IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 11.13

THIS PARALEGAL IS A COMPUTER
An unprecedented finder of precedents
P. 28

A SPINAL SOLUTION
Treating paraplegia with electronics
P. 44

NODAL NIGHTMARE
Getting lost on the semiconductor road map
P. 24

THE HIGHS AND THE HAZARDS OF BITCOIN
Read about them at SPECTRUM.IEEE.ORG

Beam Me Up

LASERS WILL LIGHT THE WAY FORWARD FOR BMW'S i8 ELECTRIC SPEEDSTER—AND EVERYTHING ELSE, TOO P. 34

IEEE

RF DESIGN: Simulation results show the electric field distribution on top of the microstrip lines of a Wilkinson power divider. The S-parameters show input matching at 3 GHz and evenly divided power at the two output ports.

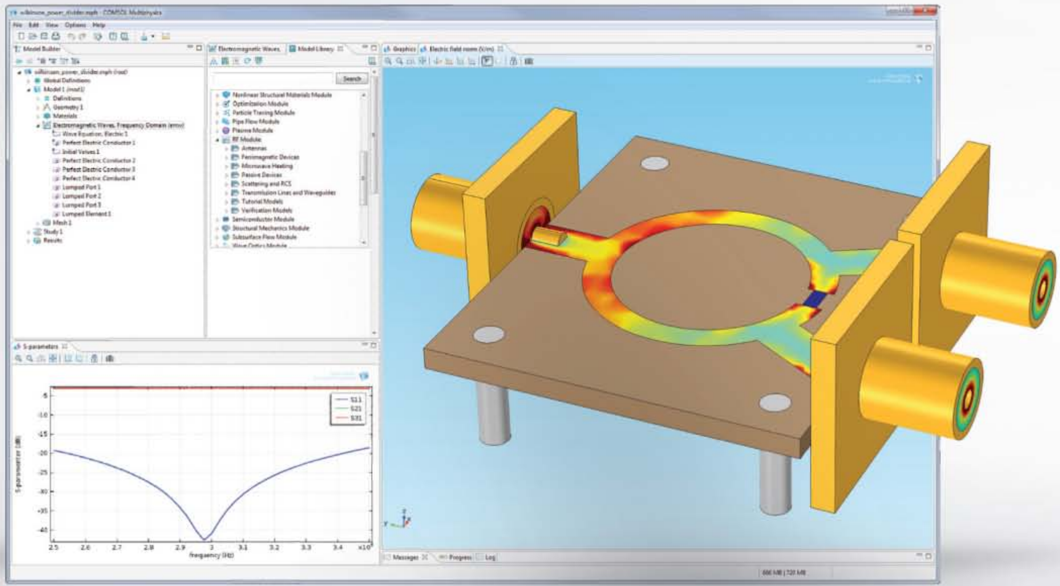

Verify and optimize your designs with COMSOL Multiphysics®

Multiphysics tools let you build simulations that accurately replicate the important characteristics of your designs. The key is the ability to include all physical effects that exist in the real world. To learn more about COMSOL Multiphysics, visit www.comsol.com/introvideo

Product Suite

COMSOL Multiphysics

ELECTRICAL
AC/DC Module
RF Module
Wave Optics Module
MEMS Module
Plasma Module
Semiconductor Module

MECHANICAL
Heat Transfer Module
Structural Mechanics Module
Nonlinear Structural Materials Module
Geomechanics Module
Fatigue Module
Multibody Dynamics Module
Acoustics Module

FLUID
CFD Module
Microfluidics Module
Subsurface Flow Module
Pipe Flow Module
Molecular Flow Module

CHEMICAL
Chemical Reaction Engineering Module
Batteries & Fuel Cells Module
Electrodeposition Module
Corrosion Module
Electrochemistry Module

MULTIPURPOSE
Optimization Module
Material Library
Particle Tracing Module

INTERFACING
LiveLink™ *for* MATLAB®
LiveLink™ *for* Excel®
CAD Import Module
ECAD Import Module
LiveLink™ *for* SolidWorks®
LiveLink™ *for* SpaceClaim®
LiveLink™ *for* Inventor®
LiveLink™ *for* AutoCAD®
LiveLink™ *for* Creo™ Parametric
LiveLink™ *for* Pro/ENGINEER®
LiveLink™ *for* Solid Edge®
File Import *for* CATIA® V5

COMSOL

FEATURES_11.13

IEEE SPECTRUM

On the Cover **Photo-illustration for *IEEE Spectrum* by Smalldog Imageworks**

BMW

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

DEPARTMENTS_11.13

IEEE SPECTRUM

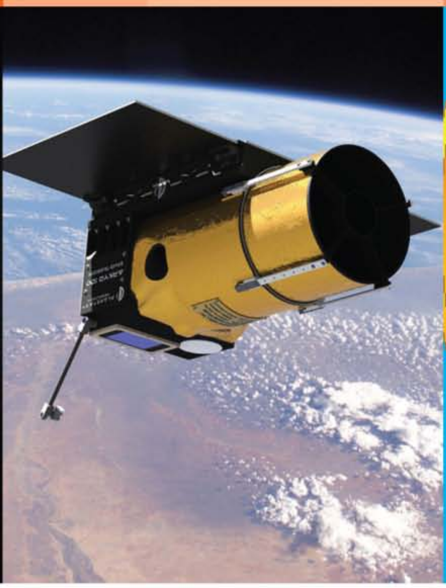

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- Automated Design Testing Procedures Using Software Simulation Tools—5 November
- Wireless Power Transfer and Microwave Energy Harvesting—7 November
- EMC Simulation in the Design Flow of Modern Electronics—14 November
- Traveling Wave Tube Design With Simulation—21 November
- **NEW PRODUCT RELEASE LIBRARY** http://spectrum.ieee.org/static/new-product-release-library
- **MASTER BOND WHITE PAPER LIBRARY** http://spectrum.ieee.org/static/masterbond-whitepaper-library
- **IBM SYSTEMS ENGINEERING RESOURCE LIBRARY** http://spectrum.ieee.org/static/ibm-systems-engineering-resource-library

The Institute

Available 08 November at theinstitute.ieee.org

- **UNCOVERING OCEAN MYSTERIES** Learn about the contributions made to oceanography by the many incarnations of England's Royal Research Ship *Discovery*, the first of which set sail in 1901.
- **HISTORY OF THE MAGNETIC COMPASS** The IEEE History Center traces the origin of the instrument that helped seafarers transport goods, explore new territory, and navigate the world's oceans.
- **TRAUMA ALERT** Roozbeh Ghaffari, an IEEE member, helped develop the CheckLight, a cap for athletes to wear under their helmets to detect serious head injuries. The product is now sold in Reebok stores.

FROM LEFT: FIA; PLANETARY RESOURCES; QUICKHONEY

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2013 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 50, issue No. 11, International edition. 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 the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

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

BACK STORY_

Driven

SOME PEOPLE LOVE CARS. Lucas Laursen isn't one of them.

He drives a car maybe three or four times a year and does the rest of his commuting on the wonderful public transit system in Madrid, where he lives. He hasn't had his own car since his teens. And until *IEEE Spectrum* sent him to the Frankfurt Motor Show, in September, to report on an electric Formula One race car, he'd never even set foot in such a venue [see "Electrifying Formula One," in this issue].

So why on earth did we send this guy to Frankfurt? Because many of the latest innovations on display there were not about driving. They were about being driven.

Laursen grew up near Los Angeles, where people spend a lot of time crawling along in traffic. Stuck in beach traffic on a recent visit home, he decided to start reporting on technology that makes driving less of a "waste of brainpower–a waste of time and energy."

It turns out that there's plenty to report: Lane tracking, self-parking, head-up displays, and even "moose avoidance" are all becoming affordable, Laursen found [see "The Race to Get Your Hands Off the Wheel" on *Spectrum*'s website].

Seeing all these technologies demonstrated at sites spread out around the show's 144 hectares took a bit of driving. Laursen didn't do any of it, of course; members of the press were chauffeured around. Naturally, he asked the chauffeur if all this autonomous driving technology had him worried. The driver just shrugged and said, "No." During the awkward silence that followed, Laursen had time to think about another implication of autonomous driving. Journalists could lose their most dependable source of quotes: the anonymous cabbie. ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. 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, Dataflow is in *IEEE Spectrum*, Vol. 50, no. 11 (INT), November 2013, p. 60, or in *IEEE Spectrum*, Vol. 50, no. 11 (NA), November 2013, p. 84.

11.13

IEEE SPECTRUM

EDITORIAL / ADVERTISING CORRESPONDENCE
IEEE Spectrum
733 Third Ave., 16th Floor
New York, NY 10017-3204
EDITORIAL DEPARTMENT
TEL: +1 212 419 7555 FAX: +1 212 419 7570
BUREAU Palo Alto, Calif.; Tekla S. Perry **+1 650 328 7570**
ADVERTISING DEPARTMENT +1 212 705 8939

PHOTOGRAPH BY **Lucas Laursen**

CONTRIBUTORS_

Wojciech Mazurczyk, Krzysztof Szczypiorski & Józef Lubacz

Mazurczyk, Szczypiorski, and Lubacz [from left], who wrote "The Spy Who Skyped Me," [p. 40] in this issue, are professors at the Warsaw University of Technology, in Poland. In 2002, as part of the Network Security Group at WUT, they founded the Stegano.net project to develop new ways to smuggle data through networks, in an effort to learn how to thwart such attempts. After many years spent anticipating evildoers and their machinations, Szczypiorski says his favorite saying comes from Indiana Jones: "Nothing shocks me. I'm a scientist."

Steven Cherry

The creator of *IEEE Spectrum*'s award-winning podcast, "Techwise Conversations," Cherry says the series was started "on a dare," back in 2010. Cherry recently left *Spectrum* to become a director at TTI/Vanguard, a forum for senior-level executives that meets five times annually to explore emerging and potentially disruptive technologies. But as a contributing editor for the *Spectrum* podcasts, he continues to interview intriguing guests, such as computer scientist Chenyang Lu [see Q&A, p. 21].

Tam Harbert

A freelance journalist based in Washington, D.C., Harbert specializes in technology and business. In our May 2013 issue, she traced a patent's tortured path through the U.S. legal system. In this issue, she profiles Lex Machina, a start-up that's applying data analytics to the black art of intellectual property litigation [see "The Law Machine," p. 28]. As data analytics automate white-collar professions like law and journalism, Harbert has begun to think a lot about how to stay ahead of the machines.

Lawrence Ulrich

Laser headlights may sound like an impossibly over-the-top car gadget, but veteran auto writer Ulrich saw them with his own eyes, at an equally over-the-top place: the FIZ, BMW's secretive skunk works in Germany [see "Whiter Brights With Lasers," p. 34], where he had to surrender his passport just to get in. Ulrich is used to adventure, having played in a rock band before switching to journalism. He covers the automotive world from his base in Brooklyn, N.Y.

Emily Waltz

As part of her research for "An Electrifying Awakening" [p. 44], freelance journalist Waltz traveled from her home in Nashville to Louisville, Ky., to visit the Frazier Rehab Institute. There she met one of the paralyzed patients involved in a groundbreaking study on spinal stimulation. His commitment to the project impressed her: "He had to stand at his support rail for an hour straight as they kept changing the stimulation configurations," she remembers. "It was tedious work."

PHOTO-ILLUSTRATIONS BY **Gluekit**

11.13

The New Global Hiring Hall

The free flow of STEM labor around the world is the wave of the future

In the annual debate over how many foreign engineers should be permitted to work in the United States, one idea is predictably absent: reciprocity.

Top engineers from around the world often find appealing jobs and pay in the United States, but the reverse isn't necessarily true for equally talented U.S. workers seeking engineering jobs abroad. To ensure an even playing field, why not promote a new global ethos of employing U.S. engineers, scientists, and software programmers anywhere in the world?

Promoting a single global employment standard for engineers would help defuse the tense debate over how many foreign engineers should be permitted to work in the United States. Linking work visas for foreign engineers in the United States reciprocally with work visas for U.S. engineers abroad would also reflect the new cosmopolitan reality. Around the world, engineers share common practices, values, and modes of expression–far more than do physicians, lawyers, or teachers. Walling in each nation's engineers is unwise, especially when the walls arise mainly to please domestic political constituencies.

To be sure, many U.S. engineers will never want employment outside the country. It would be hard to imagine them working in South Korea, for instance, where promotions are few and de rigueur camaraderie means long hours with colleagues both on and off the job. India, most flagrantly, makes the hiring of foreign engineers virtually impossible. And when German companies advertise for engineers in the United States–which they do frequently–they insist on fluency in German as well as English, even though as an everyday matter in most engineering environments, German is no longer essential.

Nevertheless, despite the cultural barriers–and in some cases much lower pay–some U.S. engineers would welcome the chance to work in Australia, China, Southeast Asia, the Middle East, thriving parts of Africa, or Western Europe, particularly if the alternative is unemployment. In those places, critical labor shortages in some technology sectors are exacerbated by rigid restrictions against qualified U.S. workers.

Things may be changing, though. "In China and Malaysia, they are *dying* for foreigners to show up and 'dump their brains' into the local ecosystem," says Michael Zielenziger, an Asia expert with Oxford Economics, a London consulting firm.

Five factors are behind what may one day become an unstoppable worldwide desire for American engineers:

1) The hegemony of English in engineering workplaces means that a U.S. worker might have a significant advantage in Shanghai or Munich.

2) Wage convergence means engineers' salaries are rising in some parts of the world and stagnating in others. Because of lower living costs in some emerging cities–and the tax holiday for Americans on much of the income they earn abroad–the actual value of a foreign paycheck may be higher than expected.

3) Half of the 10 fastest-growing economies in the world are in sub-Saharan Africa, where the availability of homegrown engineers is small and the pay for foreigners is surprisingly competitive. This year, some U.S. citizens joined an IBM research unit in Nairobi.

4) Under pressure from the U.S. government, some countries, including China, Germany, and Singapore, are lowering their barriers to imported engineers.

5) Engineers worldwide are becoming more able to compete. Many engineers in China, Germany, India, and elsewhere have some U.S. education and work experience. If there ever was a need to protect local engineers from U.S. rivals, the need is diminishing.

For engineers around the world, competition with their counterparts in the United States is intensifying in unexpected ways. While the conversation is mainly about what happens when the world's best engineers go to the United States, the odds are growing that the ultimate showdown will occur, face-to-face, in faraway places–between footloose Yanks and local engineering heroes. –G. PASCAL ZACHARY

G. Pascal Zachary is the author of *Endless Frontier: Vannevar Bush, Engineer of the American Century* (Free Press, 1997). He teaches at Arizona State University.

DIMITRI OTIS/GETTY IMAGES

NEWS

2.13 SECONDS: NEW RECORD FOR AN ELECTRIC CAR TO REACH 100 KM/H

CHARGER: The Spark-Renault SRT_01E, an electric racer, will compete in 2014 after a demo race later this year.

FIA

ELECTRIFYING FORMULA ONE

Racing exposes the best and worst about EVs

Formula One drivers have always needed physical stamina to endure crushing turns and long races. Now they'll need to be good sprinters, too. That's because in lieu of a tire-change pit stop, drivers of the new all-electric Formula E series will sprint from one car to another midrace–twice.

The car swap will allow pit crews to recharge the batteries. The relay-race aspect of Formula E harks back to the Pony Express, but the car itself, unveiled in September at the Frankfurt Motor Show, in Germany, looks ahead to the possible future of electric cars.

During the inaugural 2014-2015 season, all 10 teams will use the same base car, called the Spark-Renault SRT_01E. The car is a Frankenstein's monster of sorts, with a brand-new chassis and tires, a battery based on the Formula One system that recovers energy from braking, and motors evolved from the McLaren P1 hybrid supercar.

The challenges of designing an everyday EV are a far cry from those of relaying a pair of 800-kilogram race cars at up to 225 kilometers per hour. But the series should generate spin-offs, says Kirsty Andrew, head of commercial operations for Williams Advanced Engineering, a Formula E supplier. »

DATA BYTE

200 kW Amount of power Formula E cars can use when passing other racers

"I think motor sport is always a good place to stress-test your innovation and your ideas because it is a harsh environment, and it's a very public environment, and it's repeated frequently," she says.

Like most people involved in the series, Andrew was coy when asked to put numbers on the car's specifications. Series rules limit the motor to 134 kilowatts (180 horsepower), but they will allow drivers to use a "Push to Pass" mode of 200 kW (270 hp) at certain points, much as Formula One drivers are allowed to alter their cars' aerodynamics for short stretches.

The relative importance of the power train versus the aerodynamics may be the most important difference between the Formula E series and Formula One racing, both of which are run by the Fédération Internationale de l'Automobile. The main innovations in Formula One for the past couple of decades have been aerodynamic: Engineers have tried everything from attaching dynamic spoilers to strategically releasing exhaust over car surfaces to reduce drag. Given its lower speeds and identical cars—at least in the first season—Formula E's competition will center on how the teams manage their electric resources.

SHOWING OFF: The Spark-Renault SRT_01E made its debut at the Frankfurt Auto Show in September.

DANIEL ROLAND/AFP/GETTY IMAGES

"The problem is to have a very precise prediction of the state of charge and state of health of the battery," says electrical engineering researcher Peter Spies of the Fraunhofer Institute for Integrated Circuits, in Nuremberg, Germany, whose control system designs have attracted the attention of a leading Formula Student electric car team called EVE. (Formula Student is an international university-level engineering and motor sports competition run by the Institution of Mechanical Engineers.) His team is using new microcontrollers to build onboard monitoring instruments that in the past could be used only during pit stops. Similar fine-tuned control systems should enable Formula E teams to decide just when to pull over for a charging, he says. By season two of the series they may be able to charge using a wireless system from Qualcomm.

At least judging by what goes on in Formula Student racing, the year-to-year progress in Formula E may outpace that of Formula One. ETH Zurich engineering graduate student Julian Hügl, whose Formula Student electric car beat the gasoline-powered competition this summer, says, "What we have seen in electric racing is that the development is much, much faster compared with combustion cars." He credits rapid advances in control systems.

Another driver of innovation will be necessity. The series is cost-capped, meaning that teams will have to invent rather than spend their way out of at least some of the technical bottlenecks. Indeed, McLaren Electronics has had to develop high-grade switching components to handle the high voltages and power densities of the Formula E electric motor, says its managing director, Peter van Manen. It has also invested in thermal design and avoided exotic materials for cooling its motors, to ensure that the lessons learned will apply to other EVs. "The same approaches you develop to deal with the [racing] challenges you can then reengineer into something which is a little bit smaller in terms of performance, a little bit more cost-effective," van Manen says.

Make that a lot smaller: The SRT_01E is even chunkier than its Formula One brethren: The fairings around its 200-kg battery and the electric motor bulge behind the driver. Organizers predict the car will take 3 seconds to reach 100 km/h, compared with the 2.13-seconds record achieved by a modified Formula Student car in September. The SRT_01E's peak speed of 225 km/h lags well behind that of Formula One cars, which can exceed 360 km/h. The lower speeds may be Formula E's "most problematic thing," says Hügl.

The real measure of success for racing cars, Hügl says, is the show they put on, beginning in late 2013 with demonstration runs: "If I'm interested in racing, I don't really care about the drivetrain of the car. I want to see wheel-to-wheel racing."

—LUCAS LAURSEN

THE WORLD'S MOST POWERFUL MRI TAKES SHAPE

Medical researchers expect unprecedented resolution

An MRI scanner equipped with a superconducting magnet strong enough to lift a 60-metric-ton battle tank will offer unprecedented images of the human brain when it comes on line a little more than a year from now, say its builders.

The imager's superconducting electromagnet is designed to produce a field of 11.75 teslas, making it the world's most powerful whole-body scanner. Most standard hospital MRIs produce 1.5 or 3 T. A few institutions, including the University of Illinois at Chicago and Maastricht University, in the Netherlands, have recently installed human scanners that can reach 9.4 T. Superconducting magnets used in the Large Hadron Collider, which last year was used in the discovery of the Higgs boson, produce a field of 8.4 T.

The development of the scanner, known as INUMAC (for Imaging of Neuro disease Using high-field MR And Contrastophores), has been in progress since 2006 and is expected to cost €200 million, or about US $270 million. The project reached a key milestone this summer with delivery of more than 200 kilometers of superconducting cable, which is now being wound into coils that will produce the scanner's magnetic field.

"We're pretty proud of having met all the requirements, plus given them a little extra," says Hem Kanithi, vice president of business development at Luvata, in Waterbury, Conn., which built the superconductor.

Standard hospital scanners have a spatial resolution of about 1 millimeter, covering about 10 000 neurons, and a time resolution of about a second. The INUMAC will be able to image an area of about 0.1 mm, or 1000 neurons, and see changes occurring as fast as one-tenth of a second, according to Pierre Védrine, director of the project at the French Alternative Energies and Atomic Energy Commission, in Paris. With this type of resolution, MRIs could detect early indications of brain diseases such as Alzheimer's or Parkinson's and perhaps measure the effects of any methods developed to treat those illnesses. It would also allow much more precise functional imaging of the brain at work than is currently available. "You cannot really discriminate today what is happening inside your brain at the level of a few hundred neurons," Védrine says.

High-field MRI could also allow scientists to explore different methods of imaging. Most MRI machines rely on imaging the nuclei of hydrogen atoms, but stronger scanners might gain useful physiological information by looking for weaker signals from sodium or potassium nuclei.

Improved superconducting wire is key to making such a powerful machine. The

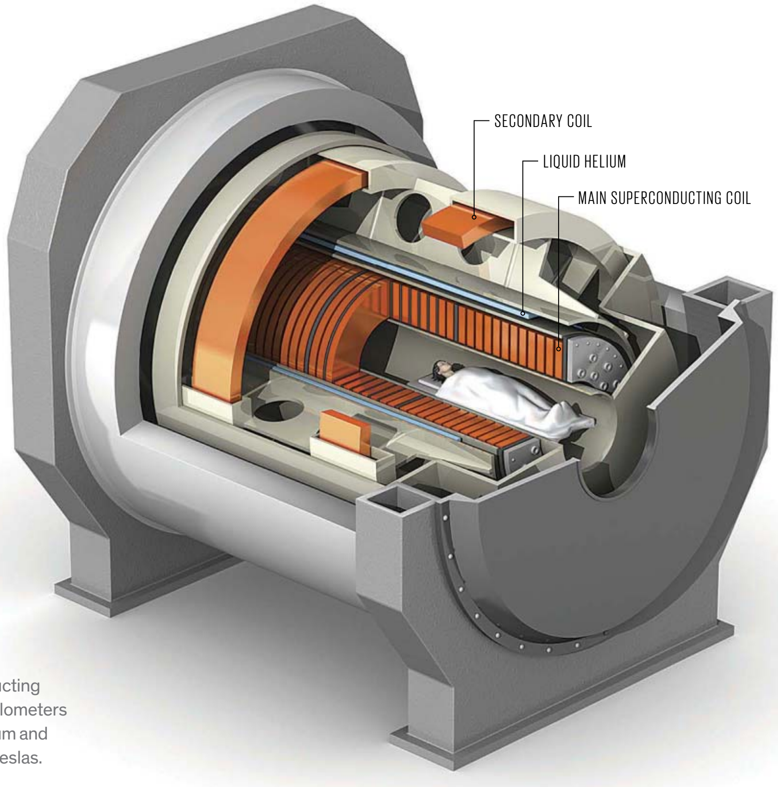

THIS ONE GOES TO 11: The main superconducting coil for the INUMAC imager is made from 170 kilometers of niobium-titanium. When bathed in liquid helium and charged, it produces a magnetic field of 11.75 teslas.

NEWS

ILLUSTRATION BY Emily Cooper

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

wire in the INUMAC magnet is made from niobium-titanium, a common superconductor alloy. But it will experience some uncommon conditions as part of INUMAC. To reach the required field strength, the electromagnet must be able to carry 1500 amperes at 12 T and be cooled by superfluid liquid helium to 1.8 kelvins. That requires specialized manufacturing and precise control of the dimensions of the wire, allowing it to be coiled so the cables are aligned to within a few micrometers of precision. "We are pushing the superconducting material niobium-titanium very close to its limits," Védrine says.

Another material, niobium-tin, can produce magnetic fields stronger than 20 T, but it was passed over for the job because it's more expensive than niobium-titanium and very brittle, making it difficult to wind.

Ultimately, Luvata produced 170 km of wire for the main superconducting coil. The company made another 58 km for two secondary coils, which will produce an opposing magnetic field to shield the area outside the machine from stray magnetic fields.

Instead of winding the wire into one long coil, as is standard in systems with lower fields, engineers are using a "double pancake" design, in which the wire is coiled into two reels that are spliced together, one on top of the other. The whole magnet will consist of 170 of these double pancakes connected in series. Védrine explains that this reduces the chances for error: Making a mistake in the winding phase using a single helical coil could ruin the whole magnet. However, a miswound pancake can simply be swapped out for a new one. The design provides space for the liquid helium bath to reach all of the coil and keep the temperature low, and it also allows engineers to place the best-performing coils at the center of the system, which improves the precision of the magnetic field.

The inner diameter of the magnet will be 90 centimeters, wide enough for a human body. Patients getting scanned will lie entirely inside the machine, but the region where the field is precise enough to get maximum resolution will be only 22 cm long. "The very good field region is only in the middle of the magnet," Védrine says. While patients could be situated so that other body parts would fit inside that region, "first we are concentrating on the brain," he says. It would take an even more massive machine to enlarge the high-resolution area.

Védrine expects to deliver the fully assembled magnet by September of next year. Other parts of the imaging system will then be added in and around the magnet, followed by about three months of testing. "Probably we'll have the first images by the beginning of 2015," he says. —NEIL SAVAGE

> The electromagnet must be able to carry 1500 amperes at 12 teslas while cooled to 1.8 kelvins

THE CACHE MACHINES

Many-core processors need the flexibility of software to use their onboard memory best

A new process for managing the fast-access memory inside a CPU has led to as much as a twofold speedup and to energy-use reductions of up to 72 percent. According to its designers, realizing such stunning gains requires a big shift in what part of the computer controls this crucial memory: Right now that control is hard-wired into the CPU's circuitry, but the substantial speedup came when the designers let the operating system handle things instead.

The CPU uses high-speed internal memory caches as a kind of digital staging area. Caches are a CPU's workbench, whether they're holding onto instructions a CPU may need soon or data it may need to crunch. And from smartphones to servers, nearly every CPU today manages the flow of bits in and out of its caches using algorithms built into its own circuits.

But, say two MIT researchers, as computers and portable devices accumulate more and more memory and CPU cores, it makes less and less sense to leave cache management entirely up to the CPU. Instead, they say, it might be better to let the operating system share the burden.

In itself, this idea is not completely new. Some of IBM's Cell processors, as well as Sony's PlayStation 3–which runs on Cell technology–allow their applications and OS kernels to fiddle with low-level CPU memory management. What's new about the MIT technology, called Jigsaw, is its middle-ground approach, which enables software to configure some on-chip memory caches but without requiring so much control that programming becomes a memory-management nightmare.

"If you go back six or seven years, you'll see that everybody was complaining that they launched the PlayStation 3 and nobody could program it well," says Daniel Sanchez, the assistant professor at MIT's Computer Science and Artificial Intelligence Laboratory and one of the inventors of Jigsaw.

