IEEE SPECTRUM

FOR THE TECHNOLOGY INSIDER | 09.15

How Much For Those ZZZZZs?

An amazing algorithm lets Airbnb put a price on any room, any time P. 26

CLICK HERE TO ACCESS THE SEPTEMBER ISSUE OF the institute

IEEE

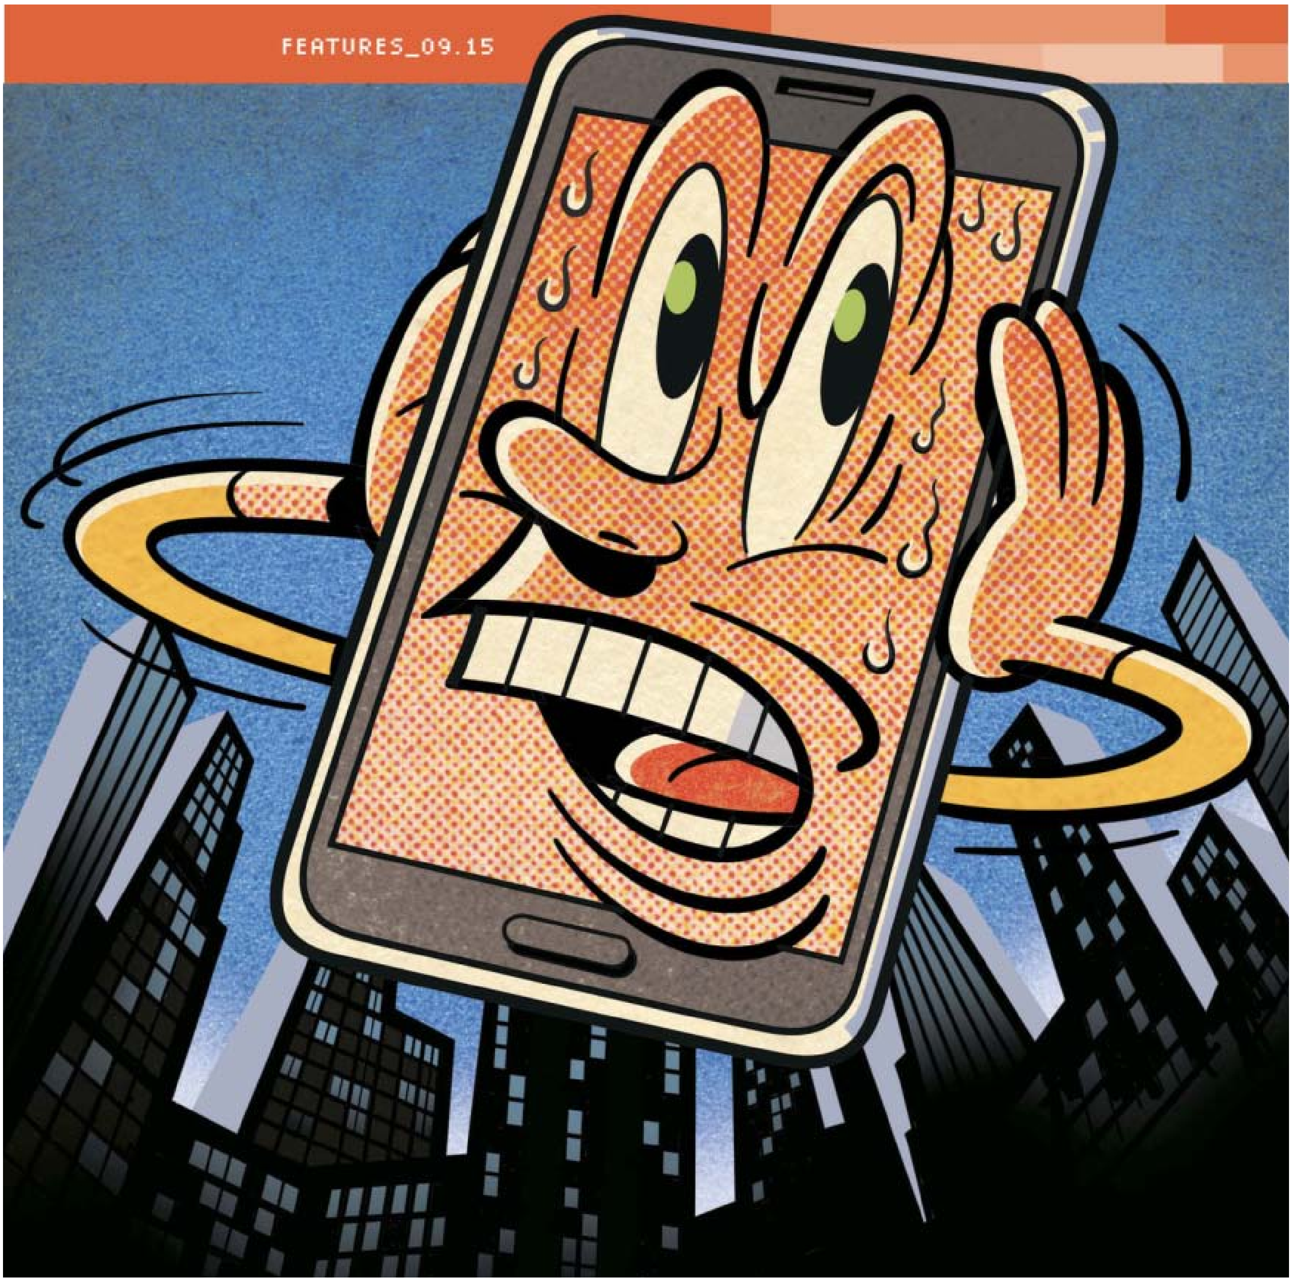

On the Cover **Photograph for *IEEE Spectrum* by The Voorhes**

Illustration by Lou Brooks

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

DEPARTMENTS_09.15

IEEE SPECTRUM

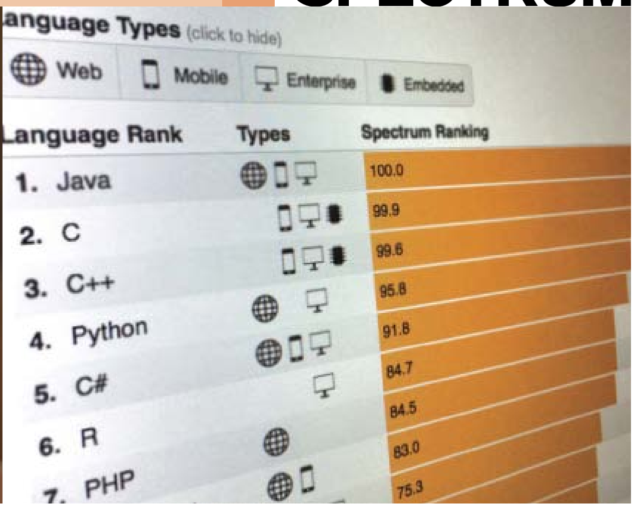

ADDITIONAL RESOURCES

Tech Insider / Webinars

Available at spectrum.ieee.org/webinar

- EMC Simulation of a Motor Control—24 September 2015
- Top-Ranking Online Master's in Engineering Degrees Designed for Working Professionals—Available On Demand
- Product Line Engineering: The Key for Using Strategic Reuse to Tame IoT Product Development Complexity—Available On Demand
- Simulating Graphene-Based Photonic and Optoelectronic Devices—Available On Demand

- **MASTER BOND RESOURCE LIBRARY** http://spectrum.ieee.org/static/master-bond-resource-library
- **NEW PRODUCT RELEASE LIBRARY** http://spectrum.ieee.org/static/new-product-release-library
- **WHITE PAPERS** http://spectrum.ieee.org/whitepapers

The Institute

Available at theinstitute.ieee.org

- **READY TO LAUNCH** Engineering and entrepreneurship often go hand in hand. In this special issue on startups, we present IEEE's efforts to attract more entrepreneurial types to the organization and to support members' ventures through online resources, networking events, and more. We also spotlight several members who founded innovative tech companies.

- **GETTING SCHOOLED ON STARTUPS** Milton Chang, an angel investor whose companies have never failed, shares advice with entrepreneurs on finding a niche market, capitalizing on an idea, and getting funding.

- **HUMBLE BEGINNINGS** Learn how multibillion-dollar companies Hewlett-Packard and Medtronic got started.

FROM LEFT: MOTUS GLOBAL; DON WEST; RANDI KLETT

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2015 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 52, issue No. 9, 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_

PETE HAEFF

Catching the Wave

JACK COPELAND AND ANDRE A. HAEFF collaborated on this month's feature "The True History of the Traveling-Wave Tube." How these two men, who have very different backgrounds and live 11,000 kilometers apart, ended up writing the article together is an interesting story in itself.

Copeland [above left], a Distinguished Professor of Arts at the University of Canterbury in Christchurch, New Zealand, developed an interest in computability theory, which eventually led him to study the history of computing, including the hardware of early electronic computers. While researching a primitive form of computer memory called the Williams tube, Copeland stumbled on a patent for another kind of computer-memory tube invented by A.V. Haeff, a name he didn't recognize. Seeking to learn about this mysterious tube, Copeland began an investigation of Haeff's work.

"I discovered he hadn't just invented the memory tube–he'd invented all these other things, including the traveling-wave tube," says Copeland. But he pieced together very little about Haeff's personal life until a student brought him a CD containing music by Haeff's brother, Alexei. That in turn led Copeland to Haeff's son, Andre, a retired physician who by good fortune had preserved his father's papers. The collection occupies "cubic meters, a completely pristine archive," Copeland says.

Copeland and Andre Haeff exchanged lengthy e-mails, which eventually evolved into a series of articles about the elder Haeff's life and work. The two coauthors met for the first time in Santa Monica, Calif., close to the beach where Haeff senior gained the inspiration for the traveling-wave tube by watching surfers. To understand how chasing breakers could inspire the invention of an early microwave amplifier, though, you'll have to read their article. Cowabunga! ■

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, And Now This... is in *IEEE Spectrum*, Vol. 52, no. 9 (INT), September 2015, p. 56, or in *IEEE Spectrum*, Vol. 52, no. 9 (NA), September 2015, p. 76.

