master idea. With computers sliced up into anywhere between 5 and 100 flexible, versatile virtual machines, users can claim exactly the computing capacity they need at any given moment. Adding more units or cutting back is simple and immediate. The now-routine tasks of cloning virtual machines and distributing them through multiple data centers make for easy backups. And at a few cents per CPU-hour, cloud computing can be cheap as dirt.

SO WILL ALL computing move into the cloud? Well, not every bit. Some will stay down here, on Earth, where every roofing tile and toothbrush seems fated to have a microprocessor of its own.

But for you and me, the days of disconnecting and holing up with one's hard drive are gone. IT managers, too, will surely see their hardware babysitting duties continue to shrink. Cloud providers have argued their case well to small-time operations with unimpressive computing needs and university researchers with massive data sets to crunch through. But those vendors still need to convince Fortune 500 companies that cloud computing isn't just for start-ups and biology professors short on cash. They need a few more examples like Netflix to prove that mucking around in the server room is a choice, not a necessity.

And we may just need more assurances that our data will always be safe. Data could migrate across national borders, becoming susceptible to an unfriendly regime's weak human rights laws. A cloud vendor might go out of business, change its pricing, be acquired by an archrival, or get wiped out by a hurricane. To protect themselves, cloud dwellers will want their data to be able to transfer smoothly from cloud to cloud. Right now, it does not.

The true test of the cloud, then, may emerge in the next generation of court cases, where the murky details of consumer protections and data ownership in a cloud-based world will eventually be hashed out. That's when we'll grasp the repercussions of our new addiction—and when we may finally learn exactly how the dream of the Internet, in which all the world's computers function as one, might also be a nightmare. —SANDRA UPSON

NO. 7

Drone Aircraft

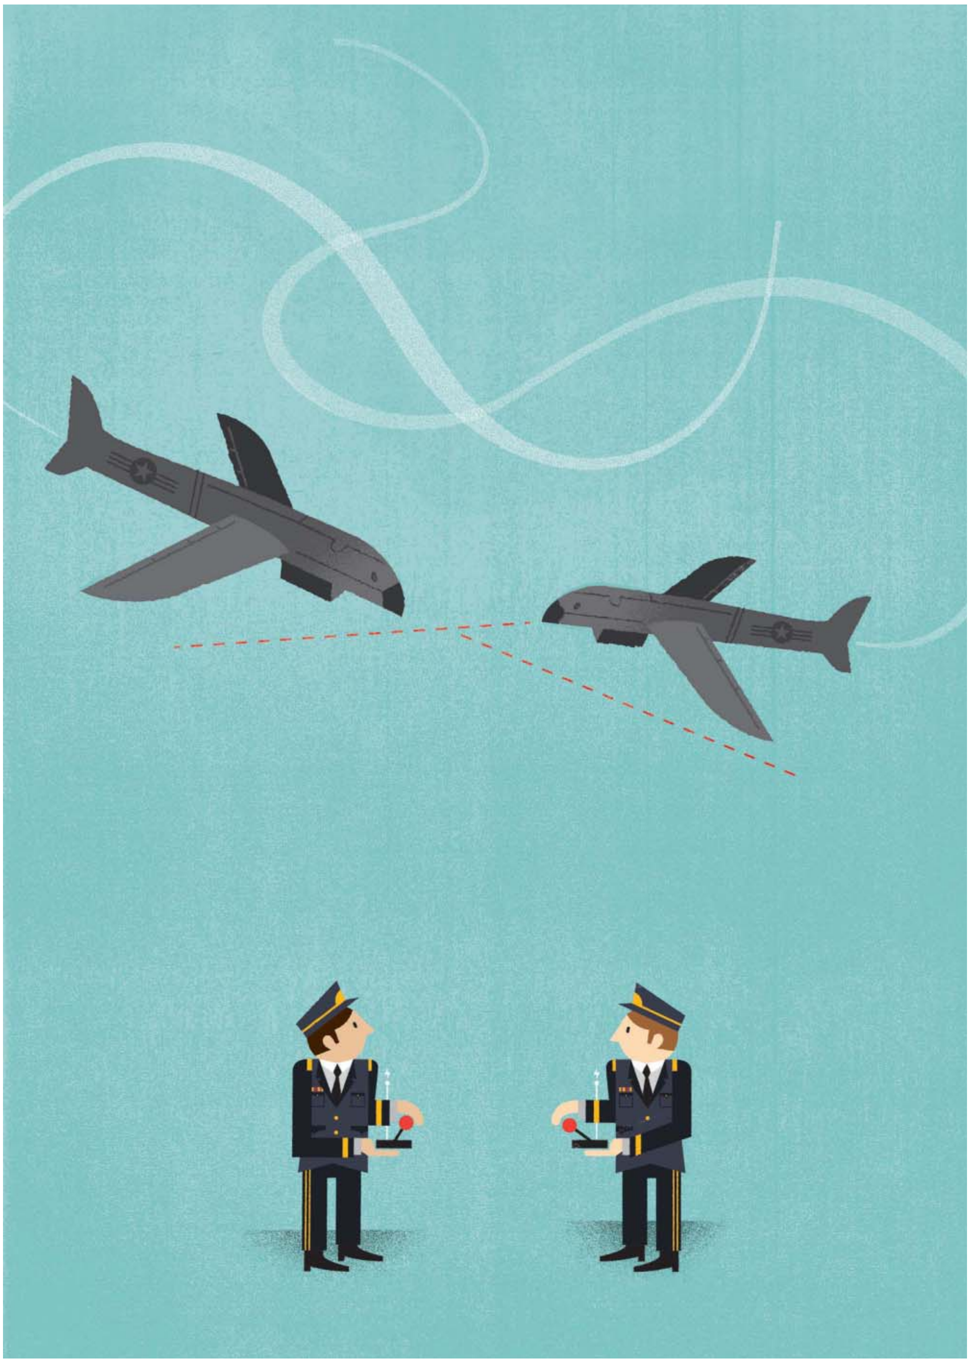

FRANK CHIMERO

HOW THE DRONES GOT THEIR STINGERS

Unmanned aerial vehicles come of age

CRUISING SILENTLY overhead, an unmanned Predator aircraft uses its infrared camera to pinpoint the telltale muzzle flashes from a sniper's rifle. The plane's operators, located half a world away, then unleash a Hellfire missile from under its wing, using a laser mounted beneath the craft's nose to guide the munition into the very window the sniper had been shooting from.

Such missions represent a technological tour de force, but they've played out so often over the past few years that they no longer make headlines. What might be news, though, is just how far back the roots of this stunning 21st-century military technology reach.

The first demonstration of a remotely piloted vehicle took place in May 1898 at the Electrical Exposition in New York City's Madison Square Garden. It was less than a month after the outbreak of the Spanish-American War, which, as history buffs may recall, was sparked by the mysterious explosion and sinking of the U.S.S. *Maine* in Havana harbor. So when

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

COMPANY TO WATCH: PRIORIA ROBOTICS, GAINESVILLE, FLA.
Prioria Robotics is building small UAVs that autonomously use video to steer around objects. The ability of UAVs to sense and avoid other aircraft is a prerequisite for flying in civilian airspace, and small smart vehicles, like the ones Prioria is developing, may be the first to do so.

SKY EYES: U.S. Army operators in Iraq prepare a Shadow for catapult launch [top]; a soldier heaves a Raven into the air from the perimeter of a U.S. Marine base in Afghanistan [bottom left]; and an airman at Creech Air Force Base checks over a Predator before flight [bottom right].
CLOCKWISE FROM TOP: U.S. ARMY; ETHAN MILLER/GETTY IMAGES; JOHN MOORE/GETTY IMAGES

the renowned inventor Nikola Tesla used the exposition to demonstrate his "telautomaton"—a small boat operated remotely by radio—the military significance of his creation must have been obvious.

Or maybe not. Perhaps no one watching Tesla steer and flash the lights of his robotic boat had sufficient imagination to see how valuable pilotless vehicles could be in war. Yet that very year Tesla sent a paper to *The Electrical Engineer* magazine describing how a remotely controlled aircraft could be used as an aerial torpedo. His submission was rejected as too fanciful.

Attitudes, of course, evolved during the 20th century. As early as 1917, the U.S. Navy pursued the development of a pilotless aircraft for use against German U-boats; during the Second World War, Nazi forces filled the skies over Britain with thousands of pulse-jet-powered flying bombs; and Israeli Defense Forces used drones against Syrian forces in Lebanon in 1982. But it was only during the past decade that unmanned aerial vehicles matured into fully controllable and reusable combat aircraft. They have also proliferated, often in miniaturized forms, providing easily deployed eyes in the sky for ground troops.

The relatively late blossoming of these vehicles can be explained, in part, by improvements in the various technologies they rely on. But an increased understanding of their utility on the battlefield also accounts for the recent upsurge in their use.

PILOTLESS AIRCRAFT began as preprogrammed drones, evolved into remotely piloted vehicles or unmanned aerial vehicles (UAVs), and are now sometimes called unmanned aircraft systems (UASs), in appreciation of the aircraft's role within a larger collection of complex equipment.

They're most commonly used for reconnaissance or sustained surveillance. But many UAVs are equipped with lasers that can illuminate a target or even mark it for automatic destruction, using laser-guided weapons launched from other platforms. And some of the larger UAVs carry missiles themselves. At the small end of the spectrum are UAVs that a single soldier can carry around and launch by hand. AeroVironment, a company headquartered in Monrovia, Calif., manufactures some of the better-known examples, including the Raven, which is fundamentally similar to the radio-controlled planes that hobbyists fly for fun, although it's much tougher and packs more sophisticated electronic gear.

So, clearly, the wherewithal to construct a small UAV has been around for a long time. Why then has the use of Ravens and other hand-launched UAVs burgeoned only in the last decade?

"The technology definitely matured," says Gabriel Torres, an aeronautical engineer and project manager at AeroVironment. He notes that the Pointer, the 1990s-era predecessor of the Raven, used nickel-cadmium batteries and could remain aloft for only 30 minutes, one-third as long as the lithium ion–powered Raven can. The Pointer's support equipment was also awkwardly bulky for a soldier to carry around. "The ground-control station was like an 80-pound box," says Torres. The craft was also limited by its rudimentary autopilot, which initially relied on an electromechanical compass. "But little by little, the technology improved," says Torres.

Strides in the fabrication of microelectromechanical systems, for example, allowed tiny gyroscopes, accelerometers, and airspeed sensors to be added to the smallest of these vehicles, along with increasingly compact and reliable GPS receivers. "The other thing that changed," says Torres, "was an appreciation of the kind of mission that could be fulfilled with this type of system." With the Raven, he explains, U.S. military planners were able to work out detailed tactics, techniques, and procedures for soldiers to use small aerial vehicles in combat. "This all became very important in 2000 and 2001," he says.

This period was also a turning point for larger UAVs, in part because of advances in computers and radio links. The more important reason to call 2001 a watershed year, however, is that it marked the very first time anyone put weapons on a reusable unmanned aircraft.

WHEN GENERAL ATOMICS Aeronautical Systems developed the Predator in the mid-1990s, it was intended solely for surveillance. Plans took a sharp turn, though, during one of its early military deployments, with NATO forces over Kosovo.