Today, CPU hardware typically controls all the on-chip caches. So those caches must be designed to handle every conceivable job, from pure floating-point number crunching (which places a small burden on caches) to intensive searches and queries of a computer's mem-

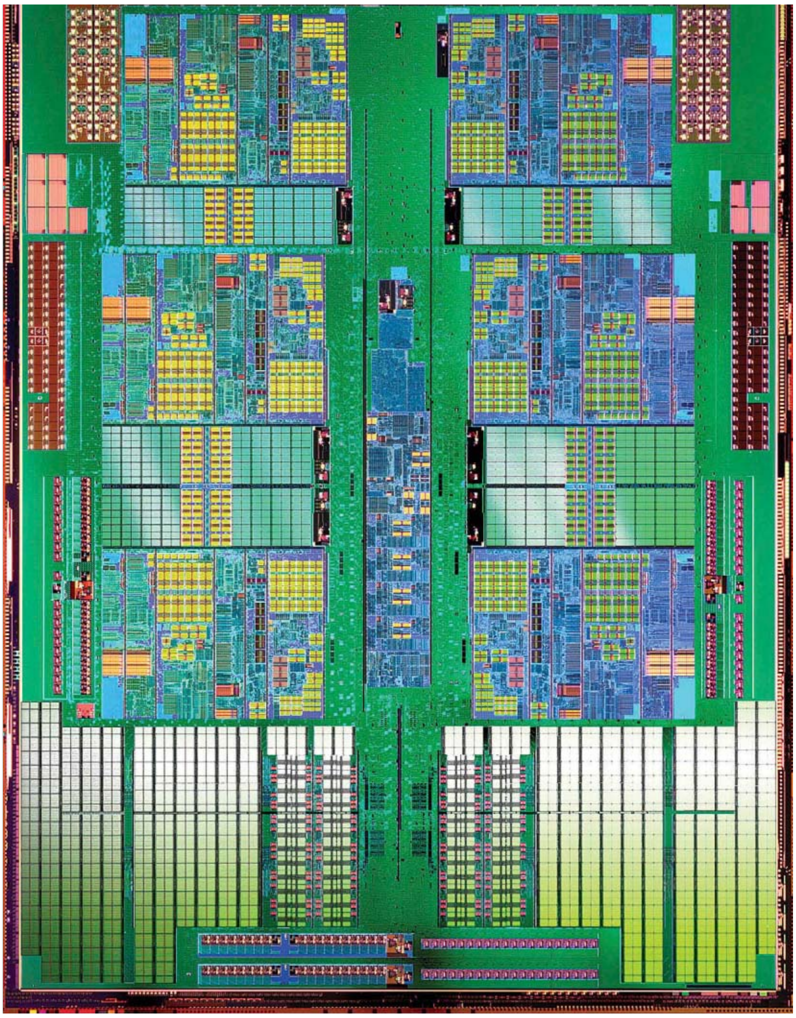

AMD

MULTICORE MEMORY: More processor cores—this AMD Opteron has six—means the computer has a harder time managing how memory moves into and out of the processor's cache.

ory banks (which can stretch their limits). Moreover, CPUs have no higher-level knowledge of the kinds of jobs they're doing. This means a self-contained numerical simulation with complex equations but little need for memory access would run with exactly the same cache resources as would a graph search, a memory-hogging hunt for relationships between stored data.

So Sanchez and his graduate student Nathan Beckmann thought, Why not let the OS trim the cache size for pure computation and swell its ranks for graph search?

The first step, they say, would be to give perhaps 1 percent of the CPU's footprint to a simple piece of hardware that could monitor in real time the cache activity in each core. Hardware cache monitors would give Jigsaw the independent oversight it would need to play air traffic controller with the CPU's caches.

Second, Sanchez and Beckmann say, the OS's kernel needs at most a few thousand more lines of code. That's not much of an addition, considering that Linux's kernel in 2012 weighed in with 15 million lines and Apple's and Microsoft's kernels unofficially contained tens of millions more than that.

One of Jigsaw's more prominent features is a software module, to be folded in with the OS, that the researchers call Peekahead. This module was adapted from the Lookahead Cache, developed more than a decade ago by Beijing computer scientists. Peekahead computes the best configuration of CPU caches based on the upcoming jobs it expects the cores to do in the coming clock cycles.

"When you let software be in charge, you have to be careful of your overhead," Sanchez says. A poorly designed cache management system, he says, might trim the cache to its optimum size and do it again every fraction of a second. But doing so taxes the CPU. And what's the point of a CPU efficiency algorithm that requires extraordinary amounts of CPU time? "The exact solution is really expensive. So we have to come up with a quick way of getting the job done so that the overhead doesn't negate the gains you get," he says.

Linley Gwennap of the Linley Group, a semiconductor consulting firm based in Mountain View, Calif., says he's impressed with Jigsaw but cautions that it's not quite ready for chip-fab prime time. "The problem is generally that a scheme that's effective on one processor may not be effective on another processor with a different hardware design," he says. "Every time the processor changes, you have to redo your software, which customers generally don't like."

Sanchez counters that software applications and utilities would remain unaffected by Jigsaw. "Only the operating system code needs to be aware of that intimate knowledge of the hardware, like the topology of the different portions of the cache," he says.

Jason Mars, an assistant professor of computer science at the University of Michigan, says Jigsaw works well as a proof of concept, which he says chipmakers might adapt as they see fit.

"The crisp novelty in this work has to do with the codesign between hardware and software," Mars says. "Much of the prior work was biased in one direction. More was expected to be done in hardware, and there was a little bit less flexibility. Jigsaw really... builds a holistic system that spans both the hardware and the software."

—MARK ANDERSON

NEWS

DATA FOR THE 31st CENTURY

New tech lets us store data for centuries, but who wants that?

Computer scientist Peter Kazansky at the University of Southampton, in England, has some words for the ages. He and a group of collaborators wrote them in quartz crystal using new optical techniques that could preserve the text for millennia. The message, which consisted of the abstract of the paper announcing the work, is stored as two types of alterations in the way quartz glass refracts light. The combination of the two allows for data-storage densities as high as 360 terabytes per disc, or more than 7000 times today's 50-gigabyte double-layer Blu-ray capacity.

There's always a catch, though. Reading the message requires an electron microscope, and the process may never provide faster access to stored data than existing technology can. This and similar over-the-horizon memory research may someday improve big-data storage, but such systems aren't an easy fit with today's data-storage needs, experts say. Improved density and durability are both helpful, but readability and the capacity to rewrite data in a different format might be more important.

Long-term data storage is one part of what Mark Watson, Oracle's director of hardware development, calls the data-storage pyramid. At the tip of that pyramid are data that users want immediate access to, such as new photos posted to social networks. For that, server farms use the latest fast-writing and -reading media, such as solid-state flash memory or spinning magnetic hard disks. Some companies are even shifting such data to their server's fastest-access memory, volatile dynamic RAM. Oracle is taking that tack to the extreme with a new database system housing a whopping 32 TB of DRAM.

But users may also need to park some data long term for legal reasons, or for future big-data analysis, without the need for frequent or fast access. For that, archivists can use storage media that are slower but much cheaper on a per-bit basis. Kazansky and his colleagues' optical innovations may fit into that layer of the data-storage pyramid.

To justify its place, a quartz media format would need to be more cost-effective than the base of today's data-storage pyramid, which is old-fashioned magnetic tape. "For archive, tape is still quite an important medium," says physicist Thomas Thomson of the University of Manchester, in England. That attitude is common among storage experts, who note that even as optical discs approach 300 GB, magnetic-tape costs continue to drop as densities, reliability, and speeds climb. One estimate puts them at 4 U.S. cents per gigabyte.

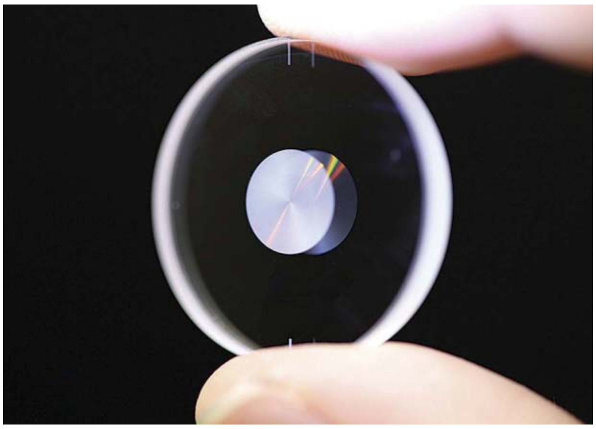

DISC OF DESTINY: Using alterations in the way quartz refracts light, we could store data in a superdense form for centuries.

One problem quartz does solve is the degradation of data in present-day media. "You're going to lose information," warns information technology consultant Michael Peterson, so archive systems must strike a balance between durability and cost. As storage media age, most archive managers copy the data to new media. Kazansky's team predicts that their data could remain legible for centuries, and the longer they can delay migrating the data, the lower the cost of archiving. "We can give all our knowledge to future civilizations," says Kazansky.

Quartz isn't the only millennial medium in the works. In 2009, researchers stored information using iron nanoparticles inside carbon nanotubes, a configuration they say could last a billion years. In 2012, Harvard University researchers encoded an entire 53 000-word book into a 50-nanogram clump of DNA.

But David Rosenthal, a digital preservation expert at Stanford University, warns that the market for such durable data is small. Instead, he argues, most archives already move their data to new media long before the original media's expiration date, to take advantage of incremental technological improvements that save space and energy. The plan to develop a 300-GB Blu-ray disc is a good example of the sort of incremental advances that help reduce archiving costs without really changing how archiving is done, Peterson says.

In fact, the way we store data, such as using standards for metadata that ensure the data's accuracy, is evolving just as fast as the hardware on which the data is preserved, says Rosenthal. "What we learn in the digital-preservation world is that it's not about the storage media; it's about the system," Peterson says.

There is one notable exception to this constant copying and upgrading: In deep space, nobody can update your archiving system. Voyager 1 recently reached interstellar space, carrying a gold-plated record and the needle to play it with, which are intended to survive the multimillennial journey to other star systems. If Voyager or any of our other starbound spacecraft inspire aliens to visit, will they find a society with thousands of years of updated digital archives? Or will it just be a few ragged bits in a long-abandoned format that read, "Look on my data, ye mighty, and despair!" —LUCAS LAURSEN

UNIVERSITY OF SOUTHAMPTON

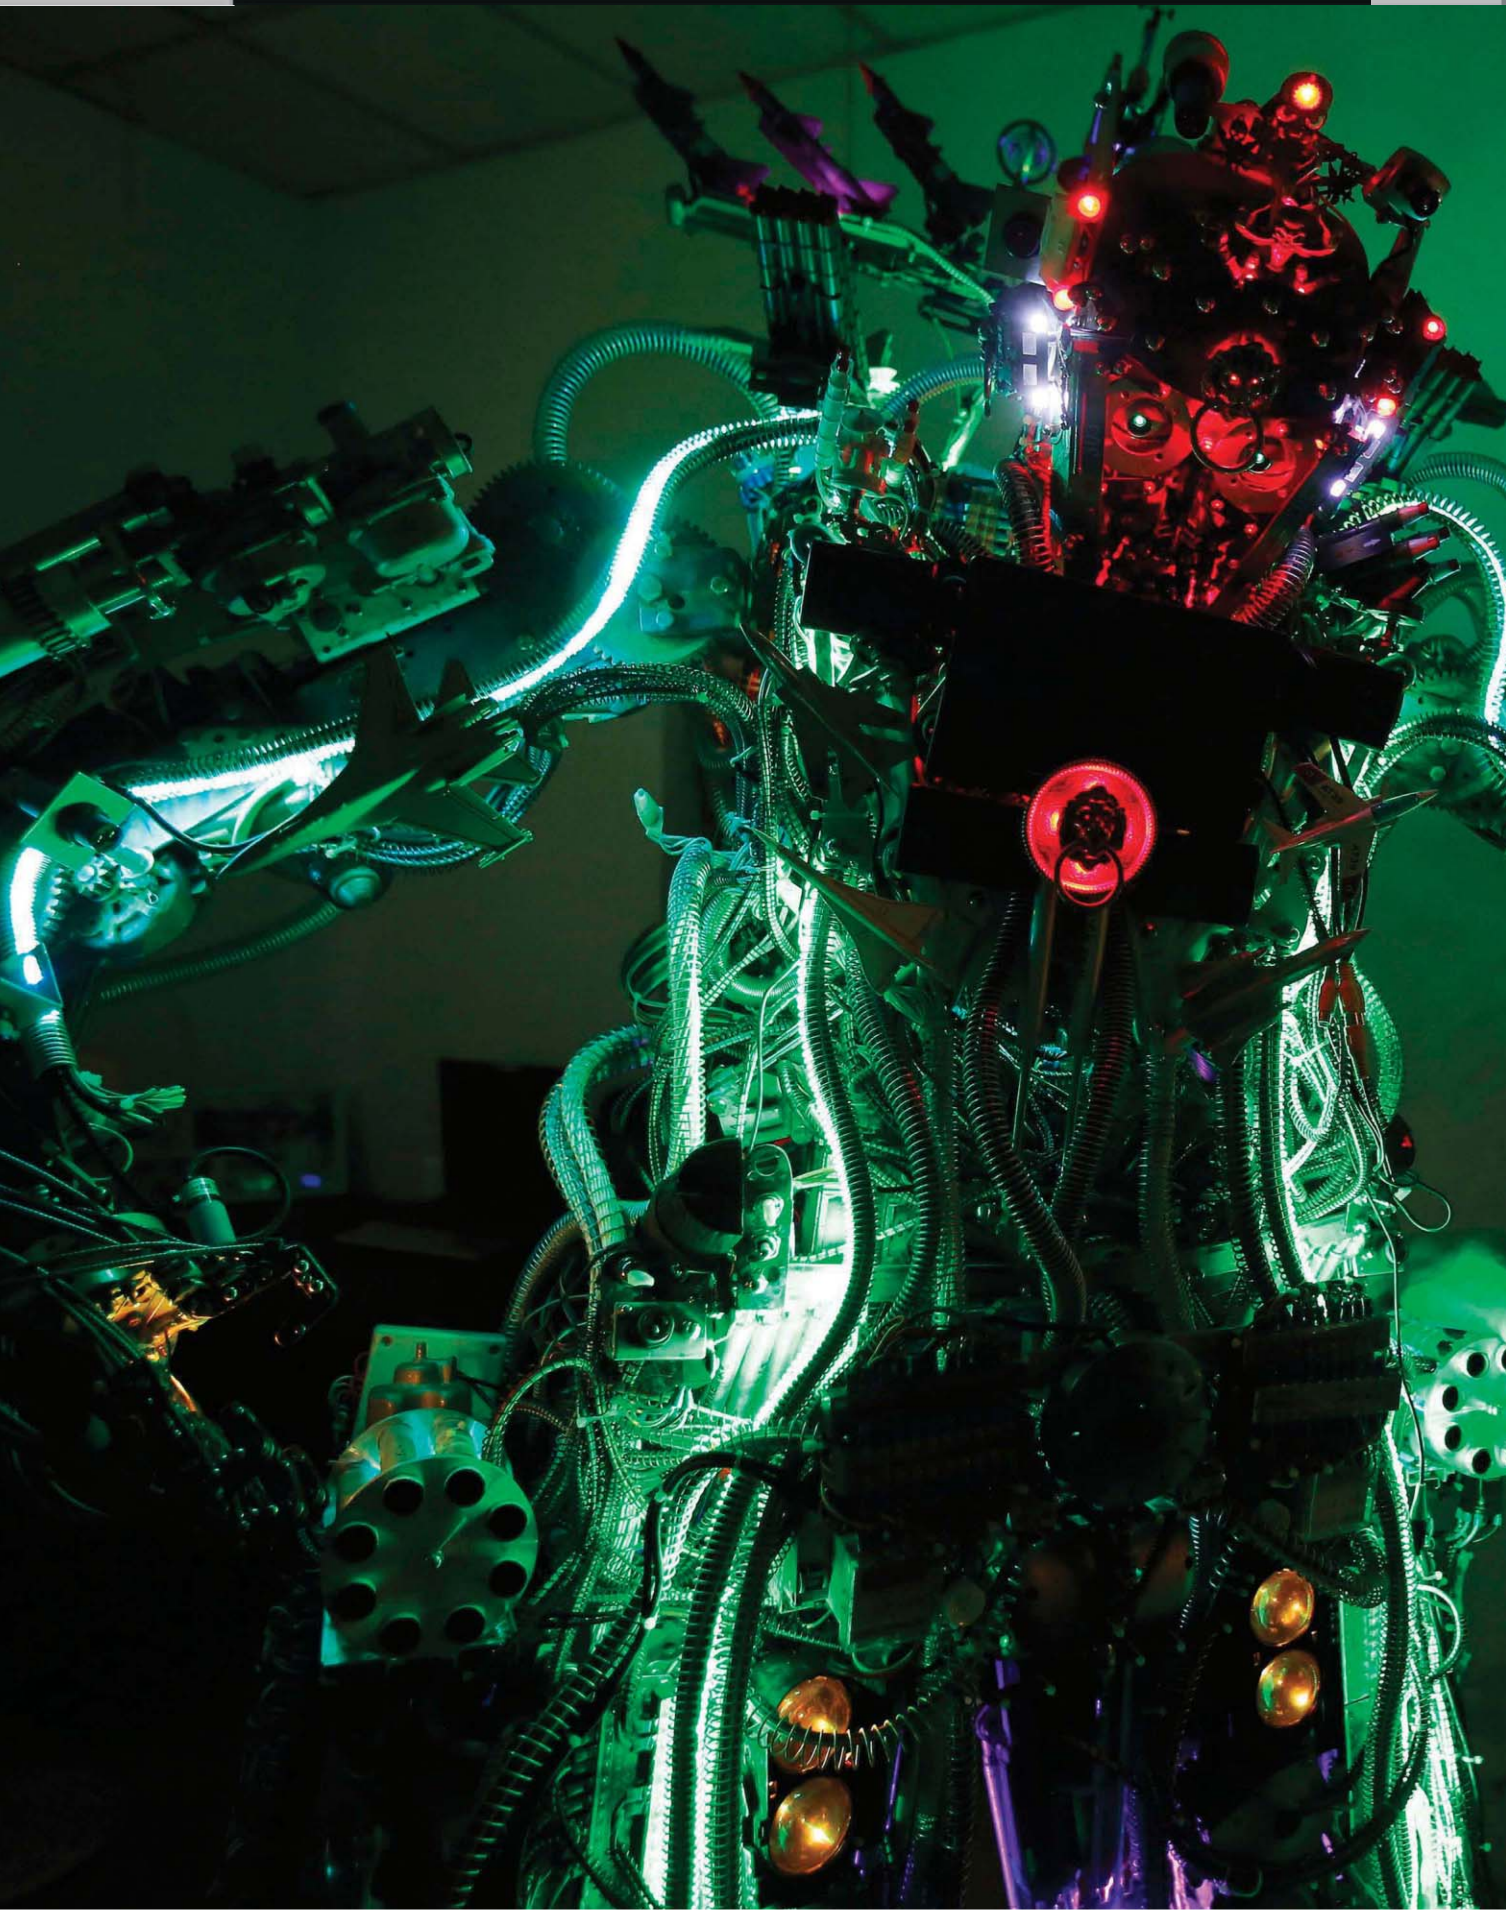

PHOTOGRAPH BY How Hwee Young/EPA/Landov

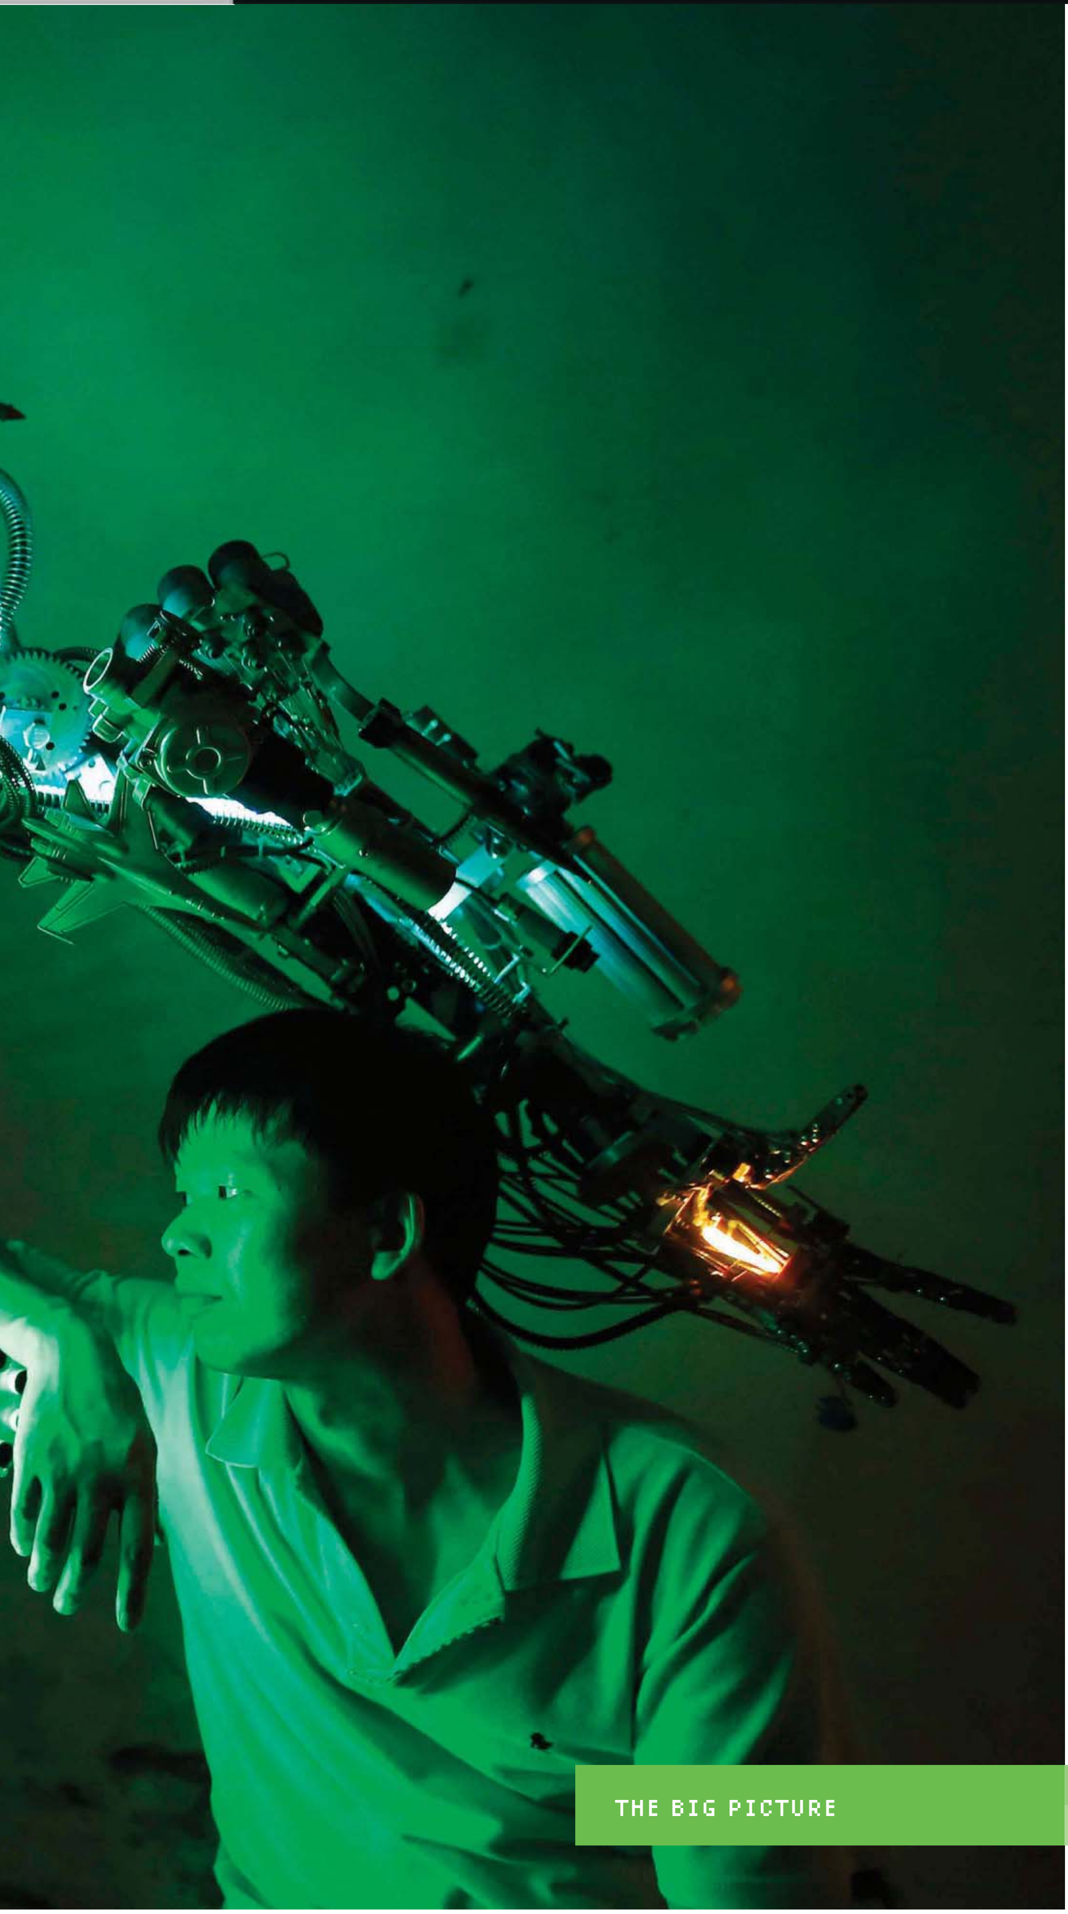

HOMESPUN HIGH TECH

THIS HUMANOID robot lights up and moves a little when you flip some switches on its back–and that's about all it does. You'd probably scoff if someone told you that building it took 150 000 yuan (about US $24 500) and 11 months of work. But what if we told you that the 2.1-meter-tall, 250-kilogram automaton was painstakingly put together with components its creator bought piecemeal from a secondhand market? And what would you do upon learning that the man behind the machine is a self-taught inventor and the son of poor Chinese farmers who couldn't afford to educate him past the fifth grade? You'd have to applaud Tao Xiangli, who learned the skills needed to make this awesome-looking robot–named King of Innovation–by taking machines apart and putting them back together.

THE BIG PICTURE NEWS

RESOURCES

50 000

THE TYPICAL NUMBER OF USERS WHO CONNECT DAILY TO ONE OF OVER 4000 TOR NETWORK RELAYS

ADAFRUIT INDUSTRIES

PRIVACY IN A BOX

A PALM-SIZE WI-FI ACCESS POINT THAT ENABLES ANONYMOUS BROWSING

RESOURCES_HANDS ON

Recent revelations about the scope of electronic surveillance performed by the U.S. National Security Agency have provoked considerable debate both inside and outside the United States. Ironically, however, over the past few years, departments within the U.S. federal government have been helping to support a project designed to combat Internet monitoring. The Tor Project permits anonymous Web browsing and publishing, and it has received funding from the U.S. National Science Foundation, the Defense Department, and the State Department. They've backed the project either for research purposes or, in the State Department's case, to assist democratic movements in countries with repressive regimes. ● Now, using the Raspberry Pi microcontroller as a platform, Adafruit Industries has released the Onion Pi Pack, a kit that creates a Wi-Fi access point with Tor software built in. [For a profile of Adafruit's founder, see "Limor Fried: Channel Your Inner Maker," *IEEE Spectrum,* May 2013.] ● For most people, using Tor involves installing a customized Web browser and the software required to anonymize Web traffic. Alternatively, users can boot their computers with a USB flash drive loaded with the Linux operating system and a Tor installation. ● But having the Tor software on a computer can be suspicious in itself. Less ominously, many companies prohibit the installation of nonstandard software. Further, if the computer normally runs OS X or Windows, using the Linux-based flash drive can make it difficult to work with some files. By offloading the Tor software to a Wi-Fi access point, users won't need to install anything on their computers, and they can work with their normal operating systems. ● Adafruit supplies the Onion Pi Pack in two versions—a long-range edition that sports a Wi-Fi adapter with a full-size 15-centimeter ▸

antenna, and a more discreet short-range edition that has a button adapter. In addition to an adapter, each kit comes with a Model B Raspberry Pi microcomputer; a case; an FTDI serial communications cable; a micro USB cord and mains adapter (to supply power); an Ethernet cable; and a 4-gigabyte SD card loaded with a version of Linux that is optimized for the Pi (but doesn't have the Tor software installed).

Putting the hardware together is trivial: Pop the SD card and Wi-Fi adapter into their slots, and then hook up the micro USB cord and Ethernet cable to the mains adapter and a router, respectively.

Getting the software running is a different story. Adafruit supplies a detailed set of tutorials online, but using them can sometimes feel like going down a series of rabbit holes. Adafruit designed each tutorial to be self-contained, so it can be used with different kits. So there's one generic tutorial for initially configuring a Raspberry Pi, another for setting it up as a Wi-Fi access point, and so on. And some tutorials direct you to complete another tutorial midway through. Some of the tutorials are also best completed using graphical-interface-based tools, so you'll probably want to hook the Pi up to a keyboard and mouse and an HDMI-capable TV for at least part of the process.

You'll also be editing a fair number of text-based configuration files. Fortunately, you can avoid a lot of error-prone typing by using the FTDI cable to connect the Pi to whatever computer you're using to read the online tutorials. Then it's mostly just a matter of cutting and pasting various alphanumeric gobbets into configuration files via an old-school terminal window.

There are also traps for the unwary during the process: For example, the British-made Pi has a U.K. keyboard set as the default. This caused a problem when I dutifully changed my log-in password as my first step and later changed the keyboard settings to "U.S." After rebooting the Pi, I had difficulty logging back in because some of the symbols I'd used for my password are in different places on a U.S. keyboard. Another example is when I tried powering the Pi using the FTDI cable alone. This feature would have been convenient, but it turns out the cable doesn't supply enough juice to sustain an Ethernet connection. It took me a while to figure out why my network had stopped working.