IEEE
SPECTRUM

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

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

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

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

SENIOR ART DIRECTOR
Mark Montgomery, m.montgomery@ieee.org

SENIOR EDITORS
Stephen Cass (Resources), cass.s@ieee.org
Erico Guizzo (Digital), e.guizzo@ieee.org
Jean Kumagai, j.kumagai@ieee.org
Samuel K. Moore (News), s.k.moore@ieee.org
Tekla S. Perry, t.perry@ieee.org
Philip E. Ross, p.ross@ieee.org
David Schneider, d.a.schneider@ieee.org

DEPUTY ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSOCIATE ART DIRECTOR Erik Vrielink, e.vrielink@ieee.org

SENIOR ASSOCIATE EDITORS
Rachel Courtland, r.courtland@ieee.org
Eliza Strickland, e.strickland@ieee.org
ASSOCIATE EDITOR
Celia Gorman (Multimedia), celia.gorman@ieee.org
ASSISTANT EDITOR Willie D. Jones, w.jones@ieee.org
SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

ADMINISTRATIVE ASSISTANTS
Ramona L. Foster, r.foster@ieee.org
Nancy T. Hantman, n.hantman@ieee.org
INTERNS Kristen Clark (Video); Lauren Young

CONTRIBUTING EDITORS
Evan Ackerman, Mark Anderson, John Blau, Robert N. Charette, Peter Fairley, Tam Harbert, Mark Harris, David Kushner, Robert W. Lucky, Paul McFedries, Prachi Patel, Richard Stevenson, Lawrence Ulrich, Paul Wallich

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia
SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

SPECTRUM ONLINE
PRODUCT MANAGER, DIGITAL Shannan Brown
WEB PRODUCTION COORDINATOR Jacqueline L. Parker
MULTIMEDIA PRODUCTION SPECIALIST Michael Spector

EDITORIAL ADVISORY BOARD
Susan Hassler, *Chair*; Gerard A. Alphonse, Jason Cong*, Sudhir Dixit, Limor Fried, Robert Hebner, Joseph J. Helble, Chenming Hu*, Grant Jacoby, Norberto Lerendegui, John P. Lewis, Steve Mann, Jacob Østergaard, Umit Ozguner, Thrasos Pappas, John Rogers, Jonathan Rothberg, Umar Saif, Takao Someya, Jeffrey M. Voas, Kazuo Yano, Larry Zhang*, Yu Zheng, Kun Zhou, Edward Zyszkowski
* Chinese-language edition

EDITORIAL / ADVERTISING CORRESPONDENCE
IEEE Spectrum
3 Park Ave., 17th Floor
New York, NY 10016-5997
EDITORIAL DEPARTMENT
TEL: +1 212 419 7555 FAX: +1 212 419 7570
BUREAU Palo Alto, Calif.; Tekla S. Perry **+1 650 752 6661**
ADVERTISING DEPARTMENT +1 212 705 8939

RESPONSIBILITY FOR THE SUBSTANCE OF ARTICLES rests upon the authors, not IEEE, its organizational units, or its members. Articles do not represent official positions of IEEE. Reader comments will be posted online, or published in print as space allows, and may be excerpted for publication. The publisher reserves the right to reject any advertising.

CONTRIBUTORS_

Lou Brooks

Brooks's best-known illustration may be the top-hatted, mustached man on the Monopoly board, which he was commissioned to redesign in 1985. The California-based illustrator relies on bold graphics that evoke comic books of the 1950s. In this issue, Brooks used a smartphone in distress to illustrate the problem of electronic noise pollution [p. 44]. The topic gave him plenty to work with, he says: "I could see right away it was my kind of meat."

Dan Hill

Hill, product lead at Airbnb, wrote the lodging-rental website's pricing algorithm, which evolved into the system he describes in this issue [p. 26]. He also cofounded Crashpadder, a home-sharing company acquired by Airbnb in 2012. Hill first did Web development to support his career as a violinist. "I sort of woke up one day and realized I hadn't really been focused too much on [the violin]," he said in a recent interview. His next thought? "I really want to spend my life working on technology and products."

Mark A. McHenry

In this issue, McHenry teamed up with Dennis Roberson of the Illinois Institute of Technology and Robert J. Matheson, who surveyed electronic noise back in the 1970s [see "Phone to Fridge: Shut Up!," p. 44]. To update Matheson's survey, they equipped a car to prospect for electronic noise. But the technician who drove the car used a low-cost, 12-volt inverter, whose own noise ruined the data. McHenry's company, Shared Spectrum, aims to tackle such interference by automating ways to share the radio spectrum.

Wolfgang Porod

Porod is a professor of electrical engineering at the University of Notre Dame, in Indiana. In this issue, he and Michael Niemier, a professor of computer science and engineering at Notre Dame, discuss their research into computing with nanoscale magnets [p. 38]. Porod, who used to focus on semiconductor devices, says working with electric charge is more straightforward: "I actually promised myself early on that I would never work on magnetism. But it turned out to have fantastic properties."

Theresa Sullivan Barger

Barger has covered business, health, the environment, and careers for *The New York Times, Family Circle,* and others. In this issue, she writes about the role that mentors play in helping women engineers navigate a male-dominated field [p. 21]. "So many women feel marginalized and undervalued that they give up on their profession," she says. "Sometimes just having a mentor who has faced similar discrimination and figured out a way to be taken seriously makes all the difference."

09.15

PHOTO-ILLUSTRATIONS BY **Gluekit**

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

The Engineered Death

How technology complicates the end of life

THIS IS NOT ONLY A PARABLE ABOUT technology and progress. This is also a true story.

When I was a boy, my mother contracted an incurable cancer called Hodgkin's disease. At a hospital in Manhattan, in New York City, ingenious medical engineers and doctors bombarded her with radiation—at far higher levels, and for far longer, than most experts at that time recommended.

Only my father knew of the severity of her illness. My mother kept that from others. I began to grasp the situation when our grandmother came to care for my brother, sister, and me while my mother stayed in the hospital. Through the prism of my 10-year-old mind, I concluded that something terrible had happened for us to endure Grandma's strange cooking and hot temper.

To our family's delight, our mother's disease vanished. She remained cancer-free. In time, she could say aloud that in the mid-1960s she was among the first Hodgkin's patients to be cured in the United States. She could also describe her anxious encounters with the bulky and fearsome radiation equipment that, guided by her innovative doctors, had saved her life.

My mother's youthful experience turned her into a lifelong adherent of the secular faith in innovation. Fifty years later, her faith was undiminished. "Obey your doctors," she often told me. "You get the best results."

This past July, I recalled my mother's creed as I sat at her feet in the intensive care unit of a Long Island, N.Y., hospital. Her 84-year-old lungs, burdened by a weakening heart, had failed. She could breathe only with the aid of a ventilator. She couldn't speak.

A specialist directed the computer-controlled ventilator as if he were playing a video game. The software was graphical, and the numerical levels were displayed in large type on a big screen. With a few keystrokes, the specialist could deliver more or less oxygen.

All of us have faith in the idea of progress—that life gets better. And most of us believe new technologies are the best indicators of progress. We believe, in short, that innovation and human improvement are one and the same. Take away the ventilator and my mother dies. Build a better ventilator and then, perhaps, she lives.

Our lives are dominated by digital technologies. Observing my mother, I became acutely aware that our deaths are also computer mediated.

My mom isn't welcoming death. She indicates through dramatic hand gestures and facial expressions that she wants doctors to keep trying to save her.

For how long, I wonder? Today is her 31st day on a ventilator. The doctors want, in their lingo, "to wean her off." Yet when they try, she struggles to breathe, her heart rate rises steadily, and the attempts are abandoned.

The doctors are doing everything they can, but what if they can't fix my mother? What if, as her CAT scan suggests, the radiation treatments that cured her of cancer also scarred her lungs so badly that now they won't work?

Unintended consequences are common features of our engineered lives. *Why Things Bite Back: Technology and the Revenge of Unintended Consequences* (1996) is historian Edward Tenner's explanation of the situation. I suspect everyone would accept the deal my mother got: 50 years of healthy life in exchange for a few months of extreme suffering. But I can't help but ask now, how many more bites can my mother endure?

I don't know when faith in technology and progress gives way to another kind of faith. Having devoted our lives to better living through engineering, when do we surrender to whatever lies outside the reach of our machines? —G. PASCAL ZACHARY

G. Pascal Zachary is a professor of practice at Arizona State University and author of *Endless Frontier: Vannevar Bush, Engineer of the American Century* (MIT Press, 1999). His mother, Rosalyn, died on 1 August.

ILLUSTRATION BY Dan Page

NEWS

1,500 MW

CAPACITY LOST TO BALTIC REGION WITH THE SHUTDOWN OF THE IGNALINA NUCLEAR POWER PLANT IN LITHUANIA

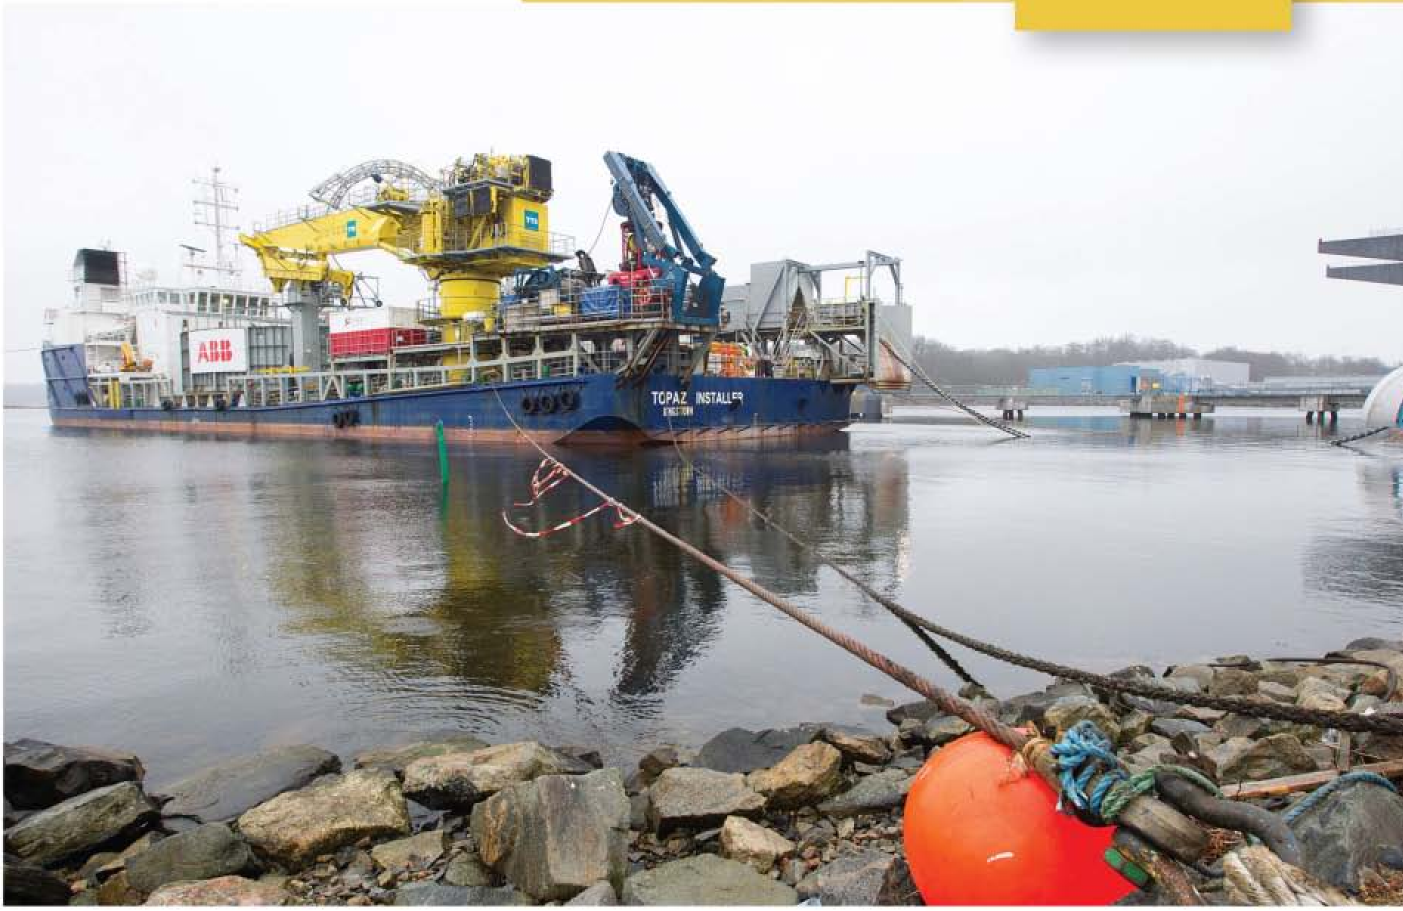

LITGRID

THE GREAT BALTIC DISCONNECT

One of the world's longest undersea HVDC cables and other new links will end reliance on Russian energy

POWER IS COMING: The *Topaz Installer* laid HVDC cable to link Sweden and Lithuania.