That's when Gen. John P. Jumper, commander of U.S. Air Forces in Europe, noticed what he called "the dialogue of the deaf." Predator operators would identify a target, say, an enemy tank lurking between buildings, and then try to guide the pilot of an attack aircraft to it by radioing verbal instructions. More often than not, this just caused a lot of confusion.

So Jumper had the Predators outfitted with laser designators, which could automatically guide missiles or artillery shells to their targets. He later pushed for the Predator to carry its own weapons, and the first instances of this UAV using air-to-surface missiles took place not long after, in the hunt for al-Qaeda members in Afghanistan in October 2001.

The following year saw the Predator's role expand to ground support, when it destroyed a machine gun bunker that had pinned down U.S. Army rangers in Afghanistan. And for the first time, a Predator fired a Stinger air-to-air missile in action, at an Iraqi fighter.

BET YOU DIDN'T KNOW

On 2 August 2010, one of the U.S. Navy's unmanned Fire Scout helicopters lost communication with operators and was fast approaching Washington, D.C., before it was brought under control.

That same year technical refinements allowed operation of a Predator to be shifted from one ground control station to another, so that UAV pilots located in combat areas could pass control to comrades stationed at U.S. bases. This is how the U.S. Air Force operates its UAVs, using pilots and crews deployed overseas to launch and recover the aircraft, while pilots at Creech Air Force Base in Nevada manage the intervening part of the missions.

This approach is decidedly different from the way the U.S. Army handles its fleet of more than 4000 UAVs, including the biggest one it flies, known as the Gray Eagle, which can carry four Hellfire missiles and looks something like a Predator on steroids.

For one, all Army UAS operators work alongside combat troops, even for vehicles that are capable of being flown over satellite links from anywhere on the globe. Also, unlike the Air Force, the Army hasn't restricted itself to using officers with piloting experience to operate its UAVs.

The Air Force initially put just seasoned aviators in that role, although it has more recently started using men and women who have spent only a few tens of hours in the cockpit—about what it would take to get a private pilot's license—to fly its UAVs. The Army, however, trains people who have never flown regular aircraft for that job.

"We recruit folks coming out of high school," says Lt. Col. Patrick Sullivan, commander of the Army's unmanned aircraft systems training battalion, located at Fort Huachuca in southeast Arizona. "We have soldiers that just came from basic training. They come to Fort Huachuca, and we train them to be UAS operators."

Army Col. Robert Sova, training and doctrine command capability manager for UAS, confirms that this approach extends throughout the Army. "Everything from our smallest systems to our largest are flown by enlisted operators," says Sova. "They don't have a pilot's license—they are not pilots. That's why we're adamant about calling them 'operators.' We have no intention of having pilots flying our unmanned aircraft."

The Army's stance on this point reflects that service's particular history and culture. But it also reflects the evolution of UAS technology. During the 1980s and 1990s, operating such an aircraft indeed required someone with considerable training and good eye-muscle coordination. It might even require two of them. The "external pilot" would handle the takeoff and landing by looking directly at the plane, while the "internal pilot" would manage operations for the bulk of the flight.

Steve Reid is vice president of unmanned aircraft systems at AAI Corp., the Hunt Valley, Md., company that makes the Shadow, a 3.4-meter-long UAV that the U.S. Army adopted for tactical use in late 1999. He explains that in 2001 and 2002, AAI added special radio equipment to allow the Shadow to make automated landings. Although engineering such capability was a considerable challenge, it relieved the Army of having to train pilots in the tricky skills needed to land these planes by eye. What's more, the automation works better. "I can tell you great stories of sandstorms rolling into Iraq—blinding sandstorms, where you can't see any aircraft out there—and Shadows land right where they are supposed to," says Reid.

Army Sgt. 1st Class Kelly Boehning, a Gray Eagle operator stationed at Fort Huachuca, has similar praise for his craft's automated landing system. "It lands perfectly, every time, without exception," says Boehning. "It takes some of the fun out, not having the stick and rudder, but it also takes the pilot error out: We don't have any incidents landing—that's where Predator's downfall is."

AUTOMATION IS INDEED a strong theme in the Army's recently published Unmanned Aircraft Systems Roadmap for the next quarter century, a document that discusses such advanced possibilities as UAVs delivering cargo to soldiers on the battlefield or flying UAV missions in coordinated swarms.

Military UAVs still have a poor safety record compared with piloted aircraft, but as trust in them mounts, there's little doubt that their use will expand and spill over into civilian applications. "There's a hunger out there in the commercial sector for this type of technology—for use in everything from UPS and FedEx flights, pipeline surveys, forestry, logging, law enforcement—just about anything you can think of that uses aircraft today could benefit from a low-cost, reliable, and safe unmanned aircraft technology," says retired Army Lt. Col. Glenn Rizzi, senior advisor to Col. Sova.

It will, of course, be a long while yet before the typical traveler will be comfortable taking an airliner with no pilot on board. But well before that, we're likely to see UAVs of other varieties flying in civilian airspace.

Understandably, aviation authorities need to police developments here cautiously so that the safety of air travel isn't compromised. But the steady advance of communication and automation technology—which is already quite sophisticated in today's airliners—will surely open the skies to pilotless aircraft of many types. As Rizzi contends, "The future [of UAVs] is only up."

—David Schneider

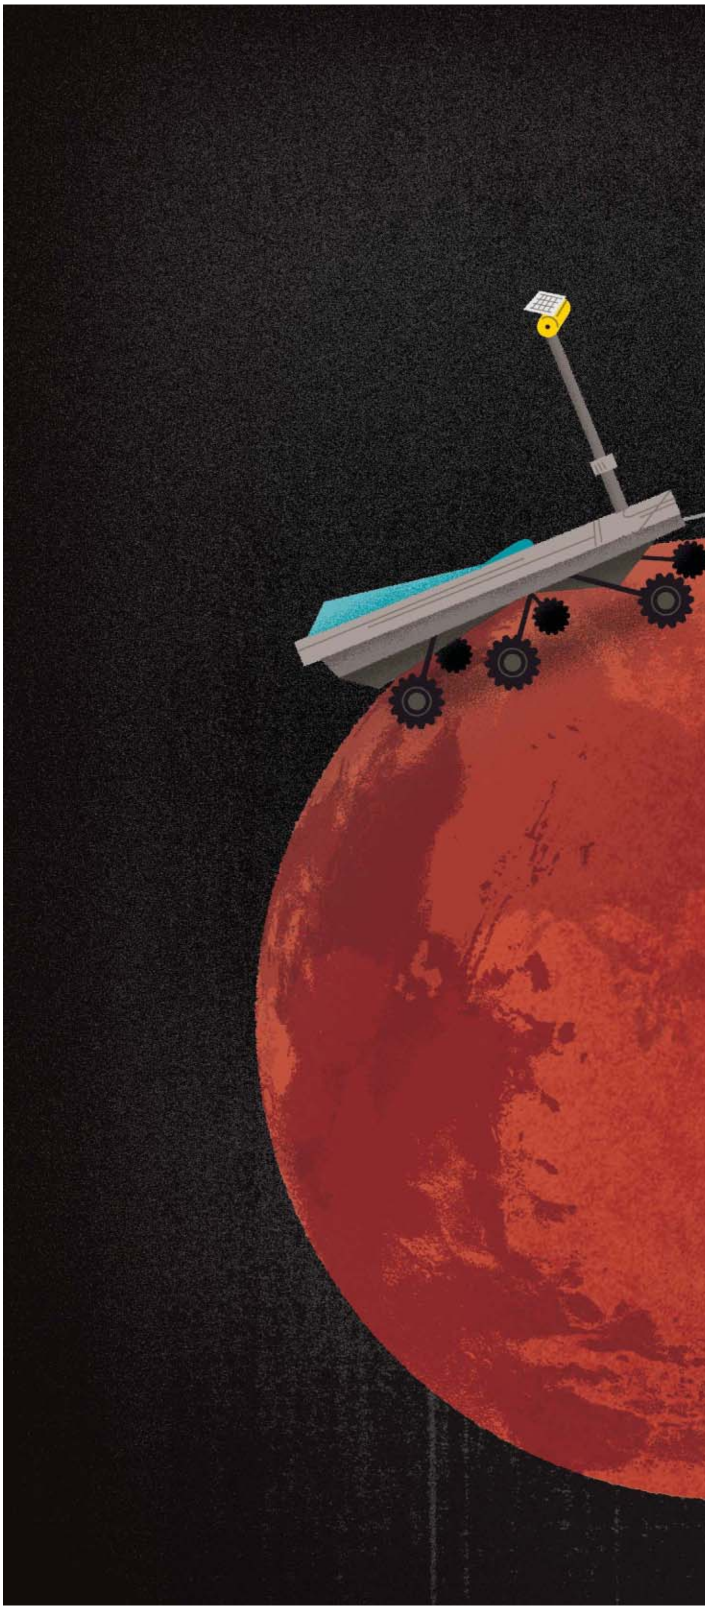

FRANK CHIMERO

Planetary Rovers

NO. 8

ARE WE ALONE?

Planetary rovers attempt to answer the most profound question in science

PETE THEISINGER stands at the back of the mission control room, his round, mustachioed face frozen in a nervous grin. Hunkered down at long rows of computer consoles, his engineers sit on the edges of their chairs. NASA's Jet Propulsion Laboratory is hanging on the brink of a jubilant victory—or a devastating failure.

Then the black-and-white images appear on a big projection screen, and the room explodes in cheers. Some 200 million kilometers from Earth, a little robotic rover called Spirit, built here in Pasadena, Calif., has awakened and called home, sending images of what it is seeing. And what it is seeing is the rocky plain of Gusev Crater, in the southern highlands of Mars.

Theisinger, the project manager of the Mars Exploration Rover mission, has only one word: "Wow."

Wow, indeed. Since that night in 2004, Spirit and its twin rover, Opportunity, which landed three weeks later, have embarked on an extraordinary journey of discovery. Designing, constructing, launching, and landing those rovers on Mars has become NASA's most thrilling and successful planetary mission ever.

Why bother to study rocks and dirt from a cold, desolate, remote world? Because the geology of Mars embodies a history that should help unravel our own, and because those Martian rocks may also hold the answer to a question we've been asking ourselves for a very long time: Are we alone?

"Finding evidence that life arose independently on another planet would be one of the most profound discoveries that humans could ever make," Steve Squyres, astronomy professor at Cornell University and the mission's chief scientist, writes in *Roving Mars* (Hyperion, 2005), his candid account of the project. Mars, he adds, "is a world that can help us learn our place in the cosmos."

In the past decade, planetary rovers have emerged as one of the most amazing exploration tools humanity has ever seen. They have also fostered scientific and technological innovations that should find applications on Earth, in areas such as

BET YOU DIDN'T KNOW
NASA's new rover, Curiosity, has approximately 50 000 parts—from microscrews to heavy actuators—or more than twice as many as an average car.

autonomous robotics, remote sensing, and materials engineering.

