THE MAGAZINE OF TECHNOLOGY INSIDERS

IEEE Spectrum

10.12

SPECTRUM.IEEE.ORG

Automation's New Face

WILL INDUSTRIAL ROBOTS BE PERSONABLE? RODNEY BROOKS IS BETTING BIG ON IT

IEEE

volume 49 number 10 international

ieee spectrum

Winner of the 2012 National Magazine Award for General Excellence

10.12

28

36

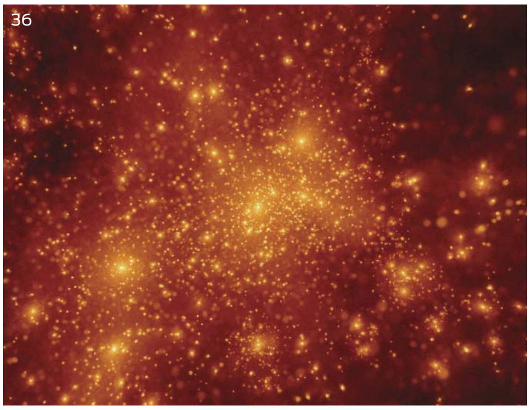

42

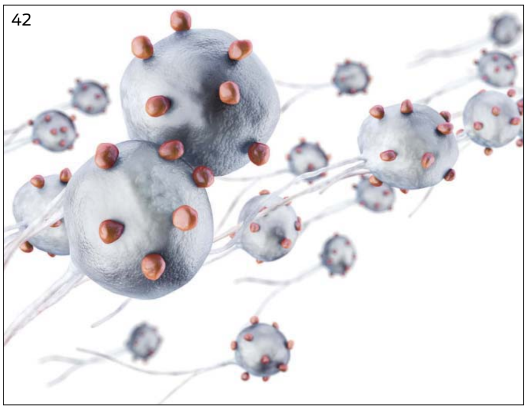

COVER AND LEFT: DAVID YELLEN; TOP RIGHT: RESIMULATION, GUSTAVO YEPES/ UNIVERSIDAD AUTÓNOMA DE MADRID; BOTTOM RIGHT: ANATOMY BLUE

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


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

volume 49 number 10 international

ieee spectrum 10.12


11

UPDATE

OPINION

DEPARTMENTS

SPECTRUM.IEEE.ORG AVAILABLE 1 OCTOBER

North Korea, IT Powerhouse

Did you know that there are a number of IT firms in North Korea—and that they accept foreign outsourcing contracts? One of them, the Korea Computer Center, has more than 1000 employees. On an *IEEE Spectrum* "Techwise Conversations" show, consultant Paul Tjia describes the unique obstacles, pluses, and minuses of outsourcing IT work to North Korea. Tjia is the author of the business guide *Offshoring Information Technology*.

THEINSTITUTE.IEEE.ORG AVAILABLE 5 OCTOBER

PROCEEDINGS FOCUSES ON REMOTE SENSING
This month's *Proceedings of the IEEE* explores novel remote-sensing methods that can help scientists better assess and mitigate the damage caused by earthquakes, tsunamis, volcanic eruptions, oil spills, and other disasters.

SCHOLARSHIP PROGRAM TAKES OFF
More than 400 budding U.S. power engineers are moving through the college pipeline, thanks to the IEEE Power & Energy Society's Scholarship Plus Initiative, which encourages students to enter the power engineering field by providing them with money and work experience.

MAKING DEVICES COMPREHEND THE SPOKEN WORD
We're used to gadgets that speak and even listen, but those that understand are still works in progress. To figure out how to make machines comprehend speech, researchers will meet in Miami from 2 to 5 December for the IEEE Workshop on Spoken Language Technology.

ONLINE WEBINARS & RESOURCES

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

4 October: TSV & Interposer: Modeling, Design & Characterization

11 October: Combined 3-D Electro-Magnetic & Spin Response Simulation of MRI Systems

17 October: Acoustics Simulation in Microphones, Mobile Devices & Hearing Aids Using COMSOL

18 October: Optimizing Data Analytics for the Smart Grid

25 October: Automotive PCBs: Efficient Signal & Power Integrity Analysis

Cutting Risk & Rework Through Requirements-Driven Quality in Complex & Embedded Systems Development

Power Converter Controller Design for Smart Grid Power Electronics

Avnet's Third Annual Virtual Power Forum

Modeling Batteries & Fuel Cells Using COMSOL Multiphysics

Quality Assurance for Medical Devices

Engineering Change to Accelerate Product & Systems Development

An MBA for Technology Professionals

Master Bond white paper library: http://spectrum.ieee.org/static/masterbond-whitepaper-library

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

CLOCKWISE FROM LEFT: BIKAS DAS/AP PHOTO; DAVID GUTTENFELDER/AP PHOTO; CHAD PURSER/ISTOCKPHOTO

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2012 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of American Business Media, the Magazine Publishers of America, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

back story

MIKE BUGDA

When Erico Met Baxter

One morning this past June, Senior Associate Editor Erico Guizzo stood facing a large 19th-century brick-and-beam building tucked into a corner of the Seaport District of Boston. It was the end of a nearly five-year odyssey.

Guizzo [left] and *IEEE Spectrum* contributor Evan Ackerman [right] were among the first journalists allowed to visit Rethink Robotics. The company, which is the brainchild of former MIT roboticist Rodney Brooks, is seeking to transform manufacturing by building robots that are adaptable, humanoid, and downright appealing.

Rethink has kept its plans completely under wraps since its founding in 2008, and it has raised more than US $60 million from several high-profile firms.

None of this was lost on Guizzo, who has long counted Brooks among his sources but whose attempts to extract information from him had gone nowhere. Cornered at conferences or accosted on the MIT campus, Brooks would remain polite but tight-lipped. Finally, early this year, he revealed that he was ready to talk—and give *Spectrum* an exclusive look inside Rethink.

Guizzo immediately called Ackerman, a Californian well-known to followers of *Spectrum*'s award-winning robotics blog, Automaton. Stepping into Rethink's top-secret lab on that June morning, Guizzo and Ackerman came face to face with Baxter, the factory bot created by Brooks and his team. In this issue's "The Rise of the Robot Worker," they describe how Baxter came to life and how Rethink plans to market it to U.S. manufacturers.

Ackerman, who credits his childhood love of science fiction for his continuing enthrallment with robots, was impressed by how Baxter uses vision to make decisions, and also by its host of safety features. "It's impressive how much Baxter is able to figure out on its own," Ackerman says. "And yes, it's also nice that it can't kill you." ❑

CITING ARTICLES IN IEEE SPECTRUM

IEEE Spectrum publishes two editions. In the international edition, the abbreviation INT appears at the foot of each page. The North American edition is identified with the letters NA. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, The Data is in *IEEE Spectrum,* Vol. 49, no. 10 (INT), October 2012, p. 64, or in *IEEE Spectrum*, Vol. 49, no. 10 (NA), October 2012, p. 80.

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

contributors

EVAN ACKERMAN, based in San Jose, Calif., is the lead contributor to Automaton, *IEEE Spectrum*'s robotics blog. He started covering robots on his own blog, BotJunkie, back in 2007. The hobby turned into nearly a full-time job, and in early 2011 BotJunkie merged with Automaton to form the top robotics blog on the Net. Ackerman cowrote "The Rise of the Robot Worker" [p. 28] with Senior Associate Editor Erico Guizzo, about start-up Rethink Robotics and its robot, Baxter.

STEPHEN CASS is *Spectrum*'s newest contributing editor, helming the magazine's careers, reviews, and DIY coverage. It's a return to the fold for Cass, a staff editor here from 2000 to 2007. In this issue, he introduces a new type of article with a report on LevelUp [p. 23], the first in a series of start-up profiles. Cass says his goal is "to celebrate the interests of techies and give them a heads-up about useful things to know."

JOSEPH A. KONSTAN and **JOHN RIEDL**, both professors of computer science at the University of Minnesota, wrote "Recommended for You" [p. 48]. As codirectors of GroupLens Research, Konstan, an IEEE Senior Member, and Riedl, an IEEE Fellow, helped create the MovieLens recommender system. The pair's scariest recommender moment came during an interview on "ABC Nightline." Just before a station break, MovieLens pitched the 1950s film noir *Sunset Boulevard* to host Robert Krulwich. Konstan and Riedl had to wait until they were back on the air to hear his verdict on the suggestion: He loved it!

SYLVAIN MARTEL, who founded and directs the NanoRobotics Laboratory at École Polytechnique de Montréal, spent many long nights trying to navigate a metal bead through a live pig's blood vessels with an MRI machine. Even the machine's manufacturer thought he was crazy, he says. But as he reveals in "Journey to the Center of a Tumor" [p. 42], his patience paid off. He is now a pioneer in the burgeoning field of medical microrobotics.

JOEL R. PRIMACK writes about simulating the evolution of the cosmos in "The Universe in a Supercomputer" [p. 36]. Primack is a physics professor at the University of California, Santa Cruz, and director of UC's High-Performance AstroComputing Center. Back in the early 1980s, pencil and paper were the tools of choice for cosmologists. "Now high-performance computers have become vital," Primack says. "They've helped transform cosmology from philosophical speculation into what's almost an experimental science."

DAVID YELLEN, this month's cover photographer, loves shoots that make great stories—like taking a portrait of Baxter the robot. Outside of photography, Yellen enjoys fishing and molecular gastronomy. In recent culinary adventures, he ate the still-beating heart of a bigeye tuna, renovated his Brooklyn kitchen, and had a dinner party where everything was cooked either *sous vide* ("under vacuum," in a low-heat water bath) or with a blowtorch.

CLOCKWISE FROM TOP LEFT: EVAN ACKERMAN; ÉCOLE POLYTECHNIQUE DE MONTRÉAL; PAUL SCHRAUB; DAVID YELLEN; ALANA RANGE

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


Enhancing entertainment
in everyday life
life.augmented
ANSYS®
Experience by STMicroelectronics. Simulation by ANSYS.

Verify and optimize
COMSOL M

ANTENNAS: The Sierpinski fractal monopole antenna shows multiband responses inherited from the nature of a fractal structure. Results show the magnitude of the electric field on the surface of the structure as well the far field radiation pattern.

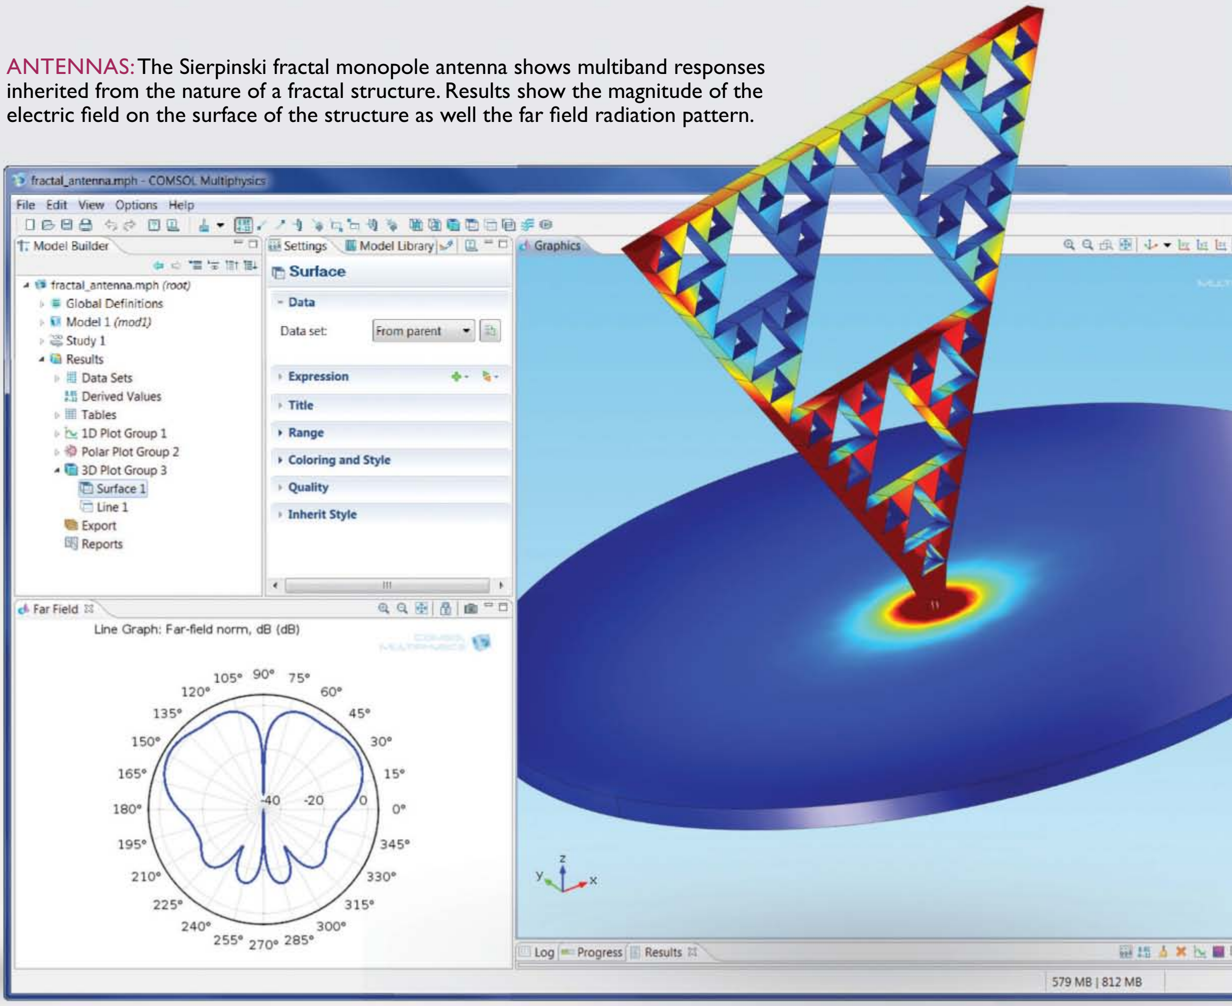

spectral lines

Obama, Romney, and the U.S. R&D Complex

MITT ROMNEY and Barack Obama, the major party U.S. presidential candidates, say little about how the United States runs its vast R&D enterprise. Their inattention is peculiar, given the annual cost: US $140 billion. However, Obama and Romney, despite many differences elsewhere, agree on the way forward: Spend more money on more ambitious projects.

This approach contributes to a troubling disconnect between research capacity, of which there is plenty, and beneficial outcomes from research, of which there are few. The candidates won't devote a presidential debate to the subject of R&D. But given the central role of publicly funded R&D to America's future, might each candidate at least agree to read a short book on the subject?

Richard R. Nelson's brief classic *The Moon and the Ghetto* is the best place to begin, in my opinion. First published in 1977, *The Moon and the Ghetto* poses a deceptively simple question, which (updated) goes like this: If government wizards can create drone assassins, why can't the same big brains, even when amply funded, create better schools, more cures for diseases, cheaper energy, and any number of other "moon shots" that would solve urgent problems?

When Nelson wrote *The Moon and the Ghetto* 35 years ago, the U.S. government had recently landed men on the moon but had failed to win the symbolic wars on poverty and cancer, or the actual war in Southeast Asia.

Of course, scientists and engineers don't guarantee winning outcomes, and even those they achieve often can't be predicted—or translated into equal benefits to all citizens. That's why Nelson highlights the importance of reforming organizational structure. He argues that the most fundamental R&D policy questions are not scientific or technical but managerial and bureaucratic.

The "key issue," Nelson insists, is "the way in which a research program ought to be managed and organized." He emphasizes the benefits to be gained from assuming what he calls "the organizational perspective." His view contrasts sharply with the conventional wisdom that budgets and objectives matter most.

By applying Nelson's concepts, the next president can tackle the urgent task of reinventing tired, bloated R&D agencies. Consider the National Institutes of Health, which spends about $30 billion a year, largely on basic science and biomedical engineering, runs dozens of research centers organized around diseases, rather than delivering improved health-care services. The agency's institutes contain many redundancies, and their work is often far removed from the ultimate end users—hospitals, doctors, and patients. Moreover, the ultimate funders of health care—the U.S. government's Medicare program, insurance companies, and individual consumers—have essentially no connection to the NIH's research activities.

Critics believe that radical reorganization of the NIH is required, which will demand presidential leadership. The same can be said for the 10 national labs run by the U.S. Department of Energy ($5 billion in costs annually). The next U.S. president could demand the consolidation of these labs—to achieve better value for money and to pursue aims clearly tied to improving the quality of American life.

To be sure, Nelson, currently an economist at Columbia University, admits that reforming federal R&D organizations can be humbling. Like other bureaucracies, R&D agencies seek to perpetuate their legacy behaviors and often become, he writes, "a powerful lobbyist of new technology for its own sake," not for the public's benefit.

So the NIH delivers a stream of "breakthroughs" that explain the origins of diseases but don't enhance everyday treatments. To improve American health requires a reorganized NIH that would also overcome the bureaucratic logic that routinely defeats reformers. That requires outstanding leadership—from the next president.

How Romney and Obama differ on R&D spending would be good for voters to know, of course. But even more important is that the two candidates possess genuine sources of wisdom about the real genius of innovation. *The Moon and the Ghetto* is a fine place to start.

—G. PASCAL ZACHARY

G. Pascal Zachary is a professor of practice at the Consortium for Science, Policy & Outcomes at Arizona State University. He writes often for IEEE Spectrum, *which published a version of this article online in July 2012.*

Correction

In "Curiosity's 1-Ton Touchdown" [Update, August 2012], we stated that distance readings from the rover's long-range radar system would be used to signal, among other things, when the parachute would be deployed. Instead, the system actually triggered the detachment of the spacecraft from the parachute and back shell.

update

more online at spectrum.ieee.org

KOLKATA CUT: Life went on during the biggest blackout in history. This barber in the eastern city of Kolkata worked by candlelight.
PHOTO: BIKAS DAS/ AP PHOTO

Blackouts Illuminate India's Power Problems

Weak links in the grid prove difficult to fix

FOR TWO consecutive days in July, India experienced blackouts that took down large portions of the country's power grid. The second outage was the largest in history, leaving more than 600 million people, nearly a tenth of the world's population, without electricity. The blackouts brought renewed attention to the country's power sector, which is struggling to supply India's growing demand. They exposed weak links in the transmission system, inadequate fail-safe systems for preventing cascading failures, and a lack of proper outage planning.

The 30 and 31 July failures may have affected more people than any blackout ever, but it's tough to argue that they were the most disruptive. "I did not even know that there was some problem in the grid, because I was working from home, and in this building we have 100 percent backup," says Sivakumaran Govindarajan, a member of the India Smart Grid Forum, who lives just outside Delhi, the largest city affected by both blackouts.

Indeed, many commercial buildings, apartment complexes, and hotels in India have diesel generators that kick in automatically when the power goes out, which happens all too often. Although there hadn't been a grid-scale failure since 2001, rolling outages are common. In July, there was a gap of about 9 percent between the country's energy requirement and the amount available, according to the Central Electricity Authority.

Even when the grid is operating at peak capacity, 300 million or more people—more than in any other nation—still don't have electricity. "The power sector in India has been very badly handled," laments S.K. Anand, a retired power engineer who

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

update

worked in Punjab for more than 40 years. Demand for electricity will only continue to grow: India's per capita electricity consumption was about 600 kilowatt-hours in 2009, leaving it far below the world average of 2700 kWh and behind regional competitors like China, where per capita consumption was around 2600 kWh.

Although new power plants are being built at a brisk clip—India added about 20 gigawatts of generation capacity in the past year—there hasn't been an equivalent investment in transmission and distribution networks, says R. Nagaraja, the managing director at Power Research & Development Consultants, in Bangalore. Increasing the network capacity will be critical, because most of India's demand is growing in the western states, while most of the suitable locations for new plants are in the east. It's politically difficult to secure right-of-way for new transmission lines, so power utilities have focused on upgrading existing corridors.

Just such an upgrade was at the heart of the July blackouts. Two days before the first outage, the Power Grid Corp. of India took off line one of the two transmission circuits near Agra, home of the Taj Mahal, to upgrade it from 400 to 765 kilovolts. India's power system is composed of five connected regional grids; four of the grids are synchronized and linked through transmission lines like the one in Agra. But because other interregional lines were already unavailable, the planned outage effectively left the remaining Agra circuit as the only link between the western and northern regional grids, according to a report issued by the committee formed to investigate the blackouts. Both outages began when this line tripped, and on the previous day, it also experienced what the National Load Dispatch Center called "a near-miss situation."

This weak link was also under a higher power load than normal. The end of July is planting season for farmers in the northern states, India's historical breadbasket. But this year's monsoon had delivered little rain, forcing farmers to rely on electric pumps to bring groundwater up from boreholes and increasing the demand for power in the northern region. Regional load dispatch centers tried to get states to reduce their consumption and generation to relieve the weak line, but the reductions weren't enough. The upgrade "should not have been done during the peak summer paddy season," says Anand.

In such a heavily loaded line, impedance at one end of the line can drop so low that it can cause a distant relay to trip, even though there is no actual fault in the system. According to the committee's report, such "load encroachment" caused the Agra line to trip prior to both blackouts, setting off a cascading failure as power swings in the system caused more circuits to trip.

But if the transmission line had proper compensation systems in place to manage reactive power, says Nagaraja, it should have been able to handle even the heavy load. The panel's report noted that flexible AC transmission systems and dynamic compensation technology could have helped prevent the trouble.

Even after the Agra line failed, it shouldn't have caused such widespread damage. The western region was able to survive both events by disconnecting some generators, and the committee's report noted that in the second blackout, after the western grid had separated itself, the remaining synchronized grid's frequency stabilized for about a minute. If some of the overdrawing regions could have been disconnected, a full blackout might have been avoided: "Trip only the one who is overdrawing," says Nagaraja, "and nothing would have happened." Relays that disconnect the line when the frequency gets too high or too low should have managed to shed enough generation or load to restore stability. India is supposed to have such systems in place, but they obviously didn't function as intended.

The blackouts have shown that even within India's large, interconnected grid, it's important to have some local generation in case of emergency. The west coast city of Mumbai, which was outside of the blackout zone, uses a scheme called islanding to dodge large outages: It has built up enough local generation to allow it to disconnect from the grid and keep essential services running. Now Delhi is considering a similar islanding scheme. The India Smart Grid Forum's Govindarajan agrees that even outside of cities, India would benefit from more such microgrids "so that we can actually manage the emergency consumption of power in a better way." —Joshua J. Romero

AMIT SOMVANSHIL/GETTY IMAGES

Intoxicam

Engineers at the University of Patras, in Greece, have come up with software that can discern from a thermal image whether a person is drunk. Alcohol dilates the skin's blood vessels, creating a particular pattern of hot spots. The system uses the differences in temperature between certain areas of the face to figure out if the person is besotted. The inventors suggest that bars could use it to determine if a patron is too pickled to purchase more alcohol.

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

US $1230 PER KILOGRAM Cost of extracting uranium from seawater using a recently improved technique. It's twice as productive as the old method but is easily outstripped by mining.

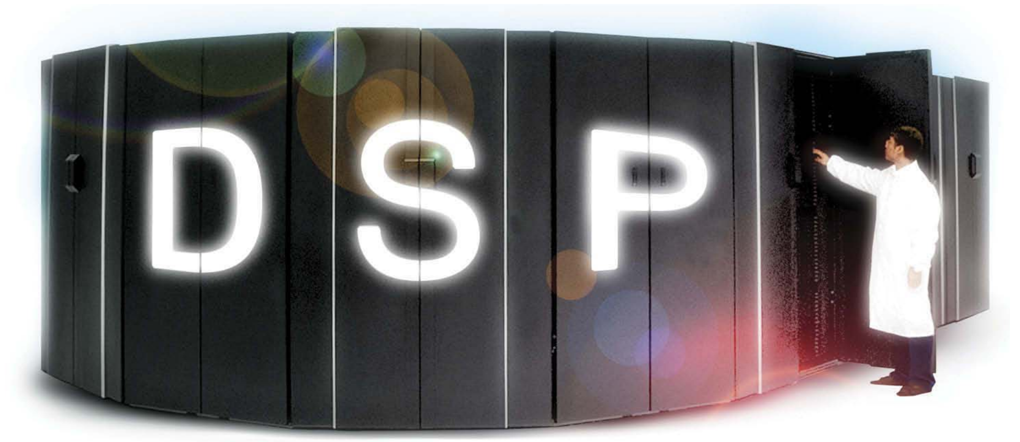

Could Supercomputing Turn to Signal Processors (Again)?

Texas team says digital signal processors could compete in high-performance computing

BUILDING HIGH-performance computers used to be all about maximizing flops, or floating-point operations per second. But the engineers designing today's high-performance systems are keeping a close eye not just on the number of flops but also on flops per watt. Judged by that energy-efficiency metric, some digital-signal processing (DSP) chips—the sophisticated signal conditioners that run our wireless networks, among other things—might make promising building blocks for future supercomputers, recent research suggests.

The DSPs that might make the jump to supercomputing come from Texas Instruments, which originally designed them for relatively modest applications. "We hadn't even thought to look at high-performance computing or supercomputers," says Arnon Friedmann, multicore business manager for TI. "It wasn't on our radar." TI's DSP chips are typically used in embedded systems, most prominently cellular base stations. For such applications, power efficiency is vital, but until recently these systems didn't require floating-point calculations, making do instead with just integer arithmetic. The advent of 4G cellular networks, however, increased the computing burden within base stations, making floating-point calculations essential.

TI engineers added floating-point hardware to the TMS320C66 family of multicore DSPs late in 2010 without appreciably slowing these processors down or increasing the power consumed. But it was only after the new chips came out that some forward thinkers at TI realized that the eight-core C6678 DSP, which can perform as many as 12.8 gigaflops per watt running flat out, might be useful for general-purpose high-performance computing.

"The question was whether we'd be able to extract that potential in the real world," says Francisco D. Igual, now a postdoctoral researcher at Universidad Complutense de Madrid. He was working at the University of Texas at Austin with engineering professor Robert A. van de Geijn when TI approached them for help. Collaborating with TI, they wrote code for the new DSP to perform general matrix-matrix multiplication, something they felt would be representative of the kind of numerical weight lifting that high-performance systems are often asked to do.

With that code in hand, the team compared the new DSP chip against some common supercomputer architectures. The DSP came out on top, at 7.4 gigaflops per watt. "We were happy, but we were not that surprised," says Igual.

But not everyone is swayed by those results. "It's very impressive, but not a fair comparison," says John Shalf of the National Energy Research Scientific Computing Center at Lawrence Berkeley National

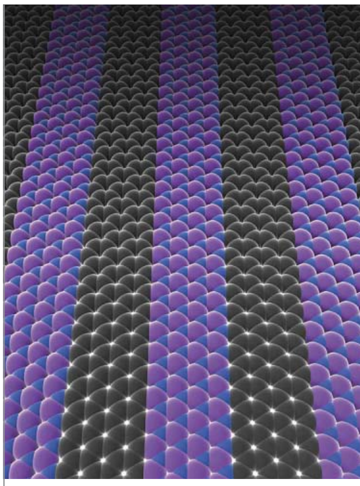

Superthin Circuits

Researchers at Cornell University, in Ithaca, N.Y., have found a way to get graphene, the nanoelectronics wonder material, to knit together with also-ran boron nitride. The result is a striped pattern of interlocking graphene conductors and boron nitride insulators that forms an array of wires just one atom thick. Sheets of the arrays could possibly be laid atop each other to form complex circuits, the researchers say. The conductive graphene can also knit together with a semiconducting version of itself; the Cornell team is going to try integrating it with the two-dimensional semiconductor molybdenum disulfide, too.

TOP: GLUEKIT; RIGHT: CORNELL UNIVERSITY

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

update

Laboratory, in California. Shalf points out that the comparison was done using single-precision (32-bit) arithmetic across the board. That's because this DSP chip can complete a single-precision operation in one clock cycle, whereas double-precision calculations take four clock cycles and thus use four times as much energy. So the number-crunching circuits in each of the competing systems tested, which are configured for efficient double-precision arithmetic, were at a disadvantage.

Texas Instruments hopes to improve the efficiency of double-precision operations on its multicore DSPs. But the energy used for double-precision calculations would, at the very least, still be double what researchers found in their single-precision tests, meaning these chips would at best be able to perform in the range of 3 to 4 gigaflops (double precision) per watt.