Former Soviet Union satellites such as Ukraine, Georgia, and the three Baltic states—Lithuania, Latvia, and Estonia—rely heavily on Russian energy, a dependence that colors the ongoing geopolitical tussle among the Russian Federation, the European Union, and NATO. But new energy infrastructure is beginning to free the Baltics from this drama, including two high-voltage direct current links to Lithuania nearing completion: the 500-megawatt LitPol Link with Poland and the 700-MW NordBalt cable to Sweden, which is among the world's longest subsea power links.

The links will give the Baltics the ability to get 100 percent of their imported power from non-Russian sources. Laying the 453-kilometer NordBalt cable was hindered by a series of nerve-wracking interventions by Russian naval ships, prompting Lithuania to lodge an official complaint with Moscow. Despite this »

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

THE BALTIC CONNECTION: Efforts to free the Baltic states from dependence on Russian energy include undersea HVDC links as well as an overland link to Poland. Synchronization with the European grid is expected later.

hurdle, both NordBalt and LitPol Link are expected to begin testing by December and should be in operation from January 2016, says Daivis Virbickas, CEO of the Lithuanian power grid operator Litgrid.

Virbickas says that LitPol Link and NordBalt represent a "new era" for Lithuania, the largest of the Baltic states. While Lithuania declared its political independence from the Soviet Union in 1990, it became more dependent on Russian energy after joining the EU in 2004, which required that it shut down the Ignalina nuclear power plant's Chernobyl-style reactors.

Overnight, Lithuania went from being one of Europe's major power exporters to heavy reliance on Russia. Since the closing of the last reactor, at the end of 2009, Lithuania has imported 70 percent of its power, roughly two-thirds of that from

LANDFALL IN LITHUANIA: The HVDC station in Klaipėda will be charged up in December.

Russia. Its domestic power plants burn natural gas imported from Russia's Gazprom. As a result, Lithuanians pay the highest energy prices in Europe.

Dependence also adds to the Baltics' sense of insecurity, which has risen in recent years as the Russian Federation has invoked power plays on other energy clients such as Ukraine and Georgia. Ukraine, whose Crimean region Russia annexed last March, now relies on Russian power imports since separatist activity has disrupted supplies from its eastern coal mines.

NordBalt and LitPol Link are, to a great extent, products of Russia's occupation of the Georgian region of South Ossetia in 2008. In 2009, the European Commission, the Baltic nations, and several neighboring states, including Sweden, forged a common action plan to end the Baltics' energy isolation and to synchronize their power grids with Europe's. That plan has been bankrolled by European financing.

Even though the lines aren't yet energized, dependence on Russia is already fading in Lithuania. Virbickas says futures contracts for electricity are

TOP: ERIK VRIELINK; BOTTOM: LITGRID

already 20 percent cheaper than they were last year, as traders factor in the impending availability of cheaper power from the EU and Scandinavia.

Natural gas options are also opening up. Last year Lithuania began operating the Baltic states' first liquefied natural gas terminal, which is receiving Norwegian gas.

The Baltics are also taking steps to synchronize with Europe's grid. Currently, the three countries' grids form the western flank of a massive interconnection that spans eight time zones and is dominated by Russia's power system. But earlier this year the Baltic transmission system operators hammered out a joint plan outlining how to desynchronize their AC grid from the Moscow-directed interconnection and synchronize it with Europe's. The specified power line reinforcements and upgraded grid controls could take a decade to complete and could cost the Baltics €1 billion (approximately US $1.08 billion), according to Virbickas.

That may be more than the Baltics can earn back from a more fluid power market. But the cost and effort is justified, says Virbickas, because it will boost supply security. He expects Lithuania to be an equal partner in the cooperative management of continental Europe's grid by the European Network of Transmission System Operators for Electricity, the Brussels-based consortium of European grid operators. In contrast, he says, the Baltic grid operators have little role in determining how the Russian grid is run.

One delicate question remains unanswered, however: the electrical fate of Kaliningrad, a Russian exclave that's sandwiched between Poland and Lithuania. Discussions on whether Kaliningrad will become an electrical island or join the European grid ended with the breakdown in Russia-EU relations last year. "It is a difficult project geopolitically because we have Kaliningrad in between," predicts Romas Švedas, formerly a Lithuanian diplomat and vice minister of energy. "It will be a lengthy process." –PETER FAIRLEY

ISTOCKPHOTO

NEWS

FBI WANTS BETTER AUTOMATED IMAGE ANALYSIS FOR TATTOOS

It's a tougher problem than facial recognition

Nothing makes a statement quite like a tattoo. And law enforcement in the United States increasingly uses them to help identify criminals and, sometimes, the victims of crime or natural disasters.

Today police take photographs of tattoos when suspects are booked, categorizing them using keywords defined in a biometric standard called ANSI-NIST-ITL 1-2011. The standard has eight main categories, such as "animal" and "plant," as well as 70 subcategories, such as "cat," "bird," "flower," and "leaf." The FBI maintains a database of tattoos as part of its Next Generation Identification Program, but searching by keyword is problematic because the categories aren't granular enough and different people often tag the same tattoo differently.

"It's very subjective as to what each person sees within a tattoo," says Eric Phillips, management and program analyst at the FBI's Biometric Center of Excellence, in Clarksburg, W.Va.

The stylized letter *D* insignia of the Detroit Tigers baseball team, for example, is easily misinterpreted, says Mei Ngan, a computer scientist at the National Institute of Standards and Technology (NIST). One person might recognize it as the team's emblem, but another might see it as just a letter, and yet another might consider it an abstract design.

That's why the FBI would prefer to use image-based tattoo recognition technology. Rather than rely on keywords, it would use an algorithm to compare and match features extracted from the image itself. The FBI turned to NIST, which issued a challenge last fall to assess the state of the art of such technology. Six organizations participated, running their algorithms against a set of data provided by the FBI. The organizations were

image-analysis and algorithm-development firm Compass Technical Consulting; the Fraunhofer Institute of Optronics, System Technologies and Image Exploitation; the French Alternative Energies and Atomic Energy Commission; nonprofit research contractor MITRE; security and identity tech firm MorphoTrak; and Purdue University.

In June, the six groups reported on how well their algorithms performed in five different types of searches. The algorithms did well in three of these searches, achieving success rates of 90 percent and above in detecting whether a given image contained a tattoo; identifying the same tattoo on the same person, over a span of time; and identifying a small segment of a larger tattoo.

The algorithms performed poorly–with hit rates as low as 15 percent–at two tasks: identifying visually similar tattoos on different people, and searching for similar tattoos across a variety of media, including sketches, scanned prints, and computer graphics.

The tattoo image algorithms are similar to those used in facial and other image-recognition technologies, says Anil K. Jain, a professor of computer science and engineering at Michigan State University, which licensed an algorithm developed by its researchers to MorphoTrak three years ago. The algorithms are all based on extracting key points in an image. However, where fingerprints have ridges and valleys, and faces have eyes and noses, tattoos have no standard features to identify and compare, he says.

(Jain recently proved that a person's fingerprints remain basically unchanged over time. Surprisingly, this assumption had gone unproved despite the heavy reliance on fingerprints in criminal justice.)

Fingerprints have ridges and valleys, and faces have eyes and noses, but tattoos have no standard features to identify and compare

Beyond highlighting the weakness in algorithms, the challenge demonstrated the need for improvement in two areas, says NIST's Ngan. First, she says, law enforcement needs to develop best practices in how to collect tattoo images: Take separate photos of each tattoo on each forearm rather than one photo of both forearms, for example, and make sure clothing or jewelry does not partially obscure the image. Second, the biometrics industry must better define what in a tattoo image it wants the technology to compare. "We need to give the algorithms a better definition of what we consider to be similar," Ngan says.

—TAM HARBERT

A WEARABLE TURNS BASEBALL PITCHING INTO A SCIENCE

Teams are keen, but some pitchers see a downside to the data

A wearable sensor that tracks strain on a pitcher's elbow is making waves in major league baseball (MLB). This season, 27 MLB teams and their minor league affiliates are trying out the device, called the mThrow, in the hope that it will help monitor pitchers' workloads, improve pitching mechanics, and prevent injuries. The device's maker, Motus Global, in Massapequa, N.Y., plans to officially launch a consumer version this month. Teams seem to like it, but some players might have reservations about sharing their data.

Injuries to professional pitchers in the United States have become epidemic. The reconstructive procedures known as Tommy John surgeries, which repair the elbow's ulnar collateral ligament (UCL), have increased among major league players from 14 performed in 2002 to 31 in 2014, according to the blog Baseball Heat Maps. Attempting to curb such injuries by arbitrarily limiting the number of throws or innings pitched hasn't been effective, says Thomas Karakolis, an expert on the subject at the University of Waterloo in Ontario, Canada. "Baseball managers should be figuring out the forces on muscles, tendons, and ligaments for each individual player and guiding them based on that," Karakolis says.

The mThrow could be the key to that. But first, Motus must derive meaning from the data it is amassing. "I don't think we quite have a grasp of all the information [the mThrow] is giving us, but I think it's the beginning of something very big," says Jamie Reed, senior director of medical operations for the Texas Rangers, whose major and minor league teams are using the device.

The mThrow consists of a compression sleeve with a small removable sensor worn in a pocket over the elbow. The sensor's accelerometers and gyroscopes

OLIVER STREWE/GETTY IMAGES

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

MOTUS GLOBAL

track arm movements, and the device wirelessly transmits the data to an app that calculates stress caused by torque on the UCL. The app also tracks a number of other metrics, such as pitch count, arm speed, release point, elbow height at release, maximum shoulder rotation, and "arm slot"—an indication of whether the throw was a side-arm pitch or over the top.

The app can gauge the amount of stress on a player's elbow joint by comparing the condition to similar cases in the company's database of pitcher workloads. The app can also compare stress levels created by, say, a fastball versus a curveball, and it can indicate when a pitcher's arm slot might be changing due to fatigue. But giving context for that—identifying the trends in workload that lead to injury—is something the company is still working on. "It's early now, but we're seeing interesting things," says Ben Hansen, chief technology officer at Motus. For example, the company is seeing a difference between players whose workloads increase and decrease smoothly and players with consistently high workloads or dramatic spikes and troughs, he says.

The mThrow offers much of the same biomechanics information that can be gleaned from motion-capture technology, mapping movements using multiple cameras that track markers placed on the body. Unlike motion capture, however, the mThrow sleeve goes anywhere the player goes. "Before, the only option was a motion-capture assessment once or twice per year. Now they get a smaller, more compact version every throw they make," says Hansen.

The Texas Rangers bought 10 mThrows, keeping a couple and giving one to each of its minor league teams. Two of the organization's players being tracked are recuperating from injury—an important use of the device, Reed says. "We are measuring their workloads so we don't stress them too much in their return to health," he says.

The Baltimore Orioles bought eight sleeves, all for the club's farm teams, says Rick Peterson, a director of pitching development for the team. "My guess is it's going to be a long time before the big league guys get into this, because they will not want to share this data," he says. "Big league players are so protective of their space. Their first question is: 'Can this be used against me?' "

That question is particularly pertinent for amateurs competing for a position through the league's draft system and for players who are looking to get a new contract. Biomechanics analysis could work in a player's favor if his mechanics are good. But if the data show red flags that could lead to a major injury, he won't be as desirable. Peterson says he already uses video footage to look for mechanics problems in potential draft picks, and biometrics data from wearables would be a likely next step.

Reed agrees. "Once we're comfortable with the Motus sleeve and the information it's giving us, I could see us using it in predraft workouts where we could measure stresses and mechanics on players before they sign," he says. "Certainly [we'd use it] after they sign to remold their mechanics."