Once configured, however, the Onion Pi is easy to use: Just power it and plug it into an active Ethernet port. Certain services are blocked to preserve anonymity, such as the "ping" command used to test connections, but in general the Onion Pi behaves like any other Wi-Fi access point. There's one big difference, though: speed.

Tor anonymizes traffic by encrypting it and sending it to a relay. The traffic then bounces around at random within the Tor network until it is decrypted at an exit gateway somewhere around the globe and passed to the Internet at large. The outside world sees the IP address of this exit gateway as your IP address. All this takes time, and the Tor network has only so much bandwidth, so attempting to download a video file, for example, can be an agonizingly slow process.

There are also limits to the veil of anonymity provided. If you transmit identifying information—such as your name—to an external server, then you've rendered the system moot. What's even more insidious is that if you use the same browser for accessing the Tor network as you do for regular surfing, you could be "clickprinted"—that is to say, identified by such things as cookies stored on your computer or the precise configuration of your plug-ins. Browser, beware. **—STEPHEN CASS**

SPACE ON $35 A DAY

CITIZEN-SCIENCE SATELLITES ALLOW ANYONE TO RUN EXPERIMENTS IN ORBIT

NANOSATISFI

IN 1957, ARTHUR FROMMER TRAVELED EUROPE on US $5 a day, and the Soviets launched Sputnik 1. Today, $5 scarcely covers an espresso at a Parisian café, but it could buy 3 hours on a satellite packed with technology that would have seemed like science fiction to Cold War rocketeers. ● In August, two small satellites were sent to the International Space Station (ISS) and will be launched into their own orbits in November. At a cost of about $35 to $45 per day of run time, purchasable in blocks of three days or more, students and hobbyists can run experiments with the satellites. These ArduSat spacecraft, from the start-up NanoSatisfi, were developed from an open-source design based on the Arduino microcontroller. Each is a cube 10 centimeters on a side and has a magnetometer, a spectrometer, temperature sensors,

SMALL SATS

a gamma-ray detector, a Geiger counter, and a 1.3-megapixel digital camera.

The ArduSats are attracting a variety of experiments and experimenters, says Chris Wake, NanoSatisfi's vice president of business development: "Computer science classes are using it to train students on open-source software and hardware, while an earth sciences class has experiments to measure Earth's magnetosphere."

NanoSatisfi keeps the price down by running 10 experiments in parallel on each ArduSat and using off-the-shelf components. The satellites have no radiation shielding and will remain aloft for just a few months before reentering the atmosphere and burning up.

NanoSatisfi isn't alone in the citizen-science satellite game. In February, the United Kingdom orbited STRaND-1 (Survey Training, Research, and Nano-satellite Demonstrator), the first satellite built around a standard consumer smartphone, the Google Nexus One. The phone was loaded with Android apps developed by university students and enthusiasts, such as an app intended to determine the satellite's position using the onboard camera. Communication with STRaND-1 was lost shortly after launch but was restored a few months later.

In April, NASA launched its own pair of Android PhoneSats, costing just $3500 apiece. During their week-long lives before reentry, the satellites took sensor readings and snapped images of Earth, beaming down the data in small packets to amateur radio operators around the world.

While building a satellite may be easier than ever, getting it into space remains fearsomely complex. For the past four years, a small company called NanoRacks has been flying science experiments into orbit on board the space shuttle and on ISS resupply missions. NASA and Japanese astronauts control

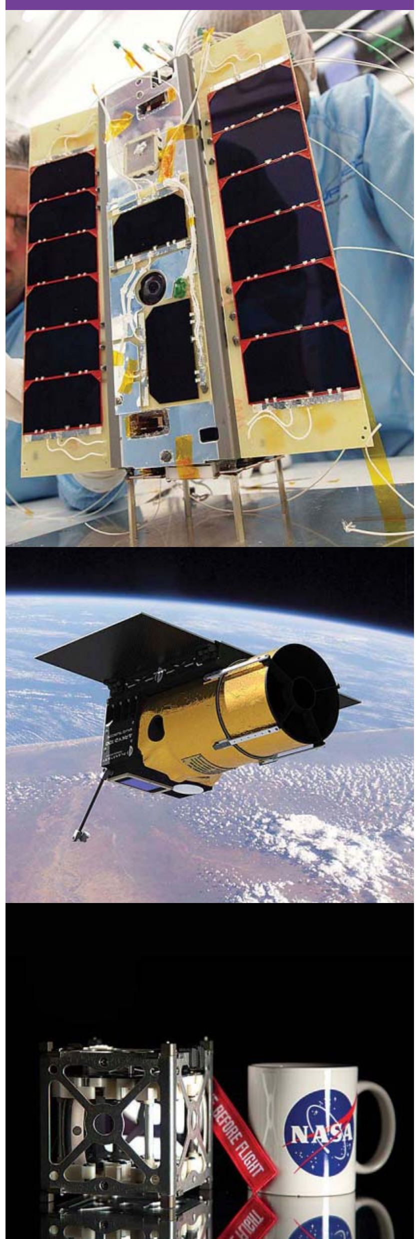

Britain's phone-based STRaND-1 [top] had a communications blackout for a few months after launch; Planetary Resources plans to sell amateur astronomers access to a small space telescope [center]; and NASA's smartphone satellite [bottom] orbited just a few days before reentry.

FROM TOP: SURREY SATELLITE TECHNOLOGY; PLANETARY RESOURCES; NASA

the experiments, many of which have been designed by high school students. Over 20 000 U.S. students have been involved, raising the necessary funds (a 30-day mission starts at $30 000) from bake sales and car washes.

Last year, NanoRacks deployed its first satellite from an ISS air lock. "Today we have over 50 satellites under contract. We're now designing our own cubesat dispenser to maximize how many satellites we can put out during a hatch opening," says Jeffrey Manber, NanoRacks' founder and managing director.

NanoRacks will handle the November launch of the aforementioned ArduSats as well as two other high-profile citizen-science satellites. Planetary Resources, a start-up hoping to prospect and mine near-earth asteroids, is planning space telescopes that armchair astronomers will use to take high-resolution photos of celestial objects for $200. Cheaper still is SkyCube, built by Southern Stars. This ultrabasic satellite will broadcast a 120-character tweet of your choice from orbit for just $1 or shoot a low-res image of Earth for $6.

"We want to democratize access to space by crowdsourcing data analysis, software, and hardware," says Peter Diamandis, cofounder of Planetary Resources. Crowdsourcing money is also part of the plan: NanoSatisfi, Planetary Resources, and SkyCube all launched their efforts on the crowd-funding website Kickstarter.

"Kickstarter validated our market areas around researchers, students, teachers, programmers, and hobbyists," says Wake. His company, NanoSatisfi, has since completed a round of venture capital funding. NanoRacks' Manber sees a bright future ahead: "I believe that by 2020, every school district in America will have had the chance to participate in space." —**MARK HARRIS**

SHAPEWAYS BRINGING 3-D PRINTING TO THE MASSES

Being in the 3-D printing business brings some conveniences—like fewer trips to the hardware store. "If we need an adapter for a vacuum hose, we don't buy it," says Ben Wilkinson-Raemer, an industrial engineer with the start-up Shapeways. "We print it."

Instead of selling 3-D printers, which typically cost upwards of US $1000, Shapeways sells printing services to designers and DIY makers. Customers first upload their designs to the company's website, and then Shapeways, which originated as a spin-off of the Dutch electronics behemoth Royal Philips in 2008, takes care of the fabrication and ships out the printed objects. Most are fabricated in lightweight plastics, but several metals and ceramics are available; objects can be as large as a night table.

To print an object, you first need a 3-D model, and many potential customers aren't familiar with modeling programs. To overcome that barrier, Shapeways has strategies for simplifying the design process: It offers online tutorials, provides apps that make it easy to create objects like smartphone cases or cookie cutters, and hosts a forum where customers can find and hire 3-D modelers.

Roughly 300 000 people have used the Shapeways site either to make or buy a 3-D doodad. "Our customers range from hardcore model-maker engineer types to suburban housewives," says Wilkinson-Raemer. About 11 000 of those customers have opened shops on the Shapeways site to sell their creations.

Shapeways' headquarters are in New York City, where the company has nearly completed a full-scale production and distribution facility in a warehouse district. Ten machines were running when *IEEE Spectrum*'s reporter visited the factory, but it could host as many as 50 printers by the time it's fully built out in January 2014.

FORM FITTING: Burlesque performer Dita Von Teese models a 3-D printed gown, produced by Shapeways, precisely tailored to her body. The mesh of the gown's "fabric" is fully articulated to permit movement.

Despite the high expense of real estate in New York, Shapeways decided to build a manufacturing plant there because the location guarantees proximity to people that the company values: customers and software developers. The city's throngs of designers and artists are all potential customers, Wilkinson-Raemer says, and the company wants to make it possible for dedicated makers to visit the factory to get a better understanding of the printers' potential. New York also has an excellent talent pool for software developers—important for a company that expects to hire more than 50 people this year, he notes.

Software is crucial to Shapeways' manufacturing process. Printing one object at a time would be impossibly slow and expensive, Wilkinson-Raemer explains. Instead, in the preproduction phase, software arranges a 3-D "puzzle" of objects that can be printed in a single run. Prices for printing vary according to size and material, but producing a plastic napkin ring, for example, formed from the letters of a custom message, will cost a customer about $9.

Michael Wolf, founder and chief analyst of the high-tech research firm NextMarket Insights, in Edmonds, Wash., thinks that Shapeways is well positioned for success. "People are fascinated by 3-D printing, but a 3-D printer is still out of reach for the average consumer," he says. The price of printers may be coming down, but Wolf says that printing services are what consumers will find affordable for at least the next five years.

Wolf credits Shapeways with expanding its market by harnessing the creativity of app developers. "They have simple apps that access this really complex printing capability," he says, citing an app called Mixee Me that guides customers in creating plastic dolls. Says Wolf: "Even my 8-year-old daughter can design a little version of herself and print out a figurine."

But all these clever tactics may not be enough to keep Shapeways in business if a deep-pocketed e-commerce company decides to enter the fray. Wolf expects serious competitors to emerge in the next year or two, naming Amazon.com as one potential rival. In proving the viability of the 3-D printing market, Wolf says, Shapeways may end up encouraging other players to jump in and take 3-D printing seriously. —ELIZA STRICKLAND

Founded: 2008; **Headquarters:** New York City; **Founders:** Robert Schouwenburg, Marleen Vogelaar, Peter Weijmarshausen; **Funding:** Recently raised $30 million in venture capital; **Employees:** 90+; **Website:** http://www.shapeways.com

SHAPEWAYS

RESOURCES_AT WORK

Q&A: CHENYANG LU

MELDING SOFTWARE MODELS WITH SENSOR NETWORKS COULD SAVE LIVES

ANOTHER BRIDGE FAILS: The collapse of the Skagit River Bridge, in Washington state, is another example of infrastructural weakness.

CHENYANG LU

IN RECENT YEARS, SEVERAL DISASTERS—INCLUDING bridge failures—have illustrated that it's dangerous to take infrastructure for granted. Chenyang Lu, a professor at Washington University in St. Louis, is researching the technology of "cyberphysical systems," which combine wireless sensor networks and simulations. He spoke with Steven Cherry for *IEEE Spectrum*'s "Techwise Conversations" podcast about how this technology could predict at least some failures.

Steven Cherry: What kind of sensors would you use to monitor a bridge, and how would you use the information they produce?

Chenyang Lu: Strain gauges would measure the strength of connection between joints. Fiber-optic sensors would detect deformation in the structures. Accelerometers would measure vibration. Then you can use model analysis, signal processing—essentially, you are doing pattern recognition, to detect if there are changes in the bridge that would be indicative of damage.

S.C.: As well as laboratory testing, you've also been working with collaborators to collect data on real-world structures, such as the Jindo bridge in South Korea. What's special about this bridge?

C.L.: Jindo Bridge is a relatively new bridge; it's a very large bridge. Jindo is an island in South Korea, so they are basically connecting the island with the mainland. My collaborators deployed hundreds of sensors on that bridge. They continuously monitor the vibrations of the bridge, including cases where there are typhoons and so on blowing by. The unique thing about their system is these are wireless sensor networks. This is especially meaningful to the United States, because of all of these old bridges, where you have to retrofit a monitoring system. That's why in recent years we've been working on wireless monitoring, as it's easier to deploy. For example, you can take a technology called IEEE 802.15.4 to create a mesh network with sensors running at very low power.

S.C.: One problem with sensor networks is that the amount of data they throw off can be overwhelming.

C.L.: That's why it's never enough just to put sensors in and collect a ton of data. I have certainly heard incidents where they have these brand-new bridges with a lot of sensors, and they accumulate a ton of data but don't know what to do with it. So that's why you have to have intelligence built into these systems. It's not just collecting data but also all these sophisticated model analyses that detect localized damage. Recently here, we've begun researching how to make prognoses, so that not only can you assess the current state of a bridge, you can also predict the remaining lifetime of the bridge.

These questions and answers have been edited and condensed. To read or listen to the full interview, visit http://spectrum.ieee.org/lu1113

STEPHEN BRASHEAR/GETTY IMAGES

PHOTO-ILLUSTRATION BY Gluekit

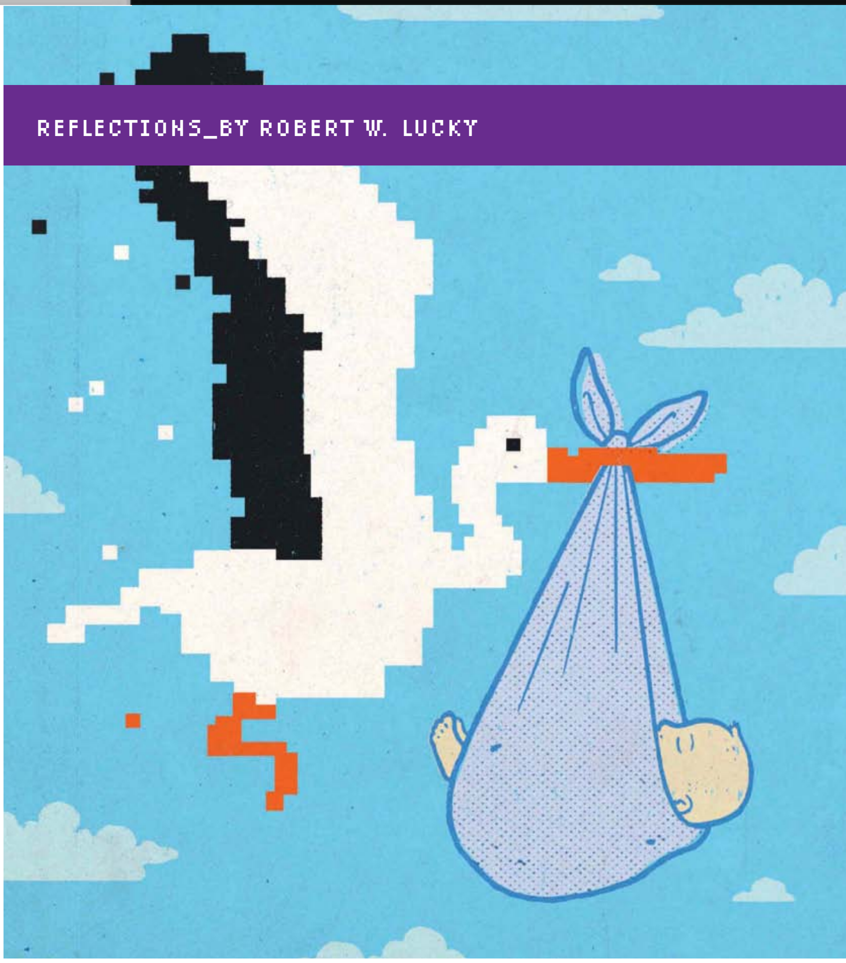

GOOD YEARS FOR TECHNOLOGY?

Is the success of an engineering career determined by birth year?

MALCOLM GLADWELL, IN *OUTLIERS* (2008), ARGUES THAT one of the factors in success is the year you were born. In technology, he reasons that the perfect year to have been born was 1955, so that at the start of the computer revolution—which he pegs as January 1975—you would have been just the right age to take advantage of the new advances. If you had been born later than 1955, you would still have been in school when everyone else was capitalizing on the new tech. If you had been born earlier than 1955, you would probably already have had a comfortable job at a place like IBM and not have been inclined to strike out on your own. • In support of this assertion he cites birth years for famous technologists. Bill Gates, Steve Jobs, Andy Bechtolsheim, Eric Schmidt, and Vinod Khosla were born in 1955. Bill Joy, Scott McNealy, and Steve Ballmer were all born only a year earlier or later. It's an impressive list, but there are enough counterexamples that I'm inclined to be skeptical. Moreover, Gladwell's measure of success is wealth, whereas we have had many engineers who through discovery or invention have changed the world without becoming rich. • Nevertheless, I have often contemplated the idea that there *are* good times—and not-so-good times—to enter the engineering profession. If we ignore the economic climate and look strictly at technology, I think there have been times when a field has been fertile and ready for harvest, and other times when a field has become overworked, with a lot of hard effort required for ever-diminishing returns.

I sometimes think of this in cosmological terms. The universe is expanding at a rate known as Hubble's constant, whereas in electrotechnology we have an expansion given by Moore's Law. In cosmology there is believed to have been an inflationary period after the big bang when space itself expanded at a tremendous rate. I like the idea of technology undergoing a similar rapid expansion at a time like 1975, when the very space in which we could work expanded explosively.

In cosmology, gravity drives contraction, while dark energy propels expansion. There are corresponding forces in technology. Contraction happens as we push against the physical limits of a technology, while expansion is augmented by the ever-increasing numbers of engineers worldwide.

Technological inflation, unlike cosmological inflation, however, occurs repeatedly. Certainly the computer revolution was one such period, but there have been other technologies that sparked transformations, such as the transistor, the integrated circuit, the Internet, optics, and wireless. Each breakthrough greatly expanded our working space. So perhaps a better analogy would be the hypothesis in evolutionary theory of "punctuated equilibrium," which describes periods of stasis alternating with periods of sudden change.

Is the current year a good year for technology or a not-so-good year? Frankly, I have no idea, because we only seem to know these things in retrospect. When some of us older engineers get together, we often reminisce about having enjoyed the "golden years." But we didn't know at the time how good we had it.

For some years after the transistor had been invented, engineering schools still taught students about vacuum tubes. The integrated circuit was invented in 1958, yet none of the IEEE Fellows writing about the future in the *Proceedings of the IEEE* in 1965 mentioned it. The Internet sneaked up on us even as we were using it. And so forth. The walls around us had moved outward—and yet we still felt confined to the smaller space to which we had become accustomed.

So perhaps in 20 years some columnist for *IEEE Spectrum* will write about what a good year 2013 was. But maybe not. ■

ILLUSTRATION BY Dan Page

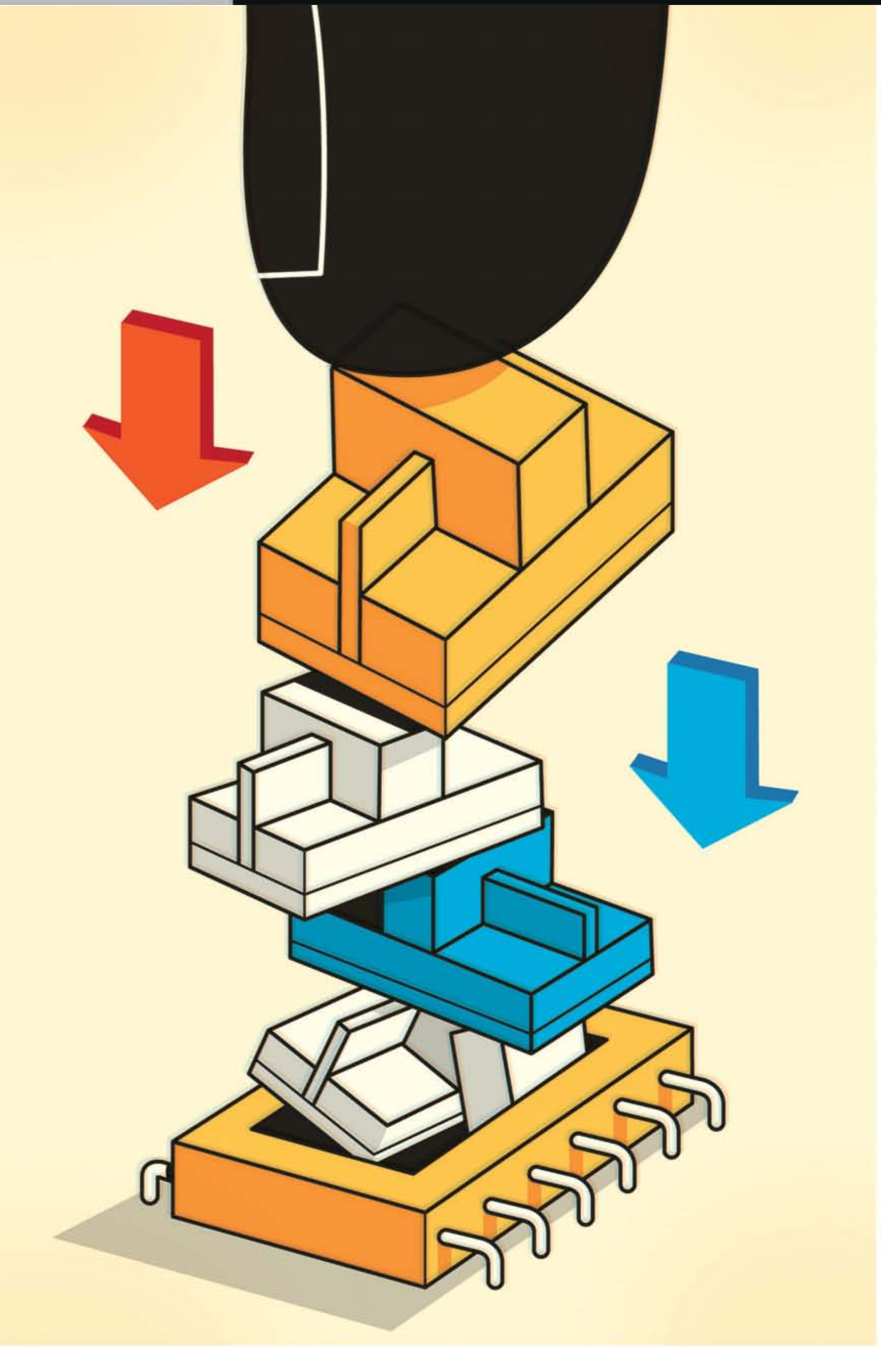

The End of the Shrink

Nobody will say that Moore's Law is over. But it's starting to get *really* complicated

BY RACHEL COURTLAND

ONE CHILLY TUESDAY EVENING LAST DECEMBER, dozens of physicists and engineers who dream up tomorrow's transistors met in San Francisco to ponder the far future. Would today's state-of-the-art switch–a three-dimensional transistor dubbed the FinFET–be able to carry chips "to the finish," a distant, possibly unreachable horizon where transistors are made up of just a handful of atoms? Or would we need a new technology to get us there?

This may all sound like the tech world's version of arguing over how many angels can dance on the head of a pin, but it actually has enormous real-world implications. The semiconductor industry pulled in revenues of US $300 billion in 2012. After decades of fulfilling Gordon Moore's prophesy of steadily doubling transistor densities (these days every 18 to 24 months), the industry is now delivering integrated circuits with transistors that are made using what chipmakers call a 20- or 22-nanometer manufacturing process. An IC fabricated with this process, such as a microprocessor or a dynamic RAM (DRAM) chip, can have billions of transistors.

Nevertheless, there on the cutting edge, the business is troubled. Each new generation of ultradense chips demands a new manufacturing process of mind-boggling industrial and technological complexity. The struggle has become so pitched that researchers are now often at a loss for words to describe the metrics of their progress.

At the December meeting, for example, Chenming Hu, the coinventor of the FinFET, began by mapping out the near future. Soon, he said, we'll start to see 14-nm and 16-nm chips emerge (the first, from Intel, are slated to go into production later this year). Then he added a caveat whose casual tone belied its startling implications: "Nobody knows anymore what 16 nm means or what 14 nm means."

SEMICONDUCTORS

ILLUSTRATION BY Harry Campbell

It's actually become a fairly common refrain among industry experts. The practice of attaching measurements to chip generations has "been hijacked by marketers to an enormous extent," one chip-design expert told me. "A lot of it's really smoke and mirrors," says analyst Dan Hutcheson of VLSI Research in Santa Clara, Calif. It's "spin," he says, that's designed to hide widening technological gaps between chip companies.

The nanometer figures that Hu discussed are called nodes, and they are, for want of a better term, the mile markers of Moore's Law. Each node marks a new generation of chip-manufacturing technology. And the progression of node names over the years reflects the steady progress that both logic and memory chips have made: The smaller the number, the smaller the transistors and the more closely they are packed together, producing chips that are denser and thus less costly on a per-transistor basis.

But the relationship between node names and chip dimensions is far from straightforward. Nowadays, a particular node name does not reflect the size of any particular chip feature, as it once did. And in the past year, the use of node names has become even more confusing, as chip foundries prepare to roll out 14-nm and 16-nm chips, custom-made for smartphone makers and other customers, that will be no denser than the previous 20-nm generation. That might be just a temporary hiccup, a one-time-only pause in chip-density improvement. But it's emblematic of the perplexing state of the field.

Moore's Law, when reflected through the steady march of node names, might seem easy and inexorable. But today a plague of intense manufacturing and design problems is forcing compromises that are sometimes sobering. And some analysts suggest that regardless of what we call the next generation of chips, the transition from old to new no longer provides nearly the kind of payoff–in cost or performance–that it used to.

"WHAT DO YOU MEAN BY 14 NM?" WHEN I ASKED AN STEEGEN THAT question at an industry conference in July, she smiled and let out a wry, knowing laugh. "Ah...what's in a name?" asked Steegen, senior vice president for process technology development at Imec, the Belgian research center. "Actually, not that much any more."

It's a state of affairs that has been nearly two decades in the making. Once upon a time, the node name told you practically everything you needed to know about a chip's underlying technology. If you trained your microscope on microprocessors made by a handful of different companies using a 0.35-micrometer process, you'd find that their products were all remarkably similar.

In the mid-1990s, when such chips were the state of the art, 0.35 µm was an accurate measure of the finest features that could be drawn on the chip. This determined dimensions such as the length of the transistor gate, the electrode responsible for switching the device on and off. Because gate length is directly linked to switching speed, you'd have a pretty good sense of the performance boost you'd get by switching from an older-generation chip to a 0.35-µm processor. The term "0.35-µm node" actually meant something.

But around that same time, the link between performance and node name began to break down. In pursuit of ever-higher clock speeds, chipmakers expanded their tool kit. They continued to use lithography to pattern circuit components and wires on the chip, as they always had. But they also began etching away the ends of the transistor gate to make the devices shorter, and thus faster.

After a while, "there was no one design rule that people could point to and say, 'That defines the node name,'" says Mark Bohr, a senior fellow at Intel. The company's 0.13-µm chips, which debuted in 2001, had transistor gates that were actually just 70 nm long. Nevertheless, Intel called them 0.13-µm chips because they were the next in line. For want of a better system, the industry more or less stuck to the historical node-naming convention. Although the trend in the measurements of transistors was changing, manufacturers continued to pack the devices closer and closer together, assigning each successive chip generation a number about 70 percent that of the previous one. (A 30 percent reduction in both the x and y dimensions corresponds to a 50 percent reduction in the area occupied by a transistor, and therefore the potential to double transistor density on the chip.)

The naming trend continued as transistors got even more complex. After years of aggressive gate trimming, simple transistor scaling reached a limit in the early 2000s: Making a transistor smaller no longer meant it would be faster or less power hungry. So Intel, followed by others, introduced new technologies to help boost transistor performance. They started with strain engineering, adding impurities to silicon to alter the crystal, which had the effect of boosting speed without changing the physical dimensions of the transistor. They added new insulating and gate materials. And two years ago, they rejiggered the transistor structure to create the more efficient FinFET, with a current-carrying channel that juts out of the plane of the chip.

Through all this, node name numbers continued to drift ever downward, and the density of transistors continued to double from generation to generation. But the names no longer match the size of any specific chip dimension. "The minimum dimensions are getting smaller," Bohr says. "But I'm the first to admit that I can't point to the one dimension that's 32 nm or 22 nm or 14 nm. Some dimensions are smaller than the stated node name, and others are larger."

"There was no one design rule that people could point to and say, 'That defines the node name.' "
—MARK BOHR, INTEL

The switch to FinFETs has made the situation even more complex. Bohr points out, for example, that Intel's 22-nm chips, the current state of the art, have FinFET transistors with gates that are 35 nm long but fins that are just 8 nm wide.