And individual chips do not a supercomputer make, Shalf stresses. Combining many of them in a high-performance computer, with its various electronic subsystems and cooling apparatus, would significantly lower the machine's overall energy efficiency.

How, then, would it compare with today's best supercomputers? The current world champion in flops is the Sequoia supercomputer at Lawrence Livermore National Laboratory, in California. With its IBM BlueGene/Q architecture, it performed just over 2 gigaflops (double precision) per watt in recent benchmark tests—similar to what you'd expect from a double precision DSP-based machine, judging by the research results of Igual and his colleagues.

But perhaps this shouldn't be so surprising, given BlueGene's ancestry. Two decades ago, TI produced a line of DSPs that, like the company's new multicore family, contained floating-point hardware. In 1998, physicists from Columbia University, in New York City, ganged thousands of them together to construct a special-purpose supercomputer for performing calculations in quantum chromodynamics, dubbing it QCDSP, for Quantum Chromodynamics on Digital Signal Processors. Texas Instruments later dropped that DSP line, but Alan Gara, one of the three physicists who pioneered the design of the QCDSP, retained some of its lessons when he moved to IBM, where he became the chief architect for the BlueGene supercomputers.

You might guess that this groundbreaking DSP-based supercomputer is now enshrined in a computer museum. But in July, when *IEEE Spectrum* caught up with Columbia physicist Robert Mawhinny, who worked alongside Gara on that project, he told us that it met a sadder fate: "We threw it out last week—literally."

—David Schneider

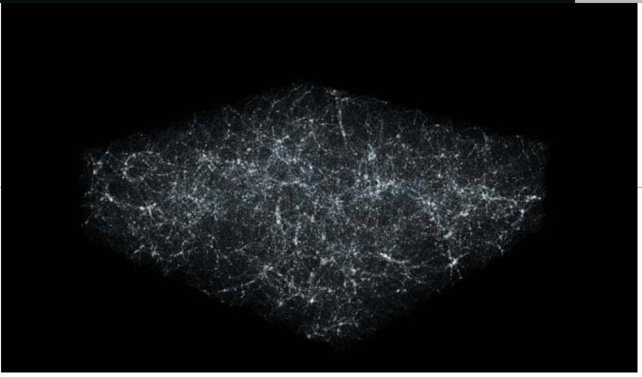

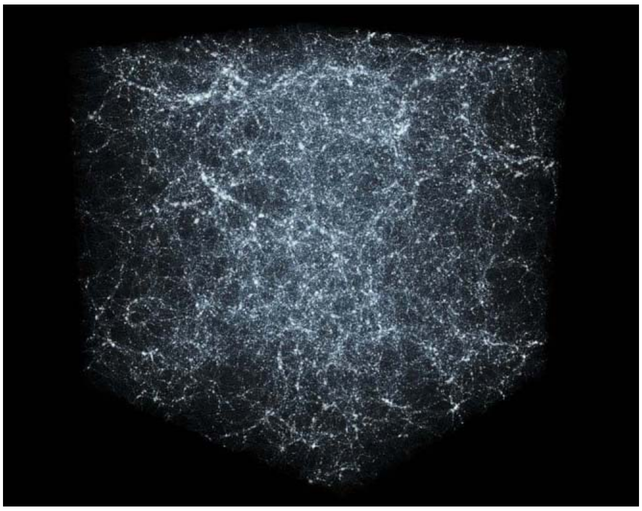

MORE, FASTER: You can explore more of a simulated universe using links that run at 100 gigabits per second [bottom] than you can using 10 Gb/s links [top]. *IMAGES: LAWRENCE BERKELEY NATIONAL LABORATORY*

SCIENCE AT 100 GIGABITS PER SECOND

U.S. energy agency lights up the first superfast science data network

The U.S. Department of Energy's high-speed science network is getting a much-needed makeover. When the new 100-gigabit-per-second Energy Sciences Network (ESnet) goes live next month, it will be the world's fastest continent-spanning science network. But this distinction will probably be temporary: Research networks around the world are in the midst of their own big upgrades.

Now more than ever, science depends on the ability to quickly and reliably move massive amounts of data over great distances. Breakthrough predictions and discoveries, such as climate change and exoplanets, rarely happen at one institution anymore. This summer's sighting of a particle resembling the Higgs boson, for instance, required sharing about 26 petabytes of data per year among more than 150 computing centers in 36 countries.

The research that ESnet users do is no exception. Traffic over the past decade on ESnet—which connects more than 40 national laboratories, supercomputing centers, and scientific instruments—has been growing at a rate of 72 percent per year, more than twice as fast as the commercial Internet, says director Greg Bell. He expects that growth will accelerate in the coming

–700 Refractive index of a new metamaterial capable of strongly bending light the wrong way. Normal materials typically have an index of 5 or less.

decade. "[Data] flows generated by the largest science networks and collaborations in the world are seriously stretching the abilities of today's 10-gigabit networks," he says.

The upgrade from 10 Gb/s to 100 Gb/s has been a multiyear project. When this generation (the fifth) of ESnet is finally finished, its engineers will have constructed an entirely new network from scratch. "We built the new advanced network in parallel with the existing one, like building a five-lane superhighway next to Main Street," says Inder Monga, ESnet's chief technologist.

To pave its new information highway, ESnet teamed up with Internet2, a consortium of U.S. universities that operates its own high-bandwidth network. The two agreed to share capacity on dormant ("dark") optical fibers that Internet2 owns. Left over from the 1990s telecom boom and bust, the fiber web spans the continent, stretching almost 21 000 kilometers if it were laid end to end. Each of its pair of thread-thin fibers can carry up to forty-four 100-Gb/s circuits, which means the shared network could be expanded to transport data at rates up to 8.8 terabits per second.

On top of the dark fiber, ESnet and Internet2 built a layer of optical technology, installing the equipment that receives, transmits, amplifies, regenerates, and combines 100-Gb/s signals. Managing its own optical infrastructure is new to ESnet. In the past, it simply leased additional 10-Gb/s circuits from commercial carriers whenever it needed more capacity. Ownership will allow ESnet to add capacity more cheaply and on its own terms, Bell says.

The last layer of technology ESnet engineers built was the routing layer—the links, hubs, and switches that ensure data packets get where they need to go.

(To reach all its users, ESnet also installed optical and routing technology on an additional 1600 km of fiber that connects centers outside the nationwide backbone it will share with Internet2.)

Engineers are now making last-minute fixes and testing for glitches before opening the network to traffic. Moving users onto the new network will take about a month, Monga says. Soon after, the old 10-Gb/s ESnet will be decommissioned, its hardware returned to ESnet's laboratory, repurposed for other research labs, or left in place for future expansions or experiments.

Although ESnet will be the first science network to make the switch over to 100 Gb/s, others will soon follow suit. Both China's and Europe's research networks have begun the migration, and Internet2 itself plans to finish its full upgrade by next summer.

—Ariel Bleicher

Lobbying 2.0

Google and Facebook are spending heavily on lobbying firms, even as they marshal their users to influence the U.S. Congress

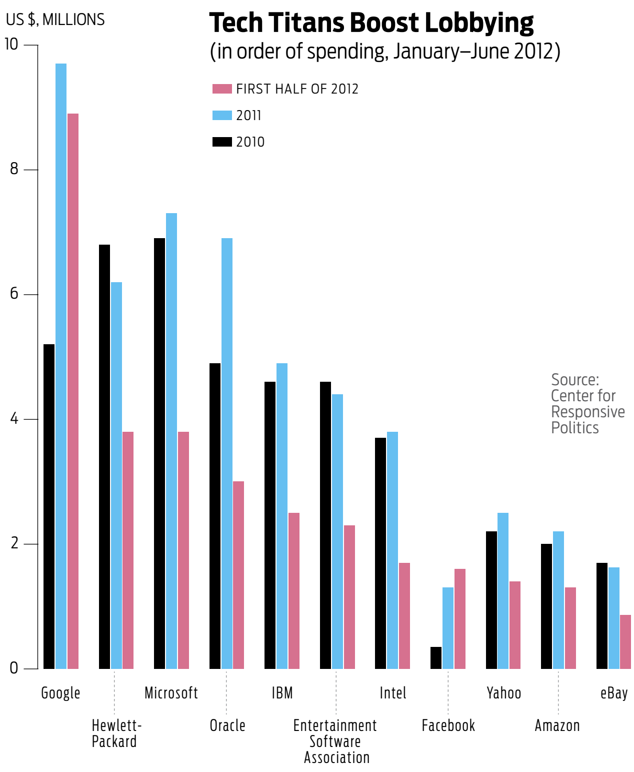

Regardless of which party controls the U.S. government after next month's elections, Washington's interest in how Internet companies use customers' data will continue to grow, and so these tech companies have recognized the need to spend more money on lobbying legislators. Unlike with some other industries, however, their investment is not limited to traditional lobbying methods. They are looking to their users to put pressure on Congress as well.

Computer and Internet companies have enlisted elite firms clustered on Washington, D.C.'s K Street in their efforts, according to Bill Allison, editorial director at the Sunlight Foundation, a nonprofit organization that promotes government transparency by putting information online. Google even hired

update

former congresswoman Susan Molinari earlier this year to head its Washington office. The companies have also been inching up their spending. Last year, the computer and Internet industry spent US $127 million on lobbying, up from $123 million in 2010 and $120 million in 2009, according to the Center for Responsive Politics (CRP). Spending is trending up again this year, at least by a few million, based on the $65 million figure for the first six months.

The most notable increases, however, are by two of the hottest U.S. Internet companies. Google's lobbying costs are on track to have nearly quadrupled in the past two years, rising to $9 million in the first half of 2012, according to tallies by the CRP. In fact, last year Google surpassed all other companies in the computer and Internet category, beating the previous top spender—Microsoft—by more than $2 million. Facebook's spending has also quickly accelerated but is still just a fraction of Google's.

But this year, tech companies also discovered new political muscle when a grassroots campaign rallied Internet users to petition Congress, leading to the defeat of the Stop Online Piracy Act (SOPA). The bill would have required websites and search engines to cut off access to non-U.S sites that the government suspects of copyright infringement. Key Internet companies—including Microsoft, Google, Yahoo, Facebook, and Twitter—opposed SOPA. But what got the most attention were constituent complaints, which were encouraged by Internet rights advocates such as Fight for the Future, Engine Advocacy, and the Center for Democracy and Technology (CDT). The tech companies rode that groundswell of public opinion; Wikipedia blacked out its website in protest. Google weighed in by covering its logo in black.

It was a watershed moment, making a huge impression on Congress, Internet firms, and Internet-rights advocates as well. Holmes Wilson, codirector of Fight for the Future, says the organization simply "stumbled across this new tactic," realizing that it could team up with Internet companies that would then appeal directly to their users. (In truth, there was a similar campaign in 1996.) If these companies "align themselves with the interests of individual users...you have this powerful coalition that can really stand up to any other industry and any other lobbyist," he says. In fact, in the wake of SOPA, Fight for the Future has formed a new group, called the Internet Defense League, to do just that. "With the Internet Defense League we want to take those tactics that were

Google and Facebook Speed Up Spending

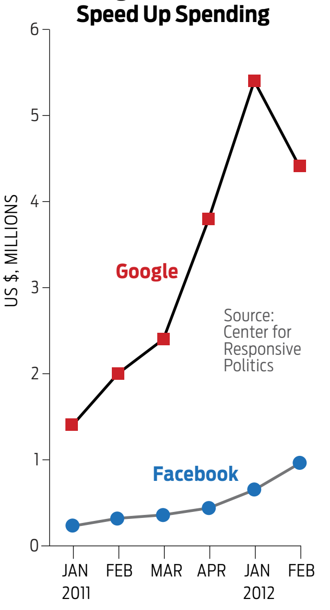

so successful in defeating SOPA and try to turn those tactics into an ongoing effort and a permanent force for defending the Internet," says Holmes.

They aren't the only ones. A new crop of organizations and partnerships is forming to try to influence Internet policy. During the third quarter of this year, several tech companies formed the Internet Association. According to several news organizations, Google, Facebook, eBay, and Amazon are among the founding members. The association, which was scheduled to launch in mid-September, said on its website that it represents both leading Internet firms and "their global community of users" and is "dedicated to advancing public policy solutions to strengthen and protect an open, innovative, and free Internet." Michael Beckerman, most recently a top aide to House Energy and Commerce Committee chairman Fred Upton, is its president.

Advocacy groups and tech companies are keeping in close contact, says Michael McGeary, political director and senior strategist of Engine Advocacy, which promotes the interests of start-ups. "There was a lot of talk during SOPA between lobbyists at big companies and some of these emerging groups, and a lot of that has kept going," he says. "We are all thinking about [how to] combine our forces on some of these initiatives."

When their interests align, such a combination "could be a very politically effective relationship," says Mark Stanley, campaign and communications strategist of CDT, which promotes policies to retain and promote Internet freedom. "Even with all the traditional lobbying money in D.C., what still actually works is millions of ordinary citizens rushing to make phone calls."

Of course, the interests of users and the interests of the companies won't always align. Consumers have been critical, for example, of the way these companies collect and use data about their users. "On privacy, they might be on the opposite side of their customers," says Allison of the Sunlight Foundation. "You won't see them using that strategy then."

—Tam Harbert

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

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

the big picture

A CHECK ON POWER

No, these men aren't playing a high-voltage game of chicken. They're using a SmartCopter ZN-2 remote-controlled helicopter to inspect a 220-kilovolt electric power transmission line. The copter, which is outfitted with special cameras, sensors, and software, lets utilities avoid both the inefficiency of having workers climb transmission towers for manual inspections and the danger of visually inspecting the lines from full-scale helicopters. One inspector who helps maintain this power line—which runs through Shouguang in China's Shandong province—says the tool has cut the time it takes to examine a 5000-meter-long section from 8 hours down to 40 minutes.

PHOTO: XU SUHUI/ XINHUA/LANDOV

hands on

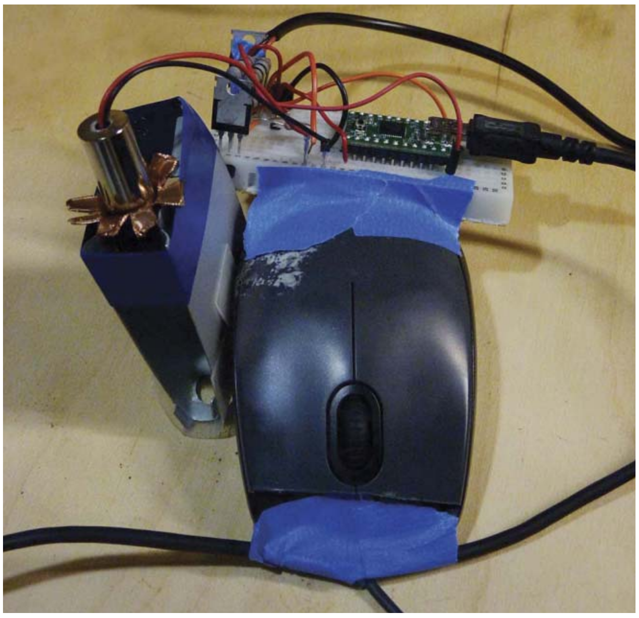

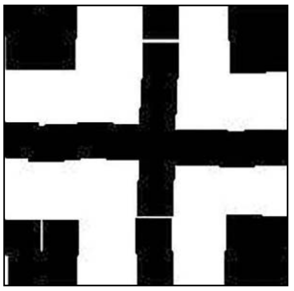

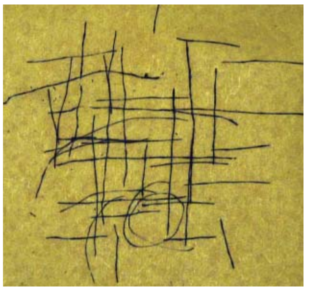

WORK IN PROGRESS: The prototype mouse engraver uses a microcontroller and a laser lashed to a US $10 mouse [left] connected to a PC. Initial testing with a checkerboard image [above left] produced almost random marks, but reducing the load on the PC leads to improvement [above right].

PHOTOS: PAUL WALLICH

LASER MOUSE

A quick-and-dirty tool for engraving large images

PREVIOUSLY, I've written about how to build a low-power laser engraver from the guts of a DVD burner [see "Laser Cuts Paper," *IEEE Spectrum*, June]. My fully automated setup was similar in spirit to that of a commercial engraver, and while effective for small areas, this approach quickly becomes unwieldy if you want to burn larger areas. You need longer rails and leadscrews, bigger stepper motors, more structure, and more space.

As I don't often require a high level of precision on larger areas, one evening I had an inspiration, and a new laser mouse was born. Normally, a "laser mouse" shines a laser onto a surface in order to track its position. My laser mouse does something different: When I move the mouse across a surface, it uses its tracking mechanism to make sure that an attached engraving laser is shining when it should be. The resulting engraving should look a little like a stone rubbing. (Note that even a small area of charred wood or paper can produce a remarkable amount of smoke.)

So how does it work? The laser mouse requires a host computer, so I coded up something ugly but functional in Lisp on my trusty basement Ubuntu Linux box (you can use the computer and language of your choice, so long as you can get access to a USB port and can read mouse-position information). The Lisp code converts a digital image into an array of brightness values; then it reads "/dev/mouse0"—Linux's standard software interface for an attached USB mouse—to decode incoming three-byte messages from the mouse into incremental *x* and *y* movements. These incremental movements are continually tallied to determine absolute position. Then my program consults the brightness array, and if the corresponding value for that position in the array is less than a threshold set when the program is started, it sends a 1 along a second USB connection to the microcontroller driving the laser. Otherwise it sends a 0.

On the mouse side, I taped a prototyping board fitted with a microcontroller to a US $10 Logitech mouse. The microcontroller I used is a Teensy, a controller slightly shorter and narrower than my thumb, that's mostly compatible with the popular but much bulkier Arduino microcontroller. The Teensy listens to its USB connection, decodes the packets being sent by my Lisp program, and, as needed, puts a high or a low voltage on an output pin that enables or disables a leftover laser driver connected to a 100-milliwatt red laser I had lying around.

Ideally, I'd like a more powerful laser to draw faster, but did I mention that even a small area of charred wood or paper can produce a remarkable amount of smoke? And even this 100-mW laser draws more power than the Teensy can supply, so I have to use an additional external power supply for the laser driver. As a consequence, instead of powering the Teensy from the USB connection as I might normally do, I also had to connect the Teensy to the external power supply to ensure a common ground. Otherwise I'd have run the risk of unpredictable voltages on the connection between the controller and the laser driver.

So far I've been using a checkerboard pattern as my test image. Initial results produced more or less random scratches, but I've proceeded to where there are recognizably denser and lighter regions. The key to improvement has been in reducing the workload on the Linux box to make it more responsive, and I'm creating a stripped-down system that doesn't run any unneeded software.

Although—or perhaps because—the laser mouse is so simple to build, it made me think about how inept modern desktops and laptops are when dealing with the outside world. I had to use an entire microcontroller just to get on and off signals to the laser driver, because there's no clean way for an ordinary mortal to tell a personal computer, "Please put 5 volts on this outside-accessible pin." And I had to carefully configure my Linux setup to allow my program access to the raw messages from the mouse. Normally programs are allowed to access the onscreen position of the mouse only within their own windows, which typically measure no more than 1800 by 1000 pixels. Because I track information from the mouse with my own code, the number of pixels in the image is limited only by the constraints of the available memory and the language, which in my case could mean an image as large as 25 000 by 20 000 pixels. (How big an actual engraving this translates to depends on a scale factor that determines how much the mouse can move and stay within the boundaries of a single pixel. In my system, this can be as small as 0.25 millimeters per pixel and as large as desired, within practical limits.)

On the positive side, one of the things I like about this project is that it is endlessly tweakable: In a single weekend you can have a basic version running with an Arduino and a prototyping board, or you can try being more ambitious by adding orientation sensing for higher accuracy, a custom mouse enclosure, hacks to draw all your laser power from your USB ports, and so forth (I've already swapped in a higher-power blue laser, for example). On the software side, you can go from simple positional lookup to complex interpolation, or even graphing code that controls the laser with no reference to a preexisting picture. (Of course, you could also burn or blind yourself, or burn down your house or apartment. If you're not sure what you're doing with a laser, don't do it.)

When I "finish" this project in a few years, I imagine that the ultimate laser mouse will look very different. The microcontroller will detect the mouse's position directly from the onboard electronics and read an attached SD card for the values of the image array. It will boast a couple of lithium batteries for power and accept additional special parameters via Bluetooth. Then all I'll have to do is build a mobile robot to wield it.

—Paul Wallich

tools & toys

Oldies But Goodies

When yesterday's gadgets are better than today's

SEVENTIES SOUND: This Pioneer SX-838 receiver, manufactured circa 1975, can output 50 watts per stereo channel. *PHOTO: TIM WHYTE*

TECH ENTHUSIASTS live in the grip of reverse nostalgia, forever pining for 18 months hence. After all, another way to state Moore's Law is "They don't make them like they're going to." But there are some electronic devices that were once made better or cooler. And there's a bustling retro-electronics subculture busy cataloging, chronicling, and collecting these old-school gems.

At the top of the stack is a world of 1970s-vintage high-end audio that Tim Whyte, based in Carmichael, Calif., caters to on his website, Classicaudio.com. To feel the difference from modern audio systems, he says, "all you've got to do is pick one of these things up. They're actually made out of wood, metal, glass. These things were a year's worth of mortgage payments back in the day."

During the golden age for audio equipment—the 1970s—emerging brands like Marantz, Nakamichi, Pioneer, and Sansui vied for the American leisure market with systems that were wildly overengineered. Whyte attributes this excess partly to tight standards imposed by the U.S. Federal Trade Commission on advertising claims about the watts-per-channel ratings for hi-fi amplifiers.

tools & toys

CLASSIC LOOKS: A 1974 Marantz receiver [above left] is more robust than similarly priced modern receivers. For those who prefer film to digital, this Bell & Howell movie camera [above] will accept Super 8 film, while a real historical oddity [left] is the Fisher-Price PXL-2000, which recorded images captured by a CCD onto audio cassette tapes. The *Gauntlet* arcade game's cabinet [opposite] allowed four people to play simultaneously.

CLOCKWISE FROM TOP LEFT: TIM WHYTE; SSPL/GETTY IMAGES; BRYAN KUNTZ/NYCDREAMIN ARCHIVES; SSPL/GETTY IMAGES

Amplifiers "had to do one-third of the rated power for 30 minutes and not blow up, which is considered a ridiculously hard test by today's standards," says Whyte, who has been working full-time since 2007 refurbishing and selling classic stereo components. "As a result, these things were built like tanks."

Today's home stereo components with four-figure price tags will likely outperform these veterans, but Whyte argues that in the range of US $400 to $600, refurbished 1970s amplifiers and preamplifiers sound as good as or better than any new equipment. And the classic components, built for a lifetime of use, can be expected to outlive their 21st-century equivalents. They're also often more easily repaired. Owners looking to replace a vacuum tube, for example, will find a flourishing market for tubes, often still in their original packaging, from companies like VacuumTubes.net and Vacuum Tubes Inc.

Outside of stereos, retro electronics is more catch-as-catch-can. There are no magic decades, for instance, of film cameras or televisions that put present-day video technology to shame. But there are individual product lines and quirky one-offs that can make the hunt back in time worthwhile. And there is no better chronicler of vintage tech on the Internet than the blog Retro Thing.

The site's two editors, Bohus Blahut and James Grahame, each has his own favorite categories and gadgets. Grahame notes that vintage Super 8 film cameras and Soviet-era still cameras are coming into their own for their value as purely mechanical, clockwork-wonder gizmos. Specialty websites like USSRPhoto.com and SovietCams.com detail the camera models to seek and the ones to avoid. Old Super 8 cameras are more widely available on sites like eBay. Although getting film developed is not as easy as it used to be, once again the Internet provides: Grahame recommends online film developing houses such as Dwane's Photo, in Kansas, and Spectra Film and Video, in California.

Blahut, a Chicago-based filmmaker and TV producer, enjoys early electronic video technology. He says he spent years hunting for the Fisher-Price PXL-2000 (PixelVision), a 1987 toy camcorder that recorded wildly wonky and distorted images to audio tapes. (Though he could have paid a premium and bought a PixelVision on eBay at any time, Blahut says he much prefers shopping for vintage tech in its native habitat—thrift stores and garage sales.) There are a surprising number of active PixelVision users, and there have even been PixelVision film festivals in recent years—

you can view some entries online at the blog PXL This.

"There are a lot of benefits to the retro lifestyle," Blahut says. "If you're into video games, you could go into any thrift store and buy [a PlayStation 2 console] for $10. And you can get as many games as you want for a buck. You just have to get over the cachet of having the latest and greatest thing."

For those who want to go even further back in time, plenty of early arcade games in their original cabinets are available on eBay and sites like the Vintage Arcade Superstore. Although it's possible, with emulators such as the Multiple Arcade Machine Emulator, or MAME, to run the actual code of classic titles like *Asteroids* and *Space Invaders* on modern personal computers, the original software was designed hand in hand with the cabinets and controllers. (The truly ambitious can travel down the path trod by *IEEE Spectrum* contributing editor Paul Wallich in his July 2011 article, "Building Your Own Arcade Game.") Games like *Missile Command*, *Battlezone*, and *Spy Hunter* had one-of-a-kind features like outsize trackballs, unusual screen arrangements, or gearshift levers that can't be replicated with generic modern hardware. Depending on the condition and title, prices range from the mid-hundreds to low-thousands of dollars. —Mark Anderson

A version of this article appeared online in September.

start-up

LevelUp: Pay-by-Phone Innovator

Its mobile payment strategy emphasizes simplicity and security

In a few years, paying for goods and services with a smartphone could be as commonplace as swiping a credit or debit card today. According to the analyst firm Gartner, by 2016 the global value of mobile payment transactions will reach US $617 billion. But for now it's still early days, with a hodgepodge of technologies and business models vying for mindshare and market share, such as Google Wallet, Isis Mobile Wallet, and Orange's Quick Tap [see *IEEE Spectrum*'s special report "The Future of Money," June].

Boston-based LevelUp entered this arena last year with a free phone app that customers link to a debit or credit card. To pay at a store, the app generates an onscreen QR code (a type of high-capacity bar code). Holding the screen up to an in-store reader completes the transaction. No credit card information is transmitted. Instead, starting with the initial QR code scan, a series of tokens is passed from the customer, through the merchant and LevelUp's servers, and finally to one of LevelUp's payment processors: Braintree Vault and Bank of America. The tokens are then paired with credit card account information, and the charge is made.

The hook for customers is automated discounts: Typically, the first time users make a purchase from a merchant with LevelUp, they receive an instant credit from the merchant, usually in the range of $1 to $3 (merchants have the option of choosing to participate in this program). Subsequent purchases can then be tracked as part of a loyalty program determined by each merchant; for example, every $50 or $100 spent with a given merchant could trigger an immediate reward of, say, a $5 credit.

On the merchant's side, the appeal is the absence of per-transaction fees. With credit cards, these fees normally run to a few percent of each purchase. With LevelUp, merchants are billed only when a customer redeems a credit. Then, in addition to the cost of the credit itself, merchants are charged 35 cents on the dollar—for example, a $2 initial credit would cost a merchant a total of $2.70.

LevelUp claims that it can avoid charging per-transaction fees because of the system's low processing costs, which it attributes in part to the security of the token-based approach. LevelUp's stated fraud rate is 1 percent of that for conventional credit card use. The company is the

start-up

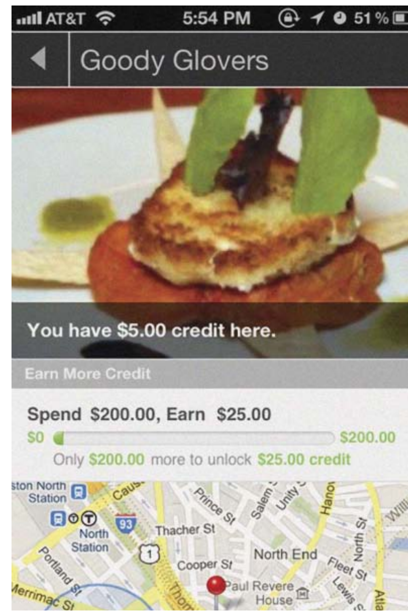

PAY HERE: An app on a customer's smartphone generates onscreen payment bar codes and records progress toward a reward [above]. LevelUp's reader [left] is based on a 3G smartphone and uses its camera to record the bar codes. *PHOTO: STEPHEN CASS; SCREEN SHOT: LEVELUP*

LevelUp
http://www.thelevelup.com
Founded: 2011
Headquarters: Boston
Founder: Seth Priebatsch
Employees: 148
Funding: $41 million *(includes capital raised by Scvngr)*

brainchild of Seth Priebatsch, who rose to prominence as the founder of Scvngr, an early leader in "gamification." The concept is to use gamelike elements to persuade people to participate in nongame activities, such as visiting a physical location and sharing that visit via social networks. The LevelUp technology was developed after Priebatsch saw that the original Scvngr mobile app was "really good at getting people to do things at places but wasn't really driving transactional revenue for businesses," says Christina Dorobek, LevelUp's vice president of partner development. "LevelUp is a tool to connect people to places, but through a transaction rather than social media."