Above all, these robotic explorers have demonstrated that unmanned missions offer formidable rewards, with immensely smaller costs and risks than manned ones. Manned missions will surely remain on NASA's agenda; human boot prints on extraterrestrial soil are too powerful a draw to relinquish. But the success of the Mars rovers has proved that before we send humans, we ought to send robots.

NASA is not alone in advancing rover technology. The European Space Agency, in a joint mission with NASA, is building a next-generation Mars rover, called ExoMars, for launch in 2018. And other countries and even private companies have dreamed up rover plans of their own.

Right now, however, JPL is commanding all the attention. A new rover, the Mars Science Laboratory—named Curiosity in a contest—is scheduled to launch later this year. Compared to its golf-cart-size predecessors, it's a monster of a machine, the size of a Mini Cooper, weighing in at 900 kilograms, equipped with a nuclear power supply, and carrying 10 scientific instruments of unprecedented sophistication, including an advanced analytical system for detecting organic molecules. The mission: Determine whether conditions for life existed on Mars and were preserved—and if they were, find a sample.

MONSTER ROVER: The size of a small car, Curiosity is equipped with a nuclear power supply and 10 scientific instruments.

ERICO GUIZZO

NASA IS LIKE a planetary system where personalities, politics, budgets, and schedules revolve around each other in erratic orbits. Every now and then, these celestial bodies align and a promising mission becomes possible. That's how the idea of sending rovers to explore the geology of Mars emerged in the mid-1990s.

Every space mission builds on its predecessors. Before Spirit and Opportunity, a Soviet program called Lunokhod put an eight-wheeled, solar-powered rover on the moon in 1970 and another one in 1973. Controllers on Earth steered the rovers and operated their cameras and instruments in near real time.

But as interest shifted from the moon to more distant parts of the solar system, flyby probes, orbiters, and landers came to dominate the scene. Rover missions seemed too daunting—or merely unnecessary, in the case of a ball of gas like Jupiter.

In the mid-1970s, NASA's Viking program put sophisticated landers on Mars to search for signs of life. They didn't find any, but then again, they were pretty much looking at their feet. That's the main drawback of landers: They're stuck in one place. You could have the most intriguing rock sitting in front of you, but if you can't get to it, you'll never know what secrets it might hold.

Rovers made a comeback in 1996, as part of NASA's Pathfinder program. Pathfinder was essentially a lander mission, but it carried a small rover named Sojourner in its belly. The rover successfully roamed around, snapping photos and analyzing the chemical elements of rocks.

The results led NASA to start planning a more ambitious rover mission. In 1999 alone the agency lost two spacecraft, the Mars Climate Orbiter and the Mars Polar Lander, and it desperately needed a successful mission.

It was good timing. JPL had been developing a number of navigation and robotic technologies that it could put to use. After a series of solicitation bids, review panels, and approved, canceled, and newly revived proposals, one idea stood out in particular. Could NASA use Pathfinder's successful landing approach to place not a lander but a rover on Mars? This approach would rely on slowing the spacecraft as it entered the Martian atmosphere, using rockets and a parachute, and then deploying air bags to cushion the touchdown.

JPL engineers believed it could work. They would strip the Pathfinder lander down to its basic structure, cram a bigger and better-equipped rover inside, and keep the rest essentially the same. The mission got the green light in mid-2000, and the clock started ticking.

COMPANY TO WATCH: MALIN SPACE SCIENCE SYSTEMS, SAN DIEGO

Malin Space Science Systems, a small outfit based in San Diego, builds some of the most sophisticated cameras in the space business. As the eyes of two Mars orbiters, the company's cameras took hundreds of thousands of photos of the planet's surface. Now they will do the same on the ground, mounted on Curiosity, which will carry three science cameras, this time able to capture not only still but also moving images.

The launch window—when the orbits of Earth and Mars would be best aligned—was just three years away. That wasn't a whole lot of time to build a complex planetary rover. Yet NASA told JPL to build two, to double the chances of success.

Cornell's Squyres led the science team building the instruments that Spirit and Opportunity would take, while Theisinger headed the engineering team working on the design of the rovers. It didn't take long for them to come to a terrifying realization.

Though things looked good in preliminary studies, when it came time to refine their designs the teams discovered they had underestimated the size the rovers would eventually become. The rovers were too big for them to reuse the Pathfinder landing system. The EDL team—responsible for entry, descent, and landing—scrambled to redesign the parachute and air bag systems. The first tests failed badly, with prototype chutes and air bags blowing to pieces.

Other problems cropped up. Instruments that had been working misbehaved inside the thermal vacuum chamber that simulated the Martian atmosphere. An error in a telecommunications module deep inside Spirit forced engineers to reopen its electronic guts. In *Roving Mars*, Squyres describes one of the commands the software team sent to the rovers quite often during development: SHUTDOWN_DAMMIT.

At one point, when Spirit was already inside a Delta II rocket awaiting launch, engineers discovered they had accidentally blown the rover's main fuse during the final assembly. The problem nearly made NASA administrators kill the entire mission.

But in the end, the engineers did what engineers do best: They solved problems, one after another, with solutions that were sometimes ingenious and other times just good enough.

The machines they created—each cost some US $400 million, including launch and operation—are beautiful pieces of technology. Their solar panels unfold like origami. Their so-called rocker-bogie suspension systems allow each wheel to move up or down independently so the vehicles won't tilt excessively. And their software lets them receive navigation instructions and then drive autonomously, avoiding big boulders and stopping before cliffs.

With panoramic cameras, a microscopic imager, a rock abrasion tool, and three different spectrometers, the rovers made many discoveries about the geology and mineralogy of Mars. Among the most important was convincing evidence that Mars once had lots of water, an ingredient essential to life as we know it.

Expected to last 90 days, the rovers have worked for seven years. Spirit is stuck in sand and probably hasn't survived the last Martian winter, but Opportunity is still roving around. These robots answered many questions and also raised fresh ones. And that's why NASA is going back—this time with a bigger rover.

FEW ENGINEERING projects compare to building a spacecraft in terms of cost, complexity, and risk. These are one-of-a-kind machines that will face the harshest conditions—crushing accelerations, extreme temperatures, radiation storms—while stuffed with sensitive instruments and moving parts. And of course, once they leave the launchpad, there's no recalling them. The number of organizations that can build jet fighters and nuclear reactors is small. Fewer still can build spacecraft. JPL is one of these.

Tucked in a small campus in the San Gabriel foothills, JPL is NASA's lead

HONORABLE MENTION

Flat-panel TVs

LCD by TKO

Underdog LCD went from desktop to wall mount

Engineers have known for decades that getting the cathode ray tube out of the picture would allow TV sets to have bigger screens while making them as wispy as runway models. But plasma displays were too expensive and power hungry to break out of the high end of the TV market, while LCDs, despite their success in desktop computing, were even further out in the cold.

LCDs suffered from blurred images, ghosting, low contrast ratios, and colors that varied dramatically at different viewing angles. The game changed with several key innovations, notably NEC Corp.'s Overdrive technique, which doubled the standard voltage supplied to a pixel while using half the pulse width (speeding up response times from around 20 milliseconds to about 2 ms today).

Meanwhile, the success of LCDs in the desktop market created economies of scale, lowering the cost per unit. Together these developments induced electronics manufacturers to invest heavily in large-screen LCD technology for television. Energy efficiency and contrast ratio were dramatically improved when LED backlights began to replace fluorescent tubes. By the 2007 holiday season, cheap, svelte LCD TVs had elbowed aside plasma and CRT.

Now children under 15 will give bulky CRTs the same quizzical look they give manual typewriters and pay phones. *—Willie D. Jones*

MCKIBILLO

facility for the exploration of the solar system. Hardware built here has flown to the moon, Venus, Mars, Mercury, Jupiter, Saturn, Uranus, and Neptune.

Late last year, when I visited Theisinger, he proudly showed me photos and decals on the walls of his cramped office—trophies of his past victories. The 40-year JPL veteran is again leading an engineering team, the one building the Curiosity rover.

Compared to Spirit and Opportunity, the new rover is "an order of magnitude" harder to design and build, Theisinger says.

With a massive 2.3-meter-long arm, the rover will be able to push a percussion drill against rocks to extract samples. Wheels half a meter in diameter will let it traverse difficult terrain, off-road-style. And thanks to its plutonium-238 thermoelectric generator, it will be able to operate in the winter and at latitudes farther from the equator.

These capabilities are key to the success of Curiosity's mission. To find out whether environments habitable for microorganisms ever existed, the rover will have to go to places and perform scientific experiments beyond what its predecessors could do. The goal is to better understand how the availability of water, energy, and elements like carbon evolved on the surface of Mars.

"With [Spirit and Opportunity], what we can do is imagine what might have happened," says Caltech geologist John Grotzinger, the mission's chief scientist.

WHAT A SIGHT: NASA managers at the Jet Propulsion Laboratory celebrate after receiving the first Mars images from the Spirit rover.

BILL INGALLS/NASA/GETTY IMAGES

"With Curiosity, we're going to determine what actually happened."

Curiosity will carry 10 instruments from five countries. The most important is known as SAM, or sample analysis at Mars. This set of instruments takes in a rock or soil sample and uses a mass spectrometer, a tunable laser spectrometer, and a gas chromatograph to characterize its molecular structure and isotopic composition, and also to test for the presence of organic carbon.

Getting the 80 kilograms of science hardware—Spirit and Opportunity each carried 5 kg—to fit on the rover was one of the project's biggest challenges. But perhaps even harder was the design of the landing system. The new rover is so big and heavy that air bags won't work. Like Spirit and Opportunity, Curiosity will travel aboard a capsule known as an aeroshell, and after it enters the Martian atmosphere a parachute will unfurl. But then, still plunging at supersonic speeds, the craft will fire retro rockets, decelerating to a gentle descent until it's just 20 meters from the ground. That's when the rover will detach from a supporting structure and lower itself on cables—much like a commando rappelling from a helicopter.

When the rover touches the ground, explosive bolt cutters will release the cables, allowing the aeroshell to fly away and crash at a safe distance. Steve Lee, a member of the EDL team, describes all this—the "sky crane," they call it—with a satisfied look on his face. When I suggest that this would make for a good movie, he smiles even more broadly: A downward camera on the rover would capture not only a top view of the landing site but also the entire touchdown action.

NASA knows the power that images of extraterrestrial worlds have to capture people's imaginations. So it's no surprise the rover is also equipped with cameras capable of obtaining high-definition stills and movies, which could be converted into 3-D. (James Cameron is a member of the camera team. Really.)

KNOWING WHAT THE JPL crew went through with their previous rovers, you might think they've grown too bold. Can they pull it off?

The project has already hit some major snags. A novel lubricant-free motor that NASA wanted to use failed during tests, forcing the engineers to go back to traditional designs.

Another setback involved a turbomolecular vacuum pump in the rover's SAM instrument. It had to be redesigned and retested, delaying its delivery—and leaving project managers wringing their hands.

But as before, the engineers march forward. At JPL's vast clean room, construction and testing of the rover's final components proceed at a frantic pace. Launch is scheduled for late November or early December.