It's a little harrowing to think that such data could cost a player his career for an injury that hasn't yet occurred—or may never occur. "It's an obvious concern, especially for guys in my shoes—minor league guys fighting for a chance to play in the big leagues," says Dillon McNamara, a pitcher for the Charleston Riverdogs, a minor league team, who has experimented with the Motus sleeve. "The mThrow is pretty much looking right into your elbow and seeing how it feels as you're pitching. That's pretty personal," he says. The players union, the MLB Players Association, did not provide comment to *IEEE Spectrum*.

The most useful application for mThrow may turn out to be for kids, and that's the market Motus plans to target. "Many of us feel, myself included, that some of the root causes of injuries at the professional level start at the 12-, 13-, and 14-year-old level," says Peterson.

"Kids are playing on multiple teams and playing year-round and trying to impress scouts all the time, so there's never a break for them from throwing," adds Bryan Bullington, a pitcher for the Orix Buffaloes in Japan's Pacific League. It's different as a professional, where players can have "a two- to four-month period in the off-season where they don't even pick up a ball," he says.

A wearable sensor like mThrow could address that concern by giving parents and coaches a cold, hard number indicating the stress on the child's arm and the number of pitches. "It might give someone a more concrete feel for what the arm is going through," says Bullington.

—EMILY WALTZ

NEWS

NEWS

STATE OF AUTOMATION

Autonomous vehicle license applications show California's robocar dominance

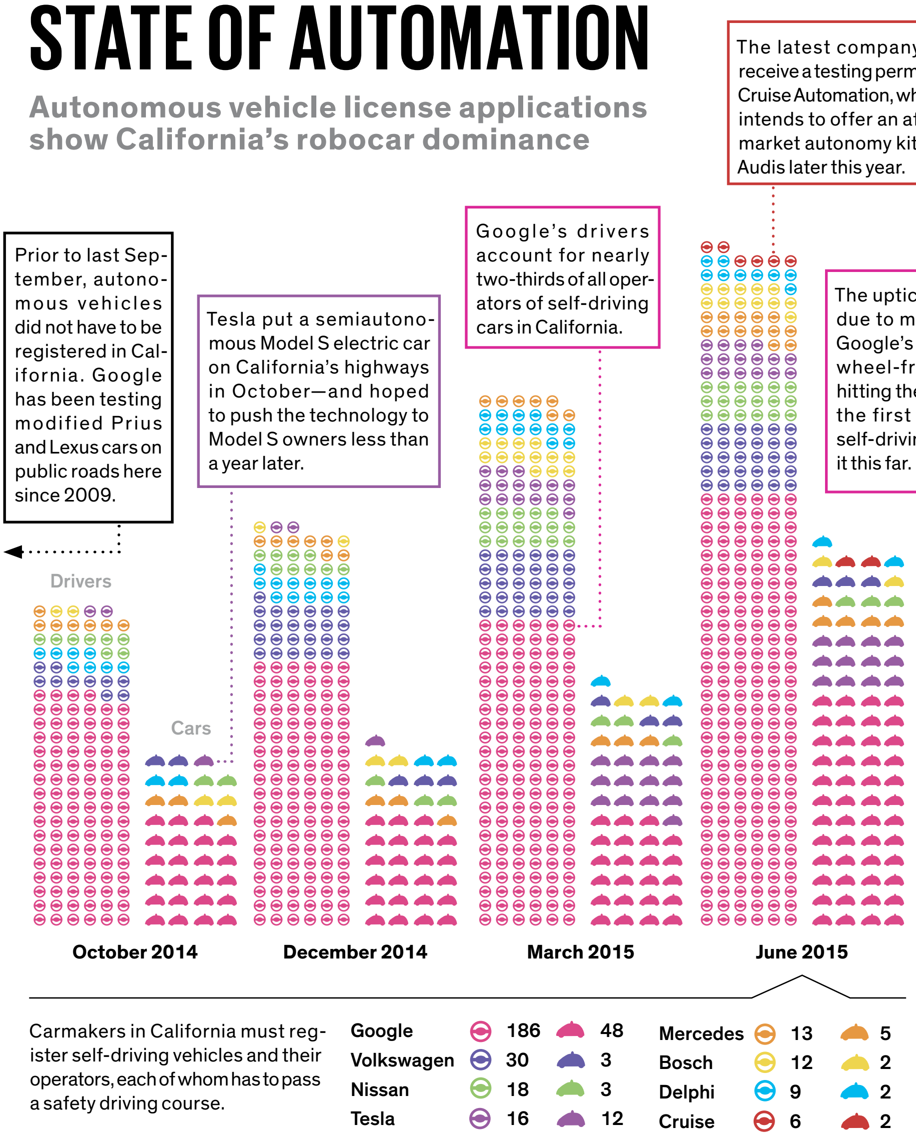

SOURCE: CALIFORNIA DEPARTMENT OF MOTOR VEHICLES

California almost certainly has more self-driving cars and operators on public roads than the rest of the world. That the state is experiencing an autonomous gold rush should not come as a surprise. For a start, its reliably mild, dry, and sunny climate is perfect for road-testing early generations of vehicles that still balk at snow, fog, and heavy rain.

Neighboring Nevada is similarly meteorologically blessed and began licensing experimental autonomous cars several years ahead of California. But the Golden State has the advantage because it's also home to Silicon Valley, where everything a prospective driverless car manufacturer needs–software engineers, hardware geeks, roboticists, venture capital–is available in a near-endless supply. Virtually every large technology company–and most mainstream automakers–have offices in the Bay Area.

But the main reason for California's supremacy in vehicular automation is the presence of a single company. While Google did not invent the self-driving car, it can lay claim to having invented the industry of self-driving cars–purchasing startups, hiring experts, and developing essential mapping and navigation technologies. From the word go, Google has remained one step ahead of the competition. It was the first to test experimental vehicles at scale, the first to move from relatively safe highways to unpredictable city streets, and the first to construct a purpose-built, steering-wheel-free, self-driving prototype. Even today, after the arrival of global carmakers and upstarts like Tesla, Google has around twice as many autonomous vehicles and drivers on California's roads as everyone else combined.

Of course, California's dominance in self-driving cars is hardly guaranteed. Mcity, a secure urban testing environment for high-tech cars at the University of Michigan, is hoping to lure traditional Detroit carmakers, and the state already has some of the most industry-friendly regulations. The United Kingdom and France opened their streets to experimental autonomous vehicles earlier this year, while the Chinese tech giant Baidu has been busy testing self-driving cars (developed with BMW) around Beijing and Shanghai.

But California still has a few tricks up its sleeve, such as its own connected and autonomous test center, a 20-square-kilometer former naval weapons site in Concord called GoMentum Station. And if, as some observers suggest, Apple is working on its own high-tech car, California should continue to boast the most roboticized roads in the world well into the future. —MARK HARRIS

INFOGRAPHIC BY Erik Vrielink

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

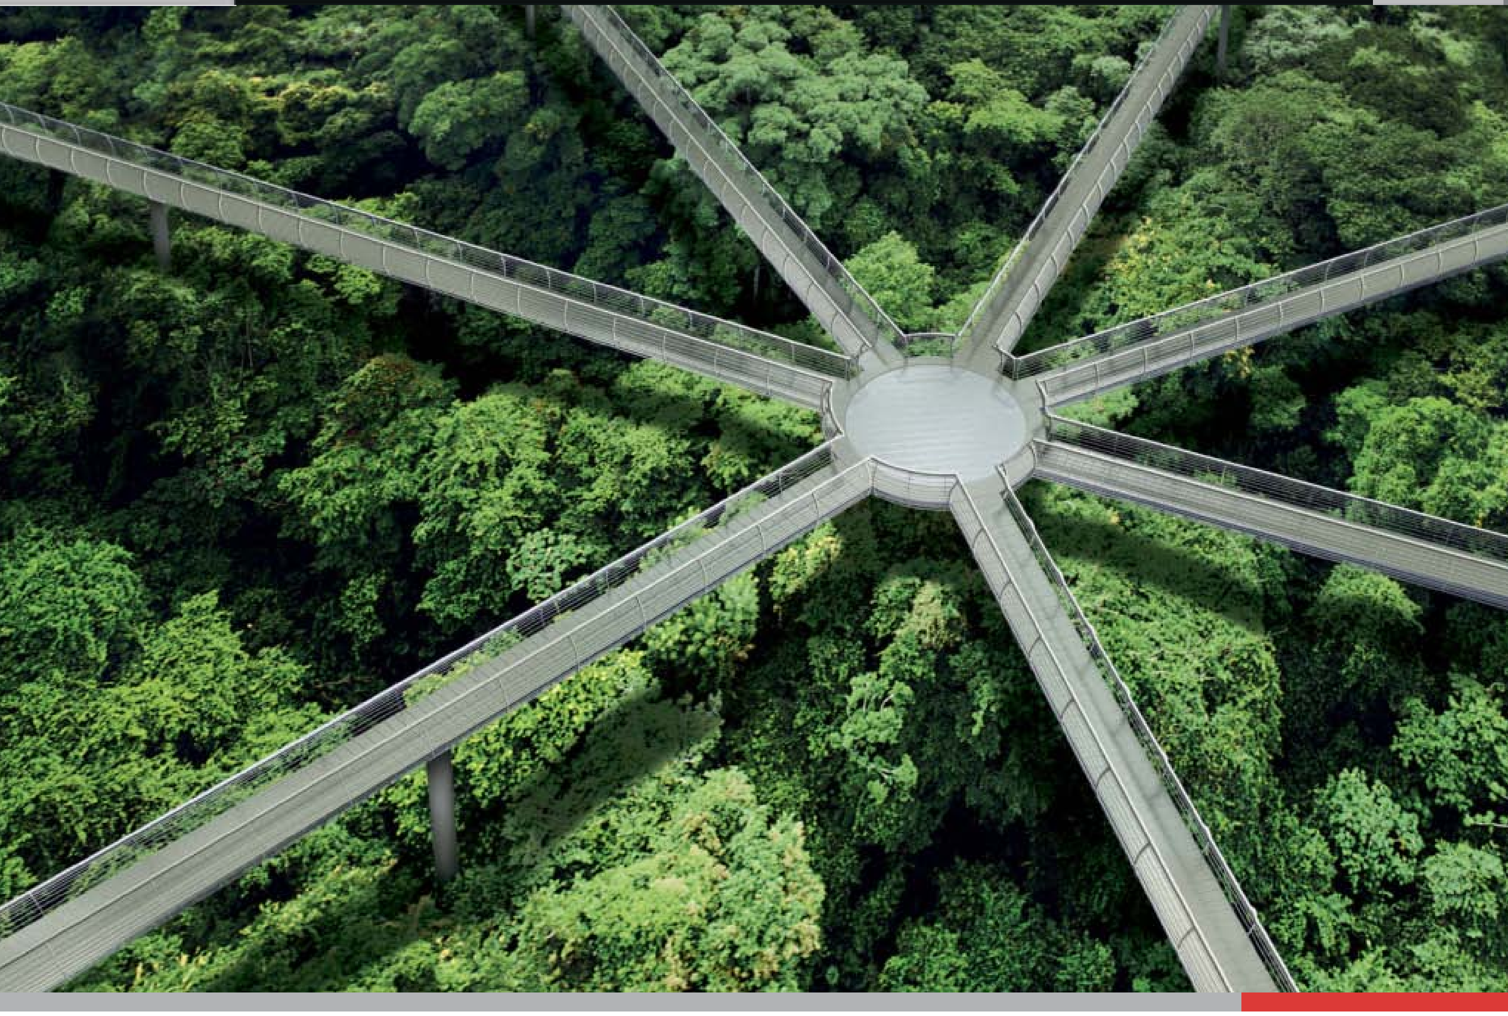

PHOTO-ILLUSTRATION BY Rex Features/AP Photo

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

UNDER-GROUND PLAY SPACE

NEW YORK CITY'S High Line Park, once an abandoned section of elevated railway on the west side of Manhattan, is now a 2.33-kilometer-long urban oasis that draws millions of visitors each year. Now a plan has been hatched to turn another out-of-use railway into a public park teeming with greenery. But this one, located on Manhattan's Lower East Side, will grow inside an unused underground trolley terminal. If things go according to plan, the subterranean park, called the Lowline, will open in 2018. But how do you get greenery to grow in a tunnel? The park's designers plan to use aboveground parabolic collectors to gather the sun's rays and direct them underground. The solar technology transmits the wavelengths of light needed to support photosynthesis, letting grass grow where the sun don't shine.