That is, of course, the view from a chip manufacturer's side. For his part, Paolo Gargini, the chairman of the International

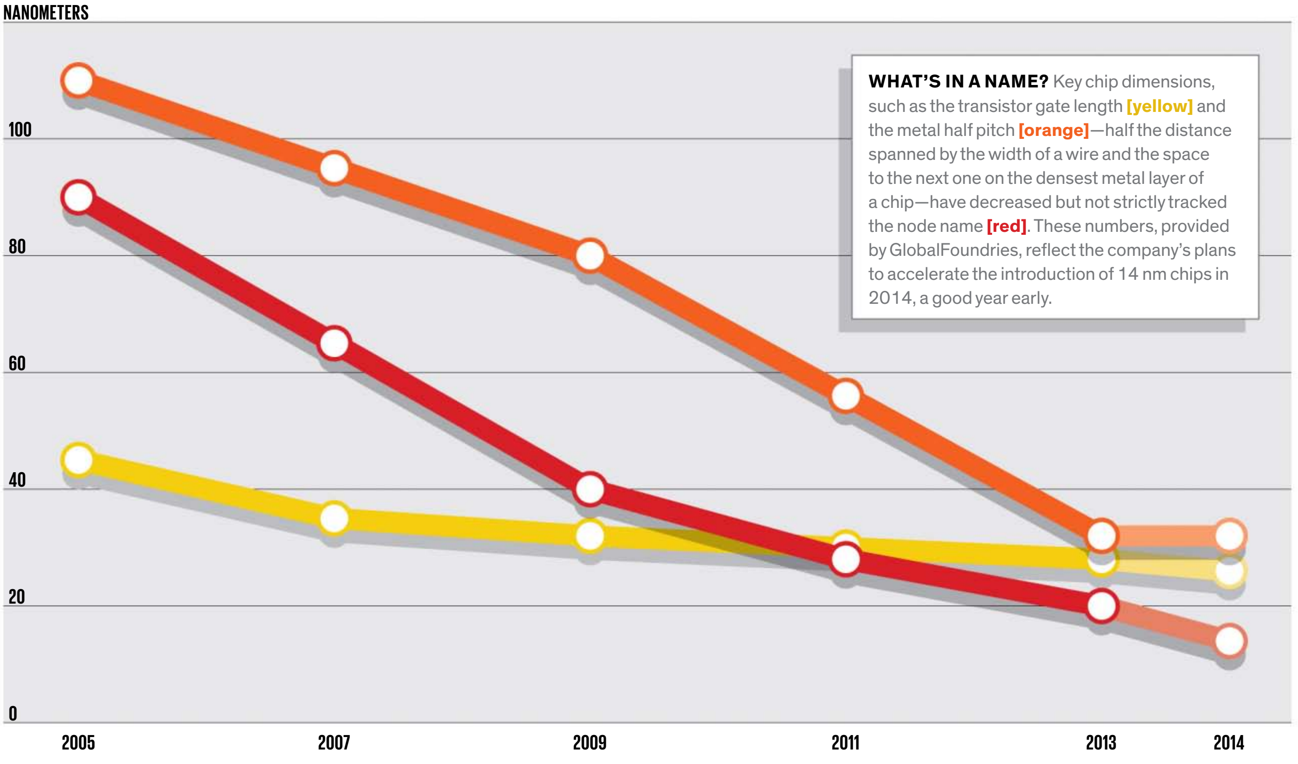

Technology Roadmap for Semiconductors, says the node is and always has been defined by the proximity of wires on the first metal layer on the back of the chip, a dimension that was reflected well in DRAM and, later, flash memory, but not in logic.

REGARDLESS OF DEFINITION, NUMBERS IN NODE NAMES HAVE continued to decline. Along with them, the distance between transistor gates and that between the closest copper wires on the back of the chip have also decreased. Both of those features help define how dense a chip can be and thus how many more you can produce on a single silicon wafer to drive down costs.

But the difficulty inherent in printing ever-finer features has now taken its toll. "When we got to around 28 nm, we were actually pushing the limits of the lithographic tools," says Subramani Kengeri, vice president of advanced technology architecture at GlobalFoundries, the world's second-biggest chipmaking foundry after Taiwan Semiconductor Manufacturing Co.

To deal with this, Kengeri and his colleagues were forced to adopt a lithographic technique called double patterning. It lets technicians pattern smaller features by splitting a single patterning step into two, relying on a slight offset between the two steps.

Intel used the technique to form transistors on its 22-nm chips, but it stuck to single patterning to make the densest metal layer. Pushing the technique to its limits, the company made wires with a pitch of 80 nm, which encompasses the width of one wire and the space to the next. By adopting double patterning, GlobalFoundries and others could push the pitch down to about 64 nm for their 20-nm chips. But that move came with a significant trade-off: Double-patterned chips take longer to make, adding significantly to the cost.

Carrying this technique over from the 20-nm node to 14 nm would mean that chipmakers would have to double-pattern even more layers of the chip. So last year, Kengeri and his colleagues announced a chip industry first: They would put a stop to the shrink. GlobalFoundries' line of 14-nm chips, which are slated to begin production in 2014, may be the foundry world's first FinFET transistors. But the company will build the new chips with the same wiring density used in its 20-nm chips. "The first-generation FinFET is basically reusing all of that and plugging a FinFET into that framework," Kengeri says. "It's really a 20-nm FinFET, in a way." Nevertheless, the company refers to these as 14-nm chips because they offer roughly a generation's-worth jump in performance and energy efficiency over its 20-nm chips.

Kengeri hopes that by putting a one-generation pause on shrinking chips and focusing on introducing 3-D transistors, GlobalFoundries will catch up with Intel, which is already shipping 3-D devices in its 22-nm chips. GlobalFoundries' 14-nm chips aren't any denser than–and therefore cost just about as much as–the previous generation, but they're still a big improvement, Kengeri says. "Our point–and our customers agree–is that as long as they see that value, they don't care what the technology is called or what is inside."

DATA SOURCE: GLOBALFOUNDRIES

SEMICONDUCTORS

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

"It is quite a controversial move," says William Arnold, chief scientist at ASML, the world's largest maker of semiconductor-fabrication equipment. "The customers of the foundries, the people who are making cellphone parts, are very skeptical of not being able to get a shrink along with a performance improvement. They're pretty vocal about saying that they're not happy about that."

THE FOUNDRIES' LATEST MOVE ASIDE, CHIPS ARE STILL MORE OR LESS doubling in density from node to node, says Andrew Kahng, a professor at the University of California, San Diego, and an expert on high-performance chip design. But for Kahng, the steady progression of node names masks deeper problems. There is a difference, he says, between "available density" (how closely you can pack circuits and wires on a chip) and "realizable density" (what you can actually put into a competitive commercial product).

The sheer density and power levels on a state-of-the-art chip have forced designers to compensate by adding error-correction circuitry, redundancy, read- and write-boosting circuitry for failing static RAM cells, circuits to track and adapt to performance variations, and complicated memory hierarchies to handle multicore architectures. The problem, Kahng says, is that "all of those extra circuits add area." His group has been scouring company specs and deconstructing images of chips for years, and they've come to an unsettling conclusion: When you factor those circuits in, chips are no longer twice as dense from generation to generation. In fact, Kahng's analysis suggests, the density improvement over the past three generations, from 2007 on, has been closer to 1.6 than 2. This smaller density benefit means costlier chips, and it also has an impact on performance because signals must be driven over longer distances. The shortfall is consistent enough, Kahng says, that it could be considered its own law.

This might be a recoverable loss. So far, Kahng says, the chip industry has made it a priority to keep up the pace of Moore's Law, ensuring that manufacturers can continue to build and release new product families while using a new process every 18 to 24 months. This means there hasn't been time to explore a number of design tricks that could be used to cut down on power or boost performance. "When you're on that kind of schedule, you don't have time to optimize things," he says. As the value of the simple shrink decreases, he says, chipmakers should then be able to revisit their designs and find chip-improving approaches they may have missed or else left on the cutting-room floor.

When will the scaling stop? Today's patterning technology, which relies on 193-nm laser light, is becoming an ever more costly challenge, and its natural successor, shorter-wavelength extreme ultraviolet lithography, has been long delayed.

Kahng says chipmakers may face a more immediate struggle with wiring in just a few years as they attempt to push chip density down past the 10-nm generation. Each copper wire requires a sheath containing barrier material to prevent the metal from leaching into surrounding material, as well as insulation to prevent it from interacting with neighboring wires. To perform effectively, this sheath must be fairly thick. This thickness limits how closely wires can be pushed together and forces the copper wires to shrink instead, dramatically driving up the resistance and delays and drastically lowering performance. Although researchers are exploring alternative materials, it's unclear, Kahng says, whether they will be ready in time to keep up with Moore's Law's steady pace.

Many people in the industry, who have watched showstopper after showstopper crop up only to be bypassed by a new development, are reluctant to put a hard date on Moore's Law's demise. "Every generation, there are people who will say we're coming to the end of the shrink," says ASML's Arnold, and in "every generation various improvements do come about. I haven't seen the end of the road map."

But for those keeping track of the road, those mile markers are starting to get pretty blurry. ■

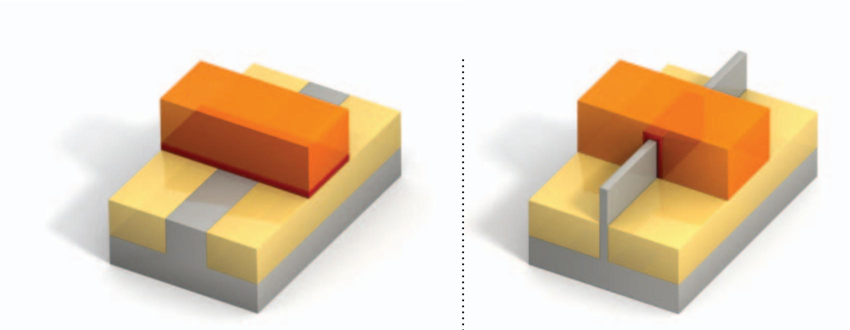

PLANAR
NODE: 20 nm // MANUFACTURER: Leading foundries // CHANNEL LENGTH: 28 nm
FIRST METAL LAYER PITCH: 64 nm

3-D
NODE: 22 nm // MANUFACTURER: Intel
CHANNEL LENGTH: 30 nm // FIRST METAL LAYER PITCH: 90 nm // FIN WIDTH: 8 nm

TWO TRANSISTORS: Chipmakers are in the process of moving from traditional planar transistors [left] to ones that pop out of plane [right]. Intel introduced these 3-D transistors in 2011, and they are now shipping widely. The leading foundries, such as GlobalFoundries, Samsung, and Taiwan Semiconductor Manufacturing Co., are in the process of ramping up production of 20-nanometer planar transistors. They will make the switch to 3-D with the next generation.

POST YOUR COMMENTS at http://spectrum.ieee.org/shrink1113

ILLUSTRATIONS BY Emily Cooper

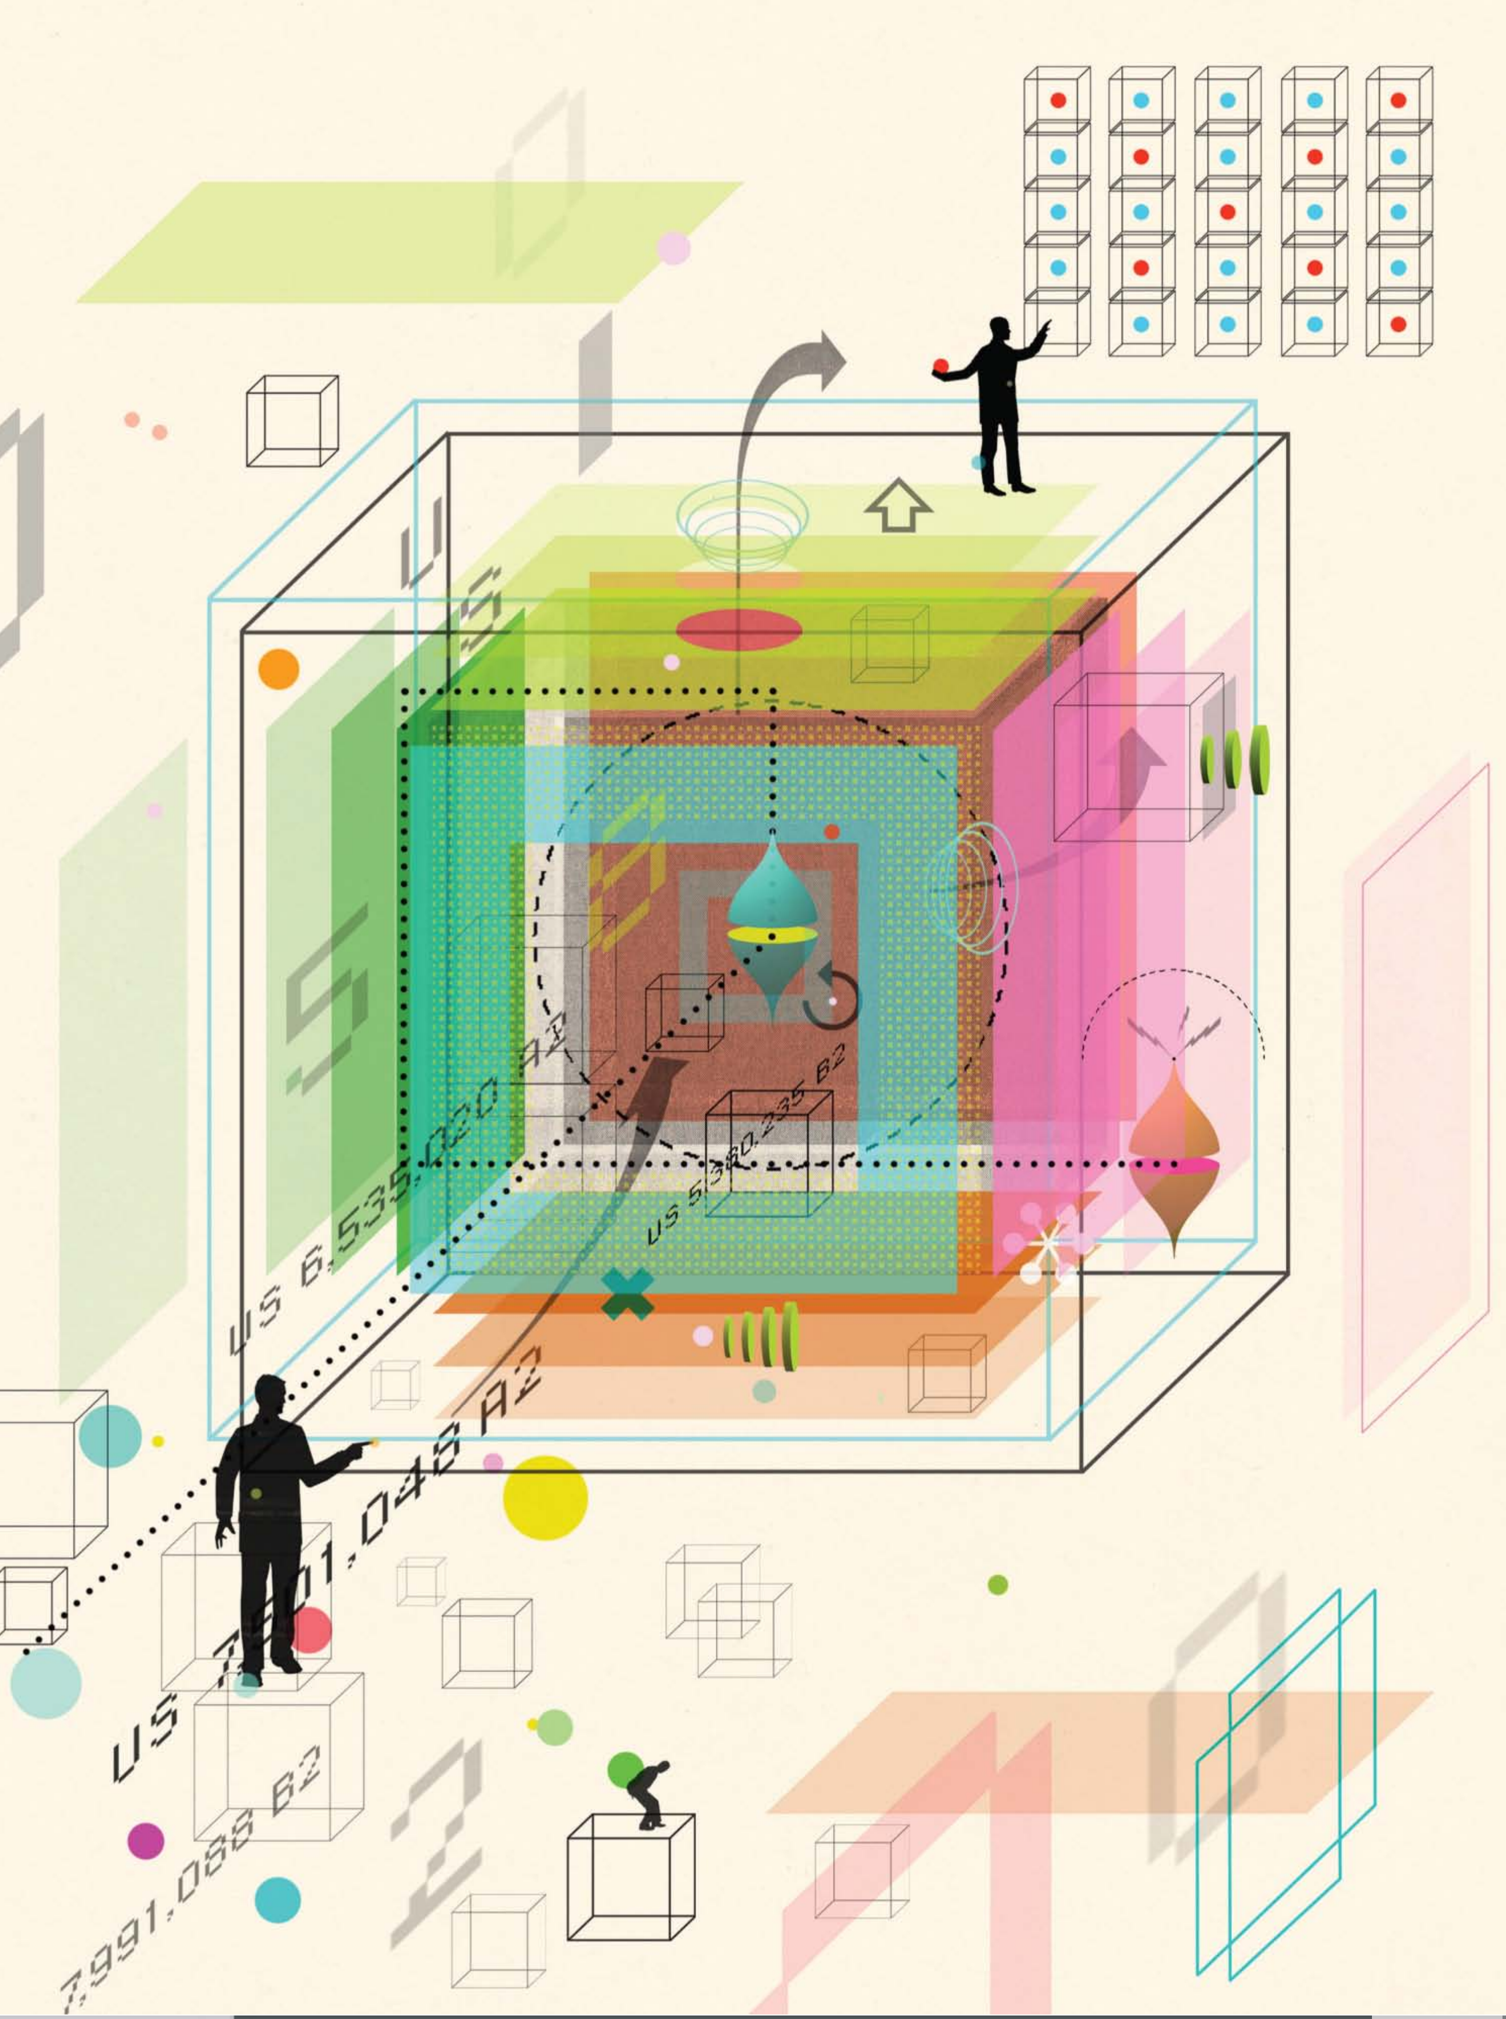
US 6,535,020 A2
7,991,088 B2

The LAW

HOW A SILICON VALLEY START-UP AIMS TO OVERHAUL INTELLECTUAL PROPERTY LITIGATION

by Tam Harbert

ILLUSTRATIONS BY MARK ALLEN MILLER

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

IN A LOW-RISE BUILDING IN MENLO PARK, CALIF., JUST upstairs from a Mexican restaurant and a nail salon, a Stanford University spin-off is crunching data in ways that could shake the foundations of the legal profession.

Here, a small group of patent lawyers and computer scientists is applying the latest in machine learning and natural-language processing to reams of documents related to intellectual property lawsuits. The result is a massive statistical database on IP litigation like nothing the world has seen before. Which attorney has the best track record in defending against semiconductor-related infringement claims? Has a particular judge ruled on cases involving patent trolls, and if so, what was the outcome? Which companies tend to go to trial, and which settle out of court? By offering up such information, the database provides corporate lawyers, law firms, and government agencies with hard numbers that will reduce the guesswork, as well as the enormous expense, of patent litigation. In short, the company is building a "law machine," from which comes its name: Lex Machina.

"Law is horribly inefficient," says Mark Lemley, a professor at Stanford Law School, director of the Stanford Program in Law, Science & Technology, and cofounder of the company. "And in some ways, it is inefficient by design." After all, lawyers get paid by the hour, so inefficiency is rewarded, says Lemley. And some are rewarded richly: Top lawyers charge north of US $1000 per hour.

Lex Machina is in the vanguard of an emerging field known as legal analytics, according to Daniel Martin Katz, an associate professor of law at Michigan State University who writes the blog Computational Legal Studies and advocates overhauling the practice of law through technology. Practitioners of legal analytics statistically parse the practice of law in search of data that can be used to augment, or in some cases replace, the more qualitative judgment of human lawyers.

"There's been a quiet transition going on in the legal world," Katz says. And that transition will shake up the legal profession. "Human reasoning, at least some part of it, is going to be replaced by machine-based prediction." If Lex Machina succeeds, there will eventually be fewer frivolous lawsuits–and maybe fewer lawyers too.

"WE'RE THE MONEYBALL OF IP LITIGATION," SAYS Josh Becker, Lex Machina's CEO. Bespectacled and unassuming, he looks more like a professor than a savvy Silicon Valley player. With law and MBA degrees from Stanford, he served as press secretary for a Pennsylvania congresswoman, worked at the Internet start-up EarthWeb/DICE and at Netscape, and founded a venture capital firm before turning his attention to Lex Machina.

Becker is also a huge baseball fan who's made a careful study of Michael Lewis's 2004 best-selling book, *Moneyball*, which tells how Oakland Athletics general manager Billy Beane used nontraditional statistics, called sabermetrics, to make judgments about players and game strategy. Looking at the numbers, for instance, Beane determined that two popular baseball plays–bunting and stealing bases–don't contribute significantly to a team's chance of winning, so he banned them. Such decisions based on sabermetrics contributed to the Athletics' making it to the playoffs in 2002 and 2003.

That approach is basically what Lex Machina is doing for law. But while baseball is known for its reliance on statistics, Becker says, law has long been a profession that is more art than science. "Some people went to law school to avoid data," he quips.

Lex Machina aims to change that. According to the company, its database covers more than 130 000 U.S. IP and antitrust cases dating back to the year 2000, including information on more than 1400 judges, 340 000 litigants, 100 000 attorneys, and 30 000 law firms. At present, it covers only the United States, but it may eventually include international patent cases as well.

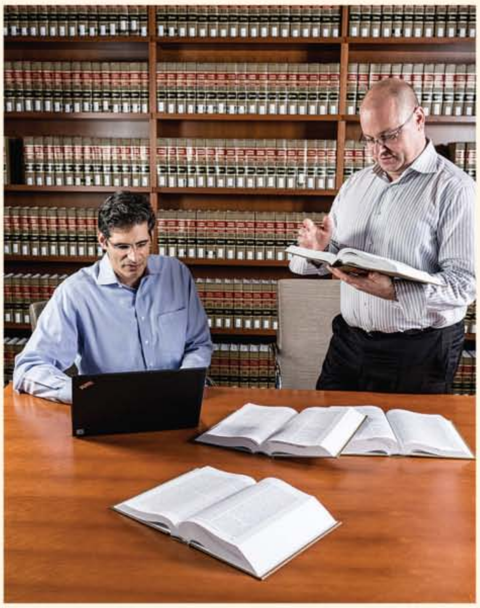

THE LAW MACHINISTS: Lex Machina CEO Josh Becker [left] and cofounder Mark Lemley aim to make intellectual property law more efficient.

With patent wars raging in every sector of the technology industry, IP litigation is big business and getting bigger all the time. The number of patent lawsuits in the United States skyrocketed between 2010 and 2012, from around 3200 filings to more than 5000, according to the Administrative Office of the United States Courts. One recent study, by James Bessen and Michael J. Meurer of the Boston University School of Law, found that defending against "nonpracticing entities"–sometimes called patent trolls–cost companies some $29 billion in 2011. Corporations are looking for a way to cut those costs.