Why go to Mars? There are many reasons, JPL engineers will tell you. But some on the team have personal motivations. Pete Theisinger says he once received a letter from a man in Ohio. The man wrote that one day he was watching TV with his young son, and they saw Spirit and Opportunity and the images their cameras had captured of the Martian landscape. The letter included a photo of the man's son. He was building a rover out of Lego blocks. —ERICO GUIZZO

NO. 9

Flexible AC Transmission

THE FACTS MACHINE

Flexible power electronics will make the smart grid smart

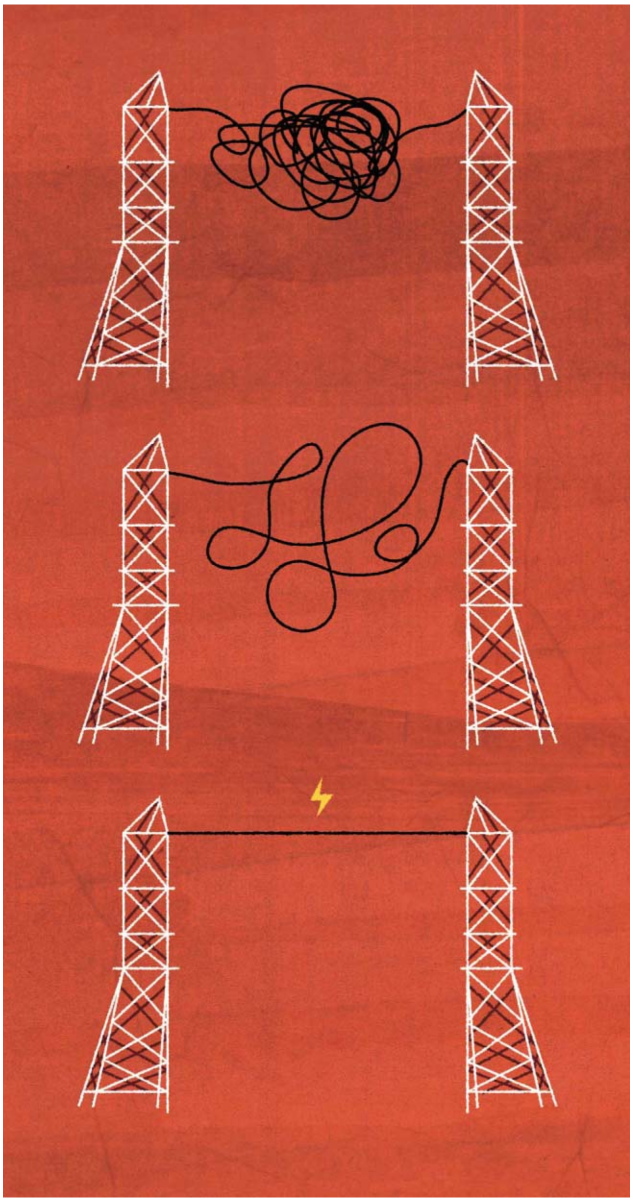
FRANK CHIMERO

POWER SYSTEMS MUST juggle supply and demand while guaranteeing glitch-free alternating current 24/7. To deliver it, engineers once had no choice but to design grids that were as passive as the Roman aqueducts, which could carry water anywhere, so long as it went downhill. But over the past decade, a confluence of innovations, regulatory change, and sheer watt-squeezing necessity has hatched a marvelous advance, one that has begun to realize the long-standing dream of pushing current where it wouldn't ordinarily go.

And it's happening not a moment too soon. These flexible AC transmission systems, or FACTS, promise to save energy in a big way by making possible the smart grid, which utilities hope will reconfigure power flows in real time, maximizing throughput and minimizing losses. They should also make it possible to smoothly incorporate wind, solar, and other intrinsically intermittent sources of energy into the grid.

The key word in the FACTS acronym is the first one: *flexible*. Modern hydraulic engineers use pumps to push water against the force of gravity, so they save immensely on bridges and tunnels in comparison with their Roman predecessors. And think about how aeronautical engineers can manipulate control surfaces from second to second to keep aloft a plane that would otherwise be only slightly more flyable than a brick. Real-time control of power systems promises similar rewards.

FACTS proved its mettle in the 1990s in demonstrations led by the power industry's Electric Power Research Institute (EPRI) and grid equipment manufacturers, such as GE and Zurich-based ABB. Over the past decade, FACTS has gone commercial and is "penetrating the network everywhere," says Claes Rytoft, chief technology officer for power systems with ABB.

Applications vary widely to fit local grid conditions and challenges.

COMPANY TO WATCH: CREE, DURHAM, N.C.
Efficient, high-temperature silicon carbide switches could slash power losses from silicon-based FACTS controllers by more than 50 percent. Cree leads a US $3.7 million project with the U.S. government's ARPA-E high-risk energy R&D fund to engineer 15- to 20-kilovolt silicon carbide power modules ready for grid-scale power flows.

"Waking giant" countries such as China and India apply FACTS to maximize the power carried by every single new transmission line they install, thus minimizing the cost of grid expansion. North American and European utilities, meanwhile, rely on FACTS primarily to tame the unruly output from wind turbines that could otherwise destabilize their congested grids. The Virginia-based consultancy NanoMarkets projects a steep growth curve for global FACTS installations, from US $330 million this year to $775 million in 2017.

THE FIRST FACTS controllers emerged in the 1970s as an improved means of balancing the two types of power that coexist on AC networks: active and reactive power. Active power, the familiar watts consumed by lightbulbs and toasters, is the product of voltage and the component of an alternating current that is in phase with the voltage. The component that is out of phase multiplied by the voltage gives the reactive power, which is measured in volt-amperes reactive or VARs (or more commonly, megavars).

Reactive power is a necessary evil: It does no work, and yet you have to add it to move active power. Reactive power results when electricity flows through an inductor or a capacitor, which causes the current to lag (the inductor) or lead (the capacitor) the voltage. When the current lags, engineers refer to it as negative reactive power; when it leads, they call it positive. It's the negative sort that tends to occur in lines, transformers, motors, and even some generators. Too much negative reactive power will cause the voltage to "sag," a condition that can damage electrical equipment. And the damage can spread: Inadequate reactive power support during peak periods is a common contributor to cascading failures—including the 2003 blackout that toppled grids from Ottawa to Baltimore.

Typically, utilities compensate for the negative reactive power caused by inductors by injecting positive reactive power into the system. Traditionally, there are two ways of doing that: by patching banks of capacitors into a circuit to convert some of its megawatts into megavars, or by tuning the generators in conventional power stations to produce current waveforms that lead voltage. FACTS got started as a more dynamic solution, and it has become increasingly relevant as deregulation has progressively turned the electricity business into a kind of promiscuous dating game, whereby supply is married and remarried frequently to match demand, sometimes on an hourly basis, and without much regard for the capabilities of the transmission assets connecting those scattered centers of supply and demand.

By means of such matchups, FACTS allows system managers to send more power over a line than it could otherwise support. The increase can be as high as 50 percent, says Ram Adapa, a technical manager for EPRI. Stability enhancement accounts for part of the boost, allowing grid operators to operate lines closer to their thermal limits.

The heart of a modern

HONORABLE MENTION

Digital 3-D

Better, Bit by Bit

With digital cinematography, 3-D finally makes sense

MCKIBILLO

Thirty years ago, Buzz Hays started his 3-D career with a lesson from Alfred Hitchcock, by helping to restore the master's 1954 *Dial M for Murder* to its original paper-glasses glory. Now he's the chief instructor at the Sony 3D Technology Center, in Culver City, Calif., where he teaches other filmmakers newly learned subtleties of stereoscopic storytelling.

Stereoscopic systems—which present a different perspective to each eye—failed in the '50s. But to critics who say, "Been there, done that," Hays responds that stereoscopy is in a "constant state of refinement," incorporating new projection technologies, brighter images, and better ways to capture footage.

Digitization was the key, says Bruno Sargeant, who manages television work at Autodesk, a computer graphics company. He says that a completely digital "pipeline" from camera to presentation has finally made 3-D filmmaking financially viable.

Joshua Greer, president and cofounder of RealD, in Beverly Hills, Calif., agrees. Theaters equipped with his company's digital projection equipment account for about 80 percent of the market share, measured in terms of 3-D films' box office receipts. His company's 3-D cinema system uses a liquid crystal filter to circularly polarize light. When combined with today's polarized shades, that means different images for the left and right eyes. Such high-quality systems, Greer says, were "just not possible before digital came along." Now business is booming: In June of 2009, RealD's light shows jumped out of 2600 theater screens; this past summer, 7500. There were 6 such films in 2008; this year, the company expects there to be 35.

Although some now point to glasses-free systems and even "holography" (which requires projecting images on mist or smoke), Hays says the next big thing will be when "mere mortals" can record and then watch their own films—including those of baby's first step—in 3-D. —*Joseph Calamia*

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

BET YOU DIDN'T KNOW

The world's most sophisticated FACTS controller keeps New York City lit when lines upstate start to max out. The system, completed by the New York Power Authority and the Electric Power Research Institute in 2003, pulls hundreds of megawatts from NYPA's congested Albany circuit and pushes it onto the Catskills line, a feat that has yet to be rivaled.

FACTS controller is an array of solid-state switches, often coupled with capacitors. Typically, the solid-state switches open to tap power from the line and charge a capacitor; then the switches fire in sequence to create a synthetic AC waveform with precisely the needed phase difference between current and voltage. That waveform is then applied to the grid. By precisely varying the phase difference, the FACTS controller can add or subtract reactive power in fine increments.

Impressive as it is, such dynamic voltage regulation is the simplest of the FACTS grid control modes. FACTS innovators went further in the 1990s by exploiting newly developed high-power semiconductor switches that could switch at frequencies higher than the standard 50- or 60-hertz AC cycles.

With relatively advanced switches, such as insulated-gate bipolar transistors, FACTS controllers could simultaneously regulate voltage and surgically remove a variety of glitches in the AC signal. One such FACTS device, the static synchronous compensator, or statcom, has played a decisive role in the more than tenfold rise in wind power capacity worldwide over the past decade. A Siemens-built statcom, for example, is stabilizing flows from the world's largest offshore wind farm, completed this past September, whose 100 3-megawatt wind turbines should feed enough energy to the United Kingdom's grid during the year to supply more than 200 000 homes.

MASSIVE WIND FARMS barely raise an eyebrow today, but Charles Stankiewicz, executive vice president at power equipment manufacturer American Superconductor Corp., says that just a few years ago they still rattled many transmission engineers, who viewed their gusty, noisy power signal as a threat to grid stability. Stankiewicz suggests that many of the more than 70 wind farms stabilized by his firm's FACTS controllers might never have been erected without FACTS technology. "It would have been one more excuse that allowed the electric utilities that maybe weren't inclined to accept wind power to basically say, we're not going to do it," says Stankiewicz.

The most advanced FACTS go beyond stabilizing a line to reducing its apparent impedance, so grid managers can actually push more power down it. This application went commercial in 1998, when Brazil commissioned a pair of 1000-kilometer, 500-kilovolt lines to link its northern grid, replete with Amazonian hydropower, to the southern grid serving its coastal population centers. A FACTS controller near the northern end of the line simultaneously drives power and damps down destabilizing feedback signals.