LevelUp is available from at least one vendor in 25 U.S. cities at press time, with larger cities such as New York or San Francisco boasting hundreds of locations that accept LevelUp. This summer it raised $21 million in funding (in total, Scvngr and LevelUp have raised nearly $41 million since 2008). Google Ventures is an investor, as is Continental Investors, a venture capital firm founded by Phil Purcell, creator of the Discover credit card.

LevelUp's stiffest competition is likely to come from banks licensing similar bar-code and token-based technologies from companies such as FIS (an established banking and payment technology provider), says Peter Wannemacher, an analyst with the business intelligence firm Forrester Research. The banks can label the technology with their own brands and exploit their existing relationships with merchants and customers to gain market share.

To prevail, LevelUp must continue to differentiate itself to customers and merchants, with such features as its reward system, in order to build its active user base. So far, if the company's claim of 200 000 active users "is accurate, that's pretty good; that's exactly where they want to be," says Wannemacher.

The start-up is also working with cash register manufacturers to develop built-in support for LevelUp's system and is demonstrating a prototype e-commerce system, in which a QR code on a customer's phone is captured via webcam to authorize online payment.

—Stephen Cass

A version of this article appeared online in August.

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


CST STUDIO SUITE 2012
Discover what happens...

technically speaking

BY PAUL McFEDRIES

From Surf to Serf

By putting the means of production into the hands of the masses but withholding from those same masses any ownership over the product of their work, Web 2.0 provides an incredibly efficient mechanism to harvest the economic value of the free labor provided by the very, very many and concentrate it into the hands of the very, very few.
—writer Nicholas Carr

IT SEEMS like only yesterday that we were extolling the virtues of such *socially produced* wonders as Linux and Wikipedia. These communal endeavors heralded a utopian age of unparalleled access to systems and information at little or no cost, and converted Web 2.0 from a mere collection of technologies to a system of liberation and empowerment.

But the road to Utopia all too often ends up detouring through the business district, and Web 2.0 has been no exception. By offering the means of production free to their users, other leviathan sites, such as Facebook, Twitter, and YouTube, have generated enormous amounts of content at almost no expense. Even better, this content is a gold mine for targeted advertising.

Over in Utopia, the "workers" who generated all those articles, photos, tweets, and videos would get a cut of the profits they helped to generate. In the business district, however, users retain their amateur status, while the companies they labor for rake in billions. Worse, contributors don't even own the content they create. The smallest of the small print in the terms of use, which you must agree to in order to get an account, states that the company can use your content as it sees fit.

To Nicholas Carr, this smacks of exploitation rather than emancipation, and he coined a term for it: **digital sharecropping**. Just like the farm laborers of old who worked the land but didn't own it, **digital sharecroppers** grow the product that earns Web 2.0 companies their profits, but they relinquish ownership. Heck, even the most put-upon sharecropper earned a share of the crop he worked so hard to cultivate; today's **digital serfs** work on their profiles, timelines, and feeds for free, with targeted ads their only "compensation."

To be sure, no one on these sites sees their uploading of posts, photos, and videos as "work." But this ad hoc content creation isn't enough for many companies, so now they're outright asking nonemployees to contribute. This was preceded by *crowdsourcing*, which involves obtaining labor, products, or content from people outside the company, particularly from a large group of customers or amateurs who work for little pay. A good example is the phenomenon of **microwork**, a short, simple task that a company **crowdsources** for a small fee, particularly to workers in the underdeveloped world.

But even **microworkers** get paid, even if it's just a few pennies. The latest trend is **unsourcing**, where companies transfer internal functions from paid employees to unpaid volunteers, particularly customers on social networks. If the task is related to product support, as it often is, it's called **social support**.

Why would people do a company's work for free? Part of the answer lies in a kind of prestige that comes from answering a large number of customer questions or solving a large number of problems. Companies coax helpers to participate by offering points for useful answers, connecting total points to specific achievement levels, and encouraging users to **level up**, for even greater prestige. (The phrase comes from the game industry, and the way concepts and ideas move from the game culture to the broader one is called **gamification**.)

Anthony De Rosa, a product manager at Reuters, calls this **digital feudalism** and laments that we "are being played for suckers to feed the beast, to create content that ends up creating value for others." What's the solution? For many it's simply keeping your content to yourself by running your own blog, website, and even Web server. But it's really more about being aware. As the writer Tim Carmody has said, "Nobody is under any obligation to allow other people to make money off of them or their data. To pretend otherwise is just silly." Unfortunately, he said that on Twitter, so I suppose they own those words now. ❑

GREG MABLY

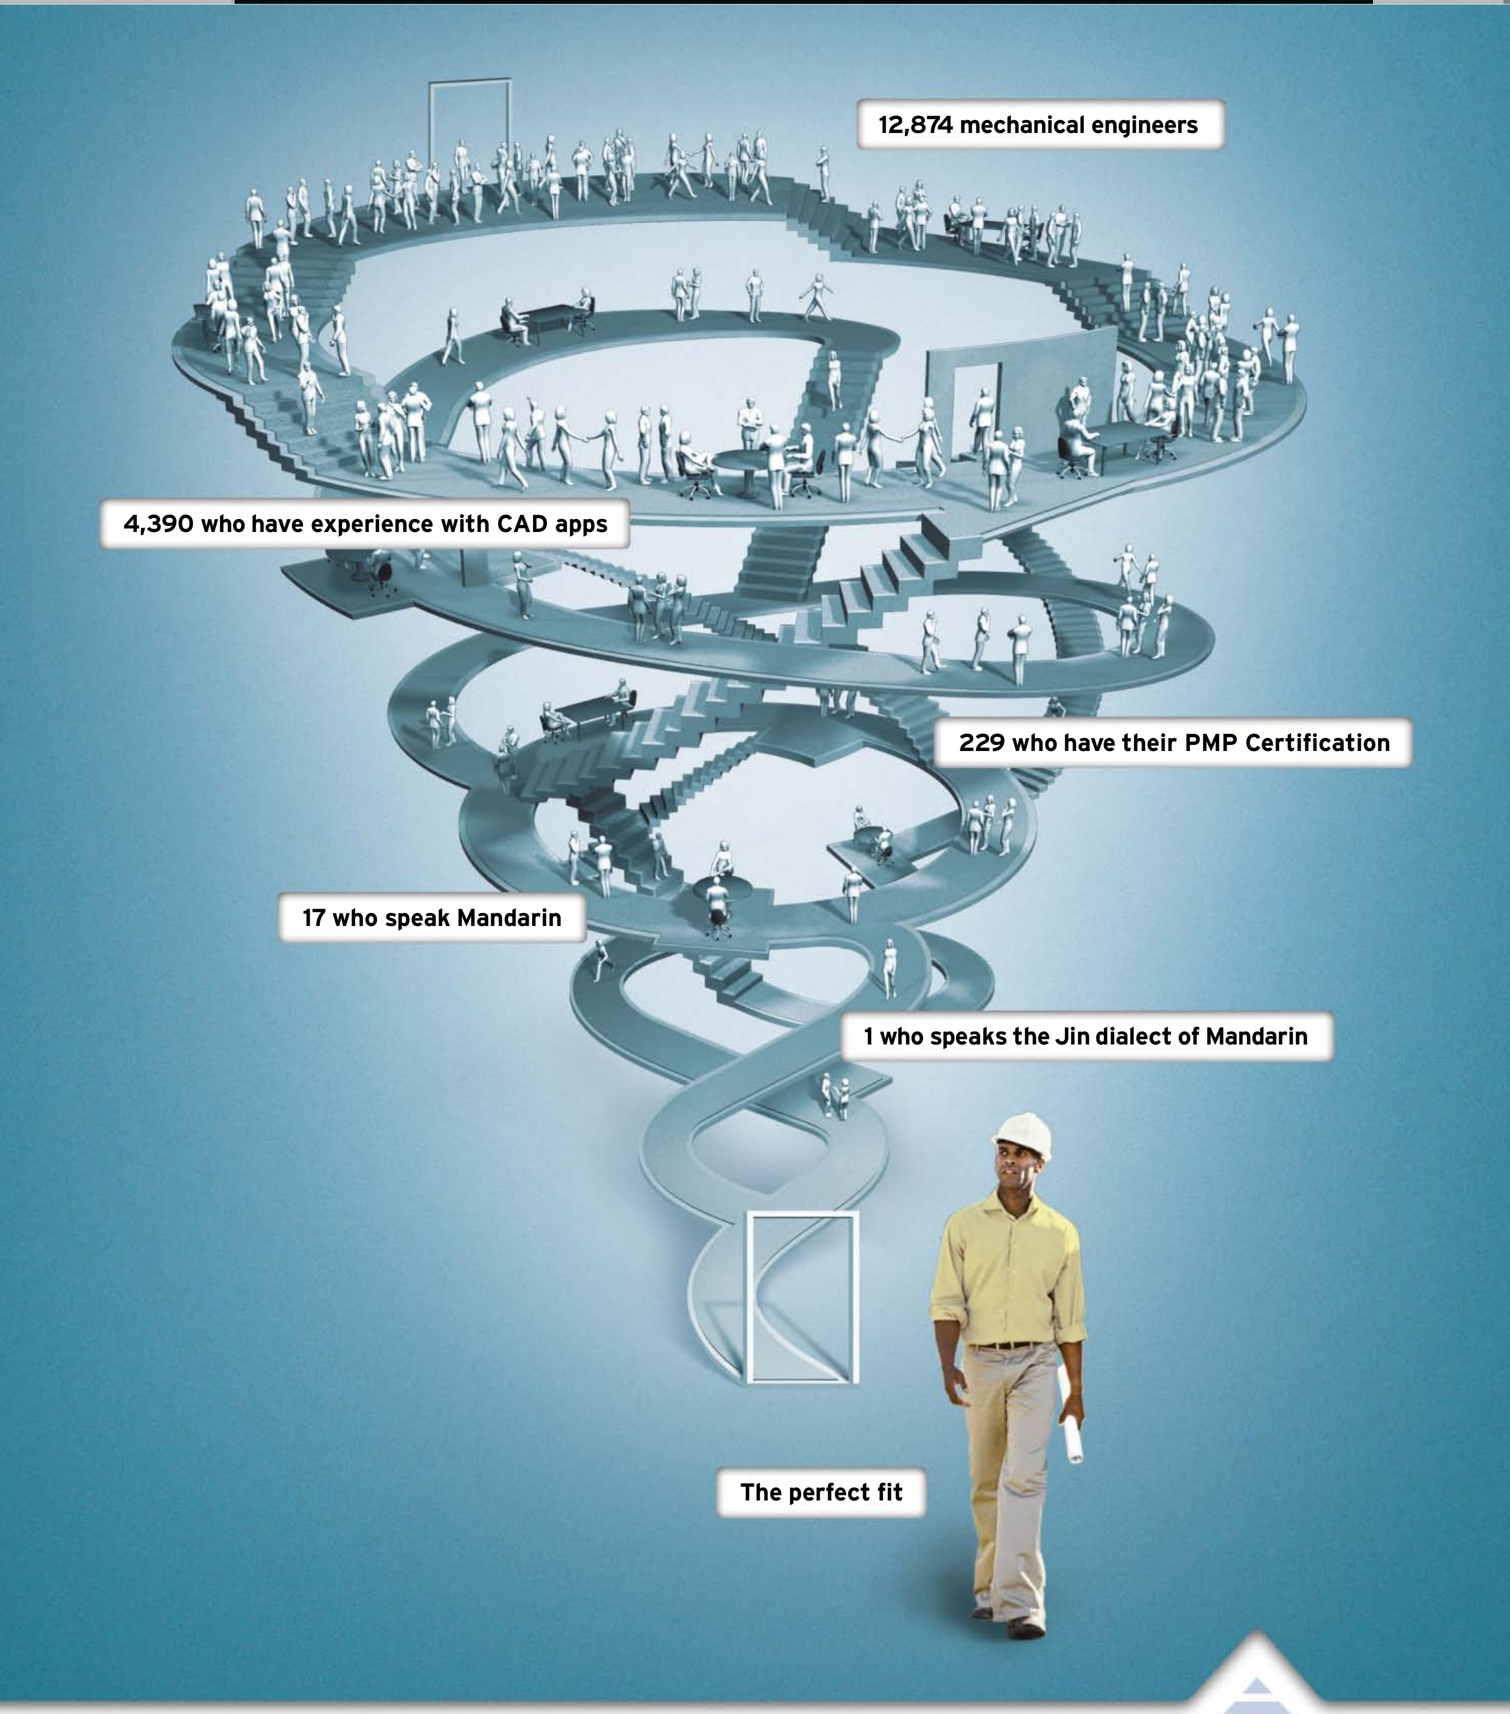
12,874 mechanical engineers
4,390 who have experience with CAD apps
229 who have their PMP Certification
17 who speak Mandarin
1 who speaks the Jin dialect of Mandarin
The perfect fit

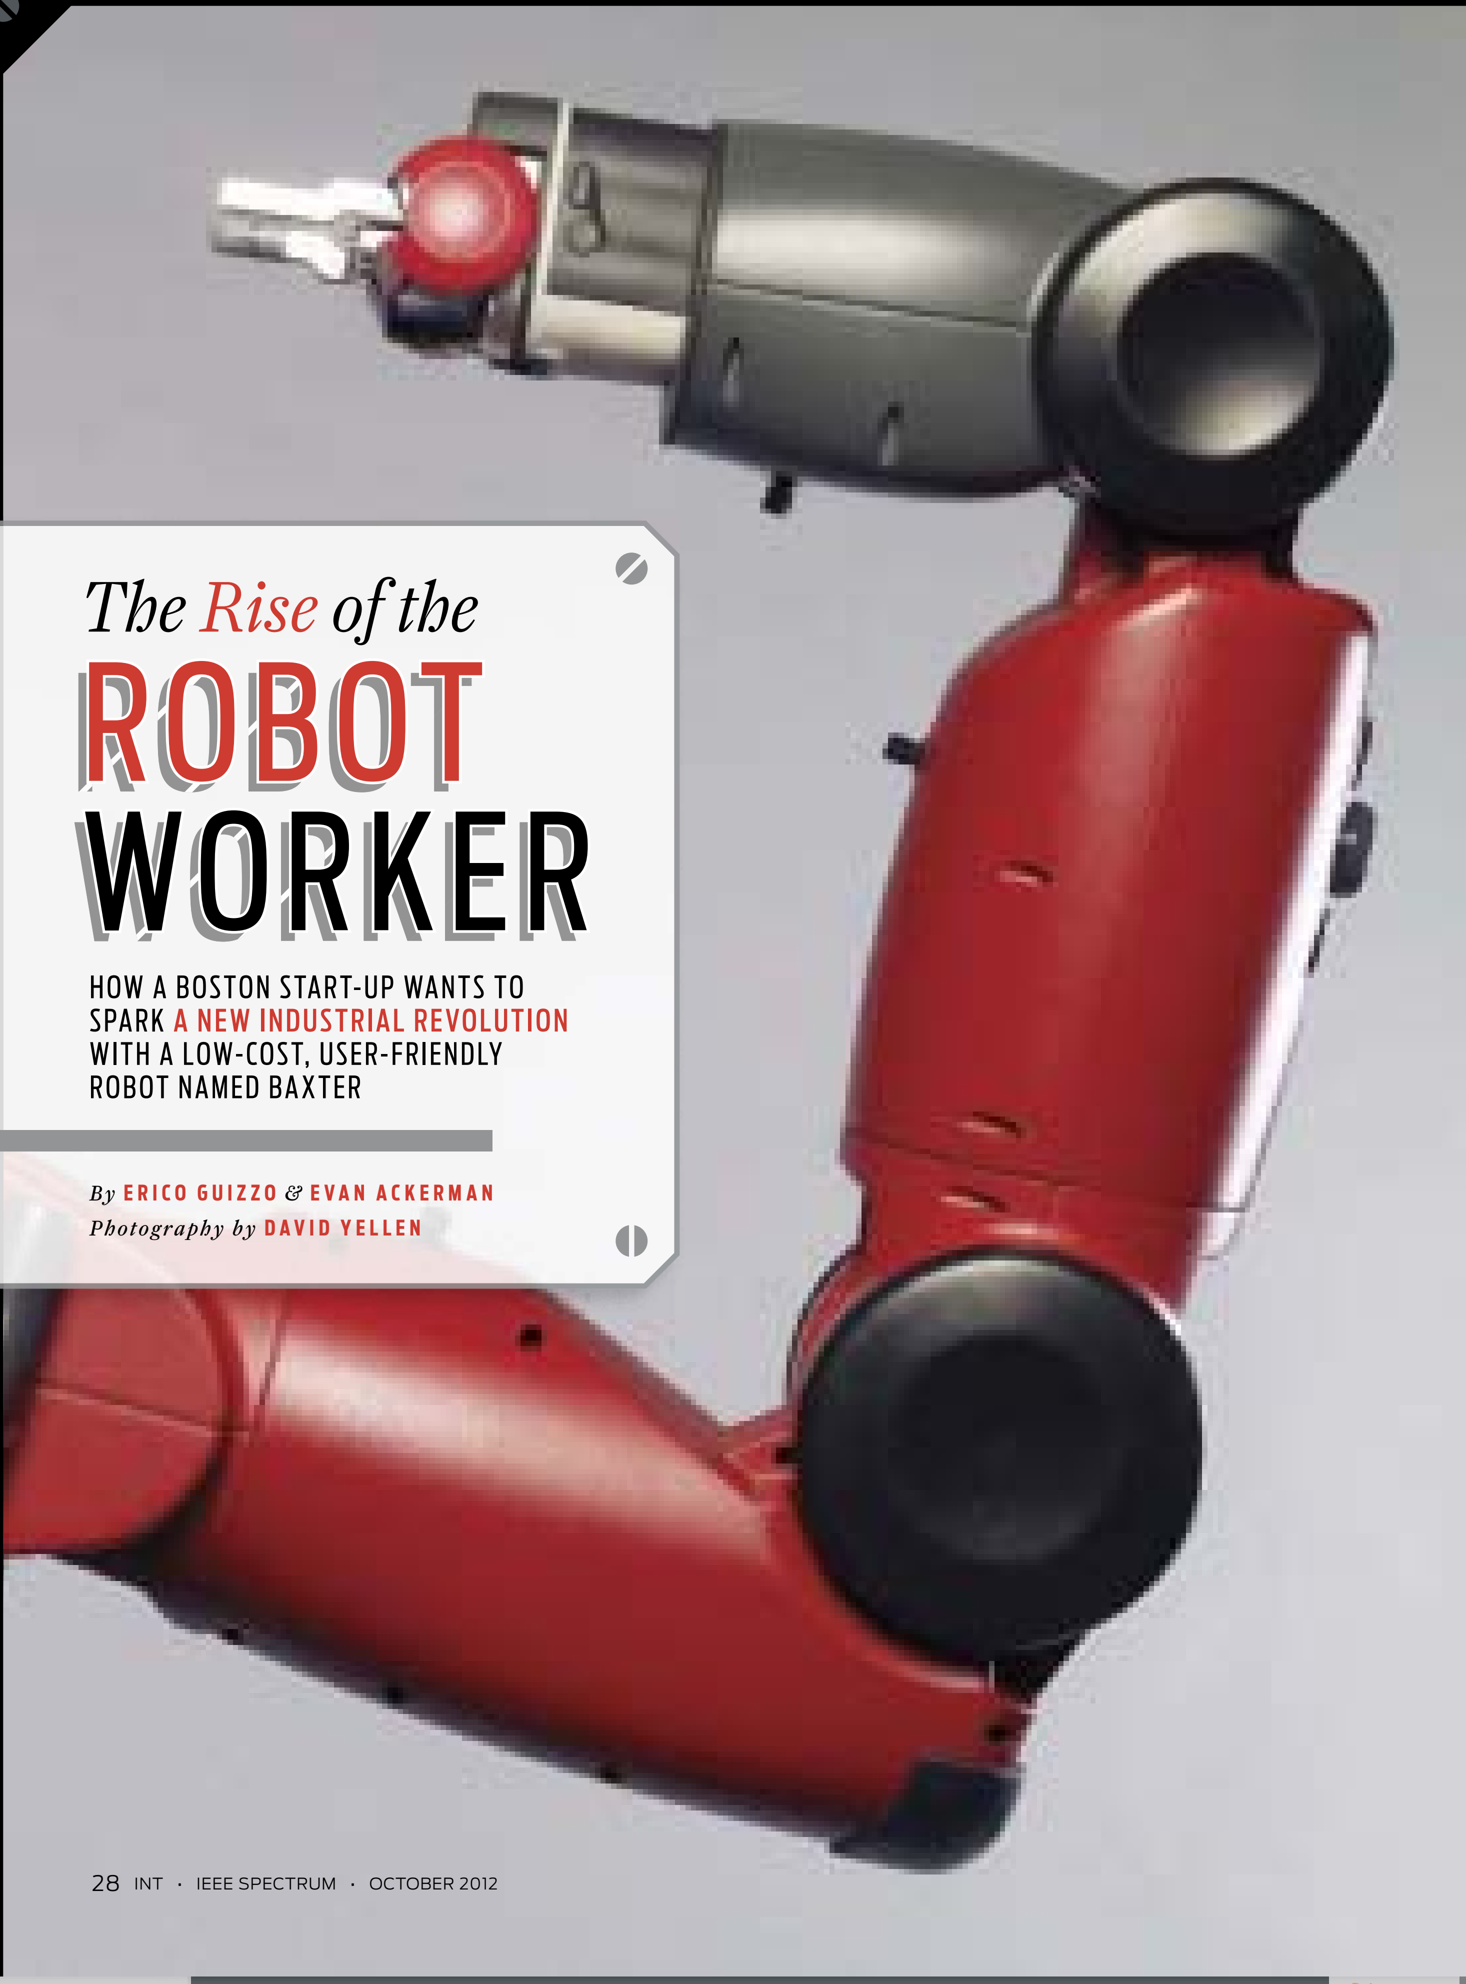

The Rise of the ROBOT WORKER

HOW A BOSTON START-UP WANTS TO SPARK A NEW INDUSTRIAL REVOLUTION WITH A LOW-COST, USER-FRIENDLY ROBOT NAMED BAXTER

By ERICO GUIZZO & EVAN ACKERMAN
Photography by DAVID YELLEN

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

MEET BAXTER: Rethink Robotics has developed a next-generation factory robot that is versatile, easy to program, and costs just US $22 000.

I N THE CENTER of an otherwise unremarkable office stand six large robotic torsos mounted on pedestals and positioned along a bench that's covered with piles of plastic widgets.

One robot methodically moves widget after widget onto a conveyor belt, the animated face on its LCD screen displaying an expression of quiet concentration. The task is mundane, but the robot is not: This is Baxter, the culmination of nearly five years of secretive development, based on the vision of Rodney Brooks, possibly the world's most celebrated roboticist. Now founder, chairman, and CTO of Rethink Robotics, the company that built Baxter, Brooks has his sights set characteristically high: to unleash a revolution in manufacturing with a friendly faced factory robot.

With US $62 million in funding from top-tier investors such as Bezos Expeditions and Charles River Ventures, Rethink Robotics has been the subject of a great deal of interest and speculation since its founding in 2008. Like others who track robotics, we'd heard the rumors: Rethink was focusing on manufacturing; its robots would be so inexpensive every factory would be able to afford one; the robots would help make workers more efficient and American factories competitive again. But nobody knew exactly what the company was up to—or if they did, they weren't talking. Rethink's robot was one of the best-kept secrets in robotics, but that secret is about to be revealed with the exclusive demonstration we're getting today in preparation for Baxter's first shipments this month.

Brooks is clearly thrilled to show off his latest creation. When he gets particularly animated, his eyebrows arch high on his forehead and his hands gesticulate wildly. You almost expect sparks to jump from his fingertips. "This is pretty cool," he says as he grabs a robot by its wrist. Baxter has thick, round arms, and you'd think they'd be heavy and rigid. But Brooks moves the arm around in a fluid motion, as if it were a weightless object. "When you hold the cuff, the robot goes into gravity-compensation, zero-force mode," he says. His eyebrows shoot up. "It's like the arm is essentially floating."

This seemingly simple feature shows just one of the many ways in which Baxter is unlike any other industrial robot. With two arms, each with seven axes of motion—or degrees of freedom, in robotic parlance—and a reach similar to that of a human, Baxter is designed to take over those simple, dumb, mindless tasks that humans hate to perform because they're

MAKING FACES: Baxter's face provides feedback to users.

so, well, *robotic*. Whereas traditional industrial robots perform one specific task with superhuman speed and precision, Baxter is neither particularly fast nor particularly precise. But it excels at just about any job that involves picking stuff up and putting it down somewhere else while simultaneously adapting to changes in its environment, like a misplaced part or a conveyor belt that suddenly changes speed.

What's more, Baxter is designed to be inherently safe. With their fast, powerful motors and hefty limbs, industrial robots are typically kept fenced off from people. Baxter's limited speed and lower weight—about 75 kilograms (165 pounds), or as much as an average adult man—mean that it can operate right alongside human workers. Brooks is so confident in this feature that he often puts his own body on the line: "I'll just walk in and have this thing bash into my head, and it's fine."

There are two other major barriers to the adoption of industrial robots that Rethink wants to overcome: ease of use and cost. As for the first, Baxter doesn't rely on custom programming to perform new tasks. Once it's wheeled into place and plugged into an ordinary power outlet, a person with no robotics experience can program a new task simply by moving Baxter's arms around and following prompts on its user-friendly interface (which doubles as the robot's face). And while a traditional two-armed robot, including sensors and programming, will typically set you back hundreds of thousands of dollars, Baxter costs just $22 000. To achieve that, Rethink designed the robot from scratch. Underneath Baxter's plastic exterior lie thousands of ingeniously engineered parts and materials that enable the robot to do what it does for the cost of a midsize car.

All these features make Baxter a potentially disruptive force in an industry ripe for disruption. Today, large suppliers and their large customers dominate industrial robotics. But almost any manufacturer can afford a Baxter (or two). In the United States alone, there are roughly 300 000 small and medium-size manufacturers, Brooks points out. So it's not inconceivable that Rethink could make humanoid robots a normal part of the manufacturing process for businesses of all sizes. And by improving the efficiency of human employees, it could make making things in the industrialized world just as cost effective as making them in the developing world.

Building robots for these small and medium-size companies "is a fantastic opportunity," says Henrik I. Christensen, a professor of robotics at the Georgia Institute of Technology, in Atlanta, who's an expert in industrial automation. (He has no ties to Rethink.) There are many tasks, he says, that don't require the speed and precision of today's industrial robots, and these tasks are begging to be automated.

But Rethink's robot invasion faces many challenges. The biggest by far will be selling the robot. For many folks, robots still evoke mental images of Robby the Robot, or worse, the Terminator. Rethink has to convince business owners that Baxter is not a gimmick, that it really can boost

BAD BOY: Rodney Brooks, who has been called the "bad boy of robotics," is back with another disruptive creation: a factory robot to help workers become more productive.

ARM & ACTUATORS
Each arm has seven degrees of freedom and can lift 2.25 kilograms (5 pounds) and perform tasks at a roughly human cadence with an accuracy of about 0.5 to 1 centimeter. The arms are powered by series elastic actuators.