Traditionally, a company that's been sued for patent infringement, or is thinking of suing because its own IP has been infringed, will hire top attorneys to pursue its case. Yet

PHOTOGRAPH BY Jonathan Sprague

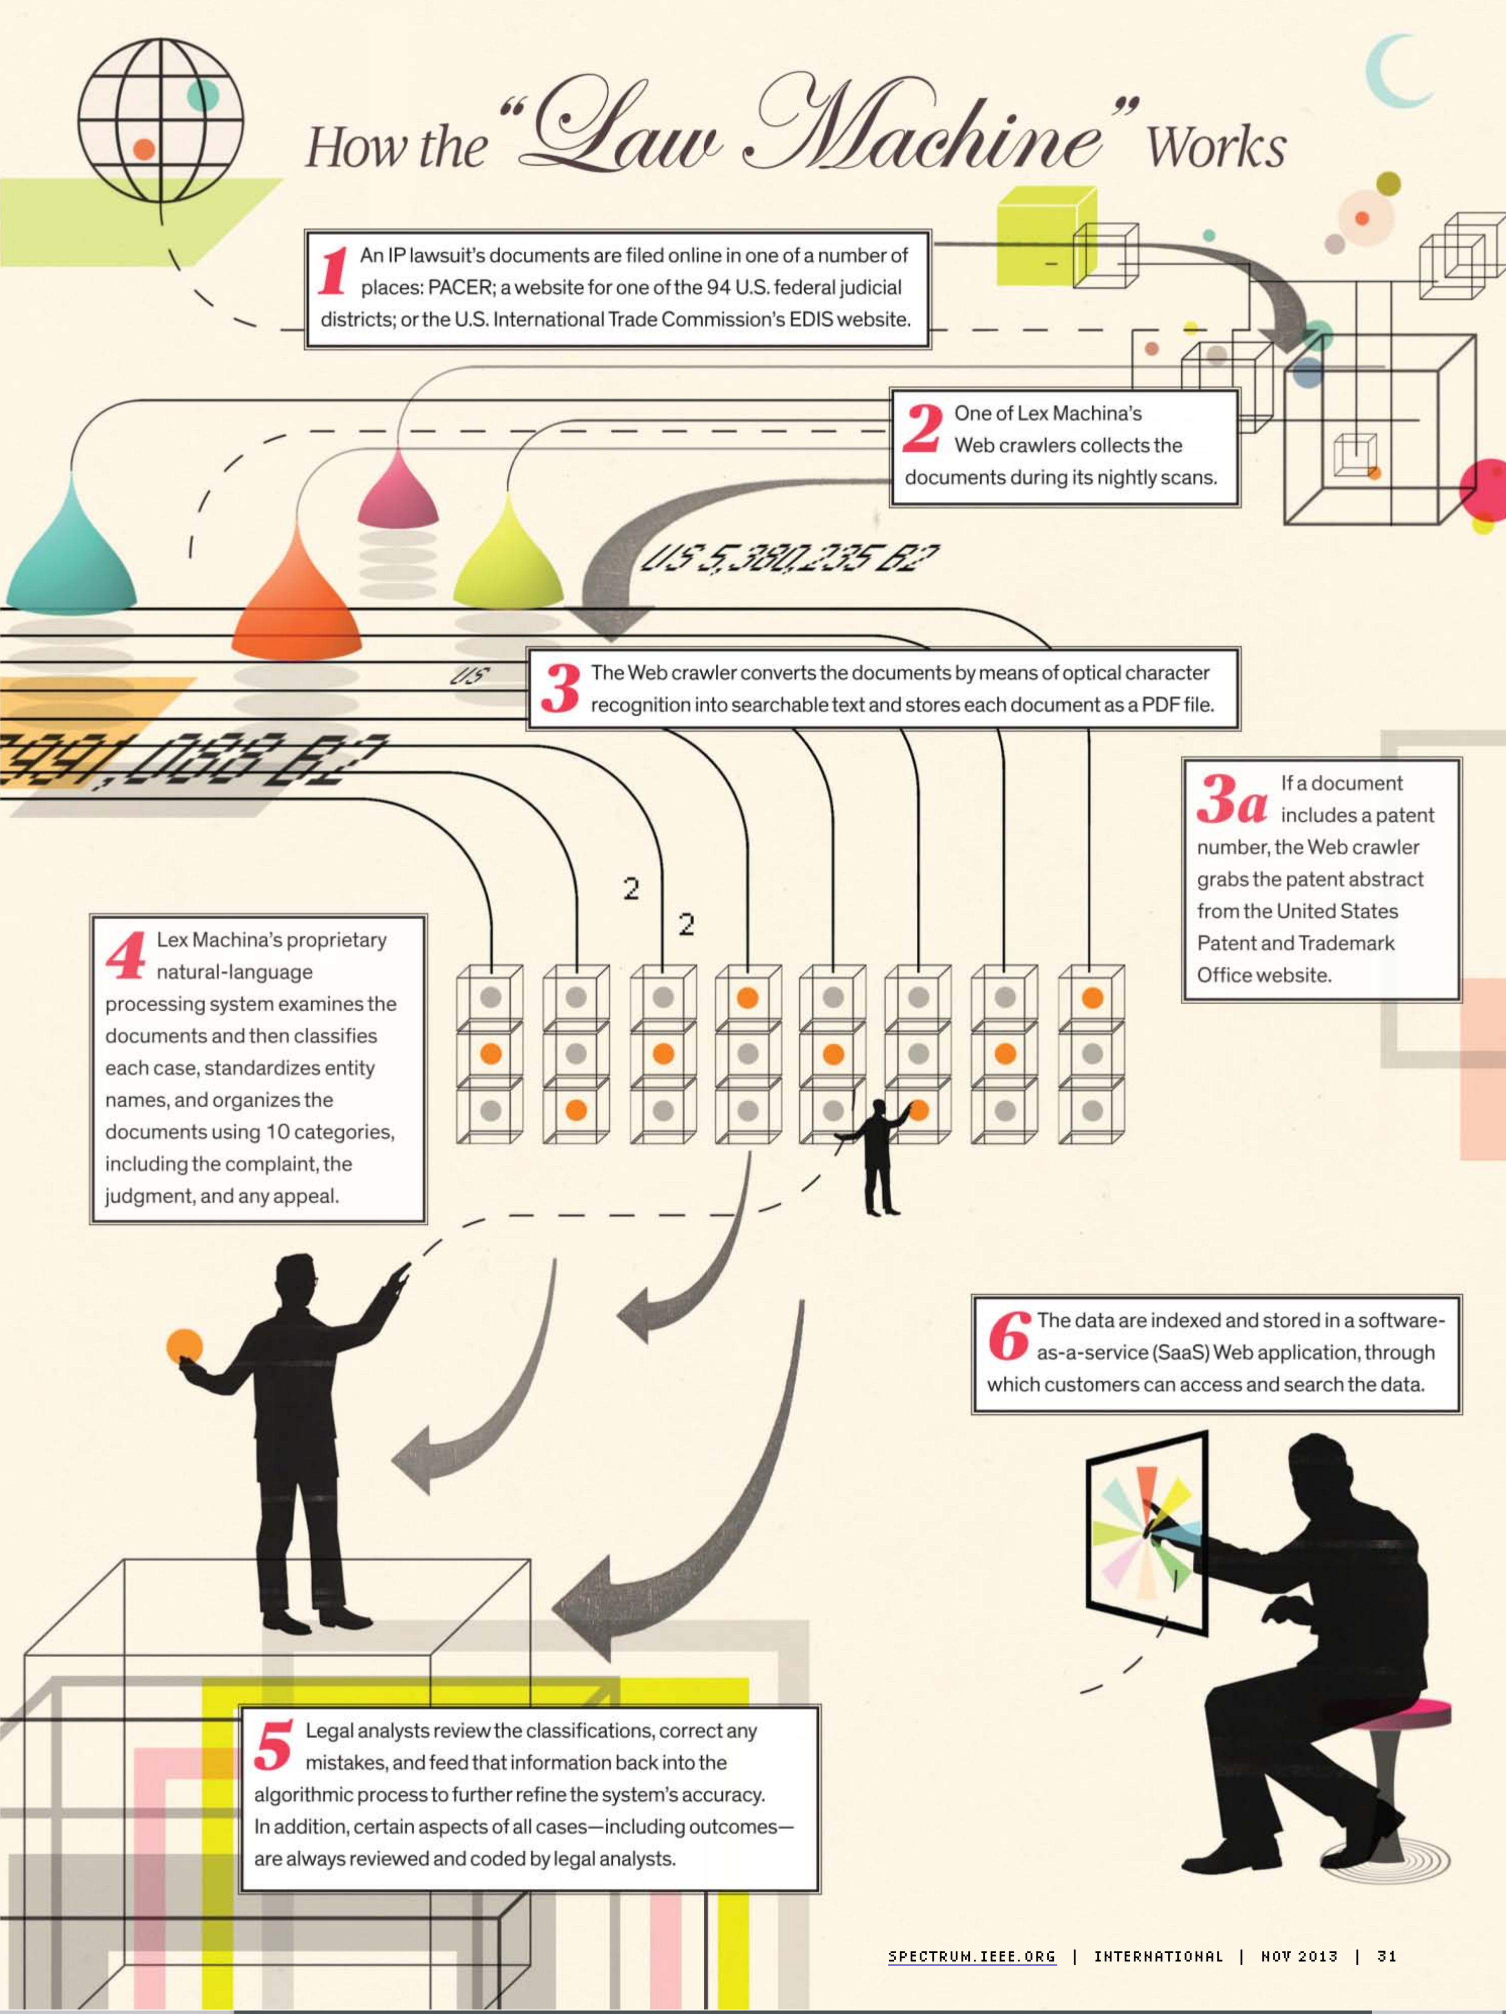

IEEE SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

the process of deciding whether, how, and even where to file such a suit is often driven by gut instinct rather than facts. Even the best patent attorney has seen maybe tens of cases that are similar to the client's. "Humans are limited. People haven't seen 10 000 cases or 100 000 cases–a human can't hold that kind of information," Katz says.

But Lex Machina can. For an annual subscription fee of around $50 000, its customers get access to 13 years of U.S. IP litigation. Just like the sabermetrics described in *Moneyball*, Lex Machina's database can aid in the formulation of broad strategy as well as the selection of players, says Becker. The company's stats reveal, among other things, which attorneys do the best against a particular patent troll, how much time and money it typically takes to fight a troll versus settling out of court, and even which judge you'd want to hear your case. The data might tell a company being sued that its peers have been settling similar lawsuits early, thereby saving money. Even if a company believes it's in the right, says Becker, a prolonged legal battle and "fighting to the death" may not make good business sense.

SO HOW DOES LEX MACHINA do what it does? It started with documents–millions of pages of legal documents that, in theory at least, are available to anyone, free of charge. In practice, though, before Lex Machina came along, there was no easy way to collectively consider that vast body of information. Figuring out how to extract relevant data from countless files and then building a comprehensive database took years of dedicated effort on the part of Lex Machina's small and eclectic team. Among its 18 employees are 6 people with law degrees, 6 with computer science degrees, and 1 who has both.

The company began as an academic research project called the Intellectual Property Litigation Clearinghouse, launched by Lemley in 2006 as a collaboration between Stanford's law school and its computer science department. As Lemley explained during an interview on the sunny terrace of Stanford Law's William H. Neukom Building, "The industry was having all these debates about how to fix the patent system, and none of them were based on actual evidence."

Lemley hoped that a law database would foster decisions based on fact rather than assumption. The tech industry was enthusiastic about the project, as evidenced by more than $3.5 million in donations from companies like Apple, Cisco, and Microsoft, as well as several law firms, the Kauffman Foundation, and Stanford Law School. Lemley recruited Joshua Walker, a cofounder of CodeX: The Stanford Center for Legal Informatics. Walker in turn hired George Gregory, then a Stanford graduate student with expertise in natural-language processing and machine learning.

Several technology developments had come together that made collecting and interpreting the raw data possible. First, the documents were already available online. In the early 2000s, all 94 U.S. federal court districts adopted electronic case-filing systems, which let parties file documents pertaining to lawsuits online and make them available through the courts' websites. Other sources of data included PACER, short for Public Access to Court Electronic Records, which gives the public online access to case and docket information from federal appellate, district, and bankruptcy courts, and the Electronic Document Information System (EDIS) of the U.S. International Trade Commission. (This last source has become increasingly important in recent years, as many companies now file patent infringement claims at the USITC in addition to the courts because USITC administrative judges have the power to bar the importation of infringing products.)

Second, the growth in computer processing power and the drop in server prices had allowed data farms to crunch terabytes of data inexpensively. And third, processes and tools for machine learning and natural-language processing had advanced sufficiently to handle the complexities of legal information. Natural-language processing, also called computational linguistics, involves developing computer algorithms so that machines can understand language. Machine learning, a branch of artificial intelligence, is about constructing systems that can learn from data. The Lex Machina team uses machine-learning techniques to identify specific legal terms and phrases and then builds natural-language processing algorithms to encode the results.

Collecting and coding all that legal data was an overwhelming task. Fortunately, researchers at the Stanford AI Laboratory were eager to take on the challenge. Christopher Manning, a professor of linguistics and computer science at the AI lab, says the project offered an opportunity to extend machine learning beyond just understanding words to understanding phrases | CONTINUED ON PAGE 50

STANFORD LAW SPIN-OFFS

Despite Stanford University's legendary spin-off history–Cisco, Google, and Hewlett-Packard, to name a few, originated in its engineering and computer science departments–Lex Machina was the first to come out of Stanford's law school. Since that happened, in 2009, there have been several other law spin-offs. "I think Lex Machina broke the ice, showing the commercial potential of collaboration between the law, business, and engineering schools," says Clint Korver, a partner at Ulu Ventures, which has invested in Lex Machina and two other law start-ups.

Like Lex Machina, many of the newcomers rely on artificial intelligence and big-data technologies:

LawGives
was founded in 2011 and has developed a platform to match people needing legal help with an appropriate attorney. Users can get legal advice for free, then pay fixed fees for common legal services; attorneys pay a fee to be listed. The platform uses machine learning to automatically interpret the questions that clients enter into the system so it can match them with the right type of attorney.

SIPX
was founded in 2012 and offers access to copyrighted material. The company is initially targeting the higher education market, where there is a lot of confusion over tracking and managing the copyrights for teaching and training materials, a problem that has worsened with the rise of online education.

Ravel Law
was also founded in 2012. It is developing a legal search technology that uses sophisticated data visualization to speed up legal search and add context and clarity to the complex Web information.

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

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

WHITER BRIGHTS WITH LASERS

BMW's laser headlights will turn night into day for its upcoming all-electric supercar

BY LAWRENCE ULRICH

PHOTO-ILLUSTRATION BY Smalldog Imageworks

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

I am standing in a twilit underground lab looking at the most powerful car headlight in the world. Its source, a blue-laser diode, is 1000 times as bright as an LED but uses just two-thirds the energy. ¶ I had to surrender my passport to see it, for I'm at BMW's highly secure, steel-and-glass research mecca, the FIZ (for *Forschungs- und Innovationszentrum*, or Research and Innovation Center), an immense complex in Munich replete with workshops, clay modelers' "caves," and a vast wind tunnel.

I'm basking in the glow of three of BMW's brightest lights: project founder Volker Levering, whose laser inspiration—a mental lightbulb, if you will—flashed on during a 2010 Christmas ski trip in the Alps; Stefan Weber, the current program leader; and Helmut Erdl, also among the technology's inventors.

"A person may not be directly aware, but you can instantly feel the difference between good and bad light," says Weber, as he switches on a wall of fluorescent panels. I certainly feel the difference: The panels simulate a sunny day above ground, right up to the 6500-kelvin color temperature that photographers consider natural daylight.

These engineers want more than mere intensity: They want a focused, high-contrast white light that mimics that of the sun. BMW's system will deliver about 5500 to 6000 K—the highest color temperature that international regulations will allow. That level is much closer than today's headlights to the cool, blue end of the scale, which helps drivers pick out objects and eases eye fatigue.

Until the dawn of semiconductor lighting, the whitest brights any headlights could manage were high-intensity discharge (HID, also known as xenon) lamps, which BMW introduced on the 1991 7 Series coupe. HID light—still an optional upgrade on many 2014 cars—is between 2800 and 3500 lumens and above 4000 K, but it's weaker, yellower, and less energy efficient than either LED or laser light.

The incandescent bulb invented by Thomas Edison in 1879 has shown remarkable staying power, not least in automobiles. The first LED headlights shone from a car a mere six years ago, when Lexus introduced them on its LS 600h L sedan. Soon, the technological torch may pass to lasers. Laser light will debut in Europe in the 2014 BMW i8, a plug-in hybrid sports car that promises 2.45 liters per 100 kilometers (about 94 miles per gallon) and a 4.4-second surge from 0 to 100 kilometers per hour (or 0 to 60 miles per hour in 4.3 seconds).

The reign of the LED headlight may be ending even before it gets properly under way, says Shuji Nakamura. He should know: Nakamura invented both the blue laser and the blue LED, which made possible the whole world of powerful white solid-state lights. And his Silicon Valley start-up, Soraa, is developing laser systems to complement its LED lighting technologies.

"The laser, we believe, is the next generation of lighting, even for general applications" such as homes, businesses, and a variety of displays, Nakamura says.

BMW'S ENGINEERS LEAD THE WAY to the *Nachtfahr*, or night-driving simulator. A cutaway interior of a 5-Series sedan, with instrument panel, seats, and steering wheel, faces a darkened simulator screen.

On a nearby workbench rests the culmination of their 2.5-year project: a pair of laser light prototypes, looking like the lethal toys of a James Bond villain. Erdl demonstrates their power, dangling an incense stick into the barely visible beam. The stick immediately begins burning, filling the lab with a cathedral's scent.

PREVIOUS PAGES: ORIGINAL PHOTO, BMW; THIS PAGE, BMW (2)

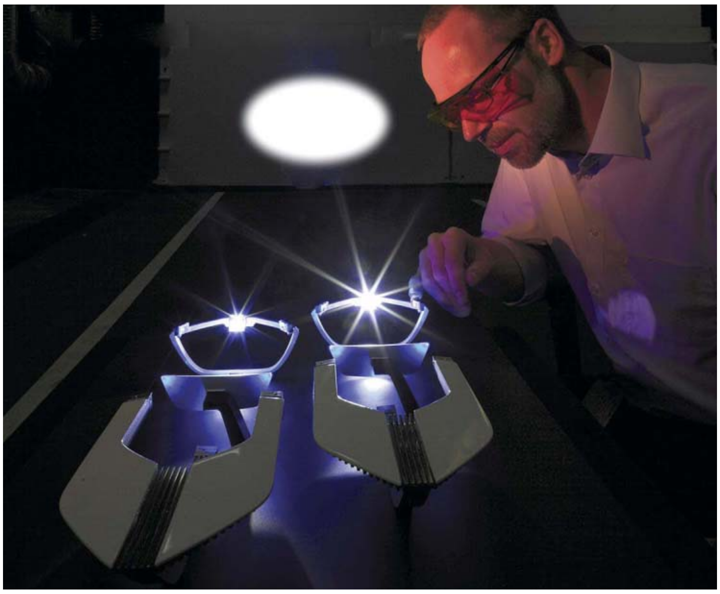

BLUE-LASER LIGHT from a tiny diode proceeds through a phosphor, which converts some of it into a wavelength in the yellow part of the spectrum. The resulting mix is a white beam that can be focused very tightly, even though it's not a laser.

But have no fear: The lasers will be safely contained, with no chance of bouncing their fierce rays off unfortunate retinas, even in the event of a collision. That's because the BMW lamps turn the intense blue beam into a tightly concentrated but nonlaser–and therefore eye-friendly–cone of white light.

The production version will have up to four Class 4 blue-laser diodes. Collimating lenses will direct their beams onto a phosphorus plate that will convert the laser beams to white light, which will bounce off secondary optics and reflect onto the road. To show me how it works, here in the night simulator lab, Erdl dips a phosphorus wafer into a blue laser beam. The wafer blocks some laser photons and lets others stream through. Among the blocked photons, some stimulate–or "pump"–the phosphorus atoms to emit yellow light. The mix of blue light streaming through and yellow emitted from within produces brilliant white light. It's the basic technique used, at lower intensities, in most white-light-emitting LEDs.

Indeed, LEDs can approach that 6000-K, white-light nirvana. And for flooding an area with diffuse white light, LEDs are great, Nakamura says.

But he and other experts agree that lasers are much better and more efficient at precisely directing light onto a distant spot. That's exactly what's required in automotive lighting (and also in movie theater projectors, which will soon incorporate laser-based lighting systems).

"Wherever you want directional, flexible applications of light, the industry is moving like a freight train in the laser's direction," says Paul Rudy, general manager of the laser division of Soraa. "It's simply the best way to direct light through a complex optical system."

The reasons are clear. At just 10 square micrometers, the laser's active light-emitting area is 1/10 000th the size of a 1-square-millimeter LED. That makes it much easier for the laser to focus and project light exactly where it's needed. The tiny chip itself, with its vastly higher current density, makes it easy to produce dazzling light without multiple, relatively bulky LEDs. And the source is already tightly focused and concentrated, so the tiny movements of the chip can be translated into the large motions of the beam. With an LED, on the other hand, "the light tends to just fan out everywhere, and it's very hard to focus it optically," Rudy says.

"With a laser, you combine the high-brightness benefits of a lamp with the reliability, long life, and efficiency of an LED," he adds. That estimated life, by the way, is 30 000 hours. The lamps could easily outlive the automobile.

LASERS ALSO BEAT LEDS WHERE IT MATTERS MOST: efficiency. It's true that LEDs are more efficient at turning electricity into light, though laser efficiency is rapidly catching up. But for overall system efficiency, it's no contest: LEDs are nowhere near as good at getting the light to where you want it to go. That intense laser, for example, can be beamed into a fiber-optic strand and lose only 10 to 20 percent of its initial energy, as opposed to what an LED could lose–up to 90 percent, experts say. A pair of the old halogen headlights drew about 120 watts from a car's battery; a couple of today's best LED headlights draw roughly 40 W. Laser light's usage is projected to drop below 30 W.

The rise of the technology can be traced back to Nakamura himself, whose first blue lasers hit the market in 2005. Low-power versions were driving Blu-ray players and PlayStations, with higher-power versions reserved for uses such as industrial welding. Engineers kept cranking up the power output in shorter-wavelength, indium gallium nitride chips, until 1-W blue lasers became available. It was a development ripe for new applications.

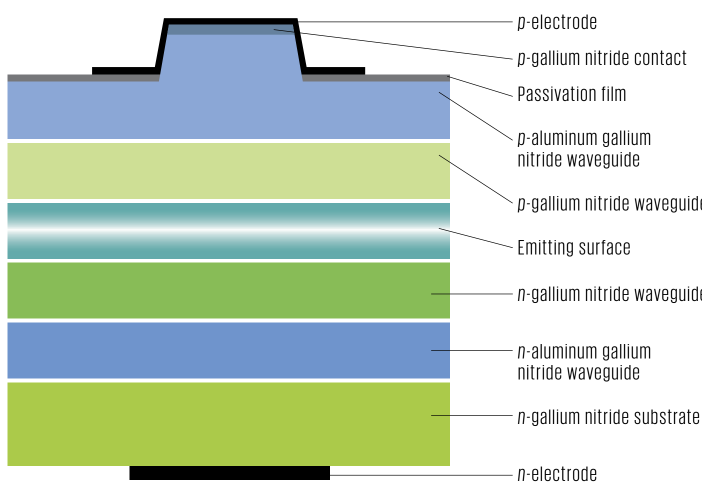

GALLIUM NITRIDE LASERS set up a resonant light wave not along their length, as do classic lasers, but crosswise, so that the light comes out from the edge. The compound semiconductor, perfected by Shuji Nakamura, has a large enough bandgap to produce the energetic waves in the blue spectrum.

DEER-SPOTTING DIODES

Even good old LEDs can do wonders with the help of an algorithm or two

NIGHT FALLS HARD at BMW's Aschheim test track, a flat, lonely outpost in the Bavarian boonies east of Munich.

"The first 30 times I saw this place, it was night," mutters BMW engineer Dominik Schneider.

I'm riding shotgun in a BMW 7 Series sedan with Schneider and his engineering partner Sebastian Orecher. As smoky clouds snuff the twilight, we drop Schneider off on a grass-choked stretch that might pass for a country road anywhere.

"Now I will turn and hunt for Dominik," Orecher says with a smile, accelerating down the rain-dampened track.

And as the BMW's headlights cut through the murk, it happens: Two secondary strobes of white light shoot far beyond the headlights' usual range, directing my eyes to targets beyond. One beam pinpoints a deer on the left shoulder and lights it up, antlers and all. The second tracks Schneider walking about on the opposite shoulder, as if it were a spotlight following a performer around a stage.

Even before the spotlights trigger, a vivid image of the road ahead—including the deer and the BMW engineer—fills the driver's instrument panel, replacing views of the digital speedometer and gauges. And as we near the light's targets, a yellow leaping-deer icon flashes on the windshield's head-up display, acoustic warnings turn increasingly acute, and the BMW precharges its brakes to prepare for a panic stop.

Welcome to a new frontier in safety and the latest step toward fully automated cars.

BMW calls its tracking beam system the Dynamic Light Spot. And as with most advances in car technology, this intelligent lighting is fused with other active systems, including the large-animal detection feature of the 2014 7 Series's night-vision unit. Such systems are critical to self-driving cars that can sense, see, and make the right decision in any situation, from a critter's crossing a road to a drunk driver's running a red light.

The "deer" in our test was actually a plastic statue, heated to the animal's body temperature. The BMW's new dual-camera night-vision system detects temperature differences of as little as 0.1 °C, making the "deer" appear crystal clear on screen, along with every sun-warmed rock, tree, and even the clouds overhead.

Real-life Bambis cause 1 million collisions, 200 deaths, 29 000 injuries, and US $1 billion in property damage each year in the United States alone. BMW's system, developed with Stockholm-based safety giant Autoliv, can identify all major deer species, along with moose, cows, horses, camels, and yes, wild boars. Those diverse animals proved tough to identify, with their different sizes, profiles, and movements. So Autoliv spent six years "training" its digital imaging system by studying wild and captive animals on five continents.

The system fuses views from a windshield-mounted near-infrared camera and a bumper-mounted long-wavelength infrared camera. These stereoscopic cameras create an image grid from which the electronic controller extracts and classifies about 150 discrete features, among them temperature, shape, velocity, and directional vectors. Each feature is classified and ranked to throw out false targets and assess the level of risk.

Small electric motors in each unit swivel the spotlight to track its target. BMW managed to fit the LED-based spotlights, which integrate with a vehicle's controller-area network bus, into the 7 Series's standard fog-lamp housing.

The system illuminates people with a steady beam and animals with a pulsing light. The idea is to avoid having an animal freeze when blinded by a harsh beam, although BMW's research suggests that animals run or halt at random no matter how they're illuminated.

The night-vision display can illuminate objects up to 500 meters down the road, accurately identifying pedestrians at about 100 meters and animals at 50 to 200 meters, depending on size and species.

BMW's studies showed that without the system, drivers spotted pedestrians only at 29 meters. The BMW takes 40 meters to stop from 100 kilometers per hour (62 miles per hour), even on slippery pavement, but with the new system, drivers recognized and reacted to people at 63 meters, providing plenty of time for a full stop. The unit also tracks oncoming cars, but it won't trigger the spotlight if there's a chance of blinding another driver in its glare. U.S. regulators want further assurances that Dynamic Light Spot won't distract other drivers. Many industry analysts suspect approval is only a matter of time.

The system should work even better when it's integrated with BMW's new laser light technology, engineers say. While LEDs can be made into spotlights, lasers would work even better because they can produce a highly concentrated beam using fewer optical components.

During some earlier road testing in Romania, Orecher says, BMW's system gave him an unexpected benefit: It spotted a warm-bodied police car lurking in a radar trap, helping him to avoid a speeding ticket. That's the kind of night owl many drivers would want looking over their shoulders. *—L.U.*

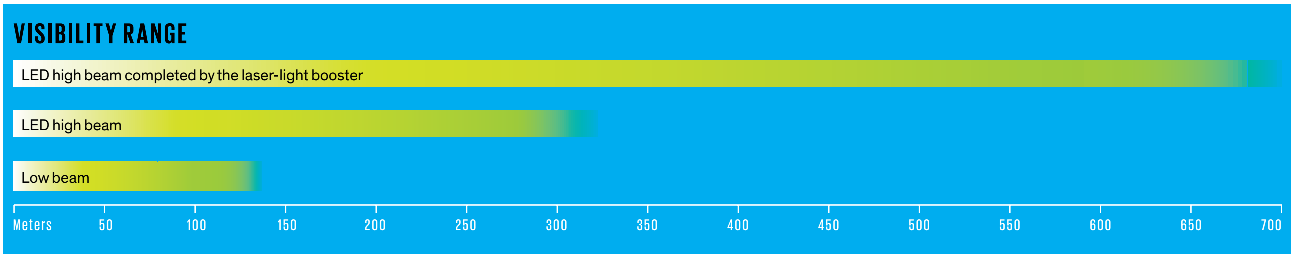

FARSIGHTED AND EFFICIENT laser headlights are able to reach beyond the range of LEDs (and older technologies, such as incandescent bulbs) because they're intrinsically bright and, more important, they can be directed to their targets with remarkable efficiency.

Digital display makers took a page from LED development, realizing they could also pump the stronger blue lasers into a phosphor to create brilliant light. The result was laser-based "lamp-free" displays, which are now being used in office, school, and home-theater projectors. Rudy calls semiconductor lighting a classic example of a maturing industry, with chip development on a steep upward curve and costs falling rapidly.

"The confluence of adjacent fields really brought laser to the forefront," he says. "And auto lighting is one awesome example of the application." Because BMW's blue-laser diodes measure just 10 micrometers, about one-tenth the length of rival LEDs, you can put them anywhere in an automobile and transmit their output light via fiber optics. That frees designers to create dramatically new headlight forms–or "eyes"–to heighten cars' personalities and also to save space and redistribute weight.

The field is blowing open, Rudy and Nakamura say. Beyond cars, projectors, and displays, expect to see uses in cellphone "pico" projector displays and future head-mounted systems, such as those now used in Google Glass. Lasers may even end up lighting our homes, offices, stadiums–you name it.

For general lighting, what makes the lasers attractive is that they can be packed much more densely on a chip than LEDs can. Laser-based lights would not only be more energy efficient on a dollars-per-lumen basis but also more flexible, able to work as spotlights or floodlights at the flick of a switch. If costs continue to fall, Rudy says, laser lights could make the leap to general use in roughly 10 years.

BMW INTENDS TO INTRODUCE the laser-based system on its 2014 i8 plug-in hybrid sports car. As with any plug-in vehicle, the i8 has a particular need to conserve battery electricity for propulsion, as well as for steering assist, entertainment, and heating and cooling. Finding small savings everywhere–even on the order of mere watts–translates directly into more miles of driving range.

Unlike slow-starting, single-brightness HIDs, the laser lights switch on in milliseconds and instantly go to 100 percent illumination. Easy to package in motorized modules, space-saving lasers "offer huge advantages for today's projection systems," Levering says.

Those innovations include BMW's new Dynamic Light Spot, a pair of bumper-mounted spotlights that are separate from the headlights. The spotlights are linked to thermal-imaging cameras and a head-up driver display that brightly illuminates large animals and pedestrians far beyond headlight range. And BMW and other luxury brands are now introducing headlights that direct a "cone of darkness" toward oncoming cars, allowing drivers to keep bright beams on to maintain peak visibility [see sidebar, "Deer-Spotting Diodes"].

The cars of the future, BMW engineers envision, will automatically adapt to roadway environments by switching among hundreds of lighting programs, if not more. Laser light, Levering says, will dovetail nicely with those developments.

A few hurdles remain. In coming months, BMW's team must still downsize the system and ensure durability outside the lab. As with LEDs, system cooling is a challenge. Directing airflow over the lamps is an obvious approach, supplemented by motorized fans or conductive materials. BMW is already trekking to Death Valley, Calif., and other locations to test performance in extremis.