The Brazilian project marked a breakthrough for AC transmission, where high impedance had previously limited lines to just a few hundred miles. The payoff was enhanced access to electricity in Brazil's hinterland, according to Stig Nilsson, who launched EPRI's FACTS program in the 1980s and now serves as a principal engineer with Exponent, a Menlo Park, Calif.–based consultancy. Had Brazil relied exclusively on lower-resistance high-voltage direct current (HVDC) technology to deliver power to the coast, it could not have tapped the line along the way to serve interior communities.

HVDC lines passing over as yet unelectrified regions of sub-Saharan Africa have inspired regional discord, and Nilsson says analogous concerns could upend proposals to build HVDC lines in the United States to carry wind power from the Dakotas to the Eastern Seaboard: "People are going to say, 'Why should I allow you to send wind power through my state that I'm not even going to get access to?' "

INDIA AND CHINA, meanwhile, have applied FACTS as a cost-cutting tool to ensure that AC lines deliver their full potential and thus reduce the number of lines required. "Instead of building three lines, they may be able to build two and put in a FACTS device," says Nilsson.

EPRI and the New York Power Authority pushed flow control to a technical pinnacle in 2003, showing that coordinated control of two FACTS controllers could pull power off one line and drive it down another, deftly guiding power around roadblocks in the grid. However, extensive application of such advanced power control to guide flows through meshed grids will require an improvement of FACTS controllers' efficiency and a greater recognition of their value. Grid experts expect both to happen in the years ahead.

Efficiency will rise as FACTS manufacturers integrate energy-efficient silicon carbide semiconductor switches that are being commercialized today for low-power applications and which EPRI's Adapa expects will be available for high-voltage application within five years [see "Company to Watch"]. High-frequency silicon insulated-gate bipolar transistors (IGBTs) used in sophisticated FACTS controllers are just 92 to 94 percent efficient and require active cooling to stay below 100 to 110 °C. The result is that a large controller consumes megawatts of power. By contrast, silicon-carbide IGBTs should be 96 to 97 percent efficient and can operate at 200 °C or higher, slashing the power "tax" for FACTS-based routing of current.

The same ideas will also make HVDC more efficient, because the latest generation of HVDC converters are close technological cousins to advanced FACTS controllers. These newer HVDC systems already make possible underwater connections that are too long for AC cables. They also allow lines to reach wind farms too far offshore for AC cables and to interconnect AC systems that are adjacent but not synchronized (those of Mexico and Texas and the two main U.S. grid systems, for instance). Silicon carbide switches could make the conversion of AC lines to DC a direct competitor to FACTS as a means of assuring power flows in the heart of meshed grids.

Meanwhile, the leading alternative to both—adding new lines—is not getting any easier. New lines can take a decade to build, if they're not NIMBYed to death by disgruntled property owners or throttled in the womb by environmental advocates, state regulators, and rival power industry players. Each new line also complicates an AC grid, adds Abdel-Aty Edris, until recently a power-electronics expert with Siemens, who helped design EPRI's power-driving FACTS control schemes. "We already have dynamics that we don't control and don't even comprehend," says Edris.

Edris's vision is for a smart grid—both flexible and capable of handling increasing flows—via pervasive application of grid controllers, whether FACTS or HVDC. "My vision is to have a grid which is fully power-electronics integrated—a grid that's as robust as an integrated circuit," says Edris. Rytoft agrees: "Power electronics is the key technology for the grids of the future." —PETER FAIRLEY

FRANK CHIMERO

Digital Photography NO. 10

THE POWER OF PIXELS

Digital photography changed not only how we take pictures but also how we communicate

TEN YEARS AGO, photography for the most part meant film. We carried rolls of it on vacation, dropped it off for processing when we got back, picked up our prints, then put them in albums or scrapbooks or, more typically, in cardboard boxes. On occasion, we thought about sending a duplicate to distant relatives, but we'd often forget. Photographs were for documenting our history, for framing, for saving.

What a difference a decade makes! The vast majority of us haven't handled a roll of film in years—it's a retro novelty at best. Digital technology has changed the very nature of photography. Digital images are free and easy and can be instantly distributed. As a result, the vast majority of photos are no longer taken to capture special moments; they're used to communicate the ordinary, with less forethought than a phone call.

Of course, digital cameras didn't simply materialize in our hands a few years ago, although it may seem like it. You could trace their history back to 1969, when the charge-coupled device (CCD) was invented at Bell Telephone Laboratories, or to 1957, when the first digital image scanner was created at the U.S. National Bureau of Standards.

Or you could start in December 1975, when Steven Sasson, an electrical engineer at Eastman Kodak Co., in Rochester, N.Y., became the first person to pick up a digital camera and take a picture.

Sasson, hired by Kodak in 1973, fresh from a master's program at Rensselaer Polytechnic Institute, in Troy, N.Y., was a fish out of water. Kodak was essentially a company staffed by chemical and mechanical engineers, but in the early '70s it started hiring a handful of EEs to develop electronic controls for cameras, like exposure systems and motor drives. One of Sasson's first assignments was to check out the new 100-by 100-pixel CCD chip developed by Fairchild Semiconductor, to see if it would be useful for Kodak.

Sasson decided that the best way to study the chip was to build it into a camera. Being an electrical engineer, he thought it would be cool to create a new, all-electronic camera, with no moving parts, rather than sticking the CCD into an existing mechanical body. He spent about a year on the effort, working on it in between other assignments, cobbling together the materials he needed from catalogs and used-parts bins. He found a tiny digital-data cassette recorder, adapted an analog-to-digital converter from a Motorola digital voltmeter, and grabbed a lens from an old 8-mm movie camera.

In December 1975 he pointed the completed prototype at a lab technician and took his first picture. He then went to his supervisor and told him that he'd turned that CCD into a working camera.

Sasson recalls his supervisor saying he'd bring some people to the lab for a demo. No, Sasson responded, it's portable. I can bring it to you. His supervisor was amazed.

Sasson started a round of demos, bringing groups of Kodak engineers and executives into a conference room, taking a quick picture of one of them, and then popping the tape out and putting it into a player to show it on a TV screen. "In 1976 we were taking pictures without

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

COMPANY TO WATCH: TESSERA TECHNOLOGIES, SAN JOSE, CALIF.

In 2008 Tessera bought FotoNation, the company responsible for a vast number of the features that make today's digital cameras so easy to use—like automatic red-eye correction, face recognition, and smile detection. The more that cameras are used for instant communications under suboptimal conditions, the more important such features become. Look to Tessera to continue to make it ever easier to get the perfect shot.

film and viewing them without paper," he says.

Sasson's project never went beyond the prototype stage. At the time, he told Kodak executives that digital cameras wouldn't catch on until they could produce images with 2 million pixels; he thought that day would come in 1990 or 1995. And the executives, he recalls, though recognizing that this would be earth-shattering for the film photography business, believed they didn't need to become too concerned because it wouldn't matter for a long time. Sasson built more cameras at Kodak over the years. Then, in 1990, he moved to the output side of digital imaging, developing color printers. His original digital camera patent, issued in 1978, expired in 1995.

Meanwhile, in 1981, Sony came out with an analog electronic camera, the Mavica. It recorded images using a television video signal, storing them on a floppy disk.

"I liked it because it woke Kodak up," Sasson recalls. "I also liked it because I knew that it was not going to succeed; it was analog, and to succeed it had to be digital."

About seven years later, Kodak created the first commercial megapixel digital camera, called the Hawkeye II Imaging Accessory. It was sold at a list price of US $23 000 each to U.S. government organizations; one camera went along on a shuttle mission in 1991.

Then in 1991, the company introduced a commercial black-and-white digital camera, the Kodak Professional Digital Camera System, later referred to as the DCS 100. In a sense, it was a step back from Sasson's prototype, because it wasn't an all-in-one device; instead, the system tethered a modified Nikon camera to 5 kilograms of electronics in a shoulder bag. It was marketed to news organizations at $20 000 to $25 000.

"Journalists just laughed at it," recalls John Henshall, then president of the British Institute of Professional Photography and a consultant with Kodak. But a few did begin using the device, because of two key features—it enabled them to immediately review the captured image on the electronic display, and it was possible to easily transmit these images by dial-up modem.

Kodak's $9995 DCS 200 in 1992 put all the electronics in the camera. Recalls photographer Stephen Johnson, "It was pretty amazing. I took my first images with it walking through the snow in Camden, Maine."

But it took Apple's marketing to finally make Sasson's vision of a handheld all-electronic consumer camera a reality. The under-$1000 Apple QuickTake 100, designed and manufactured for Apple by Kodak, came out in 1994. At its highest resolution, 640 by 480, it could store up to eight images on its internal memory. It sold only about 50 000 units, but it was a huge landmark.

"Apple legitimized the category," says Alexis Gerard, founder and president of analyst firm Future Image and of the 6Sight Future of Imaging Conference.

It came shortly after the creation of JPEG, the image compression standard that made the most of memory—still very expensive—and of Internet bandwidth. "Having a technology to crunch those huge files down to what the infrastructure could deal with was very important," Gerard says, adding, "Without JPEG, we would have had to wait another five to seven years" for the technology to catch up. And JPEG meant that digital photos taken by different cameras were fully interoperable.

Apple's second digital camera, the QuickTake 150, again from Kodak, offered JPEG compression. After that, things quickly marched forward. The Casio QV10 in 1995 had the

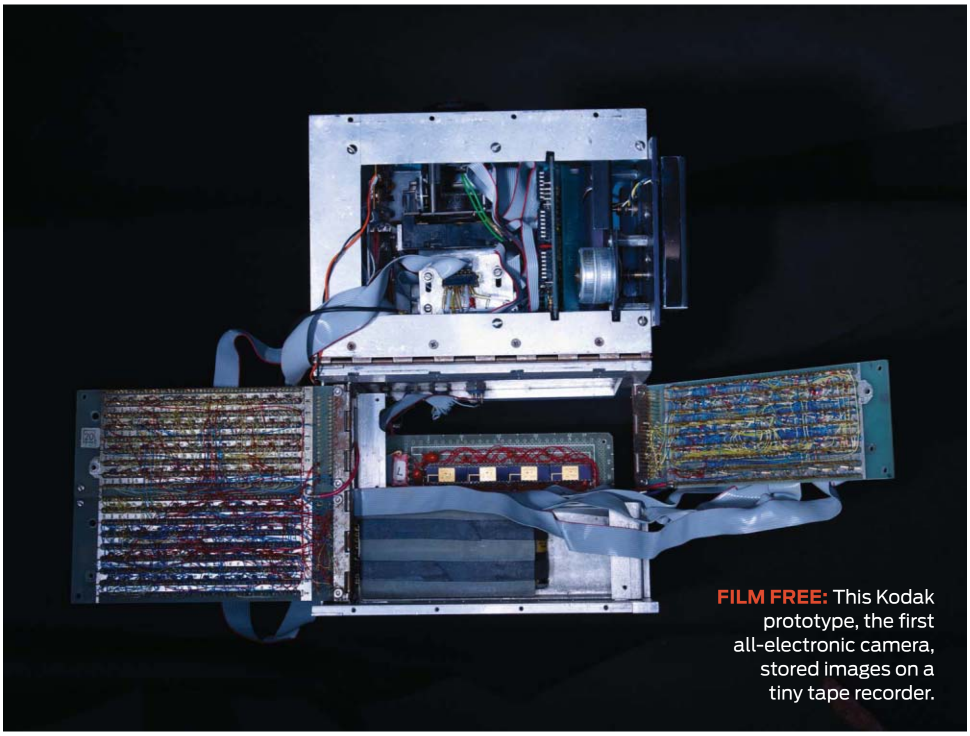

FILM FREE: This Kodak prototype, the first all-electronic camera, stored images on a tiny tape recorder.