How BAXTER WORKS

Baxter may not be superfast, superstrong, or superprecise like other industrial robots, but it's smarter. It uses vision to locate and grasp objects, and you can program the robot to perform a new task simply by holding its arms and moving them to the desired position. Baxter will even nod its head to let you know it has understood you.

productivity. Brooks isn't worried. He acknowledges that, in the beginning, Baxter will perform only a limited number of tasks. But as Rethink continues to upgrade the robot—especially its software—Baxter's full potential will become evident. And then, Brooks predicts, his new robot will "sell like hotcakes."

AFTER INTRODUCING us to Baxter, Brooks has us follow him through the maze of beige cubicles, test stands, and scattered piles of robot parts that make up Rethink's Boston headquarters, which is unusually quiet today. "We had a company outing yesterday," the 57-year-old Brooks says by way of explanation, a faint accent revealing his Australian origins. "Maybe some folks are still goofing around."

Once we get to his office, Brooks relaxes on a couch. Dressed in a button-down shirt, jeans, and red-and-blue argyle socks, he's surrounded by relics of his robotics past. One framed picture shows Genghis, a six-legged robot he built at MIT in the late 1980s, which could crawl and climb over obstacles with a remarkable, insect-like gait. Contrary to the conventional wisdom of the time, Genghis required very little programming to produce its complex behavior. It was just one of many bold, counterintuitive results that surprised the community and helped establish Brooks's reputation as "the bad boy of robotics."

In addition to his notable academic career, Brooks and two of his students, Colin Angle and Helen Greiner, founded iRobot in 1990. The start-up broke ground in two distinct areas: its Roomba vacuum cleaning robot, which became a runaway hit in the consumer market, and its PackBot military robot, used in Iraq and Afghanistan to scout for improvised explosive devices. With these accomplishments under his belt, Brooks could have retired comfortably. Instead, he decided to phase out his involvement with iRobot, gave up his tenured position at MIT, and founded Heartland Robotics, which recently rethought its name and became Rethink.

Brooks's vision for Rethink can be traced back to the time he spent in China supervising the manufacture of iRobot products. "I realized that [outsourcing manufacturing to China] wasn't sustainable, because once the cost of Chinese labor starts to go up, the appeal of doing a product there starts to go away," he says. He concluded that a fairly simple robot could do lots of those tasks, like basic material handling, packing and unpacking boxes, and polishing and grinding. Brooks's second realization was that the same advances in processors and sensors that were making PCs and smartphones better and cheaper would benefit robots, too.

From the beginning, says David "DMX" Lewis, Rethink's senior mechanical engineer and one of its earliest employees, "the vision was a two-armed robot and a head, because Rod had envisioned something that can go in and sit next to a person on the assembly line and help them out." But they decided to start with a simple design and build up from there. At first, the robot was just a single joint powered by two gearboxes salvaged from power drills. Then it grew into an arm with four degrees of freedom, a configuration known as SCARA, for selective compliant assembly robot arm. Next came an arm with six degrees of freedom, similar to a popular robot arm of the 1980s called PUMA (programmable universal machine for assembly). After much internal debate, Rethink's engineers decided that it needed a seventh degree of freedom, to better approximate the range of motion of a human arm. A team of 30 mechanical and control engineers completed the first two-armed prototype late last year and a second-generation design early this year.

As the complexity of the robot has increased, keeping the cost down has been a constant struggle. The engineers found creative ways to work with less expensive components, which often meant designing parts out of cheaper materials. (At one point they attempted to build the entire arm, including the gear systems, out of plastic, but that didn't pan out.) And by working closely with local manufacturers, they were able to get the parts they needed more cheaply and without compromising quality.

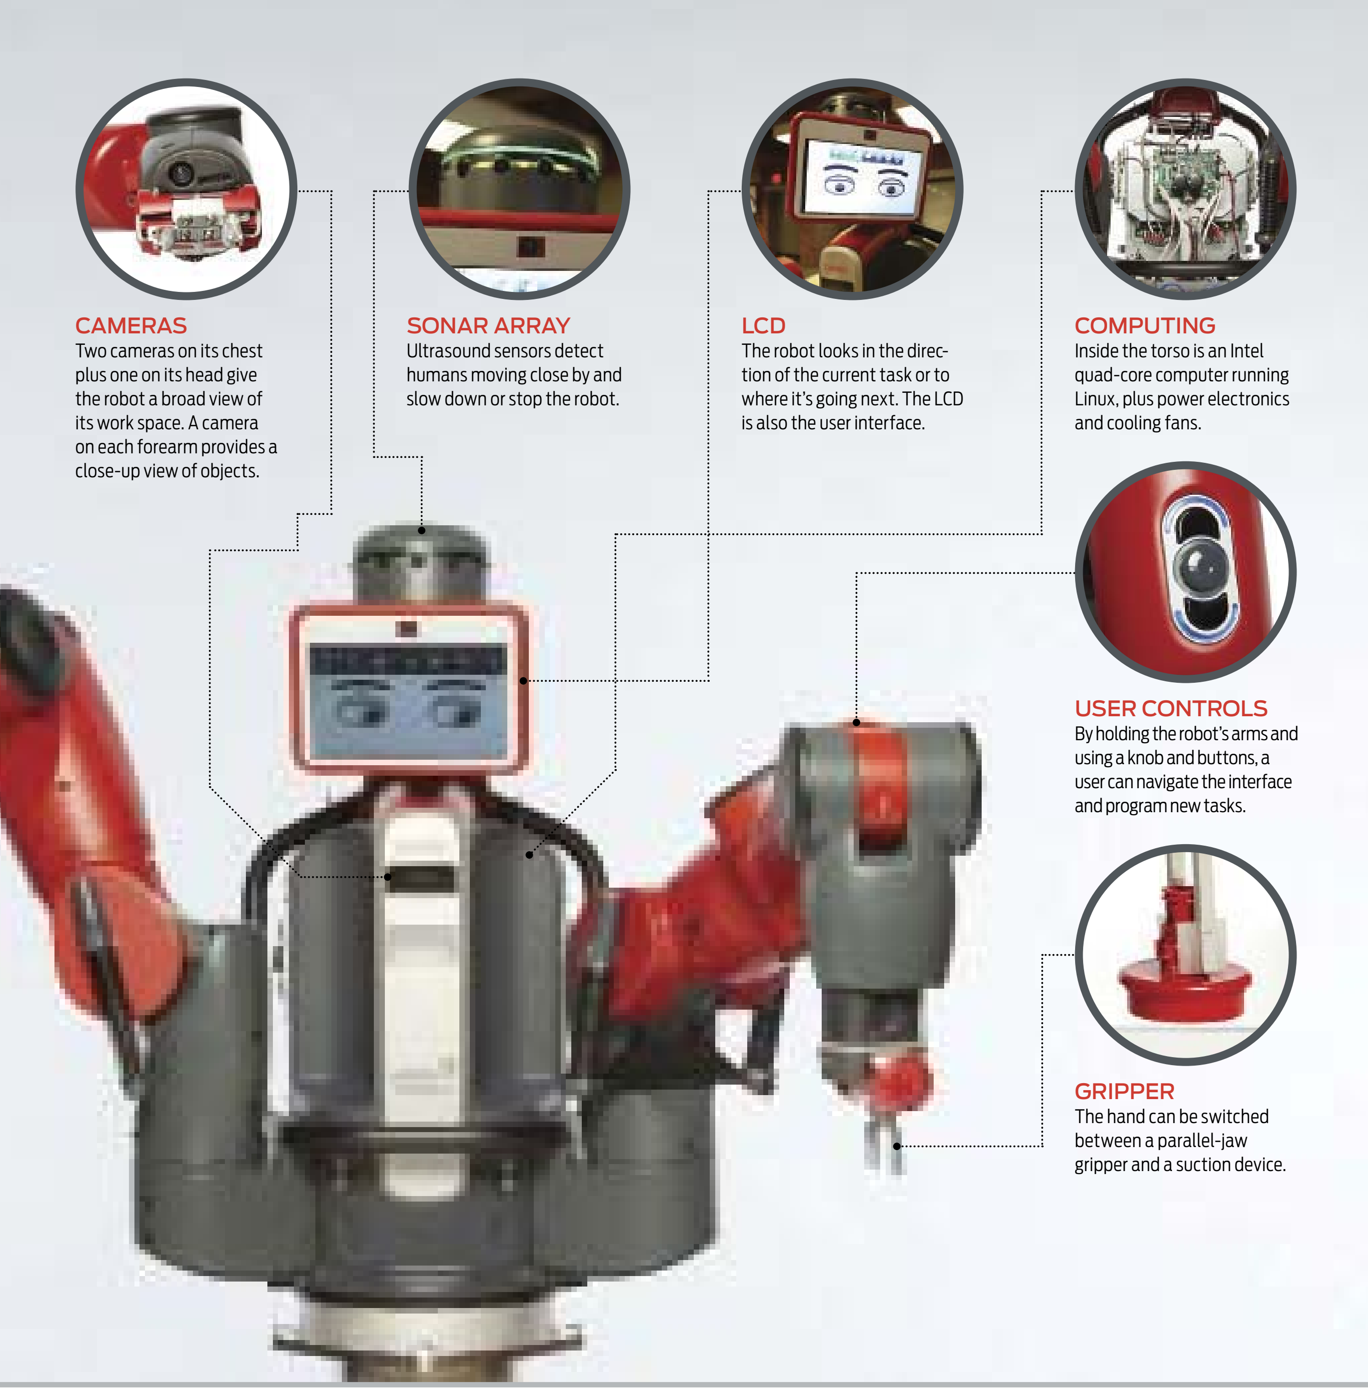

SONAR ARRAY AND LCD: ERICO GUIZZO

Brooks picks up a set of gears and waves it around. A Pennsylvania shop made them, he says, using a metallurgical process involving pressed powdered metal, which ended up being only about a fifth of the cost of traditional gears. The robot has hundreds of parts like these, and so the savings add up. Rethink used this design-for-manufacturing approach for as many parts as possible, eventually procuring 75 percent of them in the United States—a fact that Baxter advertises with a big "Made in U.S.A." sign stamped on its back.

THE OTHER way Rethink has cut costs is through software. By giving Baxter the ability to autonomously compensate for its own mechanical irregularities as well as changes in its environment, Rethink has been able to avoid having to use costly components. So Baxter may not have the most expensive motors or the stiffest bearings or the most precise gearing, but it doesn't need those things. For example, backlash in the gears—the slack in teeth couplings that may cause motion loss and vibrations—can be modeled and then adjusted on the software side so that the arms move smoothly. And the robot can use cameras to improve the accuracy of its movements, explains Elaine Chen, vice president of product development. "You don't need the robot to be able to go somewhere within 0.01-millimeter accuracy in order to pick something up," she says, "because the robot is going to see it, get a fix on it, and then pick it up."

This all gets to one of Baxter's key features: compliance. A robot is said to be compliant when it's not completely rigid and when it can sense and control the forces it applies to things. "If

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

you want a robot that's going to deal with an unstructured environment, it can't be stiff," says Matthew Williamson, Rethink's director of technology development. "You need compliance."

This idea of compliance is embodied, as it were, in Baxter's arms, which contain a mechanism called a series elastic actuator. The concept was invented and patented in the mid-1990s by then–MIT researchers Gill Pratt (now a program manager at the U.S. Defense Advanced Research Projects Agency) and Williamson. As Williamson explains, in a traditional actuator a motor drives a gearbox, which in turn drives a joint—say, the elbow of a robot arm. The motor, gearbox, and joint are rigidly attached to one another. What you get is speed and precision, which are great for a welding robot at a car plant but not for a robot that will operate around humans. If you try to hold the elbow when it's turning, for example, a traditional robot won't feel your hand and so will just keep moving.

In a series elastic actuator, however, the motor and gearbox drive a spring, and it's the spring that drives the joint. The spring makes the actuator much less stiff and also lets the device measure forces. That's because the force on the joint can be calculated from the deflection of the spring, as famously stated by Hooke's Law of elasticity. And if the control system knows how much force is being applied on each joint, the robot can use that as a feedback signal to react accordingly. So when you grab Baxter's elbow, it senses the added force and slows down, and it can even follow your motions—the "gravity-compensation, zero-force mode" that Brooks had demonstrated earlier that day.

The series elastic actuators also act as filters that help reduce friction and backlash in low-cost gearboxes, and they turn into shock absorbers if the robot ever accidentally whacks something, protecting both the robot and whatever or whomever it came in contact with. Lewis, the mechanical engineer, explains that having good force control enables Baxter to do things like pick up parts or push buttons without needing a degree of position control that would have required more expensive components. "Our robot has a certain level of softness and compliance," he says, "so it can go push a button and you don't have to worry about precision."

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/robot1012

As SOPHISTICATED as Baxter's hardware is, the "real breakthrough," according to Brooks, is "the way you program the robot." Using the word *program* to describe how you teach Baxter new tasks is perhaps overcomplicating things: It isn't so much programmed as it is simply shown what to do.

Programming a traditional robot requires writing code or running it through its paces by pushing buttons on a handheld control box called a teach pendant. It's not difficult to use a pendant, but it's time-consuming, with certain tasks requiring hours of tedious button pushing.

To show Baxter how to take an object out of a box and put it on a conveyor belt, you start by grabbing the robot by the wrist to get its attention. Baxter will stop whatever it's doing and look at you with the calm, confident eyes displayed on its LCD. You then move the arm over to the box and use buttons and a knob on the arm to navigate a series of menus on the LCD, telling the robot to use its vision to find the object. Finally, you move the arm over to the conveyor and push some more buttons to let Baxter know that this is where you want the object dropped off. Baxter even nods its head, as if to say, "I get it." Pressing the "play" button makes the robot execute the task on its own.

Despite Baxter's ease of use, how will factory owners and workers react when the machine gets rolled onto their shop floor? To find out, Rethink sent Baxter on a series of field tests, asking the folks who'll be using (and purchasing) the robot for feedback. One of them is Chris Budnick, president of Vanguard Plastics Corp., a 30-employee plastics manufacturer in Southington, Conn. Budnick met Baxter during a trial early this year. "They're really onto something with the user interface," he says. He programmed the robot to pick up a part from a conveyor belt and place it on a table in less than 10 minutes.

ARM WRESTLING:

Baxter isn't the only new robot vying for the manufacturing market. Here are some of the other contenders. (Actual prices depend on model and configuration.)

ABB
(ZURICH)
Unveiled a dual-arm concept robot designed for high-precision assembly applications, in 2011.
PRICE: not available

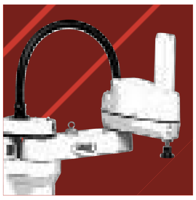

ADEPT TECHNOLOGY
(PLEASANTON, CALIF.)
Offers a variety of SCARA, six-axis, and parallel robots for manufacturing.
PRICE: not available

BARRETT TECHNOLOGY
(CAMBRIDGE, MASS.)
Its main product is the WAM, a cable-driven seven-degrees-of-freedom fully compliant arm.
PRICE: US $100 000

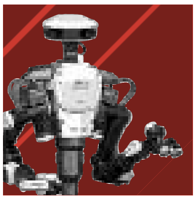

KAWADA INDUSTRIES
(TOKYO)
Its Nextage robot has two six-degrees-of-freedom arms designed for assembly.
PRICE: $85 000–$100 000

REDWOOD ROBOTICS
(MENLO PARK, CALIF.)
Founded by Meka Robotics and others, it's creating a robot for flexible manufacturing.
PRICE: not available

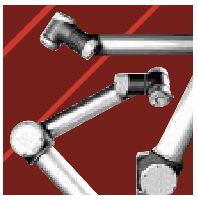

UNIVERSAL ROBOTS
(ODENSE, DENMARK)
Offers the UR-5 and UR-10 arms, designed to be affordable, safe, and easy to program.
PRICE: €20 000

FROM LEFT: ABB; ADEPT TECHNOLOGY; BARRETT TECHNOLOGY; KAWADA INDUSTRIES; MEKA ROBOTICS; UNIVERSAL ROBOTS

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

He had some quibbles, though. The Baxter prototype he saw was a little too big and slow, he says, and the robot took too long to find and grab the parts. That said, he wouldn't think twice about buying a Baxter, estimating payback in about a year.

Rethink says Baxter will get faster—and smarter. For example, it can adapt on its own to changes in position and lighting and to differently shaped objects. It's also easy to reprogram the robot: You simply get its attention and run it through the new steps. Try to do that with a traditional robot, says Mitch Rosenberg, Rethink's vice president for marketing and product management, and "you'd have to hire a team of software and hardware guys."

PART OF Rethink's plan to revolutionize manufacturing involves a business model that hasn't yet been applied to robotics but is now commonplace in the computer world. Early in the evolution of the computer, both the hardware and software were expensive and customized. The true PC revolution came when users were able to buy a standard platform that they could expand and upgrade to fit their specific needs.

And that's what Rethink is offering: future potential, embodied in the promise of frequent software upgrades that will give Baxter the ability to perform entirely new tasks. Rosenberg puts it this way: "The day you buy the robot is the day that it'll perform the least well. Over time, your investment will become more and more valuable because the software will become more and more valuable."

To that end, the company is planning to release a software development kit, or SDK, next year that will let others delve into Baxter's guts and modify its capabilities. There will also be a variety of third-party grippers for handling parts with different shapes. What's more, Baxter is relying on open platforms. It currently runs Linux and ROS (short for robot operating system), a software platform that's becoming increasingly popular within the robotics community. Rethink is now looking for people to develop a ROS-compatible SDK and is considering eventually supporting a completely open-source framework for Baxter.

But what's going to most dramatically set the robot apart is the $22 000 price tag (which includes a year of software updates and the warranty). "We want to have a product that's easy for a customer to try," says president and CEO Scott Eckert. "The whole notion of trial doesn't really exist in the traditional robotics market. You don't try traditional robots: You do an exhaustive amount of analysis and you implement and go. Our price point allows you to just go ahead and try it and see what happens." The company is confidently expecting more orders than it can fill after it starts shipping robots this month. A production software release is planned for January. Looking further ahead, Rethink will continue developing Baxter's software to add more capabilities, and eventually it'll release new hardware platforms, says Eckert.

Even so, the company is starting to feel pressure from the competition. Swiss-Swedish giant ABB has developed a dual-arm prototype, reportedly for assembly applications. Japanese firm Kawada Industries has a similar robot named Nextage. They might not cost as little as Baxter, but they will likely be able to perform high-precision tasks that Baxter can't do, like assembling electronics boards, another potentially huge market for robotic automation. Baxter isn't going to make your next iPhone, but these other robots might.

Other start-ups are also looking to enter the low-cost robotics market. A Silicon Valley company called Redwood Robotics is working on a manipulator that may become a direct competitor of Baxter. And Universal Robots, a Danish company, has a compact and compliant robot arm that sells for about €20 000 and is now used in a dozen countries in Europe and Asia. "Our robot is like a tool that people buy and use without help from an expert," says Universal Robots cofounder and chief technology officer Esben Østergaard. The concept sounds very similar to what Rethink proposes for Baxter. Østergaard says his company has already shipped robots to the United States, so Baxter is "not going to be the first robot to enter this new market."

Another challenge Rethink will have to face is the perennial question, Will the robot take people's jobs? Some people see automation as a force that needs to be checked, and not even a friendly faced robot will make them warm up to the idea of pushing even more automation into factories.

PROGRAM ME, PLEASE: Product manager Mike Bugda demonstrates how to program Baxter to perform a new task. He simply holds the arm by the cuff and moves it to the desired location. He tells the robot what to do using a knob and buttons.

To be clear, nobody at Rethink believes that armies of Baxters are going to take jobs away from humans. They may shift jobs away from the developing world if U.S. and European manufacturers bring production back home, Brooks says, but the robots are not designed to replace workers. "The robot can't do the full job of any one person, but it can do the simple cases of the simple tasks," he says. "So let the person do the stuff that they're better at, which makes the job more challenging and more interesting than some totally mindless thing."

Brooks has attended some of the field tests to see how potential customers interact with the robots. He says that when they bring Baxter onto a factory floor, some workers look puzzled, while others just laugh at the machine. But after they use Baxter for a while, people realize how the robot can help them. "And then something interesting happens," Brooks says, his eyebrows rising. "People sort of personify the robot. They say, 'It's my buddy!' " ❑

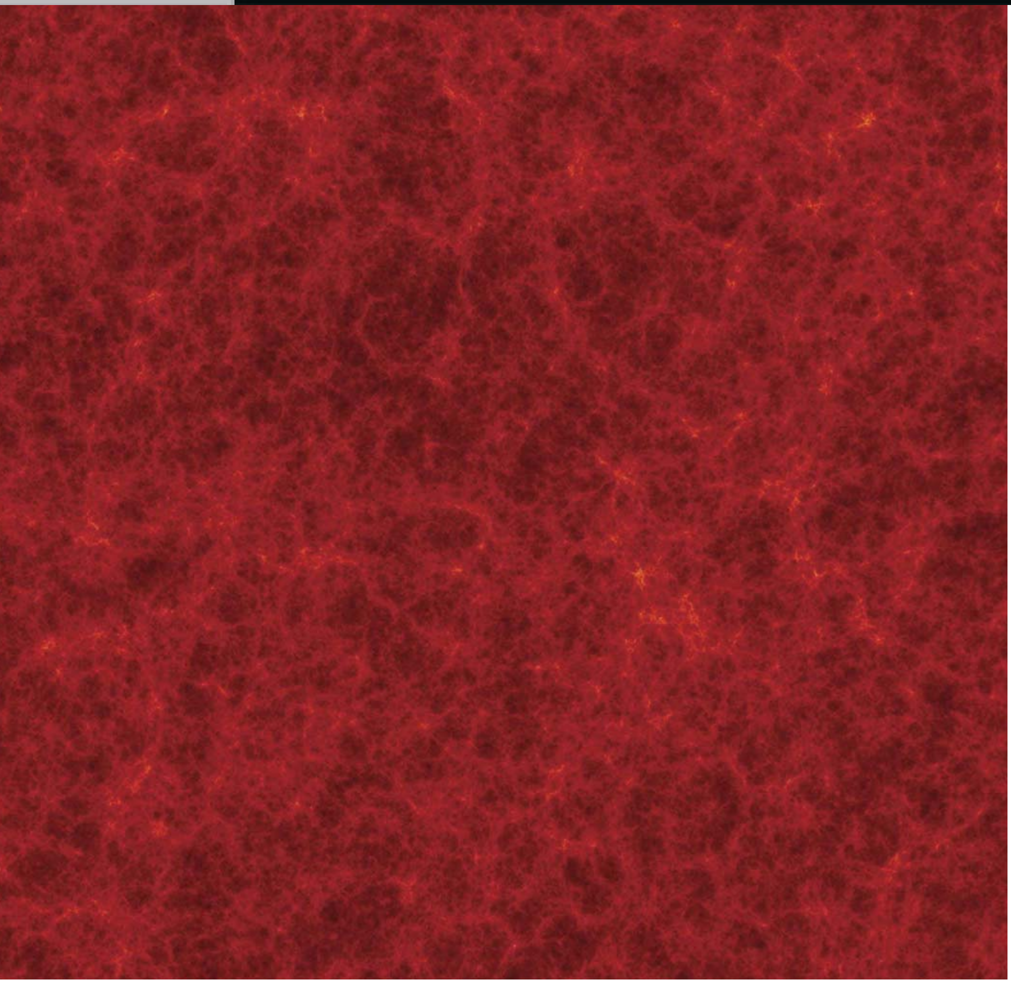

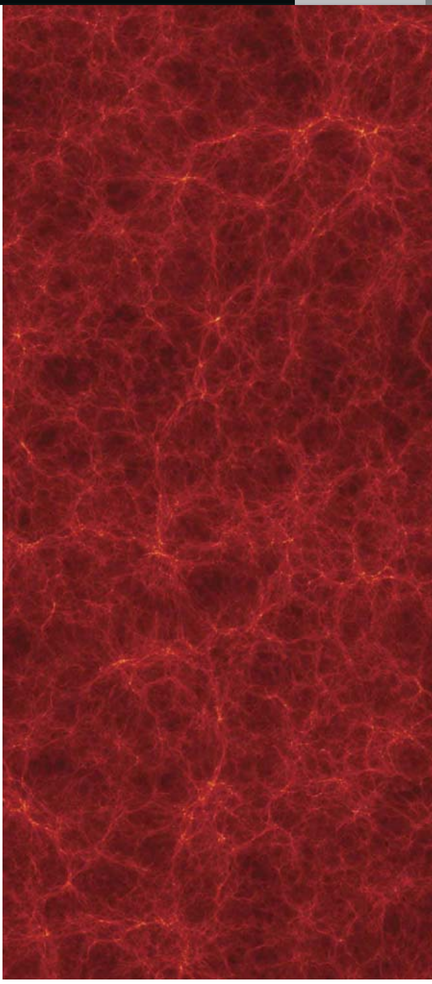

THE UNIVERSE IN A SUPERCOMPUTER

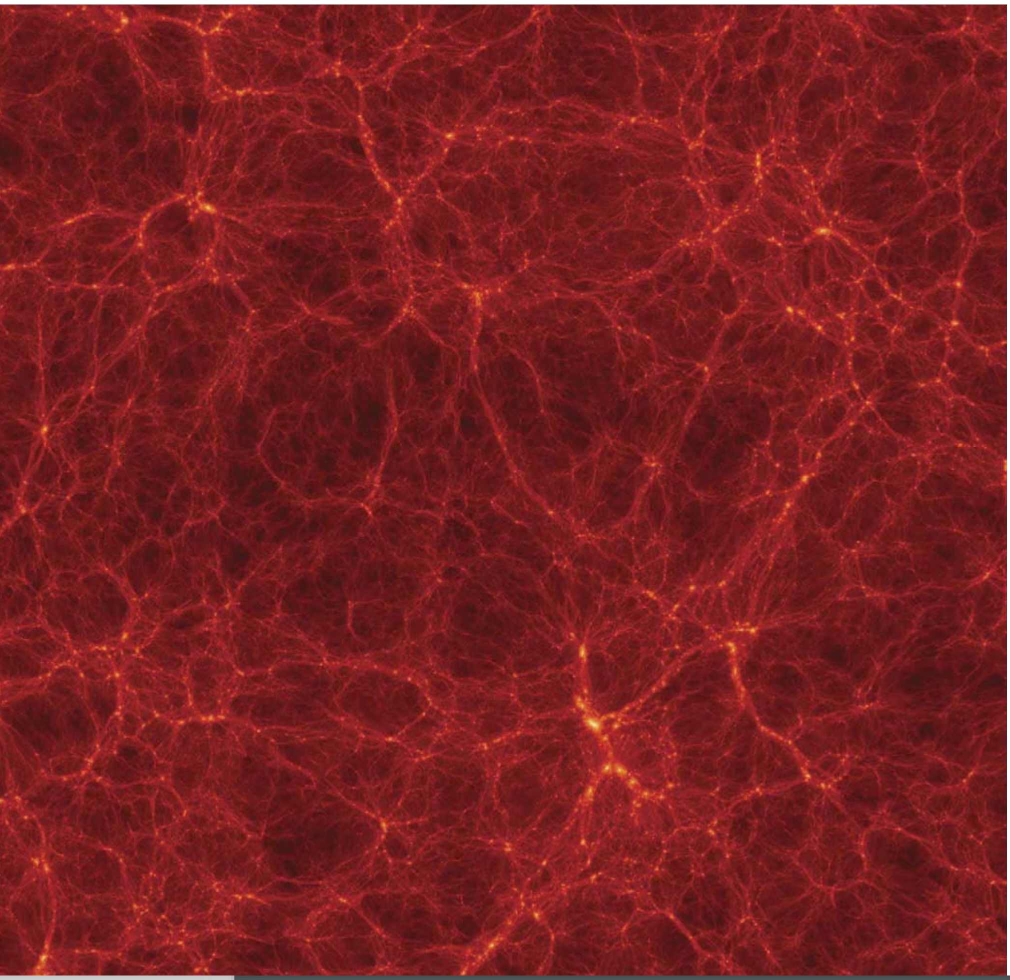

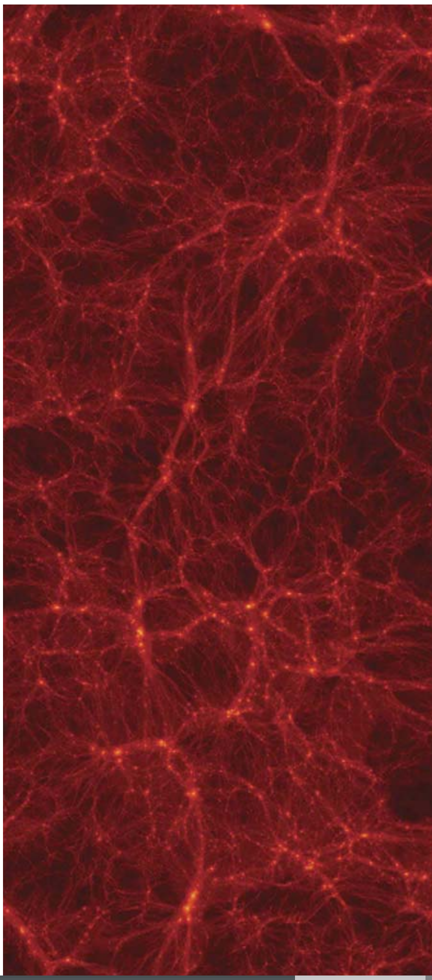

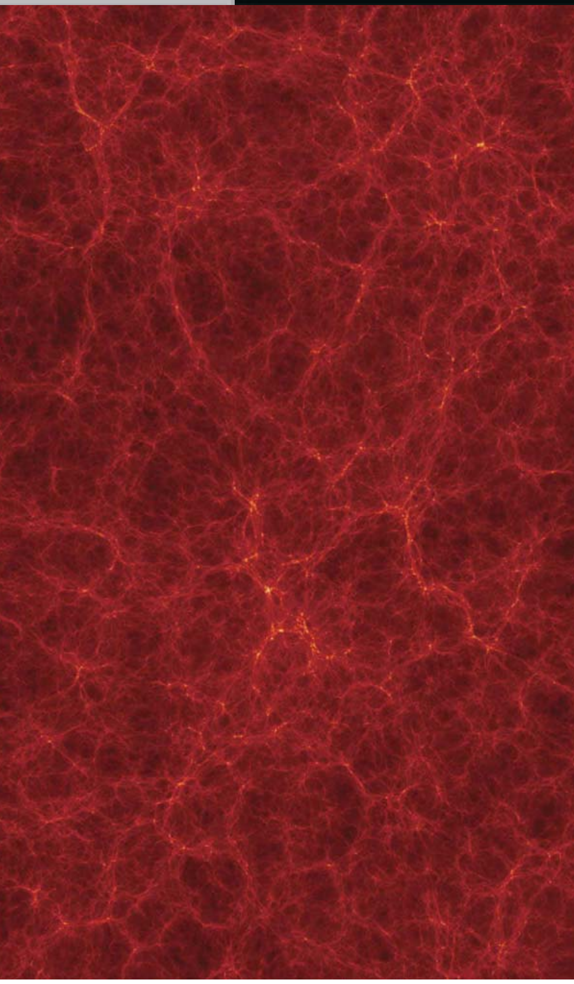

COSMIC WEB: The Bolshoi simulation models the evolution of dark matter, which is responsible for the large-scale structure of the universe. Here, snapshots from the simulation show the dark matter distribution at 500 million and 2.2 billion years [top] and 6 billion and 13.7 billion years [bottom] after the big bang. These images are 50-million-light-year-thick slices of a cube of simulated universe that today would measure roughly 1 billion light-years on a side and encompass about 100 galaxy clusters.
SOURCES: SIMULATION, ANATOLY KLYPIN AND JOEL R. PRIMACK; VISUALIZATION, STEFAN GOTTLÖBER/LEIBNIZ INSTITUTE FOR ASTROPHYSICS POTSDAM