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

RESOURCES

25 BILLION: THE NUMBER OF THINGS THAT WILL BE INTERNET CONNECTED BY 2020, ACCORDING TO GARTNER

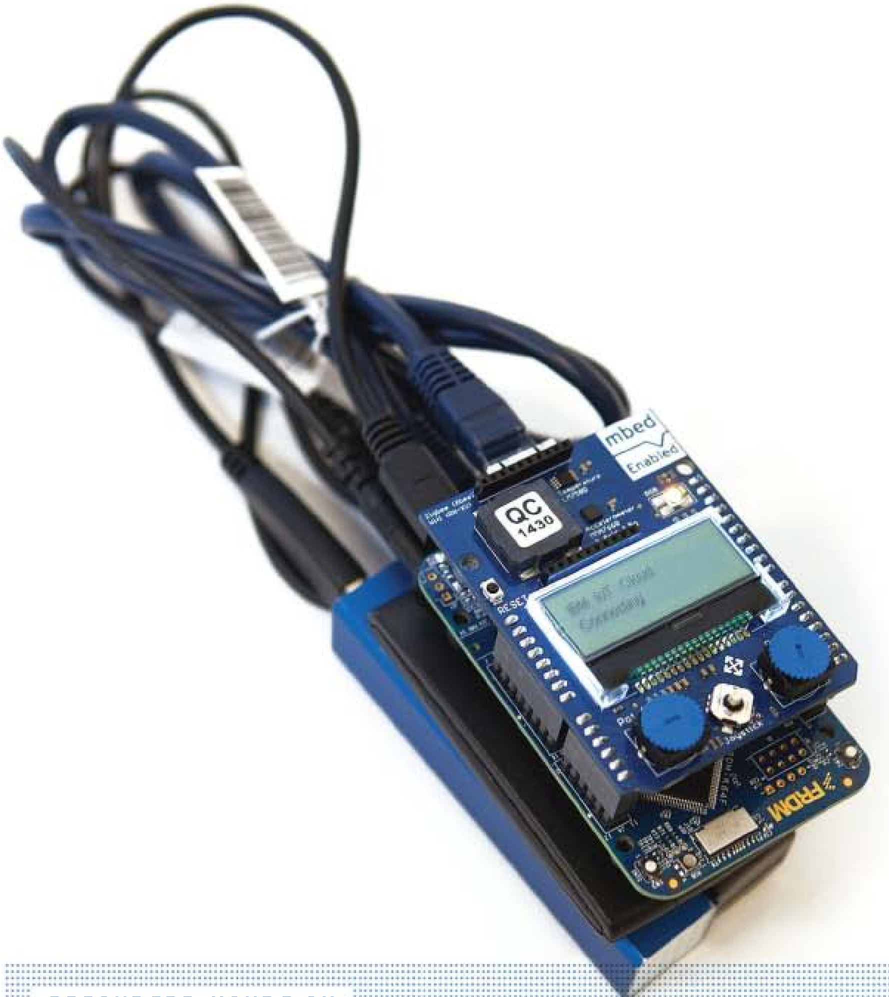

RESOURCES_HANDS ON

A KIT FOR THE INTERNET OF THINGS

IBM AND ARM TEAM UP TO MAKE EXPERIMENTING EASY

Just like pretty much every other tech outlet, *IEEE Spectrum* has been nattering on for a while now about the coming Internet of Things (IoT). This is the vision of a world where not only our PCs and smartphones are connected to the Internet but where nearly every device that runs on electricity—fridges, scales, lightbulbs—has sensors, a small glob of processing power, and a network connection. • Makers have been building IoT-style gadgets for some time, but it can be quite fiddly to glue together all the required bits of hardware and software. In a bid to make it easier for developers to get their feet wet and start experimenting with the concepts behind the IoT, IBM and ARM have teamed up to create the US $120 Mbed IoT Starter Kit, intended to be "a slick experience. . .particularly suitable for developers with no specific experience in embedded or web development." • On the hardware side, the Starter Kit consists of an FRDM-K64F microcontroller and an application shield that fits on top. Pin-compatible with Arduino boards, the FRDM-K64F has an Ethernet interface built in and is based on a chip with an ARM Cortex-M4 core that's capable of running at up to 120 megahertz (in comparison, the Arduino Mega clocks in at 16 MHz). The application shield is equipped with a temperature sensor, three-axis accelerometer, RGB LED, five-way joystick, two potentiometers, a loudspeaker, and a small 128- by 32-pixel LCD. • On the software side is IBM's Internet of Things Foundation and its Bluemix service, which lets you build and deploy apps for IoT devices. • Initial setup was as easy as promised. I pushed the FRDM-K64F and application shield together, snapped in a live Ethernet cable, and hooked up a micro-USB cable to my computer to provide power. While the ▸

PHOTOGRAPH BY Randi Klett

THE THING

Starter Kit booted up, I navigated to the URL of the IBM Quickstart page. By the time I got there, the application shield announced via its display that it was connected. I called up my Starter Kit's ID number on the shield's display, typed the number into the Quickstart page and was presented with a live visualization of sensor and other data streaming from all the shield's input channels. The whole process—from opening the box to watching data scroll by—took less than 3 minutes (including the time spent fishing around at the bottom of a drawer for the right USB cable), which has to be some kind of Hands On record.

The possibilities inherent in the Starter Kit's sensors are somewhat limited by the fact that this board is tethered by Ethernet and USB cables—for example, the length of an Ethernet cable sets sharp limits on where you might choose to monitor temperatures.

Consequently, I wanted to go wireless. There's a socket on the application shield for ZigBee and Wi-Fi daughterboards, but as most of the online examples and documentation assumed an Ethernet connection, I opted for a quick-and-dirty approach. I bought a small Netgear Ethernet to Wi-Fi adapter, set it up to talk to my Wi-Fi router, and plugged the Starter Kit in. As well as being small in size, the Netgear adapter can also run off a USB port. So I dug out a USB battery of the sort used to recharge a smartphone to provide power. A second battery powered the Starter Kit, and I was portable.

IBM's Bluemix service allows you to do more interesting things with the data coming from the Starter Kit than just watching it make simple charts. After a few minutes perusing the official online tutorial, I had the Starter Kit hooked into IBM's browser-based Node-RED editor and running the suggested sample program. This reads the temperature sensor and prints different messages depending on whether or not the temperature is above some threshold value. The editor allows you to configure and wire together graphical blocks that represent various input, function, and output nodes. This lets you build complex

THE MBED IoT STARTER KIT is a microcontroller and shield [top], to which I added an Ethernet to Wi-Fi adapter [second from top] and batteries [second from bottom]. Programs to process sensor data in the cloud are written using an online graphical editor [bottom].

FROM TOP: RANDI KLETT (3); STEPHEN CASS

applications that live in the cloud and that can communicate with IoT devices and things like smartphone apps or online databases.

But things got complicated fast when I tried to move to the next stage: sending commands from the cloud-based program to the Starter Kit. This requires the Starter Kit to be registered with the Internet of Things Foundation in order to get an organization ID, authentication token, and so forth. Once I had these credentials, I had to go to IBM's code repository, where the default program that runs the Starter Kit lives, paste details into the code, compile the code online, and download the binary file to my MacBook Pro. The Starter Kit's USB connection doesn't just provide power; it also allows the kit to look like a flash drive to a computer where binaries can be "dropped off" for installation.

In theory, transferring the updated binary should have taken just a moment of drag and drop, but in practice it resulted only in a cryptic error message. After a bit of head scratching and a helpful e-mail exchange with an engineer at ARM, it was determined that the problem probably lay in something specific to Apple's USB implementation. So I wandered through the IEEE's New York City HQ till I found a PC I could commandeer. From there, the binary installed as it should, and I could pick up the output from the newly registered Starter Kit in the Node-RED editor running back on my MacBook Pro. Getting the board to respond to basic commands sent by IBM's starter "recipe" was a matter of plugging in some more logic blocks and adding an output node, so that the blink rate of an LED on the microcontroller could be controlled by turning a potentiometer on the shield. However, getting all this to work was complicated by the scattered and sometimes inconsistent nature of the available documentation. The Starter Kit has the potential to be an impressive intro to the idea of a world in which digital smarts are seamlessly distributed among the cloud and simple pieces of hardware. All it needs is just a little more polish. **—STEPHEN CASS**

WEARABLES FOR FRAIL SENIORS

GADGETS MONITOR ELDERS' HEALTH WITHOUT INVADING PRIVACY

FASHION FORWARD: Designers of wearables are increasingly taking fashion into account, a trend that continues with CarePredict's tracker for seniors.

DON WEST

This is the final installment in a three-part series looking at exemplars of products that demonstrate how wearable devices are increasingly infiltrating everyday life; previous articles have looked at a wearable baby monitor and an activity tracker for pet dogs. In this article we look at wearables that monitor the health of elderly people unobtrusively. Rather than giving family members details about the wearers' daily activities, these systems function in the background and raise alerts only if anomalous behavior occurs.

Satish Movva lives just 10 miles from his parents and, as a dutiful son, sees them once a week. But last year, during a routine visit, he was surprised to see his father shuffling his feet. It had begun several days earlier, but his parents had not bothered to tell him. Movva knew, however, that shuffling could be a warning sign for serious health conditions such as heart failure.

Movva's father subsequently received medical treatment and the shuffling stopped, but the experience motivated Movva to research how to monitor his aging parents remotely. Not finding something suitable on the market, Movva, who had spent nearly three decades in IT, focusing on health care for the past two, founded the startup CarePredict in 2013 to develop Tempo, a wearable for seniors that will start shipping this year. Worn as a wristband, Tempo is packed with a nine-axis MEMS sensor (incorporating a three-axis accelerometer, a three-axis gyroscope, and a three-axis magnetometer), which tracks movements such as walking or lying down. Battery-powered beacons positioned around the home or a senior center, such as in the bathroom, bedroom, and kitchen, enable the system to recognize where the activity is taking place. A base station with Wi-Fi or Ethernet—and a cellular radio for backup—compresses and transfers the data to the cloud and records it in the wearer's "rhythm journal."

CarePredict can use the journal data to deduce a series of regular habits—for example, when toothbrushing occurs in the bathroom—and will log how long it takes to complete each customary activity and the time it occurred. Once a baseline is established, which takes about a week, CarePredict then monitors daily activity to detect any changes to the person's routine. In theory, a human observer could notice these changes as well, "but a human would have to be there 24 hours a day," says Movva.

Another startup, Lively, has designed a safety watch as well (also called Lively) that interacts with sensors located around the home and a central base station. Similar to Tempo, Lively links to a system that looks for changes to daily patterns and sends a reminder—for instance, if the senior has forgotten to take his medication.