DAVID YELLEN (2)

SAY CHEESE! Electrical engineer Steven Sasson took the first-ever digital snapshot in December 1975.

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

1991
Kodak DCS

1992
Kodak DCS 200

1994
Apple QuickTake 100

first built-in liquid-crystal display. Kodak introduced the first megapixel consumer color camera, the DC210, in 1997. The first camera phones appeared in Japan in 2001; they hit the United States in 2004.

"It was like a snowball rolling down a mountain; it gathered more and more snow until it blew away everything in its track," says Henshall.

The snowball also changed the essential meaning of photography. The principal purpose of photography had been to capture images for posterity. Today that is no longer primary. "Images are being used for more than just memories," Gerard says. "Many are not intended to be stored but purely to communicate information that only has value in the moment: Where did I park my car? What does this office space we're considering look like?"

Even as an art form, photography is changing, Gerard says. "If you point a camera at a good image in the real world, you will likely get a good image—there is no learning curve. For the first time we have a creative tool that people can jump into right away."

While digital photography has vanquished film, it is far from perfect. Although some people say that the cameras themselves could do a better job at matching what the eyes see—going to three dimensions, in particular—for the most part, it's not the cameras themselves that need improving. "The capability of the cameras being sold today far outstrips the average consumer's ability to use them," Sasson says. The problem is what to do with the images once you've taken them. The scrapbooks and shoeboxes of the film world are being replicated in digital forms, but they're overloaded and becoming impossible to manage. They're also not necessarily as reliable as a shoebox: Can you trust that your online photo storage company will be around in 50 years or that computers will read old camera formats?

"This is the last frontier," Sasson says. "How do you manage these images? How do you save them for 50 or 60 years, with format obsolescence, changing standards? Images are the only digital files that get more valuable the older they get." —TEKLA S. PERRY

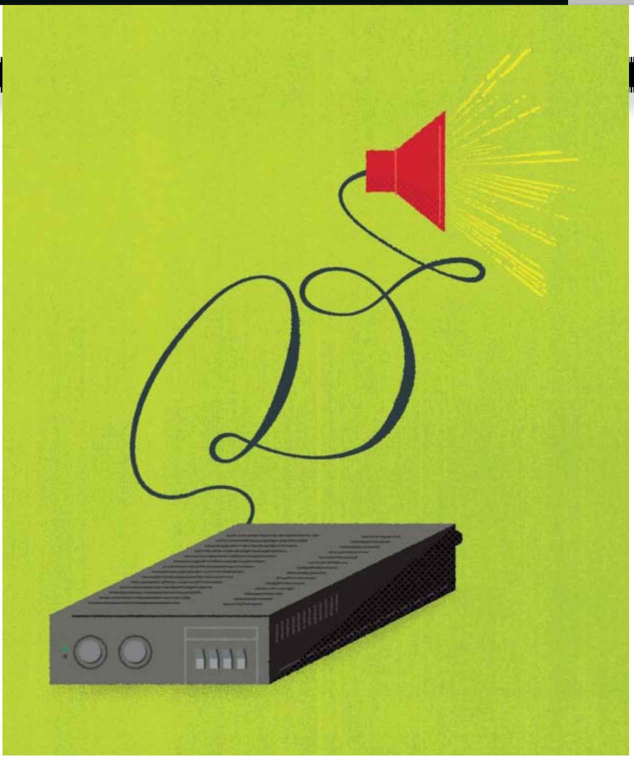

NO. 11

Class-D Audio

THE POWER AND THE GLORY

A quiet revolution is transforming audio electronics

EVEN IN THE go-go world of high tech, it's pretty rare that a technological leap delivers both markedly superior performance and stunningly greater efficiency. That neat trick happened with class-D audio amplifiers, which now dominate the market for applications in car stereos, home-theater-in-a-box systems, television sets, and personal computers.

Their success has been a long time coming. The first commercially available class-D amps came in the 1960s from the British company Sinclair Radionics (now Thurlby Thandar Instruments), but they didn't work well. Somewhat better ones came along in the 1970s and 1980s: John Ulrick designed a couple of class-D amps for Infinity Systems in the early and mid-1970s, and a decade later Brian Attwood did a series of "digital energy conversion amplifiers" for Peavey Electronics Corp. But in those days class-D theory was ahead of implementation, because the available components simply weren't good enough to produce high-quality sound.

The problem was speed: Class-D amps sample the input audio waveform

CLOCKWISE FROM TOP RIGHT: FRANK CHIMERO; JOHN HENSHALL (2); KODAK

COMPANY TO WATCH: INTERSIL D2AUDIO, MILPITAS, CALIF.

Intersil D2Audio sells products powered by its Digital Audio Engine technology, which combines class-D amplifiers with signal processing to minimize the effects of noisy power supplies or harsh automotive environments.

hundreds of thousands of times a second before amplifying it, and it wasn't until the 1990s that cheap and reliable MOSFETs became available that were fast enough to sample a waveform at such high frequencies, says Bruno Putzeys, the chief engineer at Hypex Electronics. Hypex, in Groningen, Netherlands, is a leading maker of audiophile class-D amplifier modules, which are incorporated into products sold by other firms.

"What it took was for a couple of companies to take the plunge," Putzeys adds. Those pioneers were Tact Audio, ICEpower, and Tripath Technology, all of which released their first class-D offerings in the late 1990s. Tact, now called Lyngdorf, rocked the audio community in 1998 with a US $9800 class-D amplifier called the TacT Millennium, which was designed by the Danish engineer Lars Risbo. It dazzled as much for its industrial design as for its engineering: The amp had a large volume knob with a digital display in its center.

ICEpower was founded jointly by another Danish engineer, Karsten Nielsen, who designed the company's first amplifiers and amplifier modules, and the Danish audio company Bang & Olufsen. ICEpower is best known for its combined power supply and amplifier modules, although its single biggest source of income is its MobileSound line of amplifier chips for cellphones.

Tripath, in San Jose, Calif., was the first to introduce a class-D amplifier chip, the TA1101, in 1996. It was used in Apple's celebrated Power Mac G4 Cube. Tripath's first big hit, the TA2020, in 1998, could funnel 20 watts per channel into 4-ohm speakers (see "25 Microchips That Shook the World," http://spectrum.ieee.org/25microchips0509). It was used in ministereos and early flat-screen TVs. "We were shipping millions of chips every quarter to companies like Samsung, Panasonic, Toshiba, Sanyo, NEC, Onkyo, and Sony," says Adya Tripathi, who was the company's president and CEO.

Large stocks of Tripath chips are still available. They feed an e-Bay market for tiny, ultracheap (as low as $15) amps from China and also a thriving business among hobbyists and DIYers.

Today almost all of the audio amplifier market, from cheap cellphone chips to unbelievably expensive home hi-fi, is split between class-D and class-AB amplifiers; the latter are the long-established technology, and they still dominate in home audio and also in mobile music players, including iPods, MP3s, and smartphones. Those pocket players use micropower amps, which put out fewer than 100 milliwatts per channel. Class-D chips are available for these applications, but they can't yet match the price of the class-AB chips that are available for a few cents per chip from such giants as National Semiconductor, Maxim, Texas Instruments, Sanyo, and various Taiwanese companies. Millions of those amps

FIRST CLASS: Audio amp designer Bruno Putzeys of Hypex Electronics shows that very good sound can come from small packages.

THOMAS VANHAUTE

are sold every year. And class-AB isn't even the only other competitor in this fast-growing field: Class-G amp chips share class-D's high efficiency and are also supposedly easier to integrate into a dense and highly complicated system like a smartphone.

Nevertheless, the micropower category is still a big market possibility for class-D chips, Tripathi insists. It's only a matter of time before the huge efficiency advantage of class-D tips the scales in their favor, just as class-AB amps slowly but surely displaced class-A decades ago. Some day, without a doubt, it will be a class-D world.

In the meantime, class-D amps are making steady inroads in the home market, where they can exploit to full advantage modern components that are cheap but have very high performance: MOSFETs just keep getting better and better and cheaper and cheaper. And for the technically inclined, amplifier modules and kits allow anyone with modest soldering skills to get class-D sound, highly precise and detailed, for a few hundred dollars. One of our favorite kits is the SDS-224 kit (http://classdaudio.com/index.php/sds-224-kit.html) from Class D Audio, with its robust power supply.

Class-D amps are starting to do well even in the rarefied high-end niche, where an amplifier can cost many thousands of dollars. Putzeys just completed and tested an amplifier design with total harmonic distortion of 0.001 percent at full power across the entire audio range, scalable to 2000 W and beyond. It is a specification that would have been dismissed as science fiction before class-D amplifiers. At any price.

—Glenn Zorpette

BET YOU DIDN'T KNOW

The first class-D audio amplifier offered to the public, in 1964, was an £8 kit from the British company Sinclair Radionics. Its power output was so far below the claimed 10 watts that *Wireless World* magazine, deluged with complaints, supposedly declined to take future advertising from Sinclair.