To understand the cosmos, we must evolve it all over again

By Joel R. Primack

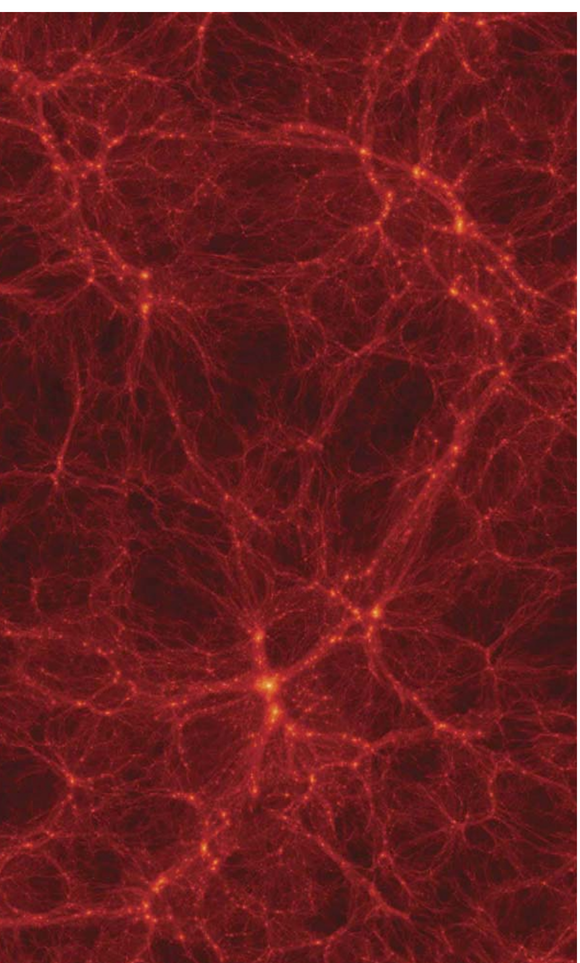

WHEN IT COMES TO RECONSTRUCTING THE PAST, you might think that astrophysicists have it easy. After all, the sky is awash with evidence. For most of the universe's history, space has been largely transparent, so much so that light emitted by distant galaxies can travel for billions of years before finally reaching Earth. It might seem that all researchers have to do to find out what the universe looked like, say, 10 billion years ago is to build a telescope sensitive enough to pick up that ancient light.

Actually, it's more complicated than that. Most of the ordinary matter in the universe—the stuff that makes up all the atoms, stars, and galaxies astronomers can see—is invisible, either sprinkled throughout intergalactic space in tenuous forms that emit and absorb little light or else swaddled inside galaxies in murky clouds of dust and gas. When astronomers look out into the night sky with their most powerful telescopes, they can see no more than about 10 percent of the ordinary matter that's out there.

To make matters worse, cosmologists have discovered that if you add up all the mass and energy in the universe, only a small fraction is composed of ordinary matter. A good 95 percent of the cosmos is made up of two very different kinds of invisible and as-yet-unidentified stuff that is "dark," meaning that it emits and absorbs no light at all. One of these mysterious components, called dark matter, seems immune to all fundamental forces except gravity and perhaps the weak interaction, which is responsible for

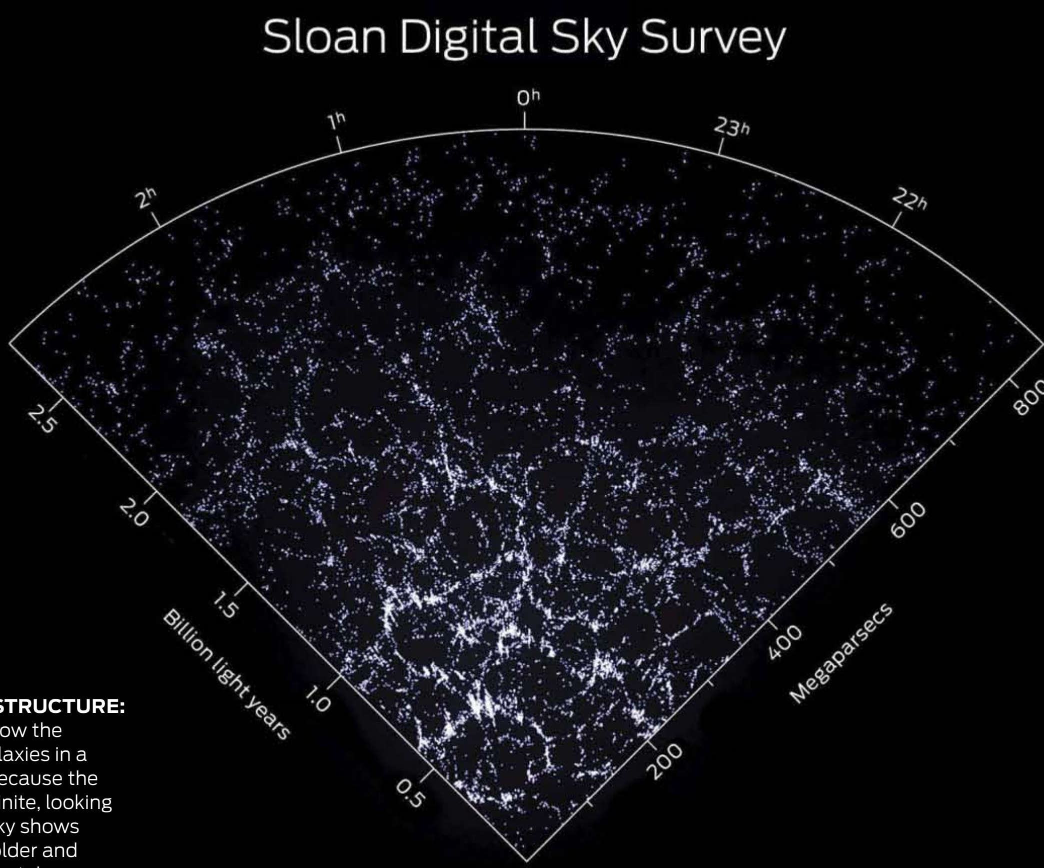

LARGE-SCALE STRUCTURE: These wedges show the distribution of galaxies in a slice of the sky. Because the speed of light is finite, looking deeper into the sky shows objects that are older and farther away. Here, telescope observations made by the Sloan Digital Sky Survey are compared with the distribution of galaxies calculated based on the Bolshoi dark matter simulation. They show a good match statistically.

SOURCES: NINA MCCURDY/UNIVERSITY OF CALIFORNIA, SANTA CRUZ; RALF KAEHLER AND RISA WECHSLER/STANFORD UNIVERSITY; SLOAN DIGITAL SKY SURVEY; MICHAEL BUSHA/ UNIVERSITY OF ZURICH

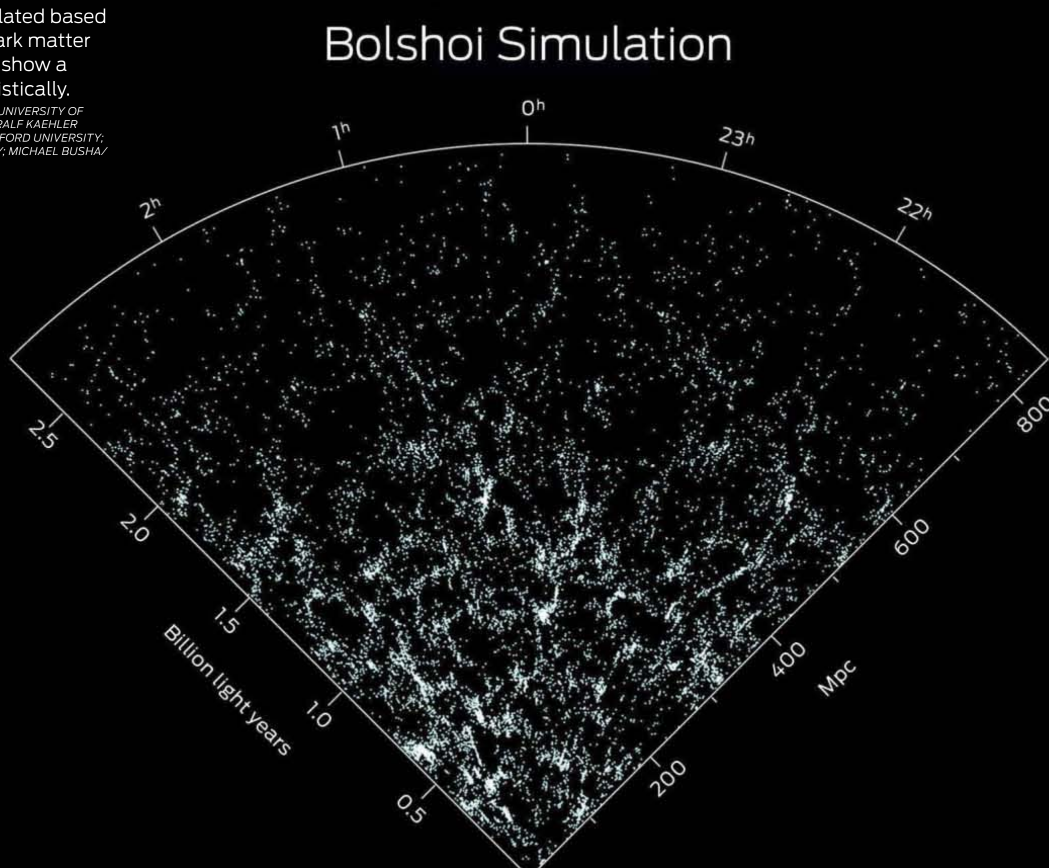

some forms of radioactivity. We know dark matter must exist because it helps bind rapidly moving stars to their host galaxies and rapidly moving galaxies to even larger galaxy clusters. The other component is "dark energy," which seems to be pushing the universe apart at an ever-increasing rate.

To identify these strange dark substances, cosmologists require more than just the evidence collected by telescopes. We need theoretical models of how the universe evolved and a way to test those models. Fortunately, thanks to progress in supercomputing, it's now possible to simulate the entire evolution of the universe numerically. The results of these computational experiments have already been transformative, and they're still only in their early days.

My colleagues and I recently completed one such simulation, which we dubbed Bolshoi, the Russian word for "great" or "grand." We started Bolshoi in a state that matched what the universe was like some 13.7 billion years ago, not long after the big bang, and simulated the evolution of dark matter and dark energy all the way up to the present day. We did that using 14 000 central processing units (CPUs) on the Pleiades machine at NASA's Ames Research Center, in Moffett Field, Calif., the space agency's largest and fastest supercomputer.

Bolshoi isn't the first large-scale simulation of the universe, but it's the first to rival the extraordinary precision of modern astrophysical observations. And the simulated universe it produces matches up surprisingly well with the real universe. We expect Bolshoi's computer-generated history of the cosmos to improve the understanding of how the Milky Way (the galaxy we live in) and other galaxies formed. If we're lucky, it might just reveal crucial clues to the nature of the mysterious dark entities that have steered the evolution of the universe and continue to guide its fate.

MOSTLY DARK: If you add up all the matter and energy in the universe, you'd find little that is familiar. The stars and gas that astronomers see in their telescopes make up just 0.5 percent of the cosmos. Just 0.01 percent of the universe is made of elements heavier than hydrogen or helium. Because of uncertainties, the numbers in this chart do not add up to 100 percent.

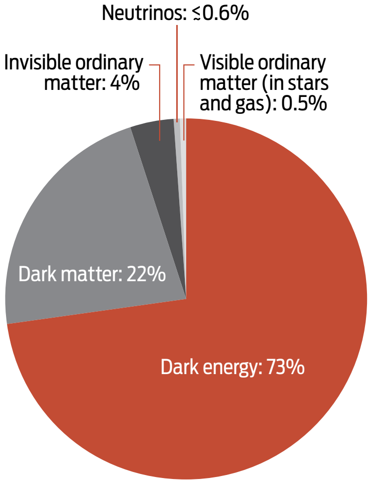

COSMOLOGY took a dramatic turn in 1998. That's when two teams, both studying light from past stellar explosions, showed that the universe isn't expanding as expected. Astronomers have been aware of the overall expansion of the universe for many decades, but most figured that this expansion must either be slowing or coasting along at a steady rate. So they were astonished to discover that the cosmos has, in fact, been expanding faster and faster for the past 5 billion years, pushed apart by a mysterious pressure that pervades space.

Although nobody knows what the dark energy that creates this pressure is, its discovery has actually been a boon for cosmology. It helped clear up a lot of old contradictions, including indications that some of the stars in our galaxy were older than the universe itself.

Nowadays, cosmologists have settled on a basic picture of a universe—one that's still full of mysteries but is at least quite self-consistent. The universe is about 13.7 billion years old and dominated by dark stuff: roughly 22 percent dark matter and 73 percent dark energy. Although dark energy strongly shapes the universe today, dark matter was more influential early on. It clumped up, producing a weblike scaffold that drew in ordinary matter and enabled the formation of all the galaxies as well as larger structures, including the galaxy groups and clusters we see today.

Telescopes on the ground and in space, like NASA's Wilkinson Microwave Anisotropy Probe (WMAP), launched in 2001, provided key evidence for developing this basic understanding of the universe. But cosmologists couldn't have arrived there without computer simulations to verify their interpretation of what they were seeing.

The most influential such simulation so far has been the Millennium Run, which was developed by a team led by Volker Springel, who is now at the University of Heidelberg. Completed in 2005, the Millennium simulation was the first to be big enough and detailed enough to model the evolution of an entire population of galaxies from their birth, 70 million or so years after the big bang, up to the present day. (Galaxies are the tracer particles in cosmological simulations. By comparing the simulated population with astronomical observations, cosmologists can see how well their models of the universe match reality.)

The Millennium simulation was a big step forward, and it's been the basis for more than 400 research papers. But that simulation had a few nagging shortcomings. For one thing, Millennium—as well as a later, smaller, but higher-resolution simulation called Millennium II—used very early WMAP results. Released in 2003, those data were from the first year of the telescope's operation, and they weren't very precise. As a result, the Millennium simulation made some predictions about galaxy populations that don't match very well with observations. What's more, although Millennium II had enough resolution to model the dark matter component of smaller galaxies (like the Milky Way's nearby companion, the Large Magellanic Cloud), it didn't simulate a large enough volume of space to enable precise predictions about the statistics of such satellite galaxies.

To help remedy this situation, my team at the University of California, Santa Cruz, partnered with Anatoly Klypin's group at New Mexico State University, in Las Cruces. Our aim was to improve on Millennium's size and resolution and to incorporate the very best knowledge of key cosmological parameters, derived in part from the latest WMAP results.

TO SIMULATE THE UNIVERSE inside a computer, you have to know where to start. Fortunately, cosmologists have a pretty good idea of what the universe's first moments were like. There's good reason to believe that for an outrageously brief period—lasting far less than 10^{-30} second, a thousandth of a trillionth of a femtosecond—the universe ballooned exponentially, taking what were then minute quantum variations in the density of matter and energy and inflating them tremendously in size. According to this theory of "cosmic inflation," tiny fluctuations in the distribution of dark matter eventually spawned all the galaxies.

It turns out that reconstructing the early growth phase of these fluctuations—up to about 30 million years after the big bang—demands nothing more than a laptop computer. That's

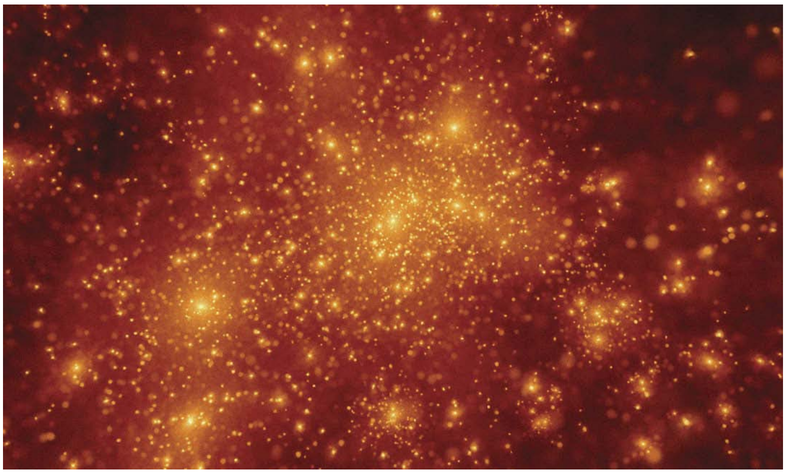

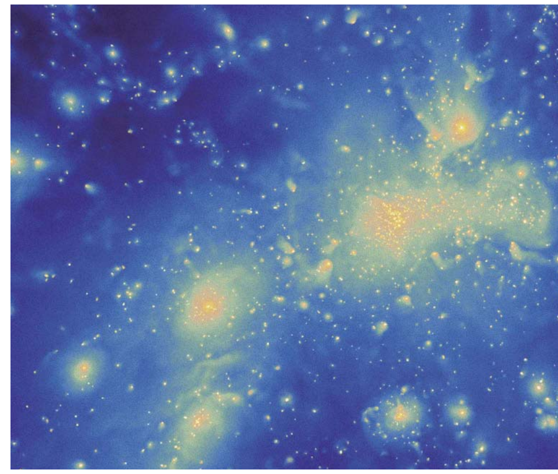

Visit the Bolshoi simulation website at http://hipacc.ucsc.edu/Bolshoi

because the early universe was extremely uniform, the differences in density from place to place amounting to no more than a few thousandths of a percent.

Over time, gravity magnified these subtle density differences. Dark matter particles were attracted to one another, and regions with slightly higher density expanded more slowly than average, while regions of lower density expanded more rapidly. Astrophysicists can model the growth of density fluctuations at these early times easily enough using simple linear equations to approximate the relevant gravitational effects.

The Bolshoi simulation kicks in before the gravitational interactions in this increasingly lumpy universe start to show nonlinear effects. My colleagues and I began by dividing up a cubical volume of the simulated universe into a uniform, three-dimensional grid, with 2048 lines running in each direction. We placed a simulated dark matter particle at each of the resulting 8.6 billion grid-line intersections.

In the real universe, there are likely to be far more dark matter particles than in a single cubic kilometer of space. But using a realistic number would overwhelm even the largest supercomputer. So we used this relatively small set, setting the mass of each simulated particle quite high so that it represents a huge amount of matter. We found that we could get pretty good results using these ultramassive placeholders, each set to about 100 million times the mass of the sun, about a hundredth of a percent of the total mass of dark matter in the Milky Way. After laying out one of these particles at each grid point, we shifted each starting location slightly to match our theoretical estimate of the primordial density fluctuations in the early universe.

Because our simulation was supposed to model boundary-less space and not a walled patch of universe, we followed the convention that's used in many video games, including that old classic *Asteroids*, where if a player's ship goes off the right edge, for example, it reemerges from the left side. In this way, we made our simulation represent a random, borderless chunk of the universe.

The last thing we did before setting the simulation in motion was to prime all the particles by assigning them initial velocities depending on the degree to which each particle was shifted away from a grid intersection. Then we flipped the switch and, like a divine watchmaker, sat back and watched what happened next.

Once the simulation began, every particle started to attract every other particle. With nearly 10 billion (10^{10}) of them, that would have resulted in roughly 10^{20} interactions that needed to be evaluated at each time step. Performing that many calculations would have slowed our simulation to a crawl, so we took some computational shortcuts. One was an algorithm called adaptive mesh refinement, which adjusts the resolution of the simulation depending on how closely particles are grouped together. If there are too many particles in a single cell, this approach divides that cell into eight smaller cells and calculates the interactions between particles with finer time and spatial resolution. This algorithm allows the simulation to dedicate most of its computational power where it's needed: in those small regions of the simulation where particles tend to cluster.

As Bolshoi ran, we logged the position and velocity of each of the 8.6 billion particles representing dark matter, producing 180 snapshots of the state of our simulated universe more or less evenly spaced in time. This small sampling still amounts to a lot of data—roughly 80 terabytes. All told, the Bolshoi simulation required 400 000 time steps and about 6 million CPU-hours to finish—the equivalent of about 18 days using 14 000 cores and 12 terabytes of RAM on the Pleiades supercomputer. But just as in observational astronomy, most of the hard work comes not in collecting mountains of data but in sorting through it all later.

TO MAKE SENSE of Bolshoi's output, the first thing my colleagues and I did was to analyze each snapshot in search of what in the real universe corresponds to galaxies. We did this by identifying distinct groupings of dark matter particles bound together by gravity. Astronomers call these bound clumps "halos" because they seem to be rather fuzzy clouds that extend in all directions beyond the visible boundaries of galaxies. In Bolshoi, these halos numbered up to 10 million at each time step. We characterized them in many different ways, measuring such properties as size, shape, orientation, mass, velocity, and rotation.

The Bolshoi simulation would be just an expensive fantasy if it didn't match observations of the real universe. The simplest way to make a comparison is to assign basic galactic properties to each of the Bolshoi dark matter halos. We know from observations that the more luminous a galaxy is, the faster its stars move. So we rank the halos by how rapidly the dark matter particles move in them, and we rank observed galaxies by their luminosity, and then we match the two distributions.

This process gives each simulated halo a likely galactic identity. We don't expect each galaxy to be a literal match, but we would hope that the overall statistics of the simulated galaxy population corresponds to what we see in the night sky.

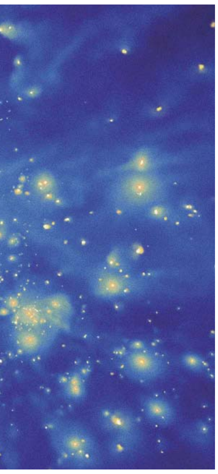

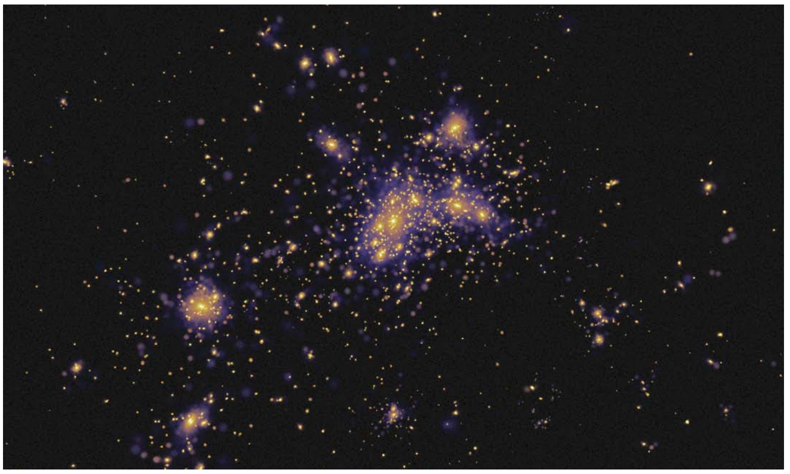

CLUSTER STRUCTURE: These three images show the distribution of matter in a cluster of galaxies from a small region of the BigBolshoi simulation. This region was resimulated, this time also taking into account the ordinary matter. The dark matter distribution [left] forms the gravitational scaffold for gas [center] and stars [right]. All the matter distributions were simulated over the age of the universe but are shown as they would appear today. The dimensions of each image are about 38 million by 50 million light-years.

SOURCES: RESIMULATION, GUSTAVO YEPES/UNIVERSIDAD AUTÓNOMA DE MADRID; VISUALIZATION, KRISTIN RIEBE/LEIBNIZ INSTITUTE FOR ASTROPHYSICS POTSDAM

One way to find out is to examine the probability of finding a galaxy within a given distance of another galaxy. To our delight, we found the Bolshoi simulation agrees quite well with astronomical observations in this regard. It did much better than the Millennium simulation, which ended up, for example, with about twice as many pairs of Milky Way–size galaxies separated by up to 3 million light-years as have actually been observed.

We performed other checks as well. For example, using the halo data from the Bolshoi simulation, Risa Wechsler's group at Stanford calculated that about 5 percent of the time, a Milky Way–size galaxy should host two satellite galaxies as bright as our galaxy's Large and Small Magellanic Clouds. We also calculated how often there would be just one such bright satellite galaxy and how often there would be none at all. When we compared those statistics with data collected by the ground-based Sloan Digital Sky Survey, we found remarkable agreement.

Yet another way to compare simulations with observations is to consider the number of galaxies with their stars and gas moving at various velocities. Here we found a good match with astronomical observations. But we also found a problem: The Bolshoi simulation ended up with many more galaxies containing relatively slow-moving stars than have actually been observed. It remains to be seen whether this is a serious issue. Perhaps it hints at an undiscovered property of dark matter. Or it could just represent astronomers' inability to detect the faintest galaxies. Deeper surveys are now under way that could help answer this question.

One of the most important products of our early analysis of the Bolshoi output—one that we expect will become an important resource for the astrophysical community—is the "halo merger tree." This tree describes the history of each halo, showing how each one formed from mergers with other halos all the way back to the beginning of the simulation.