The Lively safety watch, which is available now for US $50 with service subscription plans ranging from $28 to $35 per month, also comes with an emergency button. When pushed, the button alerts Lively's care team to dispatch emergency services. An extended Bluetooth antenna, coiled inside the watch, allows the wearer to travel more than 450 meters from home and still stay connected to the hub.

Lively is also working on a feature that will allow the watch to recognize if its wearer has fallen and automatically dispatch emergency services without the press of a button. The system will calculate the *g*-force and look at other clues, such as the amount of movement in the seconds after the accident, to figure out if help is needed.

Ultimately, seniors want to stay independent without having to don a device that screams "old age." Lively CEO Iggy Fanlo believes that this won't be the case with his company's product. "All the cool kids are getting smart watches," he says. **—ELLEN LEE**

DEFYING THE DRONE DATA DELUGE

KSI DATA SCIENCES GIVES OPERATORS CRITICAL CONTEXT

KSI DATA SCIENCES

KSI Data Sciences, based in a former tobacco warehouse in Danville, Va., is unlike most drone startups. For one thing, this company is not designing robotic hardware or offering field services. Instead it's doing something that is perhaps more important—figuring out how to handle the information collected by our robotic minions.

"Anybody can push out a video stream, but when you start wrapping it with all the telemetry data, all the sensor data, when you're trying to plot it on maps, when you're trying to make sense of the area around you, not just push out video—it gets much more complex," says the company's cofounder Jason Barton.

Soon after the terrorist attacks of 9/11, Barton started a company called EchoStorm Worldwide, which helped the U.S. military manage video from the many drones being flown in Iraq and Afghanistan. In early 2012, two years after selling EchoStorm, Barton saw the need for something similar in the burgeoning world of commercial drones. He teamed up with Jon Gaster, who had worked in data management for the Hollywood film industry, to start KSI Data Sciences. Their vision was to create a system that didn't require formal training, much less an engineering degree, to use. The result is a user interface to their cloud-based service that resembles someone's Facebook page more than anything else.

KSI's initial market targets include police, fire, and other first responders that use robotic vehicles in their work, along with utility companies that use such platforms to inspect their equipment remotely (KSI doesn't limit its scope to aerial drones, working with terrestrial and submarine remote-controlled robots as well). KSI provides its MissionCaster box to which customers feed video and associated data, formatted according to either the military's MISP (Motion Imagery Standards Profile) protocol or to the MavLink standard, which many small drones use to communicate with their ground-control stations.

KSI's hardware uploads the video and data to the Internet over an LTE cellular network. At that point anyone with appropriate credentials can log into the system and view what the drone is surveying, while also having maps and other information important to gain situational awareness near at hand—all with only a few seconds' delay.

Frank Gillett, an analyst with Forrester Research, says he hasn't heard of other companies pursuing anything similar, but he's not surprised to see this product and service come into being. "You want to be able to integrate the mobile sensor with other context," he says. Those who benefit most from that ability are remote users, say a town's fire and police chiefs, who from their offices can even open a video chat with the robotic vehicle's operators. Think of KSI's system as a GotoMeeting for drone operations.

Barton and Gaster are fully aware that their system could be valuable for more than just those supervising robots: Handling feeds from police dash- and bodycams leaps to mind. Today those videos are recorded rather than uploaded to the Internet in real time, but that could very well change in the next few years. Still, for the moment, KSI is focusing on robotic camera platforms, in part because that's where this fledging company sees its strongest competitive advantage. It probably also has something to do with the fact that messing with robots is so much fun. **—DAVID SCHNEIDER**

Name: KSI Data Sciences **Location:** Danville, Va. **Founded:** 2012 **Employees:** 8 **Funding:** US $250,000 from the Launch Place plus other undisclosed amounts

FOR WOMEN ENGINEERS, MENTORING HELPS

AT SCHOOL AND AT WORK, MENTORS IMPROVE JOB SATISFACTION

EZRA BAILEY/GETTY IMAGES

During Estefania Ortiz's first internship as a freshman computer science major, the Stanford University student felt frustrated, insecure, and unsuccessful.

But when she sought a mentor's advice three-quarters of the way through her internship, she received valuable guidance on what to expect and how to salvage the experience. Managers are not like teachers, her mentor taught her, and she shouldn't expect encouragement and regular feedback.

"This mentor was saying, 'Pick your battles. Focus on networking.' Basically, she gave me a strategy to end on a good note," says Ortiz, who begins her senior year this fall. "That conversation with my mentor changed my whole course."

It's been widely known for decades that women in STEM fields switch majors away from science, technology, engineering, and mathematics at higher rates than their male counterparts. And this pattern continues after graduation—even for those who attain master's and doctorate degrees. But research shows that having mentors—both peer mentors and more senior mentors—reduces the number of women who leave engineering and increases job satisfaction for those women who stay.

"The feelings of competence and efficacy are stronger when you are less isolated and feel more confident in your job performance," says Naomi C. Chesler, a biomedical engineer who has studied gender disparities in engineering education for two decades.

Women who enter college confident and enthusiastic can face a rude awakening when encountering professors who undervalue their contributions and a culture geared more toward competition than collaboration and cooperation. Kimberly Bryant, an electrical engineer and founder of Black Girls Code, agrees, saying mentors have been pivotal to her career since her first year in college. Her first mentor, who was a couple of years ahead of her with the same major at Vanderbilt University, in Nashville, advised her on which classes to take and how to select an advisor. As Bryant was one of the few female electrical engineering students, her mentor also helped her feel less alone.

Later, while rapidly rising in her career at Pfizer, Bryant was recruited to join an all-white-male team of directors. After encountering friction from other directors in staff meetings, she kept silent. The man who had recruited her pulled her aside, told her he wanted her perspective, and encouraged her to speak up. Having that mentor in her corner helped her share her insights, she says, even when facing pushback from other colleagues.

Women who have been successful in STEM fields report having both informal peer mentors as well as more senior mentors. While men and women in STEM place about equal value on formal mentors, women value informal mentors more, according to Marie Garland, executive director of Syracuse University's National Science Foundation Advance IT initiative for increasing the participation and advancement of women in engineering careers.

What are some guidelines for those seeking the best out of a mentor-mentee relationship? Lynn Mayo, a vice president at AECOM, says she has sought guidance from mentors throughout her 30-year career and suggests forming a relationship with a mentor before asking for career advice. "It's much easier to mentor someone when you know their strengths, weakness, and aspirations," she says. Mayo also finds it helpful for the mentee to ask focused questions, as it's much easier to give advice on a specific topic than respond to a general request for advice. Mentees should also respect the time of their mentors: Save senior mentors for the important questions; use peer mentors to gain a different perspective or to talk through a challenge.

How can young engineers find mentors, and how can more experienced engineers willing to be mentors make themselves known? Those seeking mentors should look for people they want to emulate and with whom they have a professional relationship. Mentors need not be the same gender, and a mentor could be a more experienced colleague or a supervisor. Veteran engineers can open the door to potential mentees by being approachable and offering to answer any questions.

"It's best to have several mentors," says Mayo. "When you have multiple mentors, you can approach different mentors based on specific needs." —**THERESA SULLIVAN BARGER**

"NEW JAPAN" AT 70

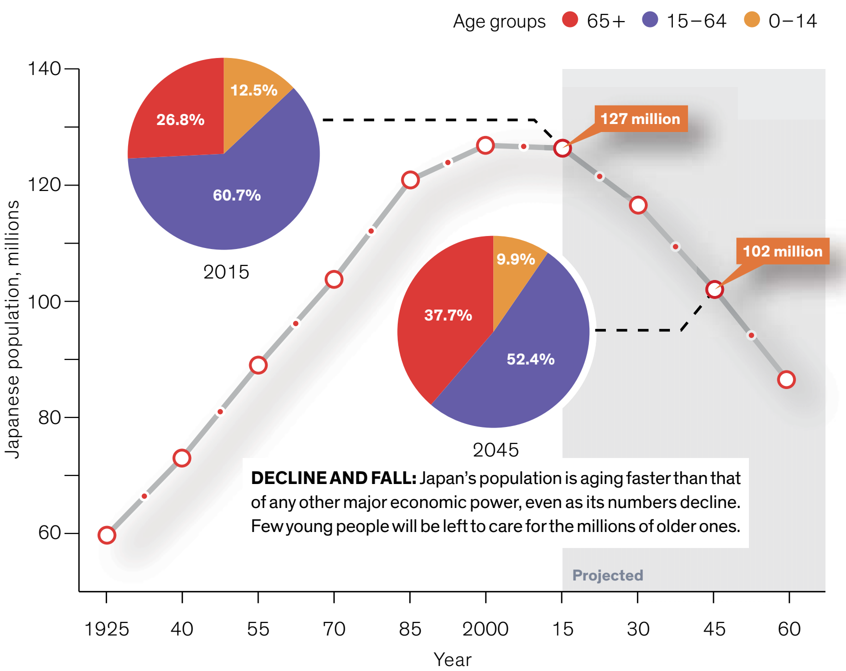

DECLINE AND FALL: Japan's population is aging faster than that of any other major economic power, even as its numbers decline. Few young people will be left to care for the millions of older ones.

SOURCE: IPSS (NATIONAL INSTITUTE OF POPULATION AND SOCIAL SECURITY RESEARCH)

ON 2 SEPTEMBER 1945, REPRESENTATIVES OF THE Japanese government signed the instrument of surrender on the deck of the USS *Missouri*, anchored in Tokyo Bay. So ended perhaps the most reckless of all modern wars, the outcome of which was decided by U.S. technical superiority even before it started. Japan lost in material terms even before it attacked Pearl Harbor: In 1940 the United States produced roughly 10 times as much steel as Japan did, and during the war the difference grew further. ● The devastated Japanese economy did not surpass its prewar peak until 1953. But by then the foundations had been laid for the country's spectacular rise. Soon its fast-selling exports ranged from the first transistor radios (Sony) to the first giant crude-oil tankers (Sumitomo). The first Honda Civic arrived in the United States in 1973, and by 1980, Japanese cars claimed 30 percent of the U.S. market. Japan, totally dependent on crude-oil imports, was hit hard by the OPEC oil price rises of the 1970s, but it adjusted rapidly by pursuing energy efficiency, and in 1978 it became the world's second largest economy. By 1985 the yen was so strong that the United States, feeling threatened by Japanese imports, forced its devaluation through the Plaza Accord. But even afterward the economy soared: In the five years following January 1985 the Nikkei index rose more than threefold. ● It was too good to be true; indeed, the success reflected the working of an enormous bubble economy driven by inflated stock and real estate prices. In January 2000, ten years after its peak, the Nikkei was still at only half its 1990 value, and only recently has it risen above even that low mark.

Once-iconic consumer electronics manufacturers like Sony, Toshiba, and Hitachi now struggle to be profitable. Toyota and Honda, global automotive brands once known for their unmatched reliability, are recalling millions of vehicles. Takata's defective air bags recently resulted in the biggest recall of a manufactured part ever. And Yuasa made unreliable lithium-ion batteries for the Boeing 787. Add to this the rapidly changing governments, the March 2011 tsunami followed by the Fukushima disaster, and worsening relations with China and South Korea, and you get a worrisome picture indeed.

But in the long run the fortunes of nations are determined by population trends. Japan is not only the world's fastest-aging major economy (already every fourth person is older than 65, and by 2050 that share will be nearly 40 percent), its population is also declining. Today's 127 million will shrink to 97 million by 2050, and forecasts show shortages of the young labor force needed in construction and health care. Who will maintain Japan's extensive and admirably efficient transportation infrastructures? Who will take care of millions of old people? By 2050 people above the age of 80 will outnumber the children.