The preproduction i8 tucks four laser lights below its stretched hood, but the showroom version will use the laser system for the high beams only. The i8's low-beam LEDs will emit 50 candelas per square millimeter at their source, compared with 580 cd/mm² for the laser high beams. That pushes visibility as far as regulations allow. The United States sets a lower limit on total overall light output than Europe does, but automakers generally design systems that can easily be tweaked to pass muster internationally.

As a further safety measure, the i8's high beams will operate only above roughly 40 km/h, to preclude the possibility of someone's staring into a static light. Photodiodes will monitor the high-power pumped lasers, switching them off if they fail during a collision or even just as the result of wear and tear. And the system automatically switches to low beams when oncoming cars are detected, as with current adaptive units.

My laser-guided tour of the FIZ complete, I head to BMW World, a dramatically billowing structure nestled among the BMW headquarters, one of its factories, and a tentlike swimming stadium built for the 1972 Olympics. There, throngs of car lovers ogle the i8 and its concept versions, snapping photos like mad, oohing and ahhing over all the carbon-fiber goodness.

It's all great stuff. But if and when you see your first BMW i8, do Herren Erdl, Levering, and Weber a favor: Spare a little love for those headlights. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/laserheadlight1113

IMAGES: BMW (2); DATA SOURCE [TOP]: BMW

THE SPY WHO SKYPED ME

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

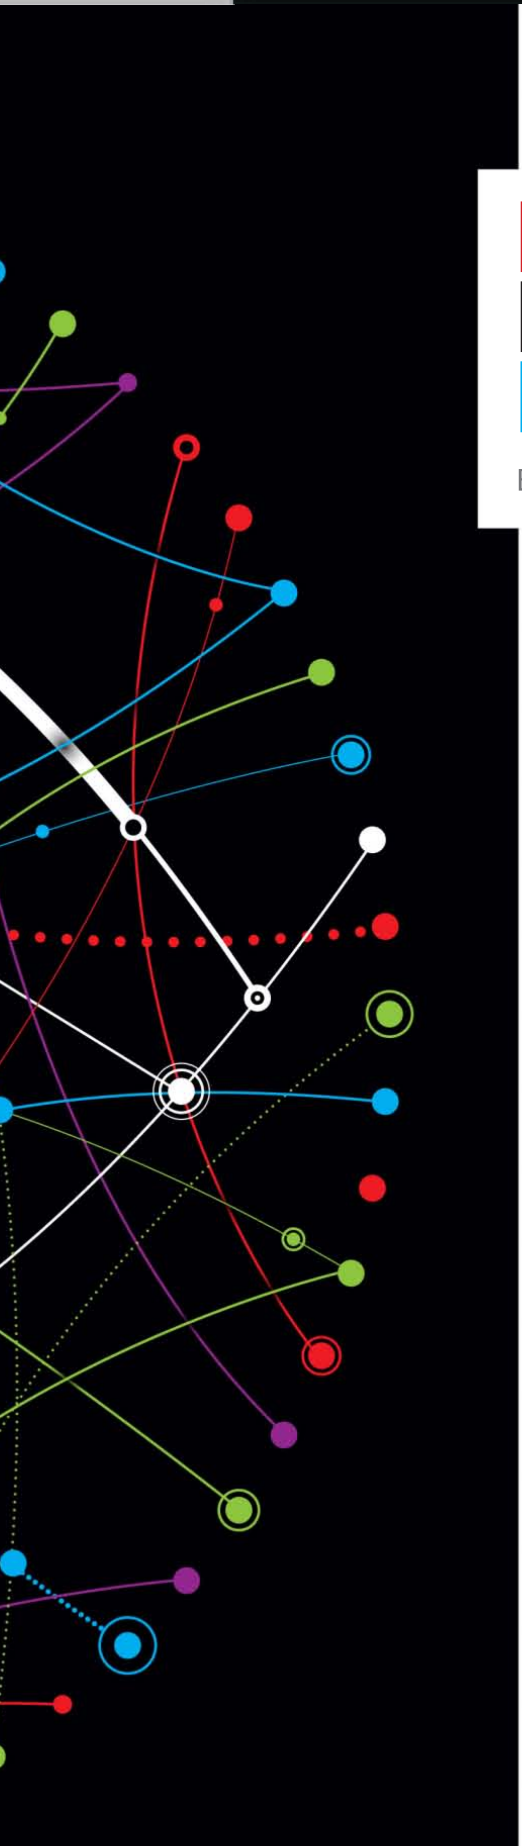

ILLUSTRATIONS BY CARL DE TORRES

FOUR WAYS TO SMUGGLE MESSAGES THROUGH INTERNET SERVICES

BY WOJCIECH MAZURCZYK, KRZYSZTOF SZCZYPIORSKI & JÓZEF LUBACZ

THEIR NEIGHBORS THOUGHT THEY WERE JUST ORDINARY U.S. residents, but secretly they were spies, sent by Russia's Foreign Intelligence Service to gather information on U.S. policies and programs. For years they thwarted detection, partly by hiding secret correspondence in seemingly innocent pictures posted on public websites. They encoded and decoded the dispatches using custom-made software. ¶ But the scheme wasn't as covert as the spies had assumed. Eventually, investigators from the U.S. Department of Justice tracked down the altered images, which helped build a case against the Russians. In June 2010, federal agents arrested 10 of them, who admitted to being secret agents a few weeks later. ¶ The act of concealing data in plain sight is known as steganography. Since antiquity, clandestine couriers have used hundreds of steganographic techniques, including invisible ink, shrunken text, and strategically placed tattoos. Picture steganography—one of the Russian spies' primary tactics—dates back to about the early 1990s. That they used such an old-school strategy is odd, particularly because doctored images can be detected and used as evidence. ¶ A more modern approach, known as network steganography, leaves almost no trail [see "Vice Over IP," *IEEE Spectrum*, February 2010]. Rather than embed confidential information in data files, such as JPEGs or MP3s, network steganography programs hide communication in otherwise innocuous Internet traffic. And because these programs use short-lived delivery channels—a Voice over Internet Protocol (VoIP) connection, for example—the hidden exchanges are much harder to detect. ¶ Network security experts have invented all of the dozens of publicly documented network steganography techniques. But this doesn't mean that criminals, hackers, and spies—as well as persecuted citizens wanting to evade government censorship or journalists wanting to conceal sources—aren't using these or similar tactics. They probably are, but nobody has tools that are effective enough to detect these techniques. In fact, had the Russian spies used newer steganography methods, they might not have been exposed so handily. ¶ As members of the Network Security Group at Warsaw University of Technology, in Poland, we study new ways to disguise data in order to help security experts design better detection software for those cases when steganography is used for nefarious purposes. As communication technologies evolve, we and other steganographers must develop ever more advanced steganography techniques. ¶ About a decade ago, state-of-the-art programs primarily manipulated the Internet Protocol. Today, however, the most sophisticated methods target specific Internet services, such as search tools, social networks, and file-transfer systems. To illustrate the range of possibilities, we present four steganographic techniques we've recently developed, each of which exploits a common use of the Internet.

BITTORRENT

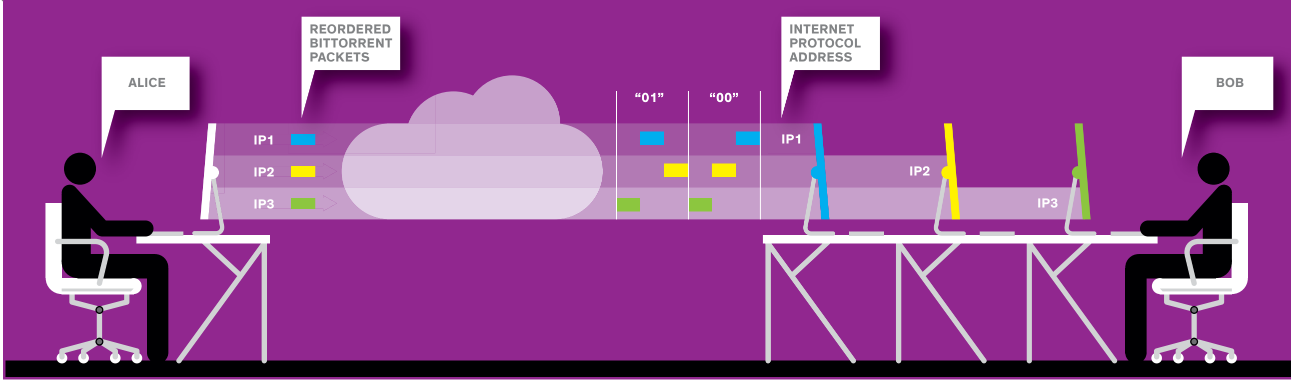

What better place to hide secrets than in one of the world's most popular file-sharing systems? The peer-to-peer transfer protocol BitTorrent conveys hundreds of trillions of bits worldwide every second. Anyone sniffing for criminal correspondence on its networks would have better luck finding that proverbial needle in a haystack.

Our group developed **StegTorrent** for encoding classified information in BitTorrent transactions. This method takes advantage of the fact that a BitTorrent user often shares a data file (or pieces of the file) with many recipients at once.

So let's say a user—we'll call her Alice—wants to send a hidden message to her conspirator, Bob. First, Bob needs to have previously established control over a group of distributed computers that all run a BitTorrent application. These are most likely computers that Bob owns or has co-opted to do his bidding. Both he and Alice need to know how many computers are in this group and what their IP addresses are.

For simplicity's sake, let's say Bob controls a group of just two computers. To initiate a transaction, he commands the computers to each request a file from Alice. In a typical BitTorrent transfer, Alice's program would transmit the data packets in random order, and Bob's computers would stitch them back together based on the instructions they contain. Using StegTorrent, however, Alice can reorder the packets to encode a specific bit sequence.

For example, if she sends a packet to computer 1 and then to computer 2, that sequence might designate the binary number 1. But if she sends a packet to computer 2 first, Bob's StegTorrent program would read the signal as binary number 0. To prevent scrambling due to packet losses or delays, StegTorrent modifies the time stamp on each packet so that Bob can decipher the exact order in which Alice sent them. Our experiments showed that using six IP addresses, Alice can relay up to 270 secret bits per second—enough bandwidth for a simple text conversation—without distorting the transfers or attracting suspicion.

SKYPE

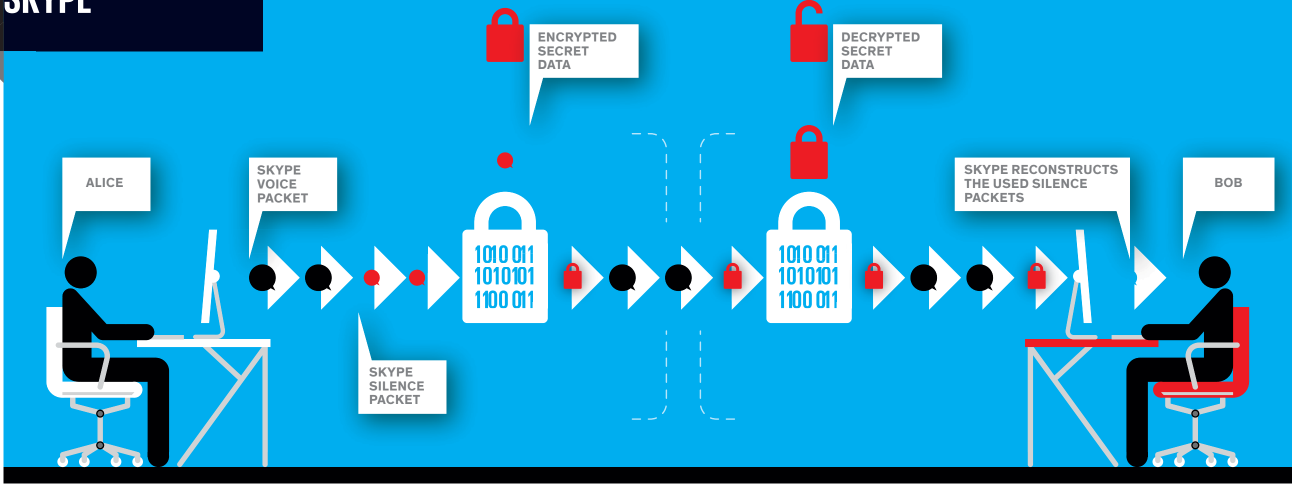

Skype, Microsoft's proprietary VoIP service, is particularly easy to exploit for steganographic purposes because of the way the software packages audio data. While the caller—Alice—is talking, Skype stuffs the data into transmission packets. But unlike many other VoIP apps, Skype continues to generate audio packets when Alice is silent. This improves the quality of the call and helps the data clear security firewalls.

But the outgoing silence packets also present an opportunity to smuggle secret information. These packets are easy to recognize because they're much smaller—about half the number of bits—than the packets containing Alice's voice.

We've developed a steganography program that allows Alice to identify the small-size packets and replace their contents with encrypted secret data. We call this program **SkyDe**, shorthand for Skype Hide. For a covert transaction to take place, the recipient of Alice's call—Bob—also needs to have SkyDe installed on his computer. The software intercepts Alice's transmission, grabs some of the small packets while letting all of the big ones pass through, and then reassembles the secret message.

Meanwhile, Alice and Bob chat away as if nothing unusual were transpiring. Bob's Skype application assumes the filched packets have simply been lost. Skype then fills the gap left by each lost packet, most likely by reconstructing its contents based on the contents of the neighboring packets. As a result, the missing silence packets sound just like all the other silence packets surrounding them.

Our experiments show that up to 30 percent of Alice's silence packets can transport clandestine cargo without causing a noticeable change in call quality. This means that Alice could send Bob up to about 2 kilobits per second of secret data—roughly 100 pages of text in 4 minutes—without arousing the suspicion of anyone monitoring their call.

A version of this article appeared online in September.

GOOGLE SUGGEST

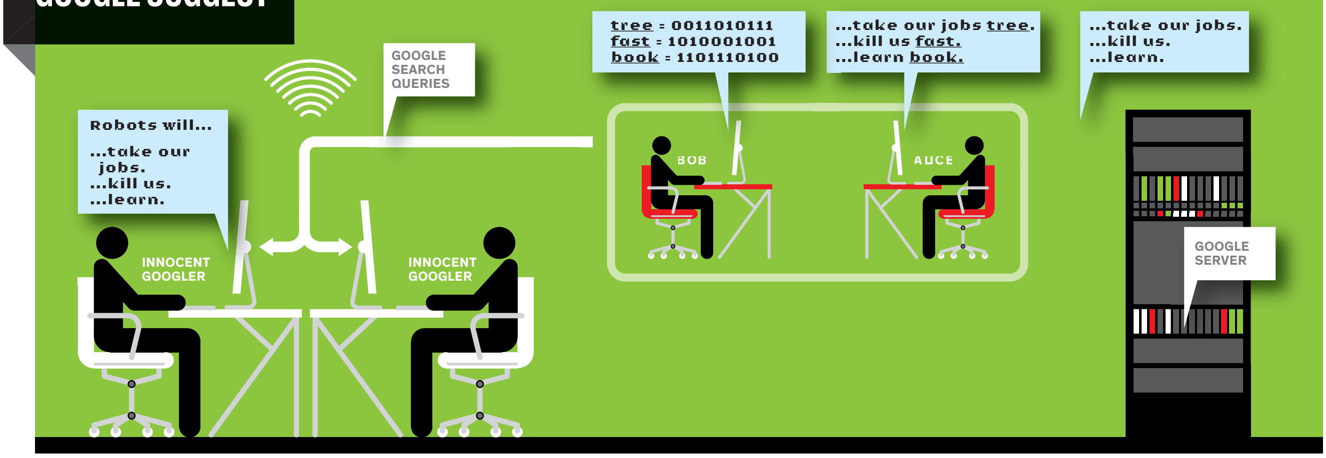

Alice can also conceal her messages to Bob—and the fact the two conspirators are communicating at all—simply by having him perform a series of innocent-looking Google searches. Our **StegSuggest** steganography program targets the feature Google Suggest, which lists the 10 most popular search phrases given a string of letters a user has entered in Google's search box.

Here's how it works: For Alice to send Bob a hidden note, she must first infect his computer with StegSuggest malware so that she can monitor the traffic exchanged between Google's servers and Bob's browser. This can be done using basic hacker tools. Then, when Bob types in a random search term, say, "Robots will…," Alice intercepts the data traveling from Google to Bob. Using StegSuggest, she adds a unique word to the end of each of the 10 phrases Google suggests. The software chooses these additions from a list of 4096 common English words, so the new phrases aren't likely to be too bizarre. For example, if Google suggests the phrase "Robots will take our jobs," Alice might add "tree," so that the phrase reads "Robots will take our jobs tree." Odd, yes, but probably not worthy of alarm.

Bob's StegSuggest program then extracts each added word and converts it into a 10-bit sequence using a previously shared lookup table. (Each of the 1024 possible bit sequences corresponds to four different words, making the code more difficult to crack.) Alice can thus transmit 100 secret bits each time Bob types a new term into his Google search box.

To send data faster, Alice could hijack the searches of several innocent googlers in a crowded hot spot, such as an Internet café or a college dormitory. In this scenario, both she and Bob would intercept the googlers' traffic. Alice would insert the coded words into Google's suggested phrases, and Bob would extract and decode them. He would pass on only the original phrases to the googlers—who would never suspect they had just facilitated a secret exchange.

WI-FI NETWORKS

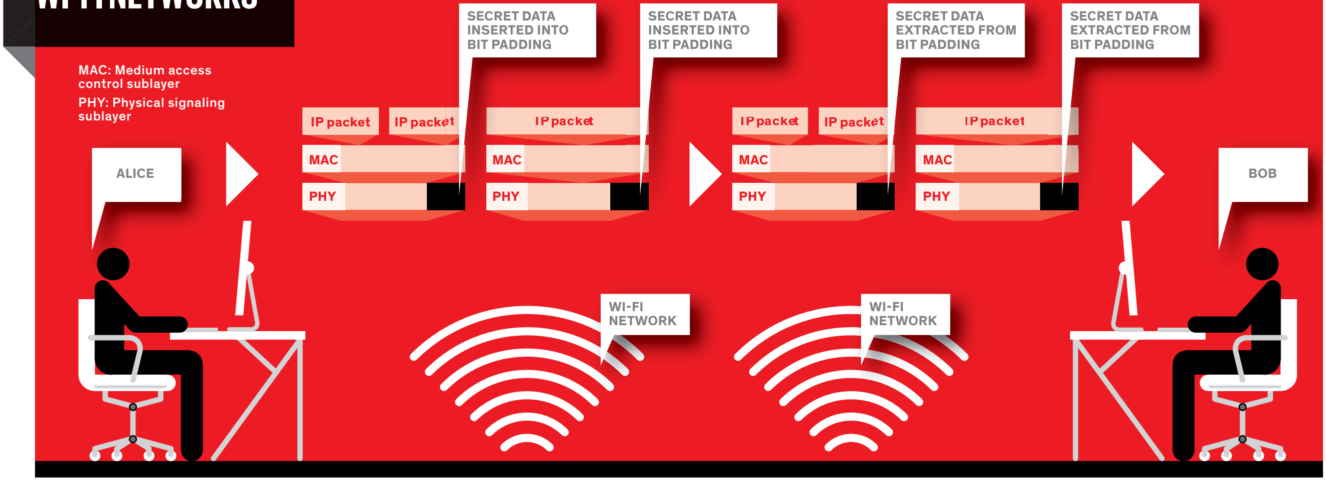

Now let's say Alice wants to secretly send video in addition to documents or text messages. In this case, she might opt to smuggle the stream in a very average-looking wireless transmission.

But not just any wireless network will do. Alice must use a network that relies on the data-encoding technique known as orthogonal frequency-division multiplexing (OFDM). Wireless standards that employ this scheme are some of the most popular, including certain versions of IEEE 802.11, used in Wi-Fi networks.

To understand how to hide data in OFDM signals, you must first know something about how OFDM works. This transmission scheme divvies up a digital payload among several small-bandwidth carriers of different frequencies. These narrowband carriers are more resilient to atmospheric degradation than a single wideband wave, allowing data to pass to receivers with higher fidelity. OFDM carefully selects carriers and divides the bits into groups of set length, known as symbols, to minimize interference.

In reality, though, a digital payload rarely divides perfectly into a collection of symbols; there will usually be some symbols left with too few bits. So OFDM transmitters add extra throwaway bits to these symbols until they conform to the standard size.

Because this "bit padding" is meaningless, Alice can replace it with secret data without compromising the original data transmission. We call this steganographic method Wireless Padding, or **WiPad**. Because bit padding is abundant in OFDM transmissions, Alice can send hidden data to Bob at a pretty good clip. A single connection on a typical Wi-Fi network in a school or coffee shop, for instance, could support up to 2 megabits per second—fast enough for Alice to secretly stream standard-definition video to Bob.

Post your comments at http://spectrum.ieee.org/steganography1113

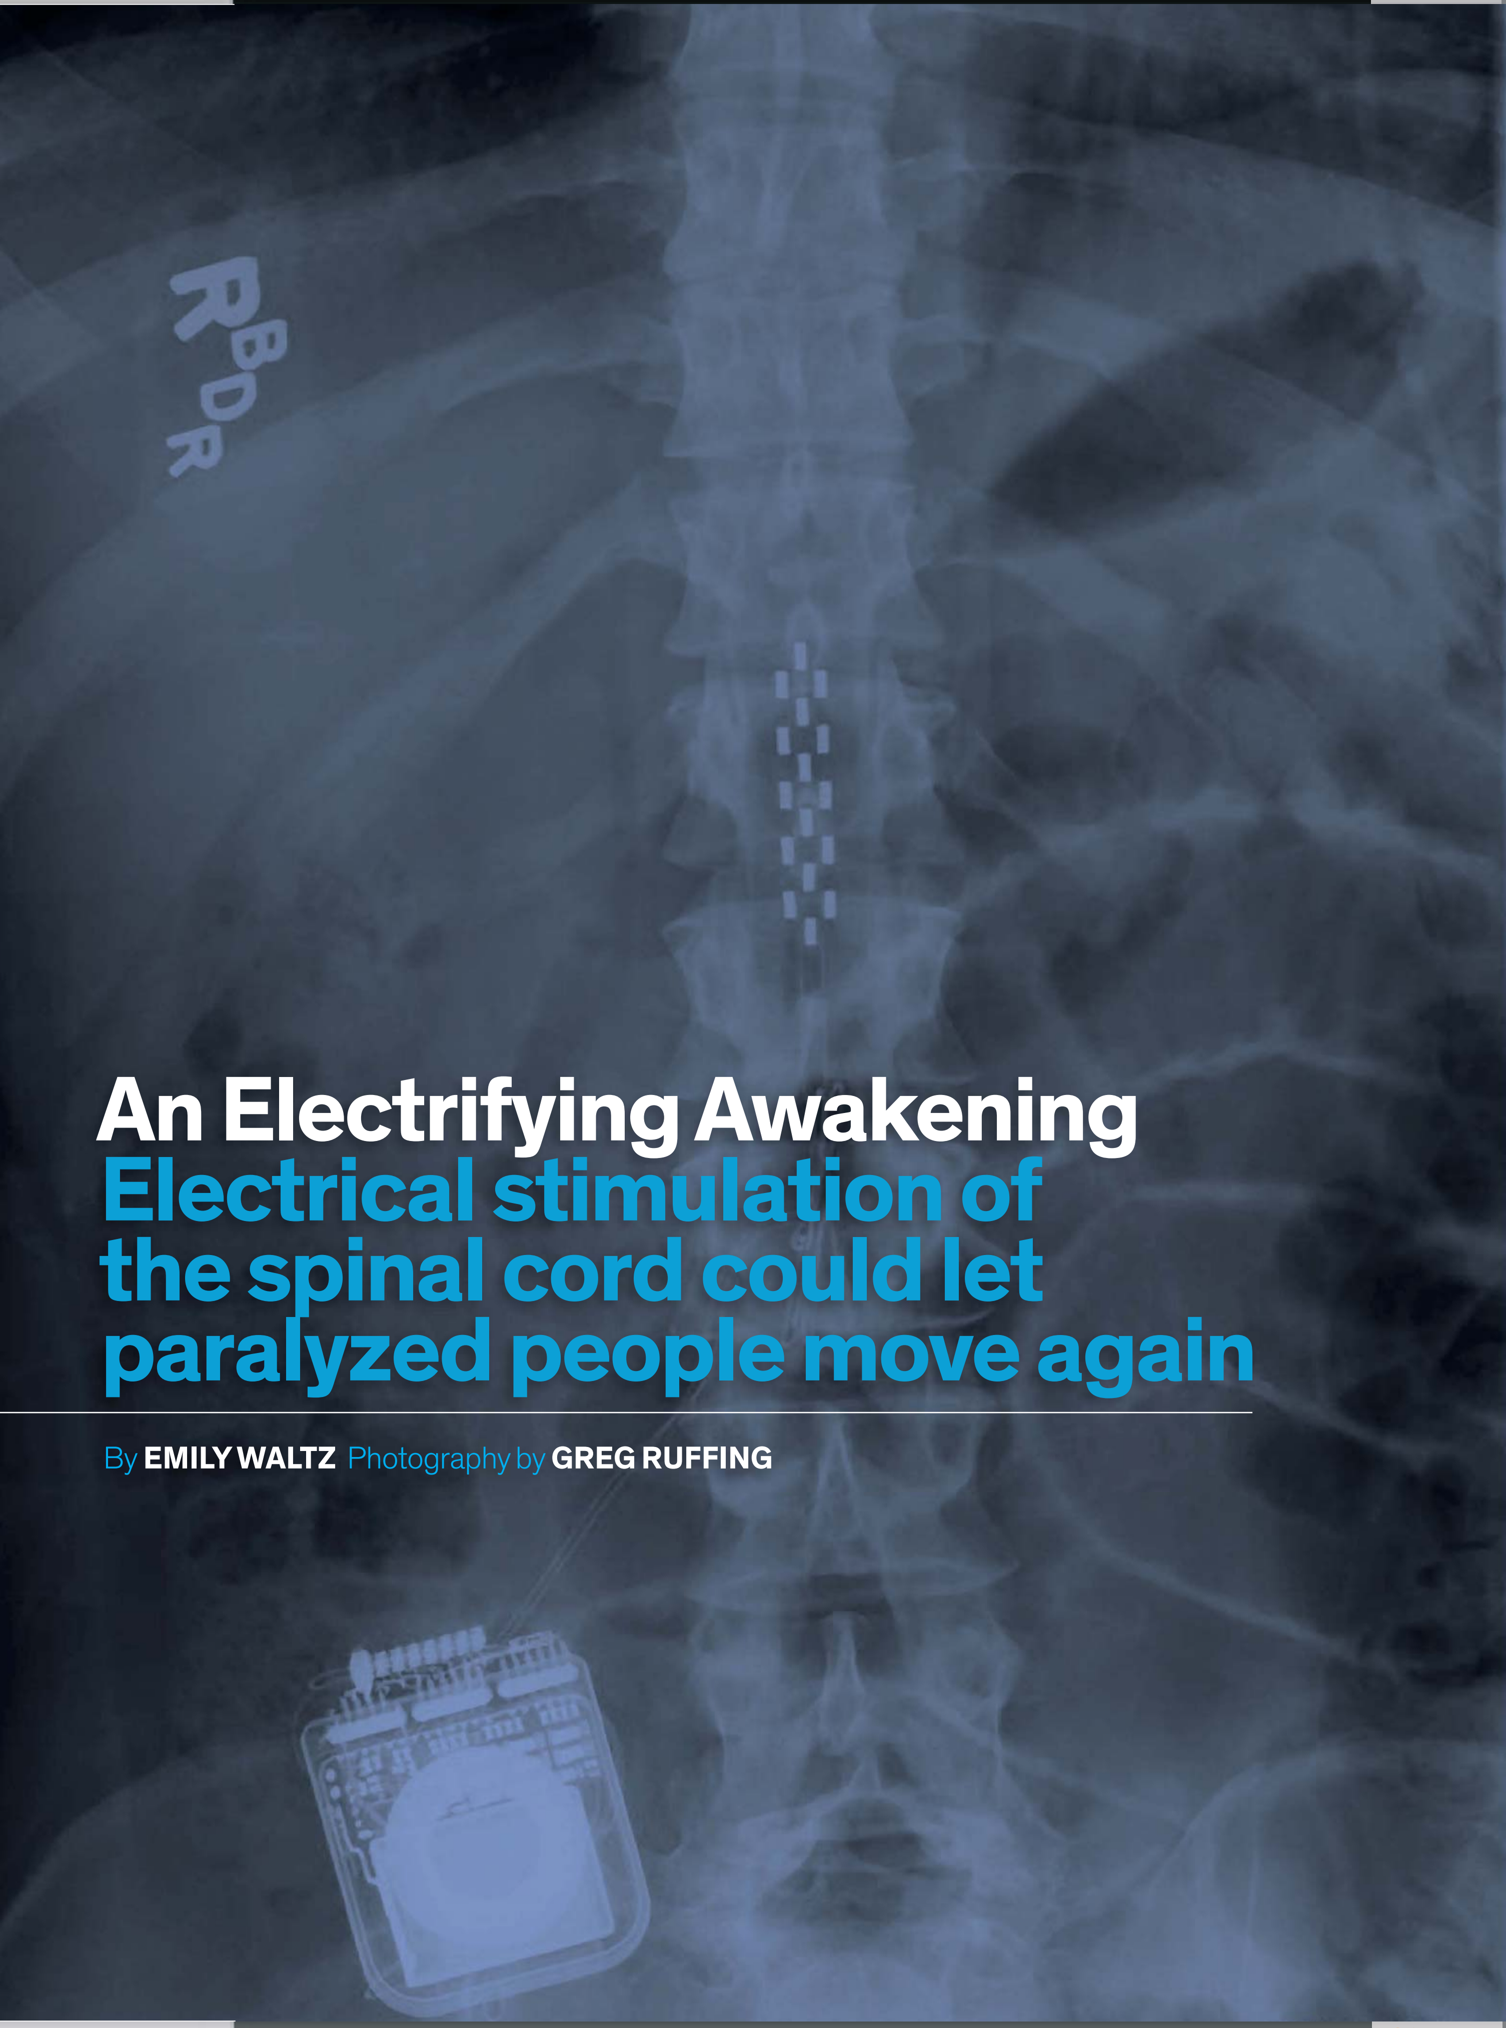

An Electrifying Awakening

Electrical stimulation of the spinal cord could let paralyzed people move again

By **EMILY WALTZ** Photography by **GREG RUFFING**

PATIENT NO. 4

Dustin Shillcox volunteered to have electrodes and a pulse generator implanted in his spine.