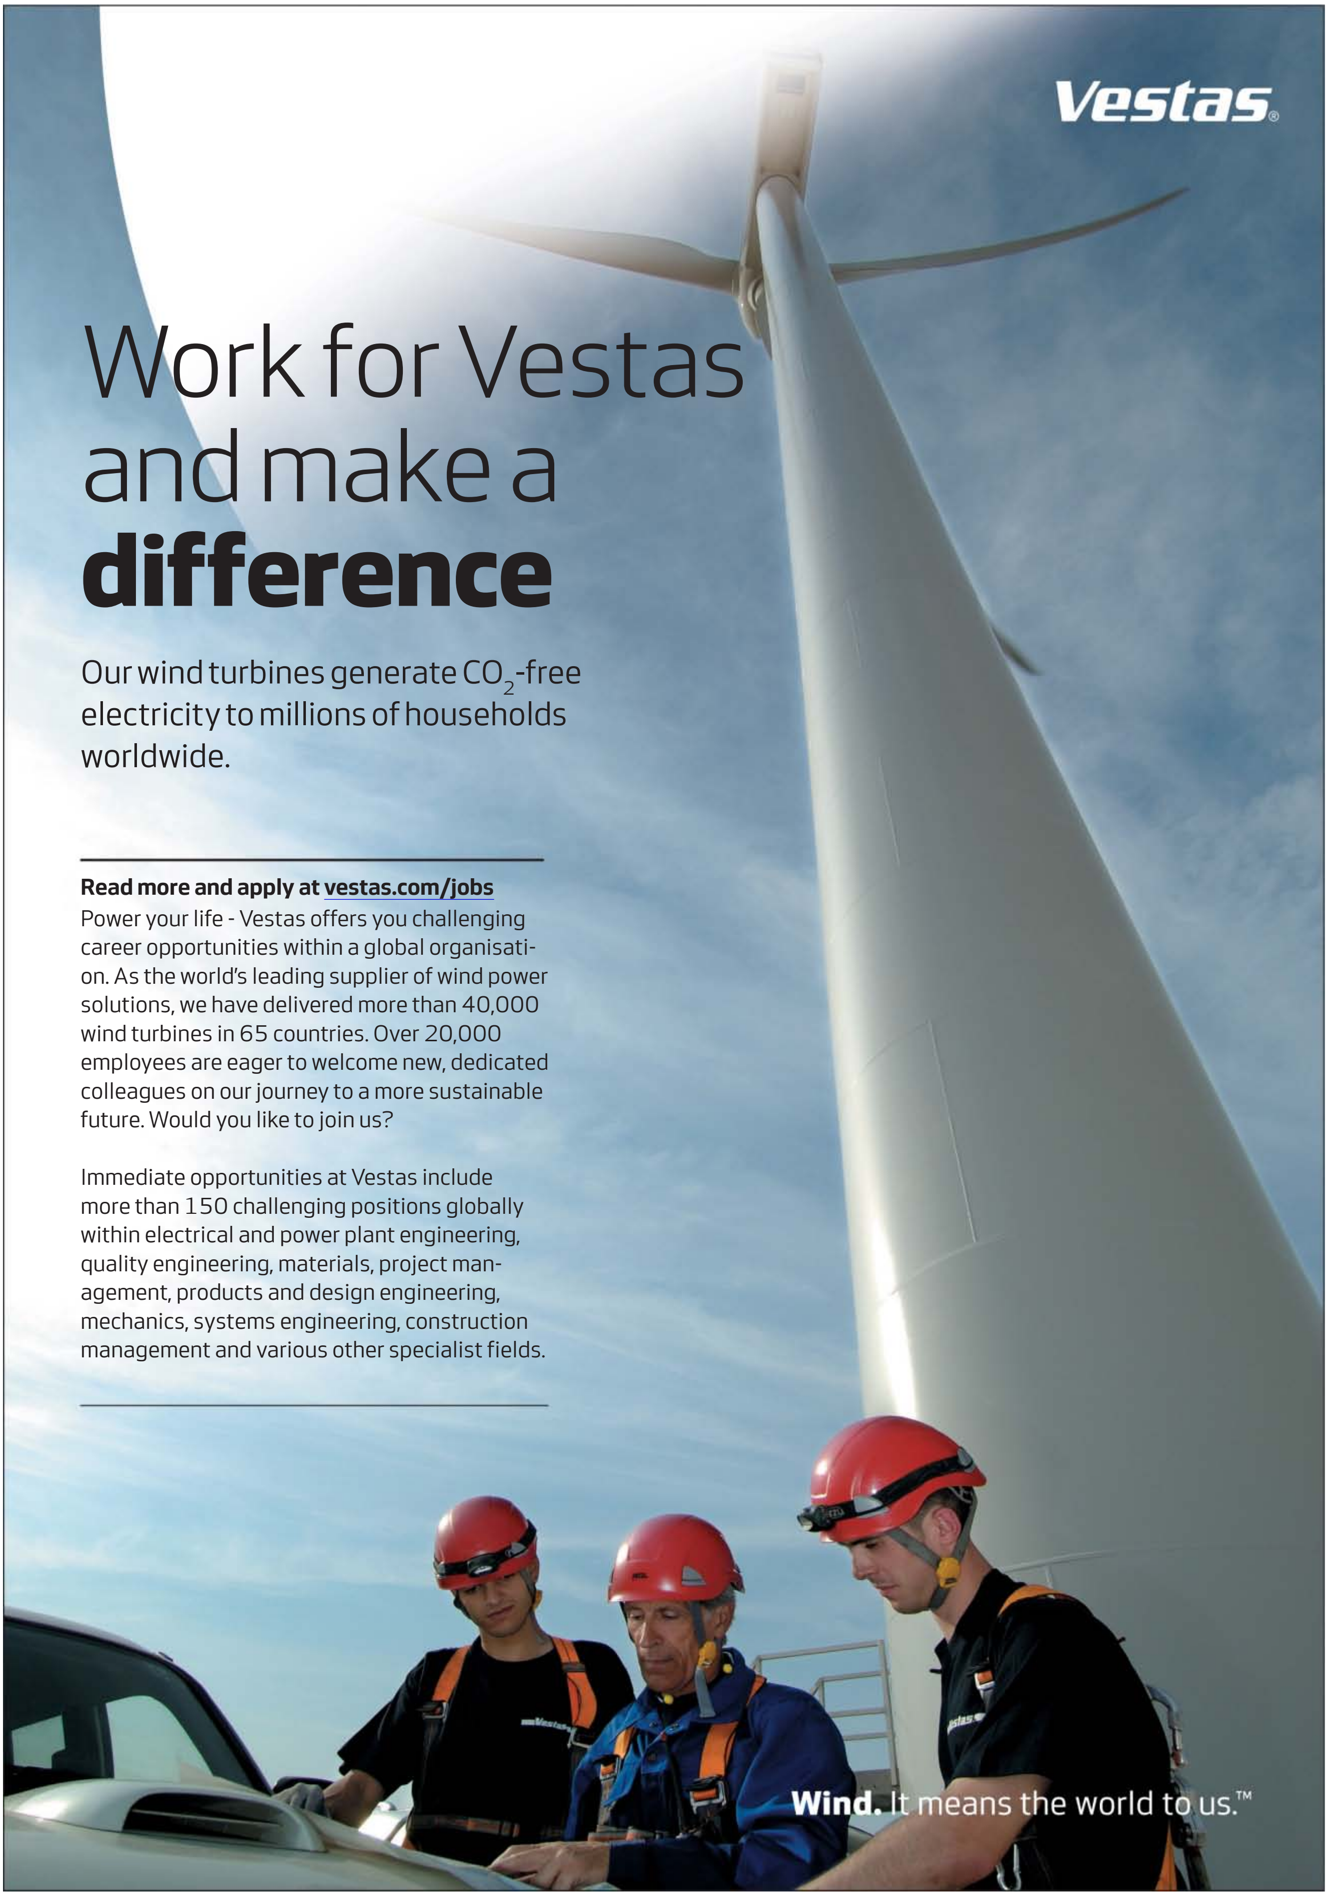
Vestas
Work for Vestas and make a **difference**
Our wind turbines generate CO_2-free electricity to millions of households worldwide.
Read more and apply at vestas.com/jobs
Power your life - Vestas offers you challenging career opportunities within a global organisation. As the world's leading supplier of wind power solutions, we have delivered more than 40,000 wind turbines in 65 countries. Over 20,000 employees are eager to welcome new, dedicated colleagues on our journey to a more sustainable future. Would you like to join us?
Immediate opportunities at Vestas include more than 150 challenging positions globally within electrical and power plant engineering, quality engineering, materials, project management, products and design engineering, mechanics, systems engineering, construction management and various other specialist fields.
Wind. It means the world to us.™

the data

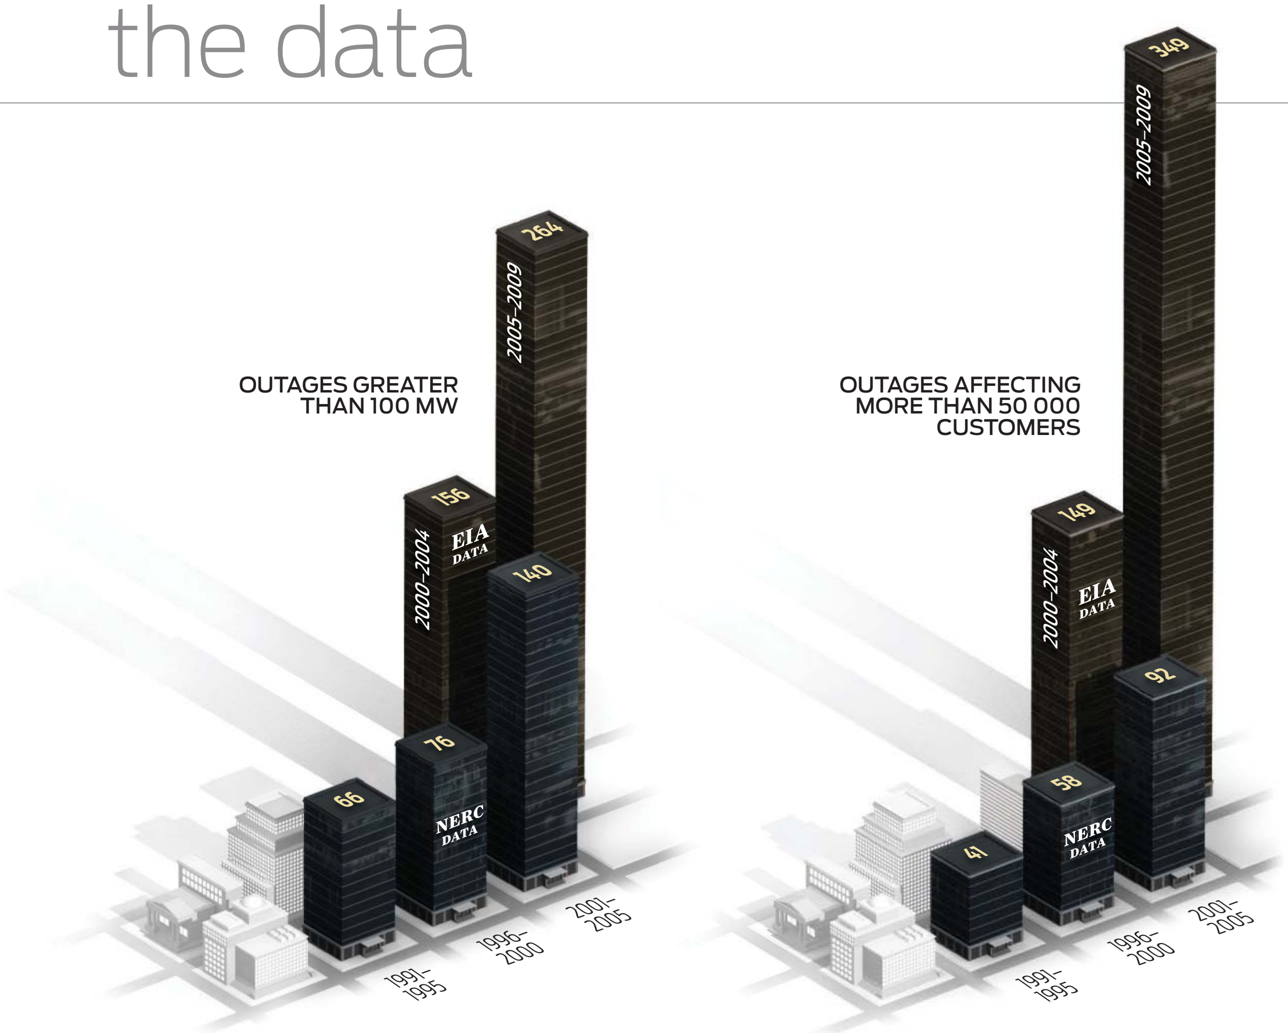

U.S. Grid Gets Less Reliable

THE U.S. electrical grid has been plagued by ever more and ever worse blackouts over the past 15 years. In an average year, outages total 92 minutes per year in the Midwest and 214 minutes in the Northeast. Japan, by contrast, averages only 4 minutes of interrupted service each year.