We can use this merger tree to feed an improved model for the formation of galaxies, one that describes in a few equations basic properties like mass, luminosity, shape, stellar age, and elemental abundances. So far, my research group and others with early access to the Bolshoi results have found very good agreement between our models and the observed properties of nearby galaxies. We are also finding that these models can help make sense of Hubble Space Telescope images that show what galaxies looked like just a few billion years after the big bang.

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/simulatedcosmos1012

AS YOU MIGHT IMAGINE, no one simulation can do everything. Each must make a trade-off between resolution and the overall size of the region to be modeled. The Bolshoi simulation was of intermediate size. It considered a cubic volume of space about 1 billion light-years on edge, which is only about 0.00005 percent of the volume of the visible universe. But it still produced a good 10 million halos—an ample number to evaluate the general evolution of galaxies.

In addition, my colleagues and I have also run a larger, lower-resolution simulation, called BigBolshoi, which models a cube 4 billion light-years across, a volume 64 times as great as that shown by Bolshoi. The BigBolshoi simulation allowed us to probe the properties of galaxy clusters: clumps of galaxies that typically span about 10 million light-years and can have masses more than 1000 times as large as the total mass of the Milky Way. And we are running another simulation on NASA's Pleiades supercomputer called miniBolshoi, to focus in more detail on the statistics of galaxies like the Milky Way.

As supercomputers become more powerful, astrophysicists would like to run a large-scale, high-resolution simulation of the universe that includes the complex interplay within galaxies of gas, stars, and supermassive black holes. These sorts of simulations could be used to create a sort of field guide that could help astronomers interpret their telescope images, by telling them what galaxies should look like at various stages of their lives. With today's supercomputers and codes, though, it is possible to do such high-resolution simulations in only relatively small regions containing one or at most a few large galaxies.

One of the biggest hurdles going forward will be adapting to supercomputing's changing landscape. The speed of individual microprocessor cores hasn't increased significantly since 2004. Instead, today's computers pack more cores on each chip and often supplement them with accelerators like graphics processing units. Writing efficient programs for such computer architectures is an ongoing challenge, as is handling the increasing amounts of data from astronomical simulations and observations.

Despite those difficulties, I have every reason to think that numerical experiments like Bolshoi will only continue to get better. With any luck, the toy universes I and other astrophysicists create will help us make better sense of what we see in our most powerful telescopes—and help answer some of the grandest questions we can ask about the universe we call our home. ❑

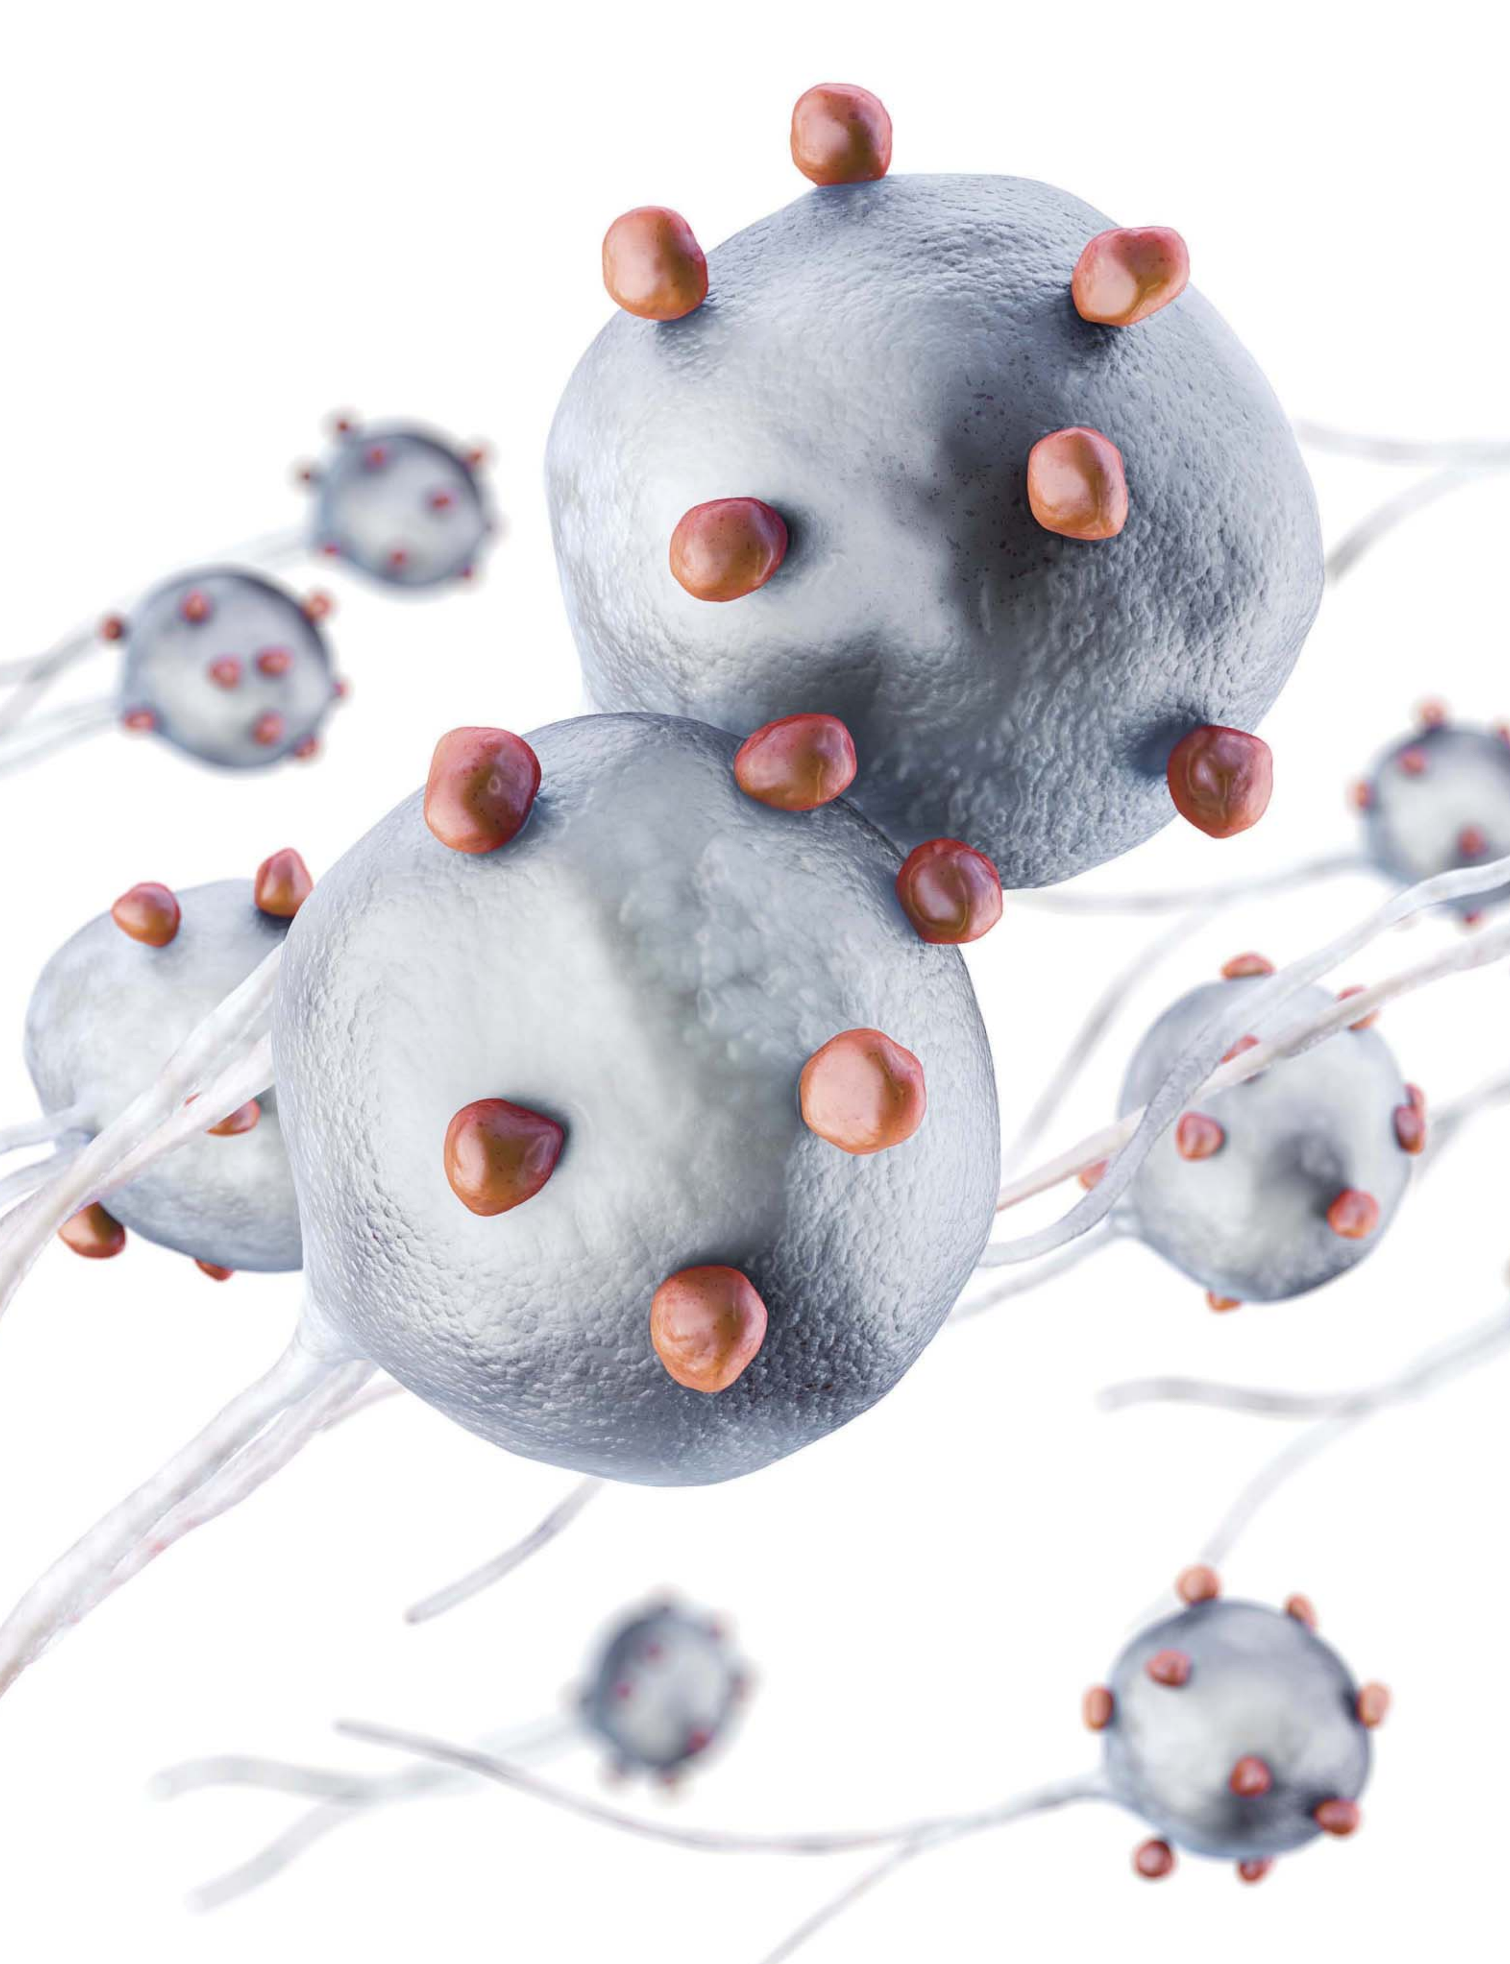

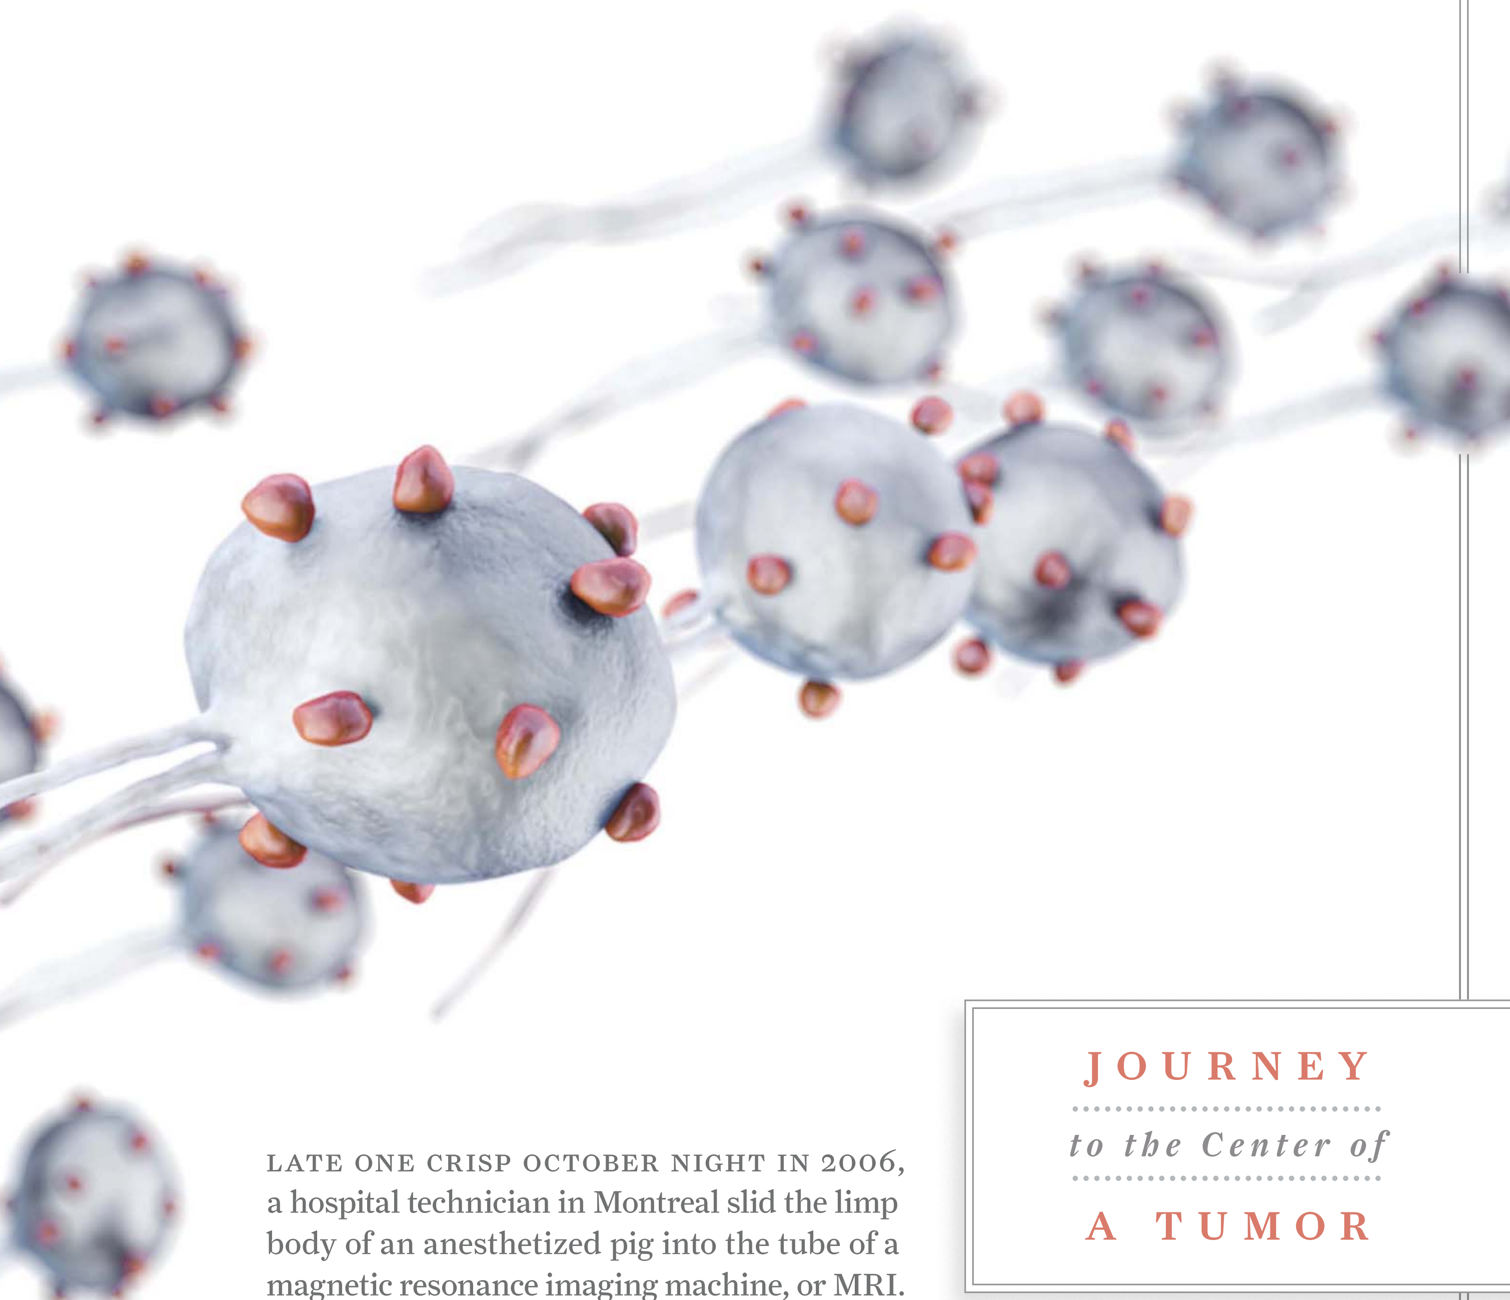

JOURNEY to the Center of A TUMOR

Magnets steer medical microbots through blood vessels

by SYLVAIN MARTEL

LATE ONE CRISP OCTOBER NIGHT IN 2006, a hospital technician in Montreal slid the limp body of an anesthetized pig into the tube of a magnetic resonance imaging machine, or MRI. A catheter extended from a large blood vessel below its neck—a carotid artery. Into the catheter, a surgeon injected a steel bead slightly larger than the ball of a ballpoint pen. ¶ In a room next door, my engineering graduate students and I held our breath. We were testing a program designed to manipulate the machine's magnetic forces, which would guide the bead like a remote-controlled submarine. Or so we hoped.

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

FANTASTIC VOYAGE

One of cancer therapy's holy grails is the delivery of drugs directly to tumors, thereby killing diseased cells while sparing healthy ones. A promising solution may be to deploy armies of carefully engineered microbot-bacteria hybrids in the body. The submarine-like bots could power through fast currents in large blood vessels, transporting their drug-packing bacteria cargo to the network of tiny vessels that lead inside a tumor.

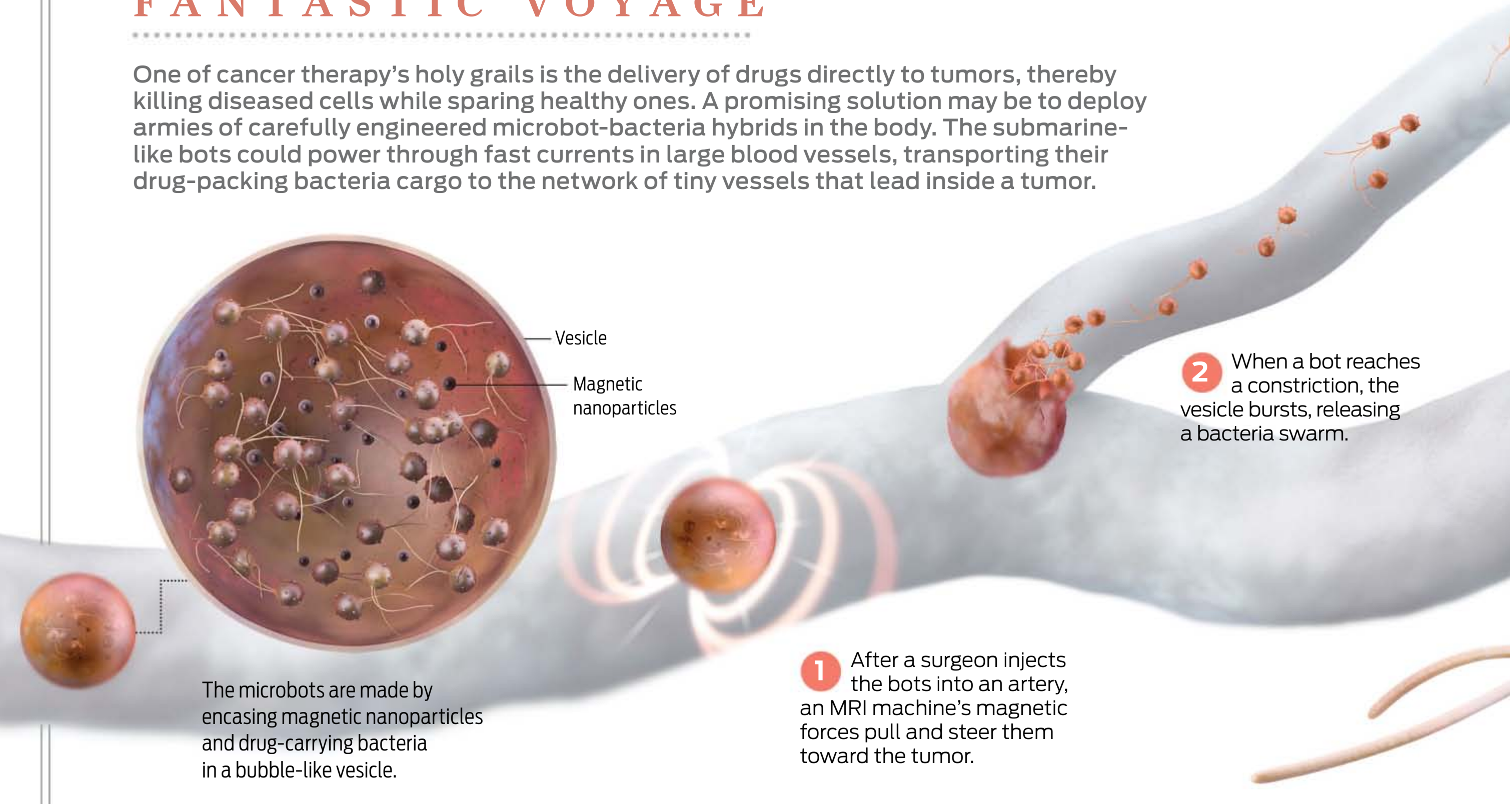

On a computer screen, the bead appeared as a square white tracking icon perched on the gray, wormlike image of the scanned artery. We stared at the square and waited. Nothing. Seconds ticked by, and still the bead refused to budge. Then suddenly the room erupted in cheers as we clapped our hands to our mouths and pointed at the screen. There, the bead was hopping up and down the artery, tagging every waypoint we had plotted.

That was the first time anyone had steered an object wirelessly through the blood vessel of a living creature. The experiment convinced us that we could engineer miniature machines to navigate the vast circulatory system of the human body. The microrobots would be able to travel deep inside the body, cruising our tiniest blood vessels to places that catheters can't go and performing tasks that would be impossible without invasive procedures.

It's easy to imagine many such tasks—delicate surgeries, diagnostic tests, and the installation of stents and other artificial structures. But the first real application will be treating cancers. Today's cancer drugs work by circulating throughout the body, killing healthy cells along with cancerous ones. Even antibody-equipped drugs designed to target cancer cells don't always hit their marks. Injecting drugs into a tumor is out of the question because the pressure would cause cells to spew from it like a volcano, spreading the disease elsewhere. So why not deploy robots to deliver the medicine?

It's a huge engineering challenge—one that will require input from many disciplines. The bots will have to contend with a complex maze of vessels churning with viscous blood. To make matters worse, today's manufacturing techniques can't build a motor and power source small enough to squeeze inside the capillaries—delicate vessels just a few micrometers thick—that feed a tumor. These vessels are too small for medical imaging tools to discern. Maneuvering through them would be like driving a car with your eyes closed.

Yet despite these challenges, my team at École Polytechnique de Montréal has made enormous progress on the problem. We have designed magnetic drug carriers as small as 50 micrometers that we can steer, like the metal bead inside the pig, through large arteries and arterioles using an MRI machine.

To reach tumors deeper in the body, we need robots too small to be powered by technology. And so we have turned to nature, harnessing swarms of swimming bacteria to serve as drug mules.

There are many problems to solve and medical trials to do. But I believe the world will start to see cancer-fighting robots on the market within the next decade. One day, microbots may even be able to hunt and kill cancer cells before they grow into tumors.

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/microbots1012

STORYTELLERS AND SCIENTISTS ALIKE have long dreamed of miniature robots that can roam the human body, sniffing out disease and repairing organs. You may recall the old science fiction flick *Fantastic Voyage* (1966), in which a submarine and its crew are shrunk to the size of a microbe and sent into the bloodstream of an eminent scientist to destroy a blood clot in his brain. Physics Nobel laureate Richard Feynman shared a similar fantasy. In a classic 1959 speech, he described a friend's "wild idea" that "it would be very interesting in surgery if you

PREVIOUS AND THIS SPREAD: ANATOMYBLUE

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

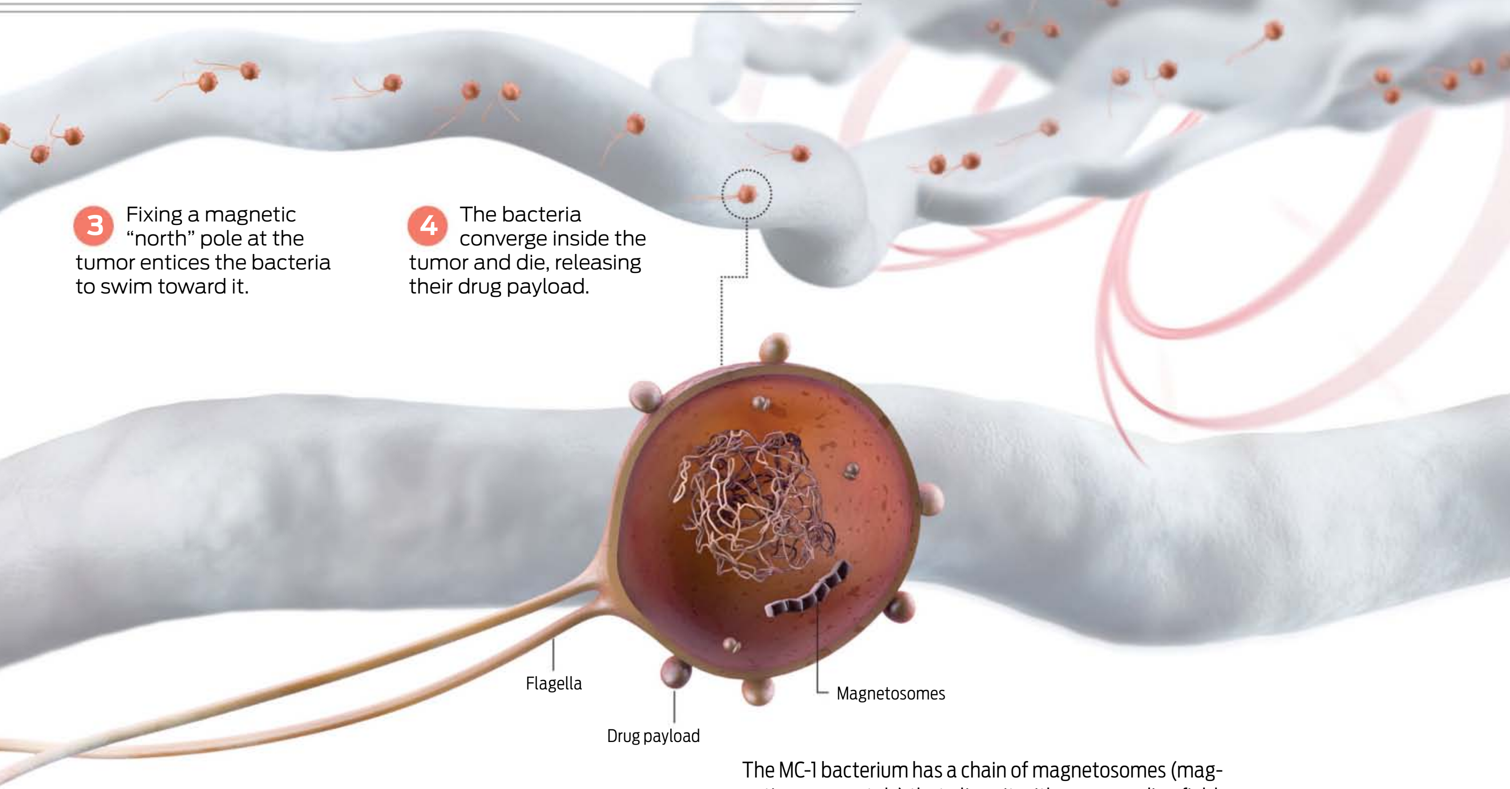

The MC-1 bacterium has a chain of magnetosomes (magnetic nanocrystals) that aligns it with a surrounding field. Anticancer drugs are cocooned in liposomes (fat-like sacs) and attached to the bacterium using antibodies.

could swallow the surgeon." As Feynman's friend imagined it, a little mechanical surgeon marches through the bloodstream to the heart and surveys its surroundings. Upon spotting a problematic obstruction or a faulty valve, it simply "takes a little knife and slices it out."