DUSTIN SHILLCOX FULLY EMBRACED the vast landscape of his native Wyoming. He loved snowmobiling, waterskiing, and riding four-wheelers near his hometown of Green River. But on 26 August 2010, when he was 26 years old, that active lifestyle was ripped away. While Shillcox was driving a work van back to the family store, a tire blew out, flipping the vehicle over the median and ejecting Shillcox, who wasn't wearing a seat belt. He broke his back, sternum, elbow, and four ribs, and his lungs collapsed.

Through his five months of hospitalization, Shillcox's family remained hopeful. His parents lived out of a camper they'd parked outside the Salt Lake City hospital where he was being treated so they could visit him daily. His sister, Ashley Mullaney, implored friends and family on her blog to pray for a miracle. She delighted in one of her first postaccident communications with her brother: He wrote "beer" on a piece of paper. But as Shillcox's infections cleared and his bones healed, it became obvious that he was paralyzed from the chest down. He had control of his arms, but his legs were useless.

At first, going out in public in his wheelchair was difficult, Shillcox says, and getting together with friends was awkward. There was always a staircase or a restroom or a vehicle to negotiate, which required a friend to carry him. "They were more than happy to help. The problem was my own self-confidence," he says.

A few months after being discharged from the hospital, in May 2011, Shillcox saw a news report announcing that researchers had for the first time enabled a paralyzed person to stand on his own. Neuroscientist Susan Harkema at the University of Louisville, in Kentucky, used electrical stimulation to "awaken" the man's lower spinal cord, and on the first day of the experiments he stood up, able to support all of his weight with just some minor assistance to stay balanced. The stimulation also enabled the subject, 23-year-old Rob Summers, to voluntarily move his legs in other ways. Later, he regained some control of his bladder, bowel, and sexual functions, even when the electrodes were turned off.

The breakthrough, published in *The Lancet*, shocked doctors who had previously tried electrically stimulating the spinal nerves of experimental animals and people with spinal-cord injuries. In decades of research, they had come nowhere near this level of success. "This had never been shown before–ever," says Grégoire Courtine, who heads a lab focused on spinal-cord repair at the Swiss Federal Institute of Technology in Lausanne and was not involved with the project. "Rob's is a pioneer recovery. And what was surprising to me was that his was better than what we've seen in rats. It was really exciting for me to see."

The report brought renewed hope for people living with paralysis. The prognosis is normally grim for someone like Shillcox, who has a "motor complete" spinal-cord injury. That level of damage usually results in a total loss of function below the injury site.

Teams of scientists have been working on transplanting stem cells for neural repair and modifying the spinal cord in other ways to encourage it to grow new neurons, but these long-term approaches remain mostly in the lab. Harkema's breakthrough, however, produced a real human success story and gives hope to paralyzed people everywhere. It presents a viable means of regaining bowel, bladder, and sexual functions, and maybe–just maybe–points the way toward treatments that could give paralyzed people the ability to walk again.

But Harkema's first experiment involved only one patient, and many researchers wondered whether the improvement they saw in Summers was an anomaly. "The next big question was, Will you ever see these things in more than one subject?" says neurobiologist V. Reggie Edgerton of the University of California, Los Angeles, a collaborator in the Louisville experiments.

The U.S. Food and Drug Administration (FDA) had given Harkema the go-ahead to try the technique in four more paralyzed people. Shillcox put his name in the pool the night he saw the news report. He was selected, and in July 2012 he packed his wheelchair into his retrofitted Dodge Journey and drove himself from Green River to Louisville to begin 18 months of experiments.

The circuitry of the lower spinal cord is impressively sophisticated. Neuroscientists believe that the brain merely provides high-level commands for major functions, like walking. Then the dense neural bundles in the lower spinal cord take over the details of coordinating the muscles, allowing the brain to focus on other things. That division of labor is what lets you navigate a party and focus on the conversation rather than on your steps. After a spinal-cord injury, damage prevents the high-level signal from the brain from reaching the neurons below. Yet those neural bundles remain intact and are just waiting to receive a signal to start the muscles working. Stimulating the lower spinal cord with electrodes can awaken that circuitry and get it functioning, astonishingly, without instructions from the brain.

It has been known since the mid-1970s that direct stimulation of the spinal cord can actually induce the legs to move as if they were taking steps, without any input from the brain. Edgerton and other

PREVIOUS PAGES, X-RAY: SUSAN HARKEMA/THE LANCET/ELSEVIER

PATIENT NO. 3: Drew Meas practices standing independently during an electrical stimulation session at the Frazier Rehab Institute, in Louisville, Ky. Neuroscientist Susan Harkema [standing at right] varies the electrical signals sent to Meas's spinal cord to control different muscle groups.

researchers have demonstrated the concept definitively in paralyzed cats, rats, and a few humans. But in most of these demonstrations, researchers were blasting a large amount of electrical current into the body to force the muscles to move. "Everyone, including us, was hung up on the idea that you have to stimulate at this high level to induce the movement," says Edgerton. What they missed was that the stimulation was essentially overwhelming the neurons in the lower spinal cord and was actually interfering with their ability to process sensory information that can help the body move on its own.

The neurons in the spinal cord don't only receive signals from the brain; they also process sensory feedback from the body as the muscles move and balance shifts. The importance of that sensory feedback gradually emerged with some animal experiments Edgerton reported in *Nature Neuroscience* in 2009. The study suggested that sensory input could actually control the motor commands produced by the spinal cord.

Harkema, a former student of Edgerton's, ran with that concept. In her experiments with Summers, she stimulated his spinal cord just enough to wake it up and then let the sensory input do its thing. "It's like putting a hearing aid on the spinal cord," says Edgerton. "We've changed the physiological properties of the neural network so that now it can 'hear' the sensory information much better and can learn what to do with it."

Harkema's group uses an off-the-shelf neurostimulation system—made by Minneapolis-based Medtronic—that's FDA approved for pain management. The system's array of 16 electrodes is surgically implanted in the epidural space next to the outermost protective layer of the spinal cord. The array is then connected to a pulse generator (which resembles a pacemaker) that's implanted nearby. Finally, the pulse generator receives a wireless signal from a programming device outside the body.

The array spans approximately six spinal-cord segments, the ones generally responsible for movement in the lower half of the body. By placing the electrodes over them, the researchers can generate a response in the corresponding muscle groups. Electrode 5, for example, is located near a segment of the spinal cord that controls hip muscles. Electrode 10 is located at the bottom of the array, over the segment that controls the lower leg.

Each of the array's 16 electrodes can be set to act as a cathode or an anode or be completely shut off. Stimulation intensities can range from 0 to 10.5 volts with pulses sent at frequencies ranging from 2 to 100 hertz, although the researchers usually don't go beyond 45 Hz. Picking the right combination of electrodes and stimulation parameters to generate a simple response in a single muscle is relatively straightforward. But generating a complex behavior like standing, which involves many muscle groups and

a considerable amount of sensory feedback, is far more difficult. Choosing the right electrode configurations for standing requires both a tremendous amount of intuition and plenty of trial and error. "That's the challenge: to create the electrical field that's going to give you the desired behavior," says Harkema.

On a Wednesday in February of this year, Shillcox arrived at the Frazier Rehab Institute in downtown Louisville for one of his first stimulation sessions. The array and pulse generator had been implanted a few weeks before. He wore Nike sneakers and black gym shorts, revealing thin legs atrophied from lack of use.

Shillcox joined Harkema and her team in a large room equipped with custom rehabilitation equipment. He wheeled himself to a three-sided stand Harkema had made out of metal pipes that she'd bolted to a piece of plywood. Researchers taped 14 sensors to Shillcox's legs. Using electromyography (EMG), these sensors would measure the electrical activity produced by his muscles and indicate how Shillcox was responding to the stimulation. Two trainers hoisted Shillcox from his wheelchair onto his feet and into the stand. Then they took their positions to keep him upright—one in front of Shillcox with both hands pushing against his knees and the other behind, steadying his hips. Shillcox held onto the stand with his hands, and a bungee cord supported him from behind.

That day, Harkema planned to test new stimulation configurations to see whether one of them would allow Shillcox to stand on his own. She took a seat in front of a screen displaying the EMG signals while two other researchers helped monitor the data from other screens. To start the session, Harkema called out the electrode settings: "1+, 2+, 3+, 9+, 14+, 12+, 13+, 6+, 7-, 8-, 4-, 10-." This configuration used 12 of the 16 electrodes, 8 of them as anodes (positively charged) and 4 of them as cathodes (negatively charged). Harkema instructed her team to set the pulsation frequency at 30 Hz and the initial intensity at 1 V and to ramp up by a tenth of a volt at a time. "Left independent," a trainer called out when the stimulation reached 1.5 V. Shillcox bore his weight on his left leg without assistance for about 30 seconds.

Harkema jotted in her lab book and instructed the team to turn off electrode 10, the one targeting Shillcox's lower leg. "Going to zero," a researcher called out. He powered down the system, punched in the new electrode configuration without electrode 10, and ramped it up again. At 2.6 V, Shillcox's knees buckled. "It shot me out," Shillcox said. The electrodes hadn't sent the signal to the legs to stand straight but had twitched his knees forward instead. The stimulation pattern and parameters weren't quite right.

Harkema tried more configurations, but each time Shillcox felt nothing until Harkema hit a particular voltage threshold, at which point Shillcox's knees would give way. After 75 minutes, on the 10th and last try, Harkema removed the bungee supporting Shillcox from behind. The muscle activity on the EMG monitors skyrocketed. He'd been balancing so perfectly with the bungee cord that he hadn't been getting enough external sensory information to activate his muscles, Harkema concluded, so there had been little input flowing back to the lower spinal cord. She instructed her team to devote the next few sessions to the last electrode pattern of the day, but without the bungee.

The technological limitations of the stimulation system make these trials unnecessarily difficult. Each time Harkema changes the configuration of electrodes, she has to turn off the electric field they generate and start over at 0 V. It's a safety feature of this off-the-shelf stimulator, but it destroys the body's neural momentum. "You can get really close, and you think the person is almost standing independently, and if you could just shift the field a little you would have it. But you can't. You have to go to zero. And then everything starts over," says Harkema. The limitation makes it especially difficult to induce a stepping motion in her patients. "It's a left-to-right problem. If we get the right leg to step, the left is doing nothing," she says.

It doesn't help that there are something like 4.3 x 10^7 possible electrode patterns she can try and that each can be tried with a

BALANCING ACT: Dustin Shillcox works on balance and core muscle control at the Frazier Rehab Institute, in Louisville, Ky. The spinal stimulation system is turned on throughout the exercises.

range of frequencies and voltages. Without an algorithm to help her choose parameters, Harkema must rely on her experience, some limited neural mapping data, and what she sees on her monitors. "I have to look at the EMG data whizzing by and then make decisions about what I can change out of these 4.3 x 10^7 combinations to get it better," says Harkema. She's gotten pretty good at making adjustments, but she acknowledges that no one can fully interpret the nuances of all that EMG data.

To do better, Harkema has enlisted the help of a handful of engineers who say they can build a stimulation system specifically for her research. At the California Institute of Technology, mechanical engineer Joel Burdick is developing a machine-learning algorithm that aims to take some of the guesswork out of choosing stimulation parameters.

The algorithm is based on statistical methods that predict the patient's likely response to stimulation patterns—even those that haven't been tested yet. The prediction part is crucial because there's no way to try out all the options: There are millions of electrode configurations, and every patient is different. And just to make things even more complicated, patients' spinal cords change during the course of the stimulation experiments. "The amount of time it would take to test that space is beyond a patient's lifetime," says Burdick. So the algorithm has to learn quickly. It must apply reasonable stimulation patterns and then use the patient's EMG responses to choose better configurations.

Burdick's team is working with Edgerton's lab at UCLA to test the algorithm on paralyzed rats. The researchers are starting simply, using just a couple of electrodes and trying to maximize the response in a particular muscle. The first step is to make sure the algorithm is making reasonable decisions. The team has also begun a small human pilot study, Burdick says.

Meanwhile, John Naber, an electrical engineer at the University of Louisville, and a team of engineers are developing a stimulation system that would give Harkema independent control of all 16 electrodes in Medtronic's array. The design would allow her to transition from one configuration to the next without shutting off the current. The team is building a new pulse generator using off-the-shelf components, and they've already written the code and roughed out a design. The challenge, Naber says, will be getting it approved by the FDA in a reasonable amount of time. "It's not like a commercial integrated circuit or product, because of the FDA requirements for human implants," Naber says.

The lingering question is whether Medtronic's 16-electrode array is the best one for Harkema's work. It was designed to treat pain, so the current diffuses rather broadly. Yu-Chong Tai, an electrical engineer at Caltech, thinks that an array with smaller electrodes arranged more densely might offer the precise stimulation needed after spinal-cord injury. The prototype he's testing in rats has 27 electrodes arranged over a 2-centimeter-long array. A human version would be similar in size to Medtronic's (about 5 cm long) but would contain hundreds of electrodes. Of course, more electrodes would mean exponentially more configuration options. "If we give them more electrodes, they will need a smart algorithm," says Tai.

Until Naber and Tai's prototypes can be approved by the FDA and Burdick's algorithm can be fine-tuned, the Medtronic system will have to suffice. That may limit what Harkema can achieve when she puts Shillcox and her other research subjects on the stimulator, especially in terms of stepping. Even so, Shillcox has reason to hope that the experiments will boost his quality of life. Rob Summers, Harkema's first subject, says his perspective on life has greatly improved since he regained bladder, bowel, and sexual functions. "This project has given me my freedom back," he says.

Research subjects No. 2 and No. 3 have completed their initial trials. Like Summers, both were able to stand while on the stimulator, as Harkema and her colleagues reported at a Society for Neuroscience meeting in 2012. The researchers have not publicly announced whether other voluntary movement and physiological functions, such as bladder control, have returned for those individuals.

Shillcox—subject No. 4—remains hopeful, but he's trying to keep his expectations realistic. "I don't want to be too optimistic, and I'm trying to be prepared for no results at all," he says. "I hope that whatever they find from this research will at least benefit other people." Shillcox will likely complete his training by the end of the year, and Harkema says she cannot yet publicly reveal their preliminary results. Whatever the medical benefits ultimately prove to be, working with Harkema as a pioneer on an experimental treatment for spinal-cord injury has boosted Shillcox's confidence around others. "I have no problem asking for help now," he says. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/spine1113

CONTINUED FROM PAGE 32 | and contexts. For instance, locating all cases related to *patent infringement* was complicated by the fact that the exact term didn't always appear in a document's text. "It was a matter of translating upwards and understanding the concept of infringement regardless of the words they were using," Manning says.

Another difficulty the researchers encountered was that each court website uses its own variant of electronic filing. They therefore had to design a Web crawler for each one. Once collected, the data then had to be standardized to account for the variations in the way courts file data. And in many instances, the data were just plain wrong; in more than half of all cases, the final decision in the case had been incorrectly coded, according to Walker, who served as Lex Machina's first CEO and is now an attorney at the law firm Simpson Thacher & Bartlett. The way the cases were tagged–as patent, copyright, or trademark infringement, for example–was also often wrong. And there were no tags for certain types of cases, such as those involving trade secrets.

The researchers had to manually sort through, categorize, and correct the data. "There were hundreds of thousands of legal judgments that had to be made" as they sifted through the information, Walker says. In total, it took the team about 100 000 hours.

Once the team had cleaned up the data and understood its many complexities, the engineers designed algorithms to automatically review each document and sort the results. Similar algorithms are used in Web searches, but interpreting legal documents requires more sophistication. "From a science perspective, the baseline [of experience] was zero," says Lex Machina's former chief technology officer, Mihai Surdeanu, who had previously worked on natural-language processing at Yahoo Labs. "There was nobody doing this."

Existing machine-learning techniques don't work very well on litigation data, he says. Even a relatively simple process like normalizing names poses a challenge. The computer has to recognize, for example, that *IBM*, *International Business Machines*, and *IBM Corp.* all refer to the same company. More problematic were law firm names, because firms sometimes change their names when they merge with other firms or when partners join or leave. Even a firm's own attorneys can get the name wrong, Surdeanu says. "One of the firms in our database has 89 different legal spellings," he adds.

The system must also be able to handle complex legal constructs. Unlike baseball, which is a numbers game, the legal world is based on qualitative information, subtle distinctions, and most of all, words. "People argue about the meanings of words and make arguments with paragraphs of text," explains Manning. The machine needs to understand phrases and strings of commonly used legal language as well as context so that it can tell the difference between, for example, the summary judgment document (in which a judge determines which party wins the case or at least certain issues in the case) and a minor procedural filing that simply mentions the summary judgment.

To help parse the legalese, Lex Machina has developed a set of rules–a sort of legal grammar for the machine. The company does this through an iterative process: A legal analyst reviews the algorithms'

results and, if necessary, corrects them, and then an engineer tweaks the algorithms [see illustration, "How the 'Law Machine' Works"].

The result is "this ontology of terms that has been developed over the years" and continues to be refined every night, when the system crawls the Web to collect the latest data, says Owen Byrd, Lex Machina's chief evangelist and general counsel. So far, the company has coded more than 6 million docket entries.

LEX MACHINA'S DATABASE IS AVAILABLE TO ANYONE who can afford its annual fee. The pharmaceutical company Impax Laboratories uses it to guide its strategy for bringing generic drugs to market. Introducing new drugs is a highly structured and litigious process, with specific time limits for each step, so knowing the history of a judge–in particular, how fast cases move through his or her court–is critical, says Huong Nguyen, senior director of IP at Impax (and an adviser to Lex Machina).

Nguyen also uses the database to look up the litigation history of the maker of a brand-name drug to find out which attorneys it uses and how successful they've been in defending their patent positions. And she uses the database to evaluate outside counsel: She can see how many cases they're working on at any given time and which cases they have won or lost, not only for her company but for other clients as well. "We have a stable of outside counsel that we go to constantly," she says. "I want to know what kind of job performance they have across the board."

John Dragseth, a principal at Fish & Richardson (the most active IP litigation firm in the United States, according to *Corporate Counsel* magazine), credits Lex Machina's database with helping him spot meaningful but otherwise hidden trends in IP litigation–and he won't give details. "If you published it, then people on the other side would know," he says.

Typically, Dragseth says, when he reviews cases with clients, "they just nod their heads." But when he starts reeling off statistics like how a particular judge tends to rule in certain types of cases, "they lean forward, put their elbows on the table, and start asking questions," he says. "Clients go crazy about that stuff."

It's not just about the bottom line, though. Lex Machina gives its data, at no charge, to courts, government agencies, academic institutions, and media outlets. That's an important part of fulfilling the mission of Lemley's original research project: improving the legal system.

"In the short term, people will think more intelligently about whether to file suit or when they get sued, how to react: What lawyer should they hire? Should they settle the case early?" says Lemley. Ultimately, he says, people will be able to make informed decisions, not just in individual lawsuits but also in shaping policy and in bringing badly needed reform to the patent system. "My hope is that once everyone has access to the data, some number of lawsuits will go away." ■

POST YOUR COMMENTS at http://spectrum.ieee.org/lexmachina1113

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

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

Joint Institute of Engineering

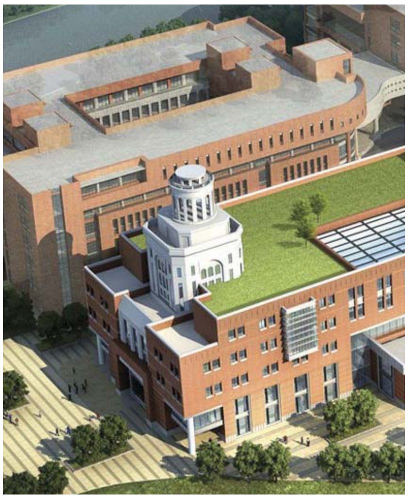

FACULTY POSITIONS AVAILABLE IN ELECTRICAL/COMPUTER ENGINEERING

Sun Yat-sen University & Carnegie Mellon University are partnering to establish the **SYSU-CMU Joint Institute of Engineering (JIE)** to innovate engineering education in China and the world. The mission of the JIE is to nurture a passionate and collaborative global community and network of students, faculty and professionals working toward pushing the field of engineering forward through education and research in China and in the world.

JIE is seeking **full-time faculty** in all areas of electrical and computer engineering (ECE). Candidates should possess a doctoral degree 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 Carnegie Mellon University to establish educational and research collaborations before locating to Guangzhou, China.

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

PLEASE VISIT: **jie.cmu.edu** for details

SHUNDE INTERNATIONAL

Joint Research Institute

RESEARCH STAFF POSITIONS AVAILABLE IN ELECTRICAL/COMPUTER ENGINEERING

SYSU-CMU Shunde International Joint Research Institute (JRI) is located in Shunde, Guangdong. Supported by the provincial government and industry, the JRI aims to bring in and form high-level teams of innovation, research and development, transfer research outcomes into products, develop advanced technology, promote industrial development and facilitate China's transition from labor intensive industries to technology intensive and creative industries.

The JRI is seeking **full-time research faculty** and **research staff** that have an interest in the industrialization of science research, which targets electrical and computer engineering or related areas.

Candidates with industrial experiences are preferred.

Applications should include a full CV, three to five professional references, a statement of research and teaching interests, and copies of up to five research papers.

Please submit the letters of reference and all above materials to the address below.

Application review will continue until the position is filled.

EMAIL APPLICATIONS OR QUESTIONS TO: **sdjri@mail.sysu.edu.cn**

SUN YAT-SEN UNIVERSITY Carnegie Mellon University

HERIOT WATT UNIVERSITY

School of Engineering and Physical Sciences - Institute of Sensors, Signals & Systems

Lecturer/Senior Lecturer/Reader in Robotics and Autonomous Systems and Signal and Image Processing

Reader and Senior Lecturer – Grade 9, £45,941 - £53,233

Lecturer – Grade 7 or 8, £29,541 - £36,298 or £37,382 - £44,607 dependent on qualifications and experience

The School of Engineering and Physical Sciences is one of the most successful engineering faculties in the UK with over £20 million annual research income.

Having recently been awarded over £7 million in the area of networked signal processing, autonomous systems and robotics, we are seeking to appoint an academic at lecturer to reader level in this area.

You will have knowledge and experience in the area of signal and image processing, robotics, computing and autonomous systems and a strong academic track record. Ideally, you will possess a successful track record of obtaining external funding, as well as previous teaching experience. A PhD in Computing, Electrical Engineering or other cognate discipline is essential, along with evidence of vision, drive and determination to succeed.

For application details see our website www.hw.ac.uk/jobs or contact the **Human Resources Office, Heriot-Watt University Edinburgh EH14 4AS** tel **0131-451-3022** (24 hours) email hr@hw.ac.uk quoting Ref EPS/01/13

Closing date: **30 November 2013.**

POSITIVE ABOUT DISABLED PEOPLE

Heriot-Watt University is a Charity registered in Scotland, SC000278

Distinctly Ambitious
www.hw.ac.uk

THE CHINESE UNIVERSITY OF HONG KONG

50 香港中文大學五十周年 50th Anniversary of CUHK

Applications are invited for:-

Faculty of Engineering
Professors / Associate Professors / Assistant Professors

(Ref. 1314/032(370)/2)

The Faculty of Engineering invites applications for several faculty posts at Professor / Associate Professor / Assistant Professor levels with prospect for substantiation in the interdisciplinary area of 'Big Data Analytics', which is a new strategic research initiative supported by the University's Focused Innovations Scheme and will complement current/planned strengths in different Departments under the Faculty. To lead the big data research initiative, senior academics in this area are particularly welcome to apply.

Currently, the Faculty is seeking candidates in the following areas:

- Theoretical, mathematical and algorithmic aspects in large data analytics;
- Large scale software systems and architecture in large data analytics;
- Application areas in large data analytics (including information analytics, network/Web analytics, financial analytics, or bio/medical analytics, etc.).

Applicants should have (i) a PhD degree; and (ii) a strong scholarly record demonstrating potential for teaching and research excellence. The appointees will be expected to (a) teach at both undergraduate and postgraduate levels; (b) develop a significant independent research programme with external funding; and (c) supervise postgraduate students. Appointments will normally be made on contract basis for three years initially, which, subject to performance and mutual agreement, may lead to longer-term appointment or substantiation later. Applications will be accepted until the posts are filled. Further information about the Faculty is available at *http://www.erg.cuhk.edu.hk*.

Salary and Fringe Benefits

Salary will be highly competitive, commensurate with qualifications and experience. The University offers a comprehensive fringe benefit package, including medical care, plus a contract-end gratuity for appointments of two years or longer, and housing benefits for eligible appointees. Further information about the University and the general terms of service for appointments is available at *http://www.per.cuhk.edu.hk*. The terms mentioned herein are for reference only and are subject to revision by the University.

Application Procedure

Please send full resume, copies of academic credentials, a publication list with abstracts of selected published papers, details of courses taught and evaluation results (if any), a research plan, a teaching statement, together with names of three to five referees, to the Dean, Faculty of Engineering by e-mail to *recruit-bda@erg.cuhk.edu.hk*. For enquiries, please contact Professor John C.S. Lui, the leader of this strategic initiative (e-mail: *cslui@cse.cuhk.edu.hk*). Applicants are requested to clearly indicate that they are applying for the posts under 'Big Data Analytics Initiative'. The Personal Information Collection Statement will be provided upon request. Please quote the reference number and mark 'Application – Confidential' on cover.

BAYLOR UNIVERSITY

The Electrical and Computer Engineering Department of Baylor University

seeks faculty applicants for three tenured/tenure-track Faculty Positions at all levels and in all areas of electrical and computer engineering. Desired areas of technical expertise include: embedded systems, cyber-physical systems, computer/network security, software engineering, sensor networks, power, and energy. Applicants seeking a senior position must have an impressive record of scholarship and sustained research funding. All applicants must have an earned doctorate and a record of achievement in research and teaching. The ECE department offers B.S., M.S., M.E. and Ph.D. degrees and is rapidly expanding its faculty size. Facilities include the Baylor Research and Innovation Collaborative (BRIC), a newly-established research park minutes from the main campus.

Chartered in 1845 by the Republic of Texas, Baylor University is the oldest university in Texas. Baylor has an enrollment of over 15,000 students and is a member of the Big XII Conference. Baylor's mission is to educate men and women for worldwide leadership and service by integrating academic excellence and Christian commitment within a caring community. The department seeks to hire faculty with an active Christian faith; applicants are encouraged to read about Baylor's vision for the integration of faith and learning at ***www.baylor.edu/profuturis/***.

Application reviews are ongoing and will continue until all positions are filled; however, applications received by January 1, 2014 will be assured of full consideration.

Applications must include: 1) a letter of interest that identifies the applicant's anticipated rank, 2) a complete CV, 3) a statement of teaching and research interests, 4) the names and contact information for at least four professional references.