I analyzed two sets of data, one from the U.S. Department of Energy's Energy Information Administration (EIA) and the other from the North American Electric Reliability Corp. (NERC). Generally, the EIA database contains more events, and the NERC database gives more information about the events. In both sets, each five-year period was worse than the preceding one.

What happened? Starting in 1995, the amortization and depreciation rate has exceeded utility construction expenditures. In other words, for the past 15 years, utilities have harvested more than they have planted. The result is an increasingly stressed grid.

R&D spending for the electric power sector dropped 74 percent, from a high in 1993 of US $741 million to $193 million in 2000. R&D represented a meager 0.3 percent of revenue in the six-year period from 1995 to 2000, before declining even further to 0.17 percent from 2001 to 2006. Even the hotel industry put more into R&D.

Investing in the grid would pay for itself, to a great extent. You'd save stupendous outage costs—about $49 billion per year (and get 12 to 18 percent annual reductions in emissions). Improvement in efficiency would cut energy usage, saving an additional $20.4 billion. —*S. Massoud Amin*

HIGH AND LOW R&D SPENDERS
PERCENT OF REVENUE, 1995 TO 2000

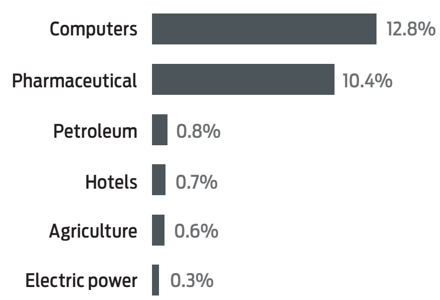

Sources: The Electric Power Monthly, *U.S. Department of Energy's Energy Information Administration; "Transmission Availability Data System Automatic Outage Metrics and Data," North American Electric Reliability Corp.; "Powering Progress: Restructuring, Competition, and R&D in the U.S. Electric Utility Industry," Paroma Sanyal and Linda Cohen,* The Energy Journal

EMILY COOPER

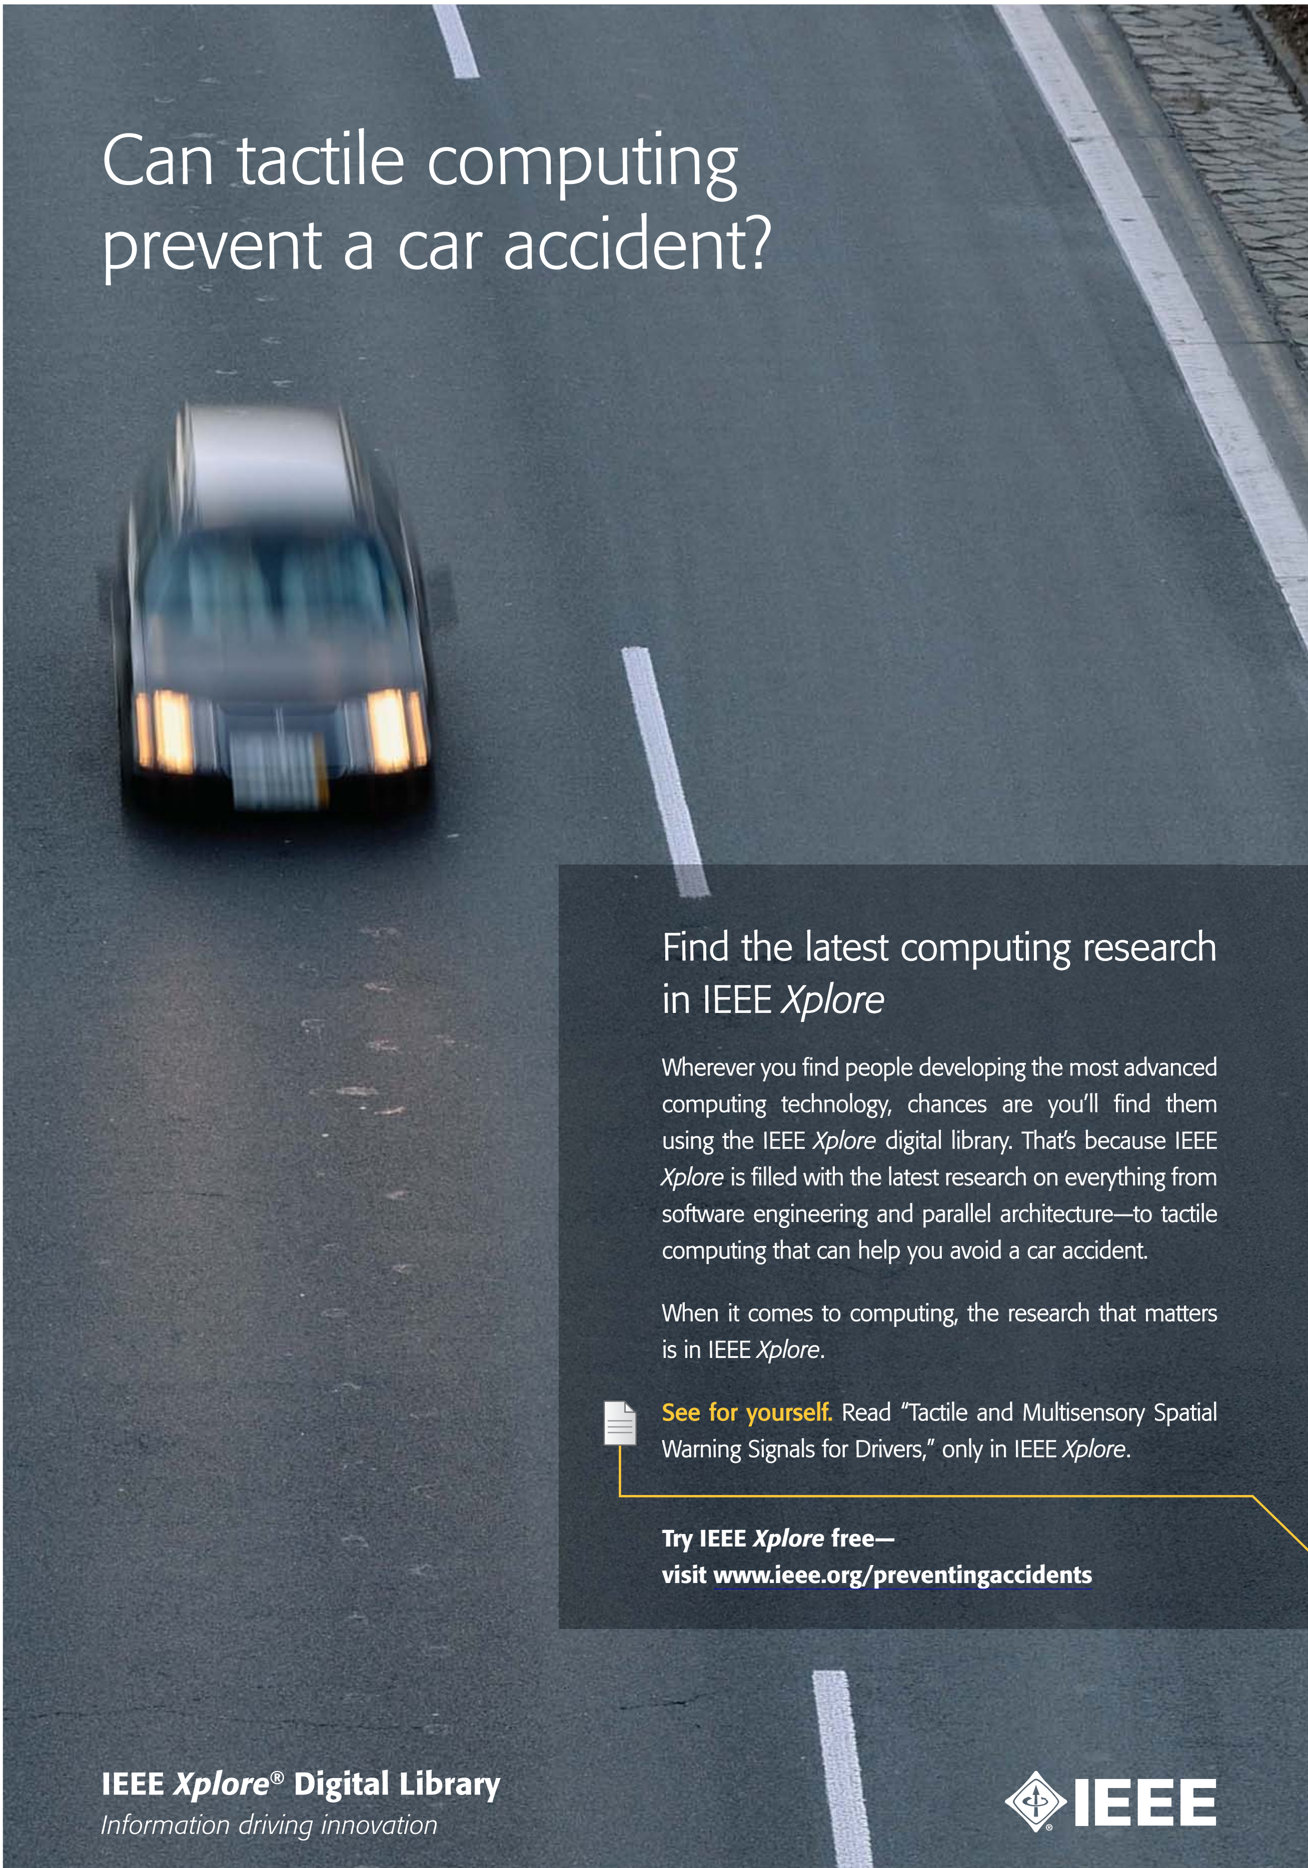
Can tactile computing prevent a car accident?
Find the latest computing research in IEEE *Xplore*
Wherever you find people developing the most advanced computing technology, chances are you'll find them using the IEEE *Xplore* digital library. That's because IEEE *Xplore* is filled with the latest research on everything from software engineering and parallel architecture—to tactile computing that can help you avoid a car accident.
When it comes to computing, the research that matters is in IEEE *Xplore*.
See for yourself. Read "Tactile and Multisensory Spatial Warning Signals for Drivers," only in IEEE *Xplore*.
Try IEEE *Xplore* free—
visit www.ieee.org/preventingaccidents
IEEE *Xplore*® Digital Library
Information driving innovation
IEEE

Figure 1: Robot Visualization
File Edit View Insert Tools Desktop Window Help
Real-Time xPC Target Spy