The problem, of course, is how to make the surgeon. Feynman envisioned an iterative manufacturing process in which ever-tinier hands make ever-tinier machines. He also considered the far-fetched dream of building things atom by atom, the way kids build toys out of Lego bricks. But he failed to offer any practical approaches. "I leave that to you," he told his audience.

Four decades later, the closest engineers had come to realizing Feynman's challenge was a diagnostic device called the PillCam. Introduced by the Israeli company Given Imaging in 2001, the disposable plastic capsule is about the size of a single piece of Good & Plenty candy and can be swallowed with water. It contains a light, a camera, batteries, an antenna, and a radio transmitter. As it passes through the gastrointestinal tract, it beams pictures wirelessly to a recording belt worn around a patient's waist. The images are then analyzed for signs of cancers, Crohn's disease, and other disorders. The procedure is less invasive than a traditional endoscopy and more sensitive than X-ray imaging.

But while the PillCam is great at inspecting the large, easily accessible cavities of the digestive system, it's much too big to travel elsewhere in the body. The device is also passive. A true medical microbot must propel and steer itself through an intricate network of fluid-filled tubes to tissues deep inside the body. And because it is so small, it must do so without the convenience of a battery-powered motor.

You might think that an obvious way to accomplish this task is to equip the robot with metallic particles and guide it with a magnet from outside the body. You wouldn't be entirely wrong. Indeed, some magnet-tipped surgical catheters work in this way. But a permanent magnet poses several problems for a smaller, tetherless robot. For instance, the magnet must be positioned quite close to the tiny particles to exert enough force to attract them. The robot could operate only near the skin and in one direction—toward the magnet. It would also be difficult to control. As it moved toward the magnet, the pull would increase, accelerating the robot and making it impossible to redirect if it went astray. And how would you know where it was going, anyway?

THE ANSWER IS THE MRI MACHINE. To understand how an MRI machine can be used for robotic navigation, you must know something about how it works. The bulk of the machine is a powerful, doughnut-shaped superconducting magnet that generates a magnetic field up to about 60 000 times as strong as Earth's. In medical imaging, the field's purpose is to align the spin of protons—the nuclei of hydrogen atoms—in the body. (A spinning proton acts as a sort of bar magnet; like a compass needle, it points in the direction of a surrounding field.)

Tucked inside the big magnet is the RF coil, which transmits radio-frequency waves. When the frequency of a wave pulse matches the spin rate of the protons, the spin "flips" direction by 90 degrees. When the pulse ends, the protons relax back to their original alignment, a lower energy state. The lost energy departs as a radio signal, which is picked up by a receiver. The density of signals gives information about the molecular

MINIBOTS *for* MEDICAL MISSIONS

Across the globe, researchers are engineering a wide array of free-roaming robots to carry out precise and delicate tasks inside the human body. Here are some of the most innovative designs.

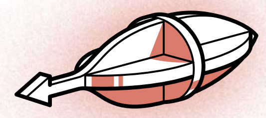

STEERABLE SURGEONS

LEAD ENGINEER: Brad Nelson, Swiss Federal Institute of Technology Zurich

BODY DESIGN: Flat nickel parts are assembled to make a 3-D tool.

POWER SOURCE: External electromagnetic coils

STEERING MECHANISM: Magnetic field gradients

TRACKING DEVICE: Microscope

SIZE (LENGTH): 1 millimeter

APPLICATIONS: Retinal surgeries, drug therapy for ocular disease

STATUS: Navigation has been tested in rabbit eyes.

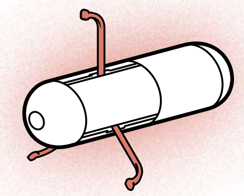

ROBOT PILLS

LEAD ENGINEER: Paolo Dario, Scuola Superiore Sant'Anna, Pisa, Italy

BODY DESIGN: A capsule contains a magnet, camera, wireless chip, and set of mechanical legs.

POWER SOURCE: DC motors, magnets outside the body

STEERING MECHANISM: Magnetic field gradients

TRACKING DEVICE: Camera, wireless telemetry system

SIZE (LENGTH): 2 centimeters

APPLICATIONS: Disease screening

STATUS: Visual inspection has been tested in a pig's intestines.

makeup of bodily tissues—distinguishing bone from blood, white matter from gray matter, tumors from healthy tissue.

Knowing where in the body the signals originate requires yet another set of coils. Sandwiched between the main magnet and the RF coil, the gradient coils generate a magnetic field that makes the main field stronger in some places and weaker in others. This variation changes the frequency of the protons' signals depending on their location in the field, which allows a computer to calculate their location in the body. By pulsing the gradient and RF coils on and off in various configurations, the machine produces a three-dimensional picture.

It's the gradient field that typically makes metal objects problematic in an MRI machine. Because the magnetic force the field exerts is uneven, it can slingshot an object through the body. That's why people with pacemakers or lodged bits of shrapnel don't get MRIs.

But there's a way to get around the projectile problem. Consider the pig experiment. First, as the metal bead slips through the catheter into the artery, only the main magnet is turned on. Its field strongly magnetizes the bead, but because the field is uniform, the bead doesn't move. In this "supermagnetized" state, a slight tug from the gradient coils will nudge the bead forward without increasing its magnetization. This prevents the bead from catapulting toward the source of the pull—as it would in the presence of a permanent magnet—and allows us to control its acceleration. It also lets us reduce the size of the bead and still move it anywhere in the body using relatively weak gradients. Then, by shifting the direction of the gradient field with precise algorithmic instructions, we can steer the bead through the artery.

The pig experiment proved the concept, but a naked 1.5-millimeter metal bead isn't much use for transporting drugs inside the microvasculature of the human body. To make a real medical microbot, we started with iron cobalt nanoparticles, which are more sensitive to magnetic forces than other nontoxic metal alloys. We coated the nanoparticles in graphite to keep them from oxidizing, and then we encased them, along with molecules of the cancer drug doxorubicin, in a biodegradable polymer sac.

We can scale these robots down to a few hundred micrometers; our gradient coils aren't powerful enough to get a grip on anything smaller than that. With beefier, custom-built coils, we reduced the size to about 50 μm. But that's the limit: Propelling even tinier robots would require gradients so strong that the huge jolt of current needed to quickly fire up the coils would begin to disturb cells in the body's central nervous system.

In 2011, we tested the ability of these microcarriers to deliver drugs in rabbits. We injected a small squad into the hepatic artery, which supplies blood to the liver and other nearby organs. As in the pig experiment, we had already plotted a simple navigation course using an MRI scan of the rabbit's interior. Our computer program then steered the robots toward their target by automatically adjusting the gradient field. We waited a few seconds, then took another scan. A distinguishable black smudge in one lobe of the liver proved that the tiny soldiers had shot through the artery, made a sharp turn at a bifurcation, and then congregated at the last waypoint. As their polymer shells degraded, they released their medicinal weapons, although in this experiment, there was no tumor to kill.

Microcarriers may be ideal for targeting liver cancers in humans because the vessel branches that lead to the liver's lobes are quite large—about 150 μm or wider. But to reach tumors hidden behind networks of smaller capillaries, such as in the

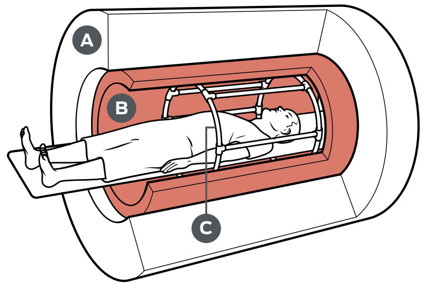

INSIDE AN MRI: To capture the body's interior, a scanner uses three sets of electromagnetic coils: a multiton superconducting magnet [a], magnetic gradient coils [b], and a radio-frequency transceiver [c].

ALL ILLUSTRATIONS THIS PAGE: JAMES PROVOST

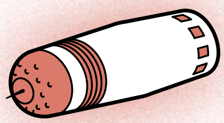

PLAQUE BUSTERS

LEAD ENGINEER: Jong-Oh Park, Chonnam National University, South Korea
BODY DESIGN: A magnetic capsule is equipped with a micro drill head.
POWER SOURCE: Electromagnetic coils
STEERING MECHANISM: Magnetic field gradients
TRACKING DEVICE: MRI and X-ray imaging
SIZE (LENGTH): 10 millimeters
APPLICATION: Removal of plaque in the arteries
STATUS: Prototype is being developed.

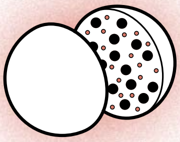

MAGNETIC MICROCARRIERS

LEAD ENGINEER: Sylvain Martel, École Polytechnique de Montréal
BODY DESIGN: Magnetic nanoparticles and anticancer drugs are encased in a biodegradable polymer sac.
POWER SOURCE: MRI machine
STEERING MECHANISM: Magnetic field gradients
TRACKING DEVICE: MRI imaging
SIZE (DIAMETER): 50 micrometers
APPLICATION: Anticancer drug delivery
STATUS: Navigation and drug release have been tested in rabbits.

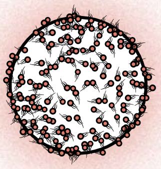

BACTERIABOTS

LEAD ENGINEER: Metin Sitti, Carnegie Mellon University, Pittsburgh
BODY DESIGN: Swarms of bacteria self-attach to polymer beads.
POWER SOURCE: Flagella of living bacteria
STEERING MECHANISM: Chemical, gas, temperature, or acidity gradients
TRACKING DEVICE: Fluorescent or X-ray imaging
SIZE (DIAMETER): 55 micrometers
APPLICATION: Disease diagnosis and drug delivery
STATUS: Chemical navigation has been tested in the lab.

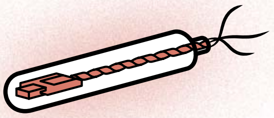

CORKSCREW SWIMMERS

LEAD ENGINEER: James Friend, Royal Melbourne Institute of Technology, Australia
BODY DESIGN: A capsule contains a piezoelectric rotary motor that whips an attached tail.
POWER SOURCE: Wireless power transmission
STEERING MECHANISM: Piezoelectric vibrations
TRACKING DEVICE: X-ray imaging
SIZE (DIAMETER): 250 micrometers
APPLICATION: Vessel navigation
STATUS: Motor has been tested using applied electric currents in the lab.

breast or the colon, we need microbots no thicker than a couple of micrometers. And because the gradient coils of an MRI machine can't be made powerful enough to propel them, they must propel themselves.

WE DON'T HAVE TO INVENT these dream machines—they already exist in nature.

The MC-1 bacterium was first discovered in 1993 in a small, salty estuary of the Pettaquamscutt River in Rhode Island. The strain, which is now cultivated in laboratories, is the model microbot. It has a spherical body about 2 µm across, which provides ample surface area for attaching molecules of anticancer drugs. It also has a pair of spinning, whiplike tails powered by molecular "rotary motors" that rocket it through water at speeds up to 150 times its body length per second. (If Olympian Michael Phelps could do that, he would swim faster than the cruising speed of a Boeing 747.)

Most important, MC-1 is magnetic. The species belongs to a group of bacteria described as magnetotactic because they have the special skill of aligning themselves with Earth's magnetic field. Though it's unclear what advantage this gives them, they perform the trick with a chain of iron oxide nanocrystals called magnetosomes that behave like a compass needle.

To guide the bacteria, we need generate only a very weak magnetic field—just slightly stronger than Earth's—pointed toward a target such as a tumor. The goal here is to direct the compass needle (the magnetosomes), not to propel the entire compass (the organism). We do this with specially designed electromagnetic coils that, in addition to being less powerful than an MRI machine, allow us to orient the field by varying the amount of current passing through each coil. The bacteria then swim toward the artificial "north" pole, navigating the twists and turns of the vascular network as if they were reeds in a river. In the hot blood of a mammalian body, the microbes survive only about 40 minutes.

We have demonstrated the potential of this system in mice. After injecting a swarm of bacteria in each animal's tail, we fixed an artificial pole on one of two tumors in the body. When we dissected and examined the tumors, we found more bacteria in the targeted tumor than in the other. But the results weren't ideal. In fact, most of the swarm never made it to a tumor at all. Many cells got swept away by heavy currents in the large vessels of the tail before they could reach the calmer capillaries.

I offer two solutions to this problem. First, swarms of MC-1 bacteria could be used specifically to treat colorectal cancers. Because the colon is easily accessed through the rectum with a catheter, the bacteria can be injected close to the tumor and so will be more likely to find it. We are now planning trials in monkeys to test whether the procedure will be safe for humans.

The second solution we are pursuing is a microcarrier-bacteria hybrid [see illustration, "Fantastic Voyage"]. We make these robots by loading fat-like bubbles, such as vesicles or micelles, with magnetic nanoparticles and drug-packing bacteria. When deployed in the body, they could be navigated using an MRI machine through large blood vessels, similarly to microcarriers, until narrower vessels block their path. At this point, the bubbles—which could be engineered to burst upon smacking into a vessel wall—would dispatch the bacteria. Then the patient would be transferred from the MRI machine into our custom coils. Guided by the coils' magnetic field, the armed microbes would swim to the tumor. And there they would die, releasing their drug cargo.

IT WILL BE SEVERAL YEARS before medical microbots of any form will be used in patients. There are still many unanswered questions. How many robots are needed to deliver enough drugs to kill a tumor? What amount of bacteria is safe to inject in humans? Will the body raise antibodies against them?

Luckily, we aren't alone in tackling these challenges. The field is growing; roboticists around the world are inventing a whole family of tiny devices for various tasks [see sidebar, "Minibots for Medical Missions"]. There is no shortage of possibilities. In my laboratory, we have already started to investigate ways of exploiting the energy properties of magnetic microbots to temporarily open up the blood-brain barrier and access tumors in the brain. Almost 50 years after Hollywood dreamed up the tale, a real-life fantastic voyage is just beginning. ❑

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


Recommended
SHOP
JUST FOR YOU
MOST
May we suggest
ENJOY
BUY
LIKE
USER

HOW ONLINE MERCHANTS PREDICT YOUR PREFERENCES AND PROD YOU TO PURCHASE BY JOSEPH A. KONSTAN & JOHN RIEDL

Recommended for You

One morning in April, we each directed our browsers to Amazon.com's website. Not only did the site greet us by name, the home page opened with a host of suggested purchases. It directed Joe to Barry Greenstein's *Ace on the River: An Advanced Poker Guide*, Jonah Lehrer's *Imagine: How Creativity Works*, and Michael Lewis's *Boomerang: Travels in the New Third World*. For John it selected Dave Barry's *Only Travel Guide You'll Ever Need*, the spy novel *Mission to Paris*, by Alan Furst, and the banking exposé *The Big Short: Inside the Doomsday Machine*, also by Michael Lewis.

By now, online shoppers are accustomed to getting these personalized suggestions. Netflix suggests videos to watch. TiVo records programs on its own, just in case we're interested. And Pandora builds personalized music streams by predicting what we'll want to listen to.

All of these suggestions come from recommender systems. Driven by computer algorithms, recommenders help consumers by selecting products they will probably like and might buy based on their browsing, searches, purchases, and preferences. Designed to help retailers boost sales, recommenders are a huge and growing business. Meanwhile, the field of recommender system development has grown from a couple of dozen researchers in the mid-1990s to hundreds of researchers today—working for universities, the large online retailers, and dozens of other companies whose sole focus is on these types of systems.

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

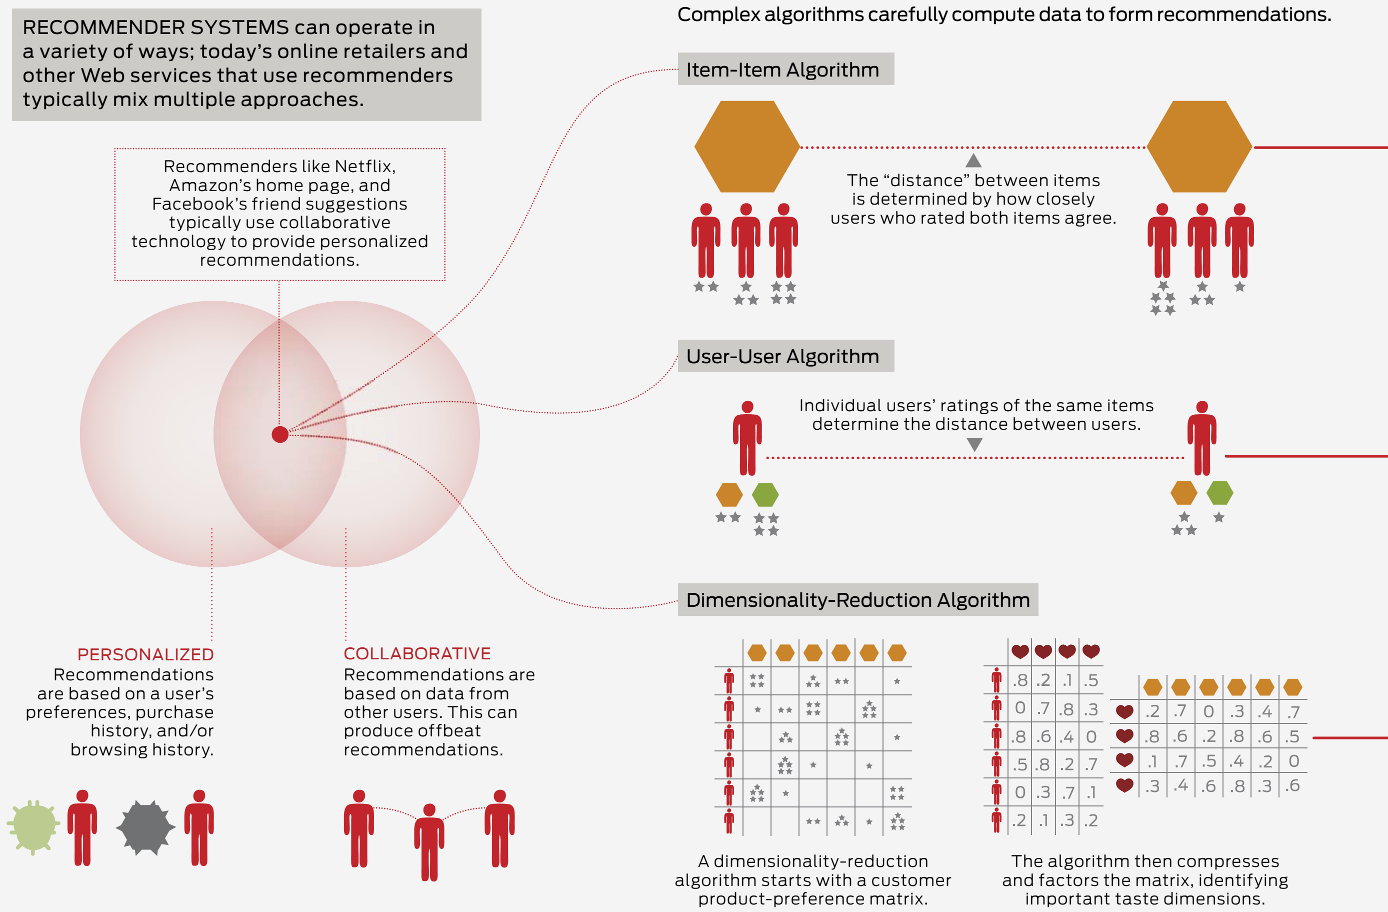

Over the years, recommenders have evolved considerably. They started as relatively crude and often inaccurate predictors of behavior. But the systems improved quickly as more and different types of data about website users became available and they were able to apply innovative algorithms to that data. Today, recommenders are extremely sophisticated and specialized systems that often seem to know you better than you know yourself. And they're expanding beyond retail sites. Universities use them to steer students to courses. Cellphone companies rely on them to predict which users are in danger of switching to another provider. And conference organizers have tested them for assigning papers to peer reviewers.

The two of us have been building and studying recommender systems since their early days, initially as academic researchers working on the GroupLens Project. Begun in 1992, GroupLens sorted through messages in Usenet discussion forums and pointed users to threads they might be interested in but had not yet discovered on their own. Several years later, we founded Net Perceptions, the leading recommender company during the first Internet boom. Our experience, therefore, gives us a lot of insight into what's going on behind the scenes at Amazon and other online retailers, even though those companies seldom speak publicly about exactly how their recommendations work. (In this article, our analysis is based on educated observation and deduction, not on any inside information.) Here's what we know.

HAVE YOU EVER WONDERED what you look like to Amazon? Here is the cold, hard truth: You are a very long row of numbers in a very, very large table. This row describes everything you've looked at, everything you've clicked on, and everything you've purchased on the site; the rest of the table represents the millions of other Amazon shoppers. Your row changes every time you enter the site, and it changes again with every action you take while you're there. That information in turn affects what you see on each page you visit and what e-mail and special offers you receive from the company.

Over the years, the developers of recommender systems have tried a variety of approaches to gather and parse all that data. These days, they've mostly settled on what is called the personalized collaborative recommender. That type of recommender is at the heart of Amazon, Netflix, Facebook's friend suggestions, and Last.fm, a popular music website based in the United Kingdom. They're "personalized" because they track each user's behavior—pages viewed, purchases, and ratings—to come up with recommendations; they aren't bringing up canned sets of suggestions. And they're "collaborative" because they treat two items as being related based on the fact that lots of other customers have purchased or stated a preference for those items, rather than by analyzing sets of product features or keywords.

Personalized collaborative recommenders, in some form or another, have been around since at least 1992. In addition to the GroupLens project, another early recommender was MIT's Ringo, which took lists of albums from users and suggested other music they might like.

GroupLens and Ringo both used a simple collaborative algorithm known as a "user-user" algorithm. This type of algorithm computes the "distance" between pairs of users based on how much they agree on items they have both rated. For instance,

FIRST PAGE (PHOTO-ILLUSTRATION) AND ABOVE: CHRISTINA BEARD

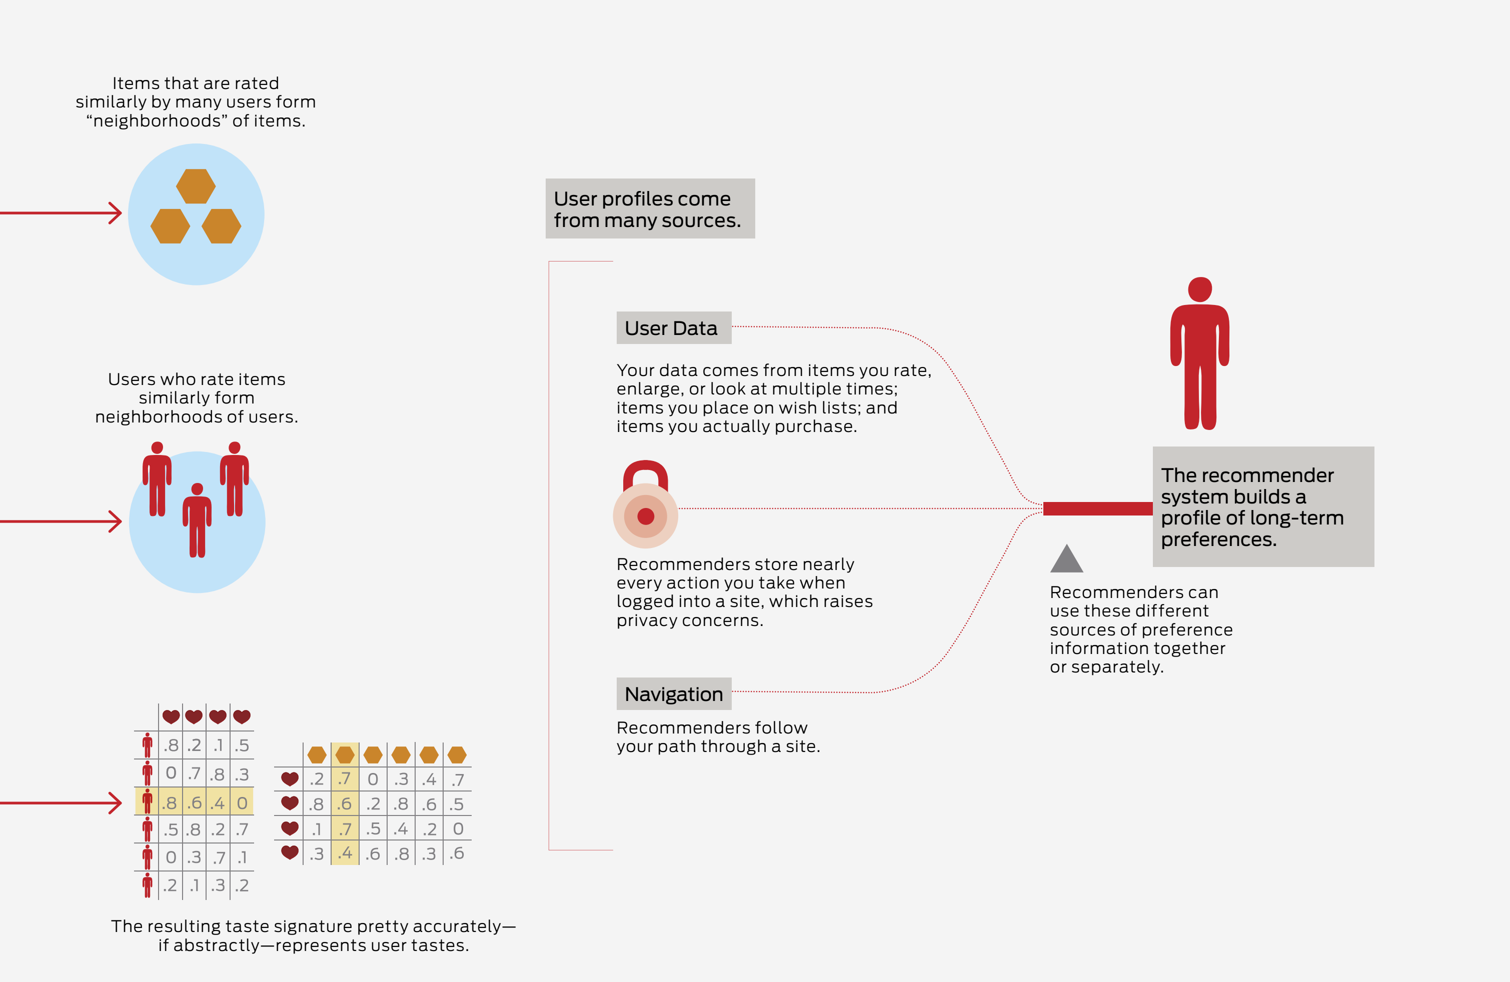

if Jim and Jane each give the movie *Tron* five stars, their distance is zero. If Jim then gives *Tron: Legacy* five stars, while Jane rates it three stars, their distance increases. Users whose tastes are relatively "near" each other according to these calculations are said to share a "neighborhood."

But the user-user approach doesn't work that well. For one thing, it's not always easy to form neighborhoods that make sense: Many pairs of users have only a few ratings in common or none at all, and in the case of movies, these few ratings in common tend to be of blockbusters that nearly everyone likes. Also, because the distance between users can change rapidly, user-user algorithms have to do most of their calculations on the spot, and that can take more time than someone clicking around a website is going to hang around.