Fortunes of all major nations have followed specific trajectories of rise and retreat, but perhaps the greatest difference in their paths has been the time they spent at the top of their performance: Some had a relatively prolonged plateau followed by steady decline (both the British Empire and the 20th-century United States fit that pattern); others had a swift rise to a brief peak followed by more or less rapid decline. Japan is clearly in the latter category. Its swift post-World War II ascent peaked in the late 1980s, and it's been downhill ever since: in a single lifetime from misery to an admired–and feared–economic superpower, then on to the stagnation and retreat of an aging society. ■

Recording Data Has Never Been This Easy!

The all new TrendCorder TRC-18 from Dewetron is easy on the eyes, and even easier on your wallet. Set-up and operation is like a streamlined app, and includes multitouch pinch and zoom.

Our powerful DejaView® capability allows you to scroll back and zoom in on any part of the data even whill still recording at full speed! Even if that data was hours or days in the past, you don't have to stop recording to examine it at any resolution. And the price?

Only $19,950 for an 18 channel system

PLUS - you can add more fast channels by plugging in another TRION module – or add dozens of slow channels using our external EPAD 8-channel bricks. The system is 100% modular and field-upgradeable.

Toll-free: +1 877-431-5166
www.trendcorder.info

Please contact us to set up a demonstration today!

- 18 analog input channels
- Wide-aspect multitouch display
- Up to 250 kS/s/ch sample rate
- 1500 V isolation on all channels!
- 8th order filters on all channels!
- Multi-touch display with pinch and zoom operation.
- DejaView® provides unlimited scroll back and viewing resolution during recording

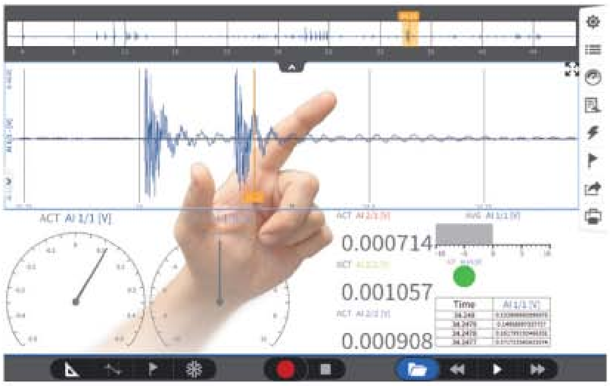

TRENDCORDER
+1 401-284-3750 www.dewamerica.com sales@dewamerica.com

TECHIES ON TV

Shows about engineers are finding an audience

A LONG, LONG TIME AGO—BACK IN THE 1980S—THERE WAS a concern within the IEEE about how to raise the public esteem of engineers. In our opinion, they didn't seem to be sufficiently appreciated. It was a time when doctor and lawyer shows were big on television, so an idea floated around some of our committee meetings: Could we get an engineer show on TV? • One distinguished colleague suggested, possibly in jest, a show entitled "L.A. Engineer," riffing off the then-hit U.S. show "L.A. Law." What a delightful fantasy that was! But I realized immediately that it would never happen, and even if it did, no one would watch it–not even us. • But times have changed. I don't think engineers worry much about the public's esteem now. We're too busy starting companies, changing the world, and stuff like that. And with no impetus from us, there is now a comedy series on the U.S. cable channel HBO entitled "Silicon Valley." It's all about engineers and computer scientists trying to start a company called Pied Piper, whose music app of the same name is based on a new data-compression algorithm. (Meanwhile, over on another U.S. cable channel, AMC, there is the 1980s period tech drama "Halt and Catch Fire," while one of the highest rated shows on U.S. broadcast television–"The Big Bang Theory"–features an aerospace engineer as one of its core characters.) • Okay, so we got a few things wrong with our "L.A. Engineer" idea, starting with the venue. Silicon Valley and its technology culture were only nascent back then. Another thing: No one ever thought our show would be a comedy. After all, this engineering is serious business. Yet "Silicon Valley" is in large part a parody, with caricatures of both the techies and the venture capitalists. Still, the caricatures aren't so broadly overdrawn that we engineers don't recognize familiar personalities and technical and business issues.

Pied Piper's engineers grapple with funding, hiring, business strategy, competition, management fads, intellectual property, and legal issues. But the one thing we don't see, however, is them actually *working*. That is, after all, boring–at least in the eyes of HBO.

Nonetheless, I got a thrill from hearing mentions of Shannon, Huffman, and Lempel-Ziv compression. Imagine this on a popular TV show! On some occasions we see "LZ" written on their backboard, and another time it's "LZW," giving mention to Terry Welch, who wrote the original compression software and improved and popularized the Lempel-Ziv algorithm.

Of course, the TV techies outdo the known lossless compression algorithms, with software that gets high marks on a fictitious metric called the Weissman score, which has since become a real measure as reported by *IEEE Spectrum* last year. This does raise some technical eyebrows, since Lempel-Ziv compression has been shown to be optimum. However, that's for a theoretical case that doesn't completely hold in practice.

In compressing written language, for example, LZ algorithms build a dictionary as they go along. However, humans bring a vast amount of understanding and context not present in that dictionary. Moreover, the underlying statistics can change dramatically, say, in going from a Shakespeare play to a medical textbook. So, yes, it is possible to get better compression than that of known algorithms.

The TV engineers frequently brag that their algorithm is lossless. Other things being equal, lossless is obviously better, but the world today largely lives on lossy algorithms like MP3, JPEG, and MPEG, where small concessions are made in fidelity to enable large gains in compression efficiency.

So maybe Pied Piper will crash and burn next season. It looks like a lot of people are going to stay tuned to find out. ■

ILLUSTRATION BY Greg Mably

MULTIPHYSICS SIMULATION

Sponsored by COMSOL

Special Advertising Section to: IEEE SPECTRUM

SEPTEMBER 2015

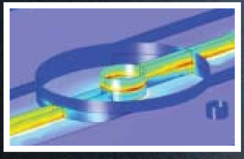

INTEL IMPROVES SIGNAL INTEGRITY AND SPEED
PAGE 18

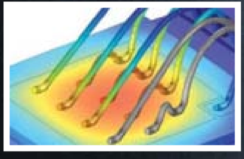

ABB EXTENDS LIFETIME OF POWER TRANSISTORS
PAGE 20

SIMULATION POWERS WIRELESS CHARGING AT WITRICITY

PAGE 8

SIMULATION APPS ARE MOVING TO THE MAINSTREAM

By **JAMES A. VICK, SENIOR DIRECTOR, IEEE MEDIA; PUBLISHER, IEEE SPECTRUM**

CHANCES ARE THAT THIS ISSUE of *Multiphysics Simulation* will have a profound impact on what you know about simulation. If you think of simulation as the exclusive domain of a select few R&D specialists, then you'll be surprised by the contents of this issue. No longer are the high costs and lack of custom simulation and analysis tools a roadblock for bringing simulation to everyone.

Easy-to-use specialized simulation apps are moving to the mainstream. A great example of a company utilizing simulation apps to improve R&D comes from Cypress Semiconductor of San Jose, CA. Perhaps best known for their smartphone touchscreen solutions, Cypress is using simulation apps to aid in the design of a wide array of consumer and industrial products. The modeling and simulation of capacitive sensors in touchscreens begins with the R&D engineers. However, instead of having to run repetitive simulations for every individual case, they are now building ready-to-use apps and distributing them to other departments. Their worldwide customer support teams can now access these apps and make use of them immediately with no learning curve and at a fraction of the cost of deploying a fully featured model. And the creation and distribution of these apps can all be done within one software environment.

This issue of *Multiphysics Simulation*, sponsored by COMSOL, Inc., offers a great way to learn about simulation application design and innovative simulation projects. If you are designing a power supply unit, first check out the article on wireless power transfer from Witricity—no cords required. Graphene is another topic that has inspired a lot of interest lately and that is now being applied in real-life environments. Within this issue, leading experts from Purdue University discuss how to efficiently and accurately simulate graphene-based photonic devices.

I hope you find this issue of *Multiphysics Simulation* inspiring! ◎ *Email: jv.ieeemedia@ieee.org*

CONTENTS

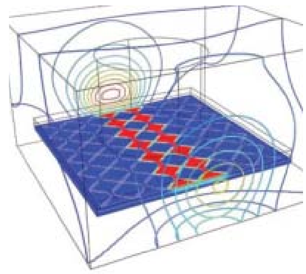

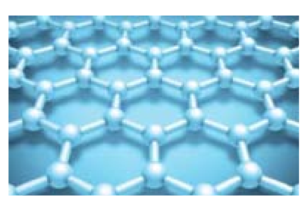

ON THE COVER: An electric vehicle charging wirelessly, parked above a charging pad. Automakers, electronics manufacturers, medical implant designers, and others are working with WiTricity to create a more efficient and consumer-friendly experience in many technology markets. See the full article, starting on page 8, for more details. Illustration is courtesy of WiTricity.

+ ONLINE: www.comsol.com/products

GETTING TOUCHY-FEELY WITH TOUCHSCREEN DESIGN

Cypress Semiconductor uses simulation to optimize touchscreen design for use in diverse products including smartphones, laptops, automotive and industrial applications, home appliances, and more.

By **ALEXANDRA FOLEY**

TO MAKE A PHONE CALL, compose a text message, or even to beat the next level of an Angry Birds™ game, we rely on being able to pick up our smartphone and interact with it without a second thought. No matter the size of our fingers, whether or not we have recently applied hand cream, or if the phone is resting on a flat surface, the touchscreen responds seamlessly, bringing just one more thieving farm animal to justice.

Engineers at Cypress Semiconductor, the leading supplier of smartphone touchscreen technologies and touch-sensing solutions, are hard at work making this possible, ensuring that touchscreen applications perform flawlessly under a variety of conditions. "And it's not just about smartphones," says Peter Vavaroutsos, a member of the touchscreen modeling group at Cypress. "Our technologies are used in smartphones, mp3 devices, laptops, automotive environments, industrial applications, home appliances, and more. For each of these uses, a different design is needed."

Capacitive touchscreens (see Figure 1, top) are by far the most commonly used method of touch sensing in the electronics industry, and consist of varying layers of transparent lenses, substrates, adhesives, and indium-tin-oxide (ITO) electrodes. Together, these elements are known as touchscreen panels (TSPs) or stack-ups. Depending on the type of product in which they will be used, each stack-up and electrode pattern is customized for its intended environment and use. A stack-up (see example shown in Figure 1, bottom) contains an LCD layer, followed by a substrate, a pattern of horizontally and vertically aligned diamond-shaped ITO electrodes, and finally an optically clear adhesive layer that bonds the glass cover onto the screen.

Angry Birds is a trademark of Rovio Entertainment Ltd.

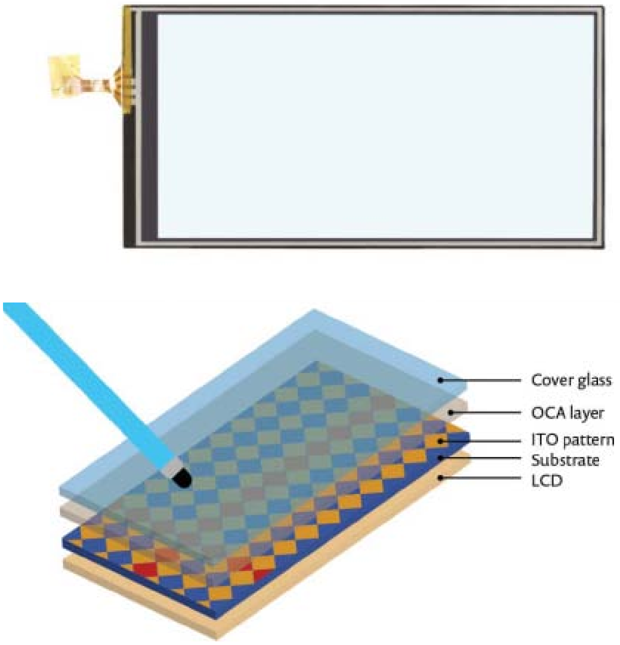

FIGURE 1: Top: Capacitive touchscreen module. Bottom: Typical touchscreen sensor stack-up where a stylus' position is detected because of the change in the electrodes' coupling capacitance.