Additional information is available at ***www.ecs.baylor.edu***. Send materials via email to Dr. Robert J. Marks II at Robert_Marks@baylor.edu. Please combine all submitted material into a single pdf file.

Baylor is a Baptist University affiliated with the Baptist General Convention of Texas. As an Affirmative Action/Equal Employment Opportunity employer, Baylor encourages candidates with an active Christian faith who are minorities, women, veterans, and persons with disabilities to apply.

NANYANG TECHNOLOGICAL UNIVERSITY

DISTINCTIVE LEADERSHIP AND RESEARCH OPPORTUNITY

with one of the world's fastest-rising universities. **NTU, Singapore**


Nanyang Technological University, Singapore, invites applications for the position of

DEAN, COLLEGE OF ENGINEERING

About the Appointment

Nanyang Technological University is seeking an accomplished and visionary academic and research leader for the position of the Dean of the College of Engineering.

Reporting to the Provost, the Dean will provide foundational vision, leadership and oversight for the strategic, academic and intellectual affairs of the College and its six constituent Schools. The Dean is expected to lead on an inter-disciplinary basis within and amongst Colleges, and build the human capital of the Engineering College with proactive recruitment of world-class faculty. Other critical responsibilities include cultivating areas of academic and research excellence including inter-College collaboration, providing leadership for securing external research funding support, and interaction with external stakeholders to raise the College's profile, academic standing and engagement. The Dean is also a member of the University Cabinet, the highest management decision-making body, and other senior leadership committees at the University.

This is an opportunity for an outstanding academician who is passionate about inter-disciplinary teaching and research, and about growing and strengthening a dynamic College, in a collaborative manner with other Colleges, within a University with a high global impact.

The appointee must have an outstanding record of academic leadership, research and teaching in a reputable university or academic institution. The appointee must demonstrate leadership, vision, and outstanding interpersonal skills to engage faculty and administrators across campus. Other essential attributes include the ability to effectively communicate with stakeholders, work cooperatively with national funding agencies and commit to faculty-shared governance. The appointee must also be passionate and committed to enhancing the existing academic strengths of the College in research, scholarship and education.

About Singapore

Singapore is a dynamic centre for the pursuit of science, technology, and innovation, and the vibrant city-state is known for its commitment to academic excellence and research. Singapore aims to be a global centre for academic excellence in higher education and research, and this strategic goal and direction has been supported by major public sector investments. Uniquely situated in an emergent and vibrant Asia, Singapore combines the eastern and western approaches to governance, education, and lifestyle. Having forged lasting and synergistic business and academic relationships with China, India, other Asian countries and beyond, Singapore sits at the crossroads of cultures and peoples.

About the University

Nanyang Technological University (NTU) is an internationally renowned research-intensive university with globally acknowledged strengths in Engineering and Business. Science and Humanities have added to the strengths of the University, and a Medical School has been set up jointly with Imperial College London. Ranked 8th in Asia and 41st in the 2013 Quacquarelli Symonds (QS) World University Rankings, NTU is the fastest rising Asian university in the QS Top 50. In this ranking, NTU's Engineering and Technology is ranked 14th in the world and 3rd in Asia. NTU is also the 2nd in the world among the young elite universities in the QS' rankings.

NTU is the 5th most-cited university for engineering research output that is among the top three universities globally (Essential Science Indicators, January 2013). The Electrical and Electronic Engineering and Computer Engineering schools have both been ranked 1st in Singapore and 4th in the world (after MIT, University of California, Berkeley, and Stanford University) by the Higher Education Evaluation & Accreditation Council of Taiwan. In the QS World University Rankings by Subject 2013, NTU is ranked within the world's top 20 for Civil Engineering (8th worldwide and first in Singapore), Mechanical Engineering (11th), Communication & Media Studies (11th), Education (13th), Electrical Engineering (14th) and Materials Science (14th). NTU was placed 32nd globally and 6th in Asia in the 2013 Financial Times Global MBA Rankings. Its accounting research in the business school has been ranked 1st in Asia and 5th position globally in the latest Brigham Young University (BYU) Accounting Research Rankings.

The University's academic and research programmes, with real-world relevance, have received strong support from major corporations and industry, in terms of research funding, industry partnerships, and global internship opportunities for the students.

A founding member of the Global Alliance of Technological Universities, NTU aims to groom active citizens of the world who can lead and manage new, complex, global challenges. The University provides a high-quality comprehensive and global education to more than 23,500 undergraduates and 9,500 graduate students. Together with the University's 3,800-strong faculty and research staff who bring international academic perspectives and depth of experience, the University's main 200-hectare residential garden campus – located at the south-western part of Singapore – is a hub for vibrant academic endeavours.

About the College

NTU's four colleges - College of Engineering; College of Science; College of Business (Nanyang Business School); and College of Humanities, Arts, and Social Sciences - comprise 12 component schools. The College of Engineering is one of the largest engineering colleges in the world, housing six constituent Schools, 14,000 students, 600 faculty and 1,600 staff. The six schools are the School of Chemical and Biomedical Engineering, the School of Civil and Environmental Engineering, the School of Computer Engineering, the School of Electrical and Electronic Engineering, the School of Materials Science and Engineering and the School of Mechanical and Aerospace Engineering.

More information on the College can be accessed at: **http://coe.ntu.edu.sg**

To apply, please send your curriculum vitae, accompanied by a cover letter, to:

The Vice-Chairman of the Search Committee, Professor S. Shankar Sastry
Dean, College of Engineering
Roy W. Carlson Professor of Electrical Engineering and Computer Science,
Bioengineering & Mechanical Engineering
University of California, Berkeley

c/o Secretary to the Search Committee
Nanyang Technological University, Level 4, Administration Building
50, Nanyang Avenue, Singapore 639798
Fax: (65) 6795 9001 Email: DEANSEARCHCOE@NTU.EDU.SG

Applications will be welcomed up to 15 January 2014

Korn/Ferry International Pte Ltd is the appointed Executive Search firm for the Dean Search.

All applications and materials submitted will be held in strict confidence.

www.ntu.edu.sg

ELECTRICAL *and* COMPUTER ENGINEERING

TEXAS A&M UNIVERSITY

Texas A&M University:

The Department of Electrical and Computer Engineering has an opening for a tenure-track faculty position at the assistant or associate professor rank. A person appointed to this position will participate in the new Center for Bioinformatics and Genomics Systems Engineering (CBGSE). The Center focuses on fundamental systems theoretic research in the general areas of bioinformatics, computational biology, genomics, and systems biology as related to health care, veterinary medicine, and agriculture, serving communities both internal and external to Texas A&M University. As a cooperative effort between Texas AgriLife Research and the Texas Engineering Experiment Service (TEES), it constitutes both an active research group aimed at fundamental knowledge and a cross-institution graduate student training program in theoretical and translational analytics. CBGSE incorporates within it the Genomic Signal Processing Laboratory, which has been a pioneer in the application of systems theory to genomics. We are looking for highly motivated faculty with outstanding demonstrated transformational research in systems-theoretic bioinformatics and computational biology, including filtering, information theory, dynamical networks, control theory, system identification, and pattern recognition (not data mining). Applicants for associate professor must have a stellar record of publication and externally funded research.

Applicants must have a Ph.D. or equivalent degree in electrical and computer engineering or related field, or have completed all degree requirements by date of hire, and must demonstrate potential for quality teaching and research leading to significant publications and funding.

The Department of Electrical and Computer Engineering currently has 72 faculty members and its degree programs have been ranked in the top 20 in recent years. Further information about the department may be obtained by visiting **http://ece.tamu.edu.**

Applications, including full curriculum vitae with a list of publications, a statement of teaching, a statement of research and the names, addresses (regular mail and email), of three references should be sent in a single PDF file, preferably electronically, to TAMU **Search@ECE.TAMU.EDU**, or in hard copy to:

Dr. Chanan Singh, Interim Department Head

c/o Ms. Debbie Hanson

Texas A&M University

Department of Electrical and Computer Engineering

TAMU 3128

College Station, TX, 77843-3128.

Texas A&M University is an equal opportunity/affirmative action employer. Candidates shall be considered solely on the basis of qualifications, without regard to race, color, sex, religion, national origin, age, disabilities or veteran status. The deadline for applications is January 15, 2014.

UNIVERSITY OF MICHIGAN-DEARBORN™

ECE Faculty Computer Engineering and Robotics Engineering

The University of Michigan-Dearborn, Dept. of Electrical & Computer Engineering (ECE) invites applications for two tenure-track faculty positions, one in Computer Engineering, and another in Robotics Engineering. Qualified individuals seeking an Assistant Professorship are invited to apply but outstanding faculty at the Associate Professor and Professor levels will also be considered. The applicant for computer engineering should have research interests that fall within the areas of embedded systems and/or high performance computing. The applicant for robotic engineering should have research interests in key aspects of robotics with a strong background in mechanical principles, electromechanics, and intelligent robots. Selected candidates are expected to establish an excellent externally funded research program in their chosen area.

Qualified candidates must have, or expect to have, a Ph.D. in Computer Engineering, Electrical Engineering, or a closely related discipline by the time of appointment and will be expected to do scholarly and sponsored research, as well as teaching at both the undergraduate and graduate levels. Candidates at the associate or full professor ranks should already have an established funded research program. The ECE Department offers several BS and MS degrees, and participates in two interdisciplinary Ph.D. programs, Ph.D. in Automotive Systems Engineering and Ph.D. in Information Systems Engineering. The current funded research areas in the department include intelligent systems, power electronics, hybrid vehicles, battery management, computer networks, wireless communications, and embedded systems.

The University of Michigan-Dearborn (UM-Dearborn) is one of the three campuses of the University of Michigan. UM-Dearborn is a comprehensive university offering high quality undergraduate, graduate, professional and continuing education to residents of southeastern Michigan, and attracts more than 9,000 students. Faculty and students have the opportunity to collaborate across all three campuses in research and scholarly activity. UM-Dearborn is located ten miles west of Detroit and thirty-five miles east of Ann Arbor. The campus is strategically located on 200 suburban acres of the original Henry Ford Estate in the Greater Detroit Metropolitan region.

The University of Michigan-Dearborn, as an equal opportunity/affirmative action employer, complies with all applicable federal and state laws regarding nondiscrimination and affirmative action. The University of Michigan-Dearborn is committed to a policy of equal opportunity for all persons and does not discriminate on the basis of race, color, national origin, age, marital status, sex, sexual orientation, gender identity, gender expression, disability, religion, height, weight, or veteran status in employment, educational programs and activities, and admissions. Inquiries or complaints may be addressed Office of Institutional Equity, 4901 Evergreen Road, Suite 1020, Administrative Services Building, Dearborn, Michigan 48128-2406, 313 593-5190. For other University of Michigan information call 734 764-1817.

Applicants should submit a cover letter, curriculum vitae including current contact information: phone number, mailing address and e-mail address, teaching statement, research statement, and a list of three to five referees to the Department of Electrical and Computer Engineering, University of Michigan-Dearborn, 4901 Evergreen Road, Dearborn, Michigan, 48128, Phone: 313 593-5420, Fax: 313 583-6336 or email to **mphicks@umich.edu.**

The University of Michigan-Dearborn is an equal opportunity/affirmative action employer.

POSITION OPEN

Toyota Technological Institute

has an opening for a tenured- or tenure-track faculty position in the Department of Advanced Science and Technology. Applications are encouraged from all relevant areas. For more information, please refer to the website **http://www.toyota-ti.ac.jp/english/employment/associate.html**

Position: Associate Professor
Research field: Advanced Functional Materials for energy, sensing or green technology.

Qualifications: A Ph.D. in a relevant field. The successful candidate is expected to demonstrate potential to develop strong and outstanding programs in the above research field. It is also necessary for him/her to supervise students, and to teach advanced and basic courses both at the undergraduate and graduate levels.

Start date: At the earliest convenience in 2014

Documents: (1) A curriculum vitae

(2) A list of research activities

(3) Copies of 5 representative papers

(4) A brief summary and future plan of your research and educational statement (within three pages each)

(5) Names of two references including phone numbers and e-mail addresses

(6) An application form (available on our website)

Deadline: January 6, 2014

Inquiry: Search Committee Chair Professor Shuji Tanaka

(Phone) +81-52-809-1775, (e-mail) tanaka_mat@toyota-ti.ac.jp

The above documentation should be sent to:

Mr. Takashi Hirato
Administration Division
Toyota Technological Institute
2-12-1, Hisakata, Tempaku-ku
Nagoya, 468-8511 Japan

Please write "Application for Advanced Functional Materials" in red on the return envelope.

South University of Science and Technology, Shenzhen, China

Faculty position in Electronic Engineering

The Department of Electronic Engineering at the South University of Science and Technology invites applications for faculty positions at all ranks of tenured and tenure-track faculty members in all areas of electrical and electronic engineering. In particular, we are in high demand of applicants in the area of microelectronics, including IC design and Device Process/Technology, Micro/Nano-Electro-Mechanical Systems, Photonic Integrated Circuits, 3D Integration and Packaging, Testing and Reliabilities, Thin Film Transistors and Displays etc. 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.

South University of Science and Technology, officially established in April 2012, is a research-intensive public institution funded by the municipal of Shenzhen, a growing international metropolitan neighboring Hong Kong. The University is accredited by the Ministry of Education, China and is a pioneer in higher education reform in China. The teaching language at the University is English or Putonghua. The choice is made by the instructor.

The University offers internationally competitive salaries, fringe benefits including medical/dental insurance, retirement and housing subsidies.

Applications including full curriculum vitae, list of publications, statement of research, and names of five referees addressed to Professor Yu Hong Yu, and should be sent by email to **ecesearch@sustc.edu.cn** as well as **hiring@sustc.edu.cn**.

Additional information is available at **http://www.sustc.edu.cn** and **http://english.sina.com/china/2012/0902/502496.html.**

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

NUS National University of Singapore

School of Computing

The Department of Computer Science, National University of Singapore (NUS), has openings for several tenure-track faculty positions. Our main focus is on candidates at the Assistant Professor level with research interests in the following areas:

- **Cyber-physical systems** • **Big data analytics**
- **Security** • **Sensor data modelling and learning**

These areas are to be viewed in a broad sense, and we are particularly interested in candidates whose research interests cut across these and related areas. We seek candidates demonstrating excellent research potential and a strong commitment to teaching. We will also seriously consider exceptional candidates in other areas of computer science. Further, we will consider candidates at senior ranks (Associate and Full Professor) who have an outstanding record of research accomplishments.

We are an internationally top-ranked department with low teaching loads, excellent facilities, and intensive external collaborations. Significant funding opportunities abound for strong candidates. The research of the faculty covers all the major areas of computer science and is well represented at prestigious international conferences and journals. The department has a thriving PhD programme and it actively strives to attract the best students from the region and beyond. More information can be found at **http://www.comp.nus.edu.sg/**.

NUS offers highly competitive salaries and generous benefits, while Singapore provides a vibrant international environment with world-class health care, excellent infrastructure, a warm climate and low taxes.

Interested candidates are invited to send, via electronic submission, the following materials to the Chair of the CS Search Committee, Prof. P.S. Thiagarajan, at **csrec@comp.nus.edu.sg**

- **Cover letter** • **Curriculum Vitae** • **A teaching statement**
- **A research statement** • **Contact information for at least three references**

Applications will be reviewed as they are received and will continue until the positions are filled. However, to ensure maximal consideration applicants are encouraged to submit their materials by December 15, 2013.

USP University of São Paulo

School of Engineering of São Carlos

Electrical and Computer Engineering Department

Brazil

The Electrical and Computer Engineering Department of the School of Engineering of São Carlos University of São Paulo (USP), seeks faculty applicants for a tenure-track assistant professor position in the research area of electronics and signal processing. Candidates must have a Ph.D. degree and will be expected to carry out research and teaching activities at the graduate and undergraduate levels.

USP is a public research university and it is consistently the top Brazilian institution of higher learning in international rankings. The city of São Carlos (population around 250,000, located 235 km from São Paulo) is home to the School of Engineering, EESC/USP.

Initial activities can be conducted in English but the selected applicant is expected to teach in Portuguese within a 2-year timeframe. Starting salary is around US$ 55,000.00/year.

For further information and details on the selection process, please contact Prof. Murilo A Romero, the Department Head, at murilo.romero@usp.br, including a letter of interest with teaching and research statements, as well as a complete CV.

College of Science & Engineering
University of Minnesota

The **University of Minnesota – Twin Cities** invites applications for faculty positions in **Electrical and Computer Engineering** from individuals with strong expertise in

(1) power and energy systems;

(2) micro and nano devices and structures; and, in support of a University Initiative on Robotics, Sensors, and Manufacturing,

(3) control and dynamical systems, robotics and automation, image processing and computer vision, and materials, devices and systems for novel sensing and actuation applications.

Women and other underrepresented groups are especially encouraged to apply. An earned doctorate in an appropriate discipline is required. Rank and salary will be commensurate with qualifications and experience. Positions are open until filled, but for full consideration, apply at

http://www.ece.umn.edu/

by December 15, 2013. The University of Minnesota is an equal opportunity employer and educator.

LECTURER

in Electrical and Electronic Engineering (EEE)

University College Cork (UCC), Ireland.

http://soe.ucc.ie

UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

Applications are invited for an Autumn 2014 start date permanent Lecturer post in EEE, School of Engineering, UCC. Lecturer topic is electrical power engineering with particular interests in electrical power systems, power electronics and control, and electric drives. Applications are sought from candidates with a Ph.D. in electrical/electronic engineering. For information on post, Contact Prof. N. A. Riza, Head School of Engineering, UCC. Email: n.riza@ucc.ie. Formal applications must be submitted by Jan.20, 2014. For application forms and further information see www.ucc.ie/hr/vacancies

IUPUI
DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING
SCHOOL OF ENGINEERING AND TECHNOLOGY
A Purdue University School
Indianapolis

Assistant/Associate Professor of Electrical and Computer Engineering,

Purdue School of Engineering and Technology, Indiana University-Purdue University Indianapolis (IUPUI), invites applications for one or more tenure-track faculty positions. The anticipated start date is August 1, 2014. Successful candidates should expect to teach undergraduate and graduate courses, supervise graduate students and conduct externally funded research. Preference will be given to applicants in the areas of computer engineering, digital signal processing and areas with applications related to "Big Data" analysis.

For a complete job description and to apply:

http://et2.engr.iupui.edu/departments/ece/about/news/index.php

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

M ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

UNIVERSITY of MICHIGAN ■ COLLEGE of ENGINEERING

Electrical and Computer Engineering
University of Michigan, Ann Arbor

The Electrical and Computer Engineering (ECE) Division of the Electrical Engineering and Computer Science Department at the University of Michigan, Ann Arbor invites applications for junior or senior faculty positions, especially from women and underrepresented minorities. Successful candidates will have a relevant doctorate or equivalent experience and an outstanding record of achievement and impactful research in academics, industry and/or at national laboratories. They will have a strong record or commitment to teaching at undergraduate and graduate levels, to providing service to the university and profession and to broadening the intellectual diversity of the ECE Division. Although the research areas of particular interest are networks and communications, computer vision, integrated circuits and optics, applications are welcome in all relevant areas of research.

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

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

Applications will be considered as they are received. However, for full consideration applications must be received by December 8, 2013.

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

Faculty positions in Robotics and in Computer Science

NAZARBAYEV UNIVERSITY

School of Science and Technology
Nazarbayev University
Astana, Kazakhstan

Nazarbayev University is seeking highly qualified full-time faculty at all ranks to join its rapidly growing programs in Robotics and Computer Science in the School of Science & Technology. Successful candidates must have an earned Ph.D. degree from an accredited university, excellent English-language communication skills, a demonstrated ability for research, and a commitment to graduate and undergraduate teaching and program development.

Launched in 2010 as the premier national university of Kazakhstan, NU's mandate is to promote the emergence of Astana as the research and educational center of Eurasia. The strategic development of this English-language university is based on the Western model via partnerships with top ranking world universities.

All suitable candidates will be reviewed. Applications are particularly encouraged from candidates with research interests that align with our thrust areas of Embedded Systems, Industrial Automation and Manufacturing **(Department of Robotics)** and Computer Systems, Intelligent Systems, Mobile Computing, Information Security, and Software Engineering **(Department of Computer Science)**.

Benefits include a competitive salary, international health care coverage, housing (based on family size and rank), child educational allowance, and home-leave travel twice per year.

To Apply: applicants should send a detailed CV, including qualifications, experience, and list of publications to **sst@nu.edu.kz**. Review of applications will begin immediately but full consideration will be given to applications submitted no later than **January 15th, 2014**.

Are You An Electrical Engineering Student Looking For An Internship?

Look No Further!

We Have Internship Opportunities for Students.

Visit us at:

www.AfterCollege.com/IEEE_Internships

IEEE Advancing Technology for Humanity

IEEE StudentJobSite

THE UNIVERSITY OF CHICAGO

The Department of Computer Science at the University of Chicago invites applications for the position of Lecturer. Subject to the availability of funding, this would be a three year position with the possibility of renewal. This position involves teaching in the fall, winter and spring quarters. The successful candidate will have exceptional competence in teaching and superior academic credentials, will carry responsibility for teaching computer science courses and laboratories. Applicants must have a Ph.D. in Computer Science or a related field at time of appointment and have experience teaching Computer Science at the College level.

The Chicago metropolitan area provides a diverse and exciting environment. The local economy is vigorous, with international stature in banking, trade, commerce, manufacturing, and transportation, while the cultural scene includes diverse cultures, vibrant theater, world-renowned symphony, opera, jazz and blues. The University is located in Hyde Park, a Chicago neighborhood on the Lake Michigan shore just a few minutes from downtown.

Applicants must apply on line at the University of Chicago Academic Careers website at **http://tinyurl.com/m3z99rw**.

Applicants must upload a curriculum vitae with a list of publications and a one page teaching statement. In addition, three reference letters will be required. Review of complete applications, including reference letters, will begin January 15, 2014, and continue until the position is filled. ***The University of Chicago is an Affirmative Action / Equal Opportunity Employer***.

THE OTHER FALLOUT FROM FUKUSHIMA

CARBON EMISSIONS RISE AS REACTORS ARE SHUT DOWN

In the wake of the Fukushima meltdown in 2011, Japan shut down all its nuclear reactors. Two reactors soon returned to operation, but now even these are closed for maintenance. Recently released estimates of Japan's energy consumption and carbon emissions for 2012 confirm that a continuing consequence of that shutdown is more fossil-fuel-based electricity generation and attendant greenhouse gas emissions. It's likely, though, that concerns about costs, not emissions, will drive a return to nuclear, as Japan's trade deficit increase has been blamed on fossil fuel purchases. —STEPHEN CASS

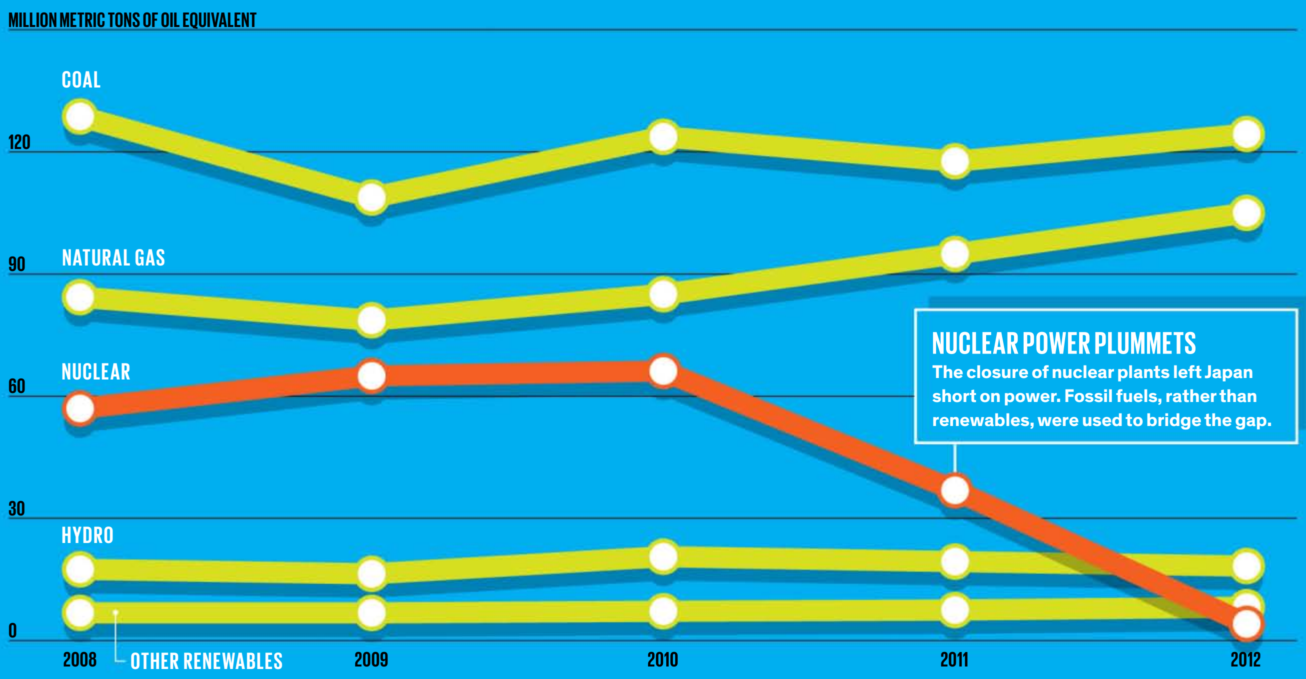

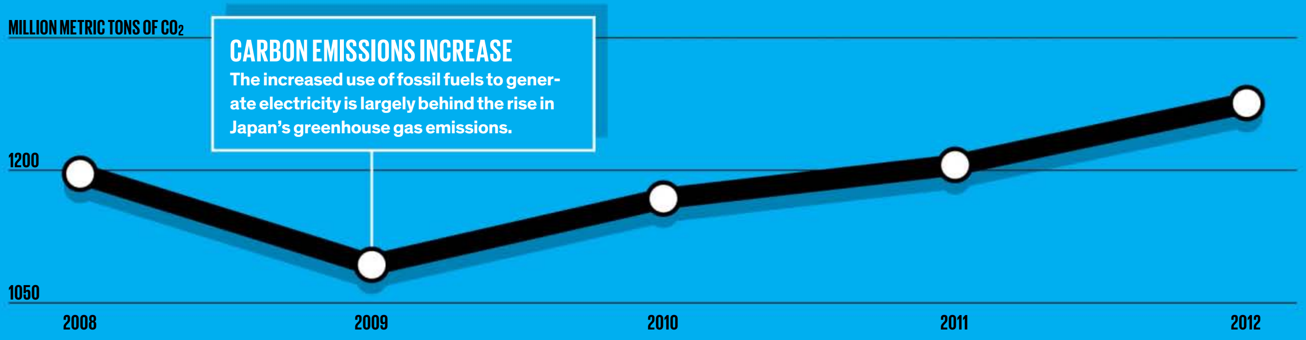

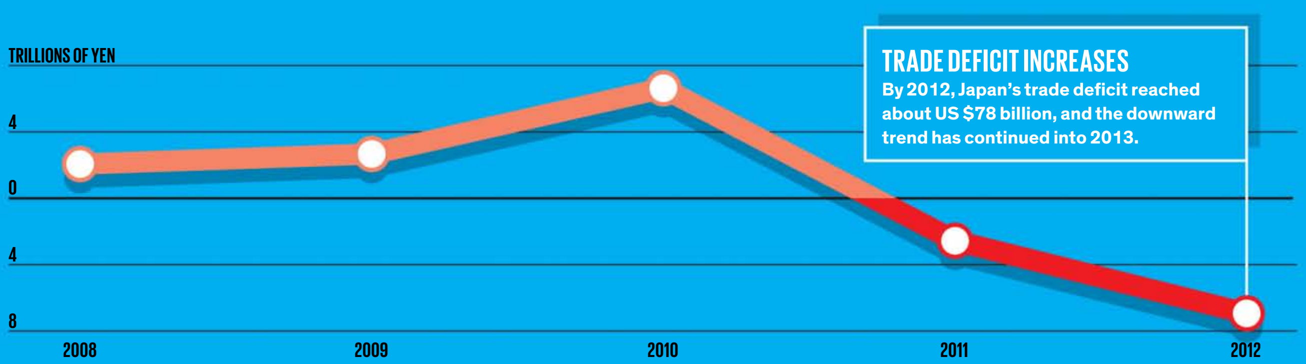

SOURCES: BP STATISTICAL REVIEW OF WORLD ENERGY JUNE 2013; REDRAWING THE ENERGY-CLIMATE MAP, INTERNATIONAL ENERGY AGENCY, 2013; HTTP://WWW.IEA.ORG/STATISTICS/TOPICS/CO2EMISSIONS; JAPAN MINISTRY OF FINANCE