So most recommenders today rely on an "item-item" algorithm, which calculates the distance between each pair of books or movies or what have you according to how closely users who have rated them agree. People who like books by Tom Clancy are likely to rate books by Clive Cussler highly, so books by Clancy and Cussler are in the same neighborhood. Distances between pairs of items, which may be based on the ratings of thousands or millions of users, tend to be relatively stable over time, so recommenders can precompute distances and generate recommendations more quickly. Both Amazon and Netflix have said publicly that they use variants of an item-item algorithm, though they keep the details secret.

One problem with both user-user and item-item algorithms is the inconsistency of ratings. Users often do not rate the same item the same way if offered the chance to rate it again. Tastes change, moods change, memories fade. MIT conducted one study in the late 1990s that showed an average change of one point on a seven-point scale a year after a user's original rating. Researchers are trying different ways to incorporate such variables into their models; for example, some recommenders will ask users to rerate items when their original ratings seem out of sync with everything else the recommender knows about them.

But the user-user and item-item algorithms have a bigger problem than consistency: They're too rigid. That is, they can spot people who prefer the same item but then miss potential pairs who prefer very similar items. Let's say you're a fan of Monet's water lily paintings. Of the 250 or so paintings of water lilies that the French impressionist did, which is your favorite? Among a group of Monet fans, each person may like a different water lily painting best, but the basic algorithms might not recognize their shared taste for Monet.

About a decade ago, researchers figured out a way to factor in such sets of similar items—a process called dimensionality reduction. This method is much more computationally intensive than the user-user and item-item algorithms, so its adoption has been slower. But as computers have gotten faster and cheaper, it has been gaining ground.

To understand how dimensionality reduction works, let's consider your taste in food and how it compares with that of a million other people. You can represent those tastes in a huge matrix, where each person's taste makes up its own row and each of the thousands of columns is a different food. Your row might show that you gave grilled filet mignon five stars, braised short ribs four and a half stars, fried chicken wings two stars, cold tofu

RECOMMENDATION ODYSSEY: An Amazon user interested in *2001: A Space Odyssey* sees suggestions from three different collaborative recommenders.

rolls one star, roasted portobello mushroom five stars, steamed edamame with sea salt four stars, and so forth.

A recommender using the matrix wouldn't really care about your particular rating of a particular food, however. Instead, it wants to understand your preferences in general terms, so that it can apply this knowledge to a wide variety of foods. For instance, given the above, the recommender might conclude that you like beef, salty things, and grilled dishes, dislike chicken and anything fried, are neutral on vegetables, and so on. The number of such taste attributes or dimensions would be much smaller than the number of possible foods—there might be 50 or 100 dimensions in all. And by looking at those dimensions, a recommender could quickly determine whether you'd like a new food—say, salt-crusted prime rib—by comparing its dimensions (salty, beef, not chicken, not fried, not vegetable, not grilled) against your profile. This more general representation allows the recommender to spot users who prefer similar yet distinct items. And it substantially compresses the matrix, making the recommender more efficient.

It's a pretty cool solution. But how do you find those taste dimensions? Not by asking a chef. Instead, these systems use a mathematical technique called singular value decomposition to compute the dimensions. The technique involves factoring the original giant matrix into two "taste matrices"—one that includes all the users and the 100 taste dimensions and another that includes all the foods and the 100 taste dimensions—plus a third matrix that, when multiplied by either of the other two, re-creates the original matrix.

Unlike the food example above, the dimensions that get computed are neither describable nor intuitive; they are pure abstract values, and try as you might, you'll never identify one that represents, say, "salty." And that's okay, as long as those values ultimately yield accurate recommendations. The main drawback to this approach is that the time it takes to factor the matrix grows quickly with the number of customers and products—a matrix of 250 million customers and 10 million products would take 1 billion times as long to factor as a matrix of 250 000 customers and 10 000 products. And the process needs to be repeated frequently. The matrix starts to grow stale as soon as new ratings are received; at a company like Amazon, that happens every second. Fortunately, even a slightly stale matrix works reasonably well. And researchers have been devising new algorithms that provide good approximations to singular value decomposition with substantially faster calculation times.

ATTENTION GLOBAL TECHNOLOGY EMPLOYERS!

IEEE Media Recruitment Services will be exhibiting with IEEE Spectrum magazine at:

electronica 2012 – the 25th International Trade Fair for Electronic Components, Systems, and Applications 13-16 November, 2012 in Munich, Germany.

Global employer representatives who visit us at the show will have the opportunity to connect with the talented IEEE job pool through a special show discount on recruitment advertising and job posts on the IEEE Job Site!

Here's how it works:

1. Visit the IEEE Spectrum booth at electronica 2012 - #A6-144
2. Speak with one of our representatives, and leave us your Human Resources contact information
3. We will forward instructions on how to receive a 20% discount on print recruitment advertising in IEEE Spectrum and one complimentary featured job post on the IEEE Job Site to the appropriate contact
4. The offer is good through 31 January, 2013

Put the "PRO" in your technology employment prospects with

IEEE Media Recruitment Services!

http://advertise.ieee.org/recruitment

ieee spectrum INSIDE TECHNOLOGY

IEEE MEDIA

BY NOW, YOU HAVE A BASIC IDEA of how an online retailer sizes you up and tries to match your tastes to those of others whenever you shop at its site. Recommenders have two other features that dramatically affect the recommendations you see: First, beyond figuring out how similar you are to other shoppers, the recommender has to figure out what you actually like. Second, the system operates according to a set of business rules that help ensure its recommendations are both helpful to you and profitable for the retailer.

For example, consider the recommender used for Amazon's online art store, which at last count had more than 9 million prints and posters for sale. Amazon's art store assesses your preferences in a few ways. It asks you to rate particular artworks on a five-star scale, and it also notes which paintings you enlarge, which you look at multiple times, which you place on a wish list, and which you actually buy. It also tracks which paintings are on your screen at the time as well as others you look at during your session. The retailer uses the path you've traveled through its website—the pages you've viewed and items you've clicked on—to suggest complementary works, and it combines your purchase data with your ratings to build a profile of your long-term preferences.

Companies like Amazon collect an immense amount of data like this about their customers. Nearly any action taken while you are logged in is stored for future use. Thanks to browser cookies, companies can even maintain records on anonymous shoppers, eventually linking the data to a customer profile when the anonymous shopper creates an account or signs in. This explosion of data collection is not unique to online vendors—Walmart is famous for its extensive mining of cash register receipt data. But an online shop is much better positioned to view and record not just your purchases but what items you considered, looked at, and rejected. Throughout much of the world, all of this activity is fair game; only in Europe do data privacy laws restrict such practices to a degree.

Of course, regardless of the law, any customer will react badly if his or her data is used inappropriately. Amazon learned this lesson the hard way back in September 2000, when certain customers discovered they were being quoted higher prices because the website had identified them as regular customers, rather than as shoppers who had entered anonymously or from a comparison-shopping site. Amazon claimed this was just a random price test and the observed relationship to being a regular customer was coincidental, but it nevertheless stopped the practice.

The business rules around these systems are designed to prevent recommenders from making foolish suggestions and also to help online retailers maximize sales without losing your trust. At their most basic level, these systems avoid what's known as the supermarket paradox. For example, nearly everyone who walks into a supermarket likes bananas and will often buy some. So shouldn't the recommender simply recommend bananas to every customer? The answer is no, because it wouldn't help the customer, and it wouldn't increase banana sales. So a smart supermarket recommender will always include a rule to explicitly exclude recommending bananas.

That example may sound simplistic, but in one of our early experiences, our system kept recommending the Beatles' "White Album" to nearly every visitor. Statistically this was a great recommendation: The customers had never purchased

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

this item from the e-commerce site, and most customers rated it highly. And yet the recommendation was useless. Everyone who was interested in the "White Album" already owned a copy.

Most recommender rules are more subtle, of course. When John recently searched for an action movie on Netflix, for instance, he wasn't offered *The Avengers*, because the blockbuster was not yet available for rental, and so the suggestion wouldn't have profited Netflix. Instead it steered him to *Iron Man 2*, which was available for streaming.

Other business rules prevent recommenders from suggesting loss leaders—products that sell below cost to draw people into the site—or conversely encourage them to recommend products that are overstocked. During our time at Net Perceptions, we worked with a client who did just that: He used his recommender system to identify—with considerable success—potential customers for his overstocked goods.

This kind of thing quickly gets tricky, however. A system that simply pushes high-margin products isn't going to earn the customers' trust. It's like going to a restaurant where the waiter steers you toward a particular fish dish. Is it really his favorite? Or did the chef urge the staff to push out the fish before its sell-by date?

To build trust, the more sophisticated recommender systems strive for some degree of transparency by giving customers an idea of why a particular item was recommended and letting them correct their profiles if they don't like the recommendations they're getting.

You can, for instance, delete information from your Amazon profile about things you purchased as gifts; after all, those don't reflect your tastes. You can also find out why certain products have been offered through the recommender. After Amazon selected Jonathan Franzen's novel *Freedom* for John, he clicked on the link labeled "Explain." He then got a brief explanation that certain books on John's Amazon wish list had triggered the recommendation. But as John hadn't read any of the wish list books, he discounted the *Freedom* suggestion. Explanations like these let users know how reliable a given recommendation is.

But profile adjustments and explanations often aren't enough to keep a system on track. Recently Amazon bombarded Joe with e-mails for large-screen HDTVs—as many as three a week for months. Besides sending him more e-mail on the topic than he could possibly want, the retailer didn't recognize that he'd already purchased a TV through his wife's account. What's more, the e-mails did not offer an obvious way for Joe to say, "Thanks, but I'm not interested." Eventually, Joe unsubscribed from certain Amazon e-mails; he doesn't miss the messages, and he has more time to actually watch that TV.

SO HOW WELL DO RECOMMENDERS ultimately work? They certainly are increasing online sales; analyst Jack Aaronson of the Aaronson Group estimates that investments in recommenders bring in returns of 10 to 30 percent, thanks to the increased sales they drive. And they still have a long way to go.

Right now the biggest challenge for those of us who study recommender systems is to figure out how best to judge the new approaches and algorithms. It's not as simple as benchmarking a microprocessor, because different recommenders have very different goals.

The easiest way to evaluate an algorithm is to look at the difference between its predictions and the actual ratings users give. For instance, if John gives the teen-romance novel *Twilight* one star, Amazon might note that it had predicted he would give it two stars, based on the ratings of other similar users, and so its recommender was off by a star. But sellers care much more about errors on highly rated items than errors on low-rated items, because the highly rated items are the ones users are more likely to buy; John is never going to purchase *Twilight*, so scoring this rating contributes little to understanding how well the recommender works.

Another common measure is the extent to which recommendations match actual purchases. This analysis can also be misleading, however, because it erroneously rewards the recommender for items users managed to find on their own—precisely the items they don't need recommendations for!

Given the shortcomings of these approaches, researchers have been working on new metrics that look not just at accuracy but also at other attributes, such as serendipity and diversity.

POST YOUR COMMENTS *online at* http://spectrum.ieee.org/recommenders1012

Serendipity rewards unusual recommendations, particularly those that are valuable to one user but not as valuable to other similar users. An algorithm tuned to serendipity would note that the "White Album" appears to be a good recommendation for nearly everyone and would therefore look for a recommendation that's less common—perhaps Joan Armatrading's *Love and Affection*. This less-popular recommendation wouldn't be as likely to hit its target, but when it did, it would be a much happier surprise to the user.

Looking at the diversity of a recommender's suggestions is also revealing. For instance, a user who loves Dick Francis mysteries might nevertheless be disappointed to get a list of recommendations all written by Dick Francis. A truly diverse list of recommendations could include books by different authors and in different genres, as well as movies, games, and other products.

RECOMMENDER SYSTEMS RESEARCH has all sorts of new ground to break, far beyond fine-tuning existing systems. Researchers today are considering to what extent a recommender should help users explore parts of a site's collection they haven't looked into—say, sending book buyers over to Amazon's clothing department rather than recommending safe items they may be more comfortable with. Going beyond the retail world, recommenders could help expose people to new ideas; even if we disagree with some of them, the overall effect might be positive in that it would help reduce the balkanization of society. Whether recommenders can do that without annoying us or making us distrustful remains to be seen.

But one thing is clear: Recommender systems are only going to get better, collect more data about you, and show up in new and surprising places. And as for you, if you liked this article, Amazon will be happy to recommend entire books on recommender systems that you might also like. ❏

Sun Yat-sen University / Carnegie Mellon University
Joint Institute of Engineering

中国

Sun Yat-sen University (SYSU) & Carnegie Mellon University (CMU) are partnering to establish the Joint Institute of Engineering (JI). SYSU and CMU are embarking on an exciting opportunity to transform engineering education in China. JI will provide world-class education and cutting-edge research in China's Pearl-River Delta region, which provides rapidly growing opportunities for future technology innovation.

JI is seeking full-time faculty, who have an interest in pursuing innovative, interdisciplinary education programs and in leading research efforts in all areas of Electrical/Computer Engineering . Candidates should possess a Ph.D. in ECE or related disciplines, with a demonstrated record and potential for research, teaching and leadership. The position includes an initial year on the Pittsburgh campus of CMU to establish educational and research collaborations before locating to Guangzhou, China.

This is a worldwide search open to qualified candidates from all countries and of all nationalities, with an internationally competitive compensation package for all qualified candidates.

Faculty Positions Available *in Electrical/Computer Engineering*

Visit **http://sysucmuji.sysu.edu.cn** *or* **http://sysucmuji.cmu.edu** *for details.*

THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

FACULTY POSITION IN **Electronic and Computer Engineering**

The Department of Electronic and Computer Engineering at the Hong Kong University of Science and Technology invites applications for a faculty position at the rank of Assistant Professor. Applicants should have a PhD with demonstrated strength in research and a commitment to teaching. Successful candidates are expected to pursue an active research program, to teach both graduate and undergraduate courses, and to supervise graduate students. Areas of research interest should include high-performance and low-power embedded systems, reconfigurable computing, FPGA or related areas.

The Hong Kong University of Science and Technology is a truly international university in Asia's world city, Hong Kong, and its Engineering School has been consistently ranked among the world's top 25 since 2004. The high quality of our faculty, students and facilities provides outstanding opportunities for faculty to develop highly visible research programs. All formal instruction is given in English and all faculty members are expected to conduct research and teach both undergraduate and graduate courses. The Department has excellent computing resources, state-of-the-art teaching and research laboratories and currently has about 40 faculty members, 800 undergraduate students and 350 postgraduate students.

Starting rank and salary will depend on qualifications and experience. Fringe benefits including medical and dental benefits, annual leave and housing will be provided where applicable. Initial appointment at Assistant Professor rank will normally be on a three-year contract. A gratuity will be payable upon successful completion of contract. Re-appointment will be subject to mutual agreement.

Applications including full curriculum vitae, list of publications, names of five referees addressed to Professor Vincent Lau, Chair of the Search Committee, and should be sent by email to **eesearch@ust.hk**. Applications will be considered until all the positions are filled.

More information about the Department is available on the website **http://www.ece.ust.hk/**.

(Information provided by applicants will be used for recruitment and other employment-related purposes.)

POLYTECHNIQUE MONTRÉAL
WORLD-CLASS ENGINEERING

COMPUTER AND SOFTWARE ENGINEERING PROFESSOR
Competition n°: 12-P-7

The successful candidate will be called upon to creatively apply their knowledge of interactive media, video games, simulation and/or virtual reality to teach undergraduate and graduate courses, supervise graduate students, and carry out innovative research projects. Researchers are provided space, technical support, modern IT environments and funds. Applicants must hold a bachelor's degree in Engineering with basic training in Computer/Software Engineering and a doctorate (PhD) in Computer Engineering, or the equivalent. The successful candidate will be a member of the Ordre des ingénieurs du Québec (OIQ), or become a member during their first contract. The candidate must also have research experience related to their area of expertise and demonstrate strong skills in multimedia. The working language is French. For more information on the department, visit **www.polymtl.ca**.

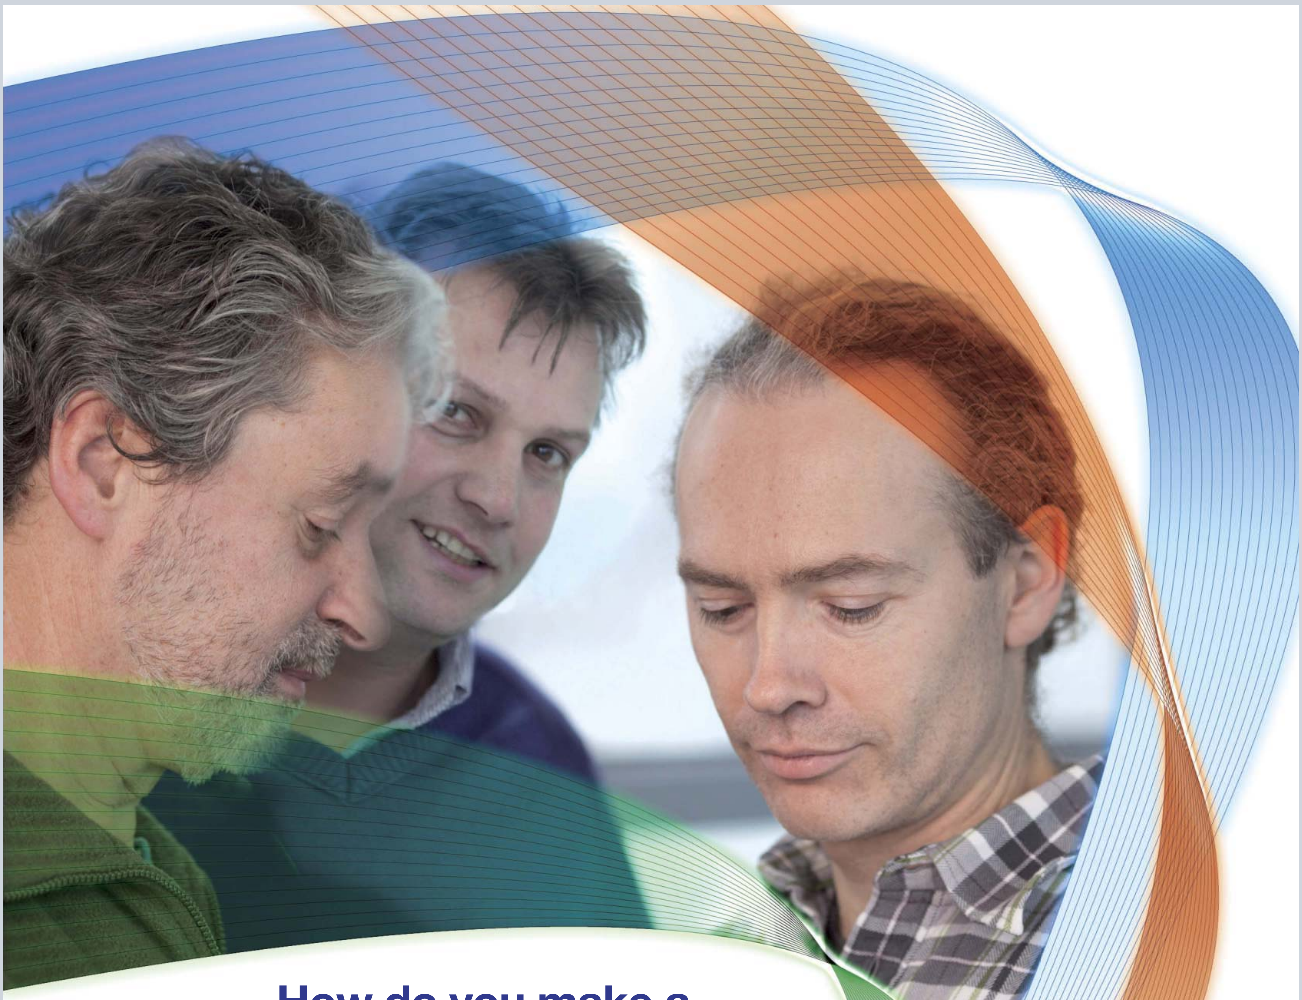

ieee spectrum Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

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

the data

Right Person, Right Place

Inventing and implementing new ideas require different networks

SINGER ELVIS PRESLEY and 18th-century potter Josiah Wedgwood are rarely mentioned in the same breath. But network science—the study of connectedness—is a world of strange bedfellows: Researchers mapping out a theory of how innovation happens have identified a key similarity between these men.

According to an analysis of network studies by John Steen of the University of Queensland, Australia, and Sam MacAulay of Imperial College London, major advances occur in two distinct phases. First, there's the idea phase—the "Eureka!" moment. Second, there's the implementation phase, when the revolutionary idea is packaged for the world at large. Each phase requires a different type of social network, and so the best people for the first phase are rarely well placed to pull off the second as well—except in rare cases, such as with Presley and Wedgwood. *—Mark Anderson*

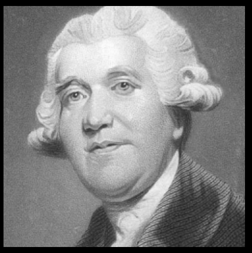

THE INNOVATORS

Here is the rarest of all innovators: one who straddles multiple worlds but who also is densely connected. Elvis Presley [left] revolutionized music because, in addition to his own talent, he had deep ties to African-American rhythm and blues and white American country and folk, coupled with access to a tightly knit network of predominantly white radio stations. Josiah Wedgwood was able to industrialize the manufacture of pottery thanks to his deep immersion in the disparate worlds of commerce, science, artisanship, and industrial design.

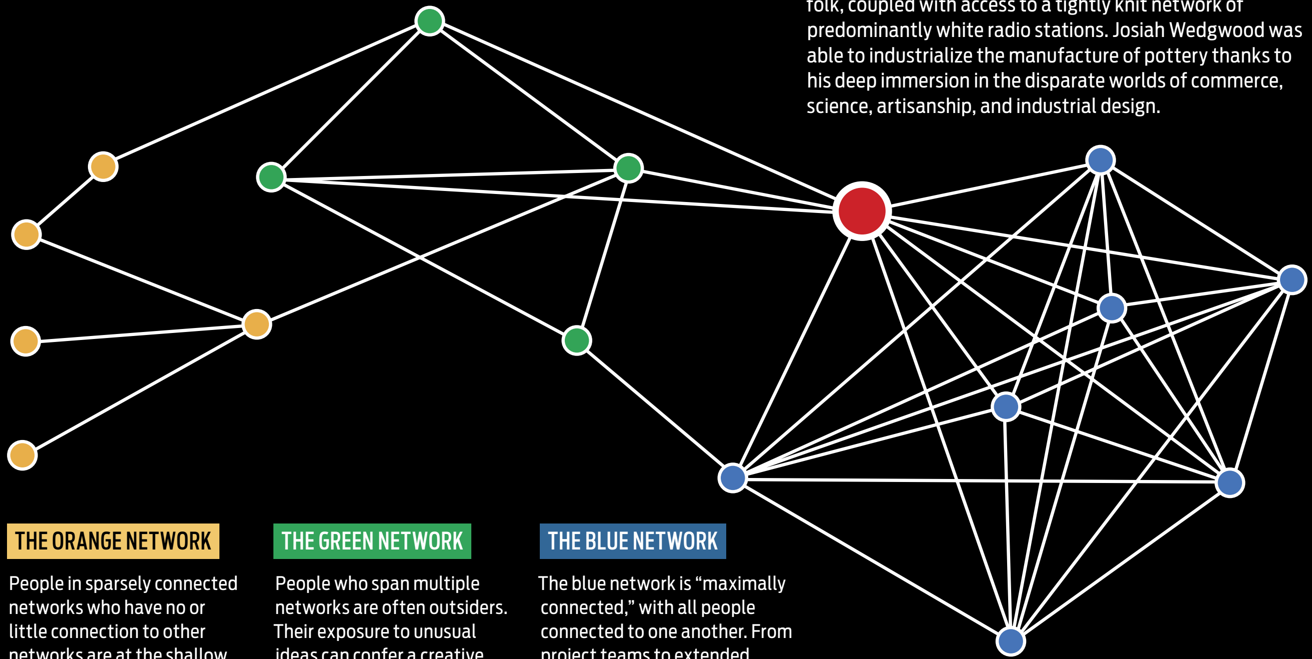

THE ORANGE NETWORK

People in sparsely connected networks who have no or little connection to other networks are at the shallow end of the innovation pool. Few connections means less exposure to new ideas and insufficient social capital to bring innovations to fruition.

THE GREEN NETWORK

People who span multiple networks are often outsiders. Their exposure to unusual ideas can confer a creative advantage that produces "Eureka" moments. But with relatively few connections, they lack the social capital needed to implement ideas.

THE BLUE NETWORK

The blue network is "maximally connected," with all people connected to one another. From project teams to extended families, such networks "form an almost ideal scaffold to get a new innovative idea off the ground," says MacAulay: Dense networks are good at mobilizing support. However, they can suffer from groupthink, and so are less likely to spawn an initial breakthrough.

For more information, see John Steen and Sam MacAulay, "The Past, Present and Future of Social Network Analysis in the Study of Innovation," **Handbook on the Knowledge Economy, Volume Two** *(Edward Elgar Publishing, 2012).*

ELVIS: MICHAEL OCHS ARCHIVE/GETTY IMAGES; WEDGWOOD: HULTON ARCHIVE/GETTY IMAGES

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


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

MathWorks®
Accelerating the pace of engineering and science

DISCOVER THE NEW LOOK AND FEEL *of* Simulink

TRY IT TODAY
visit mathworks.com

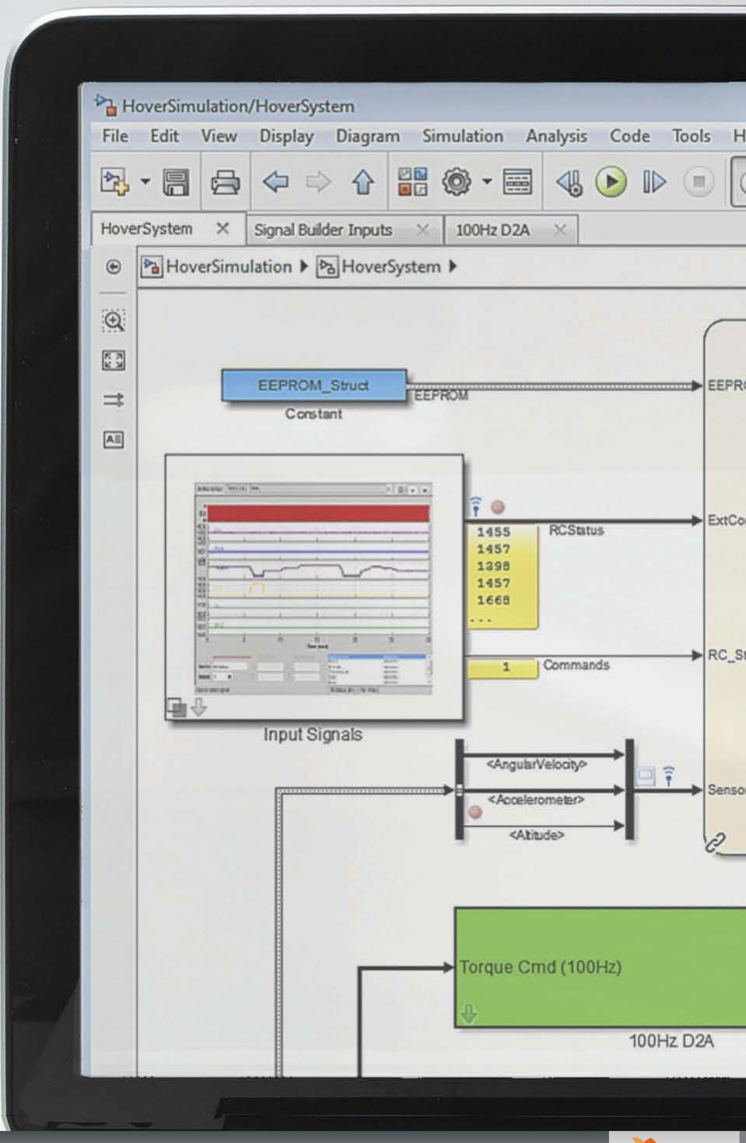

With Simulink® Release 2012b, it's even easier to build, manage, and navigate your Simulink and Stateflow® models:

- Smart line routing
- Tabbed model windows
- Simulation rewind
- Signal breakpoints
- Explorer bar
- Subsystem and signal badges
- Project management

MATLAB® & SIMULINK®