At Cypress, multiphysics simulation and simulation apps have emerged as key tools for ensuring effective product development, allowing designers to predict and optimize the behavior of numerous designs without needing to build multiple physical prototypes.

» AN ENGINEERING RULE OF THUMB

AS A RULE OF THUMB, touchscreens must track finger or stylus positions with high accuracy. This means that at any point in time, a touchscreen must not only be able to determine that it is being touched by an object of variable size, but also where, for how long, and whether the "touch object" is moving in a certain direction. To achieve this, a capacitive sensor is composed of a pattern of horizontally and vertically connected ITO electrodes, where a touch object is sensed at the grid intersection. When a finger or stylus touches the screen's surface, it distorts the electrostatic field and causes a measurable change in the coupling capacitance between the transmitting and receiving electrodes (see Figure 1, bottom).

Depending on where and how the touchscreen will be used, the stack-up components are configured in a variety of ways. "The design of a touch-

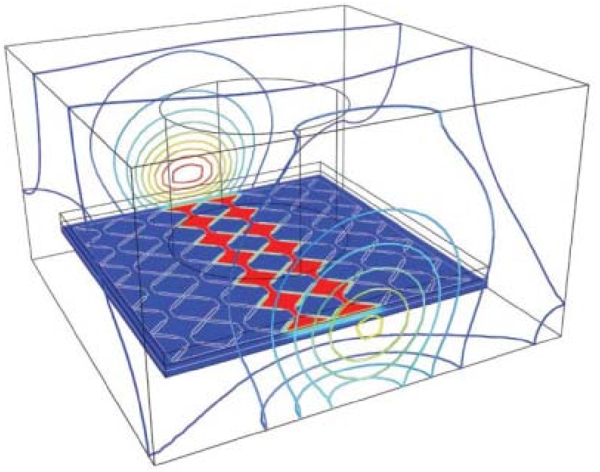

FIGURE 2: Touch sensor with electric field lines modeled using the AC/DC Module, an add-on product to COMSOL Multiphysics.

screen stack-up for the automotive industry is very different than one used in, say, a laptop," says Vavaroutsos. "My job at Cypress is to design different stack-ups for different consumer products, taking into account such things as how interactions between a horizontally mounted GPS, for example, will differ from a smartphone, which can be held and interacted with in a myriad of different ways."

Cypress R&D engineers create multiple electrostatic simulations for a particular device geometry and for many different parameters, what the team refers to as a "design box".

"Our findings from a specific design box are then used by our sales engineers and customer support team so that they can optimize certain design specifications in order to meet a customer's individual needs," explains Vavaroutsos.

Using the COMSOL Multiphysics® simulation software, R&D engineers at Cypress perform analyses to determine the electrical performance of the ITO pattern, including measuring the change in mutual capacitance between electrodes when a stylus or finger is present. In the example shown in Figure 2, floating potential boundary conditions were used in the electrostatic model, a feature that is instrumental in allowing Cypress engineers to simulate the boundaries of touch objects and any electric shielding or electrodes that are not currently being excited. Because these objects are affected by an externally applied electric field, they will be at a constant but unknown electric potential and therefore are represented as surfaces over which a charge can freely redistribute itself.

"Since the screen can be interacted with in so many different ways, in order to optimize a stack-up for use in a certain device or product, we have to run numerous electrostatic simulations in order to test different touch object positions," says Vavaroutsos. "We try to minimize effects such as when you get water on your screen and it doesn't work as well, or when you put your phone down on the table and the screen responds poorly. Simulation has been a very valuable tool for ensuring that our product responds effectively over a range of different environments and conditions, since we can single out certain factors and determine how to most effectively optimize performance."

Because COMSOL® software can be run on unlimited multiple cores and using cluster and cloud computing with no limit to the number of compute nodes, Cypress engineers are able to quickly run many simulations with virtually no limits on the size of the design boxes analyzed. "We can reduce the number of assumptions we have to employ and accurately model capacitive touchscreens by capturing changes between active electrodes in great detail while working with realistic geometry and materials," says Vavaroutsos.

Within a single design box, Cypress engineers might test different cover lens thicknesses, alter the permittivity of various layers, or change pattern parameters. Depending on the application area, a single touchscreen may be designed to have more than one electrode layer, or have different layers in a different order. For example, a design box might include a range for cover lens thicknesses from 0.5 millimeters to 1.5 millimeters. The R&D team at Cypress will model a variety of different parameter ranges in order to precisely understand a certain design, but anything outside the modeled range will remain unknown.

» TOUCHSCREEN SIMULATION APPS FOR CUSTOMIZED DESIGNS

IN ORDER TO EXTEND the usability of their models, Cypress engineers are using the Application Builder in COMSOL Multiphysics to create simulation apps based on their models. "In order to communicate more effectively

> *"Simulation has been a great tool for ensuring that our product responds effectively over a range of different environments and conditions."*
>
> —**PETER VAVAROUTSOS, R&D ENGINEER, CYPRESS**

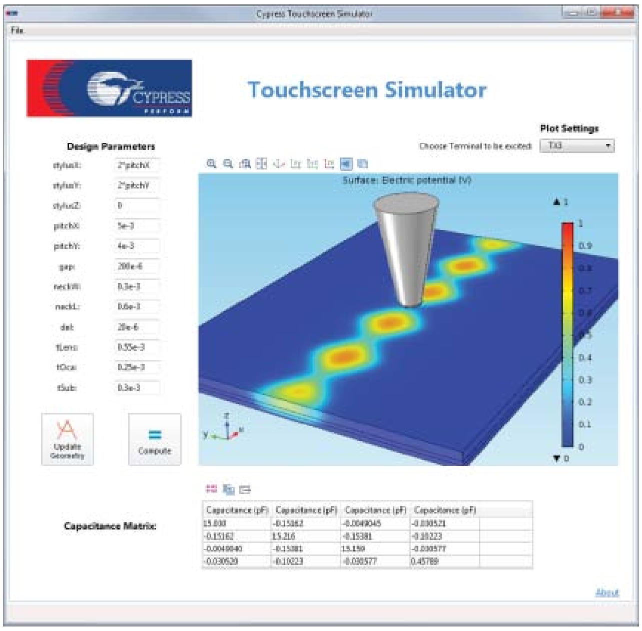

FIGURE 3: Touchscreen app created with the Application Builder and run using a Windows®-native client connected to COMSOL Server™.

with our customer support teams, we've started using the Application Builder to build simplified user interfaces over our models," says Vavaroutsos. "Before we started using simulation apps, any time a customer wanted a design that was slightly outside of the design box, we'd have to be involved again to run simulations for minor parameter changes. A lot of times, a sales engineer might try to run the simulations themselves, even though they had little experience using the COMSOL software. Not only would we have to check the simulations, but they also took up a seat on the software as well."

The example app in Figure 3 shows parameters that could be included in a touchscreen app. The finger or stylus touching the panel is represented by a conical structure. The app user can change design parameters ranging from the finger location to the thickness of the different layers in the sensor. The app then generates a report detailing the capacitance matrix, an integral piece of information for capacitive sensor design. The app can also show the electric field distribution in the sensor and a drop-down list can be used to select a solution corresponding to the excitation of different sensor traces.

Cypress is also using the COMSOL Server™ license to share their simulation apps with colleagues around the world, which allows anyone to access simulation apps using either a Windows®-based client or a web browser. "We're finding that letting our support teams have access to multiphysics simulation results is hugely helpful. We can control the parameters that the app user has access to so that we know the apps are delivering accurate results, while also letting our support engineers experiment with thousands of different design options without the need to involve an R&D engineer—or use a seat up on our COMSOL Multiphysics license."

» TEST-DRIVING AUTOMOTIVE TOUCHSCREEN DESIGNS

IN ADDITION TO TOUCHSCREENS for consumer products, Cypress also creates touchscreen designs for use in the automotive industry. For these applications, engineers experiment with different designs in response to certain automobile requirements.

"In the automotive group, our designs are more customer driven and are often created on a case-by-case basis for a specific product or customer," says Nathan Thomas, an R&D engineer working in the automotive group at Cypress. "Our design boxes are irregularly shaped, and we do more simulations that are customer-specific. For example, an automotive company might use touchscreens for different applications such as in the center console, in rear seat entertainment systems, or in overhead entertainment systems, all of which will need their own models."

Instead of creating a new model for each and every instance, the automotive group is now using apps to let field engineers test new designs that would otherwise have been outside of the design box. The apps can be used to explore special requests from customers who are interested in how varying a parameter will affect end performance. "For cases such as these, we've been using the Application Builder to create simulation apps that our field engineers can apply directly without having to go through us to create the simulation for them. While it's still a new technology, I can foresee simulation apps becoming the primary tool used by our field engineers."

» POWERING UP

WHETHER IT BE SMARTPHONE DESIGNS, automotive applications, or other industrial processes, Cypress R&D engineers can create simulation apps that allow other support engineers to experiment with designs that would otherwise have required the expertise of an R&D engineer. Through the use of simulation, Cypress engineers are delivering more customizable designs faster than ever before. ◎

COMSOL APPS BRING SIMULATION DEMOCRATIZATION

By **ALEXANDRA FOLEY AND VALERIO MARRA**

BEHIND EVERY PRODUCT LAUNCH, every technological innovation and scientific breakthrough, stand a whole team of professionals whose vision, expertise, and commitment have made the product possible. From design engineers to field technicians, to sales engineers and production managers, a wealth of knowledge and expertise surrounds every successful product.

On the R&D side, this knowledge is the work of a team of highly qualified engineers whose job it is to test, analyze, and innovate using advanced computer aided engineering (CAE) tools. With CAE, these R&D engineers build virtual physics models to gain an understanding of how the product will perform when exposed to real-world effects.

Future product updates, customizations for specific uses, or the exploration of different designs, materials, and other optimizations necessitate that simulations be conducted using the physics models developed by the R&D team. "Because of the complexity of these simulations, the engineer who created the model is oftentimes the only one who can safely make modifications and test new designs. In many companies today, this small team is therefore facilitating the work of a much larger group of people, thus creating a bottleneck," says Bjorn Sjodin, VP of Product Management, COMSOL.

How can these powerful computational tools be made available to a wider user base? Increased access to multiphysics simulation would undoubtedly lead to increased productivity and innovation, but how can this be done in a way that is cost-effective and scalable?

» *SIMULATION APPLICATIONS MOVE TO THE MAINSTREAM*

THE COMSOL MULTIPHYSICS® software and the Application Builder allow R&D engineers to share their simulation expertise with others in a way that is fast, easy, and customizable (see Figure 1). With the Application Builder, an R&D engineer can readily build an intuitive interface around any COMSOL® software model and customize it for the application's intended use. The resulting simulation app retains the functionality of the original model, but has a simplified user interface (UI) through which the app user can only edit certain inputs, materials, geometry, or other parameters, and then recalculate the new design's expected performance. Such an app can be created from within a single software environment using COMSOL Multiphysics.

Simulation apps can be used for a diverse range of tasks. For example, companies may offer apps demonstrating their product's performance to be used in place of data sheets, or may produce licensable apps as products in their own right. Within a company, a sales or support engineer might use an app to quickly analyze the expected performance, durability, and cost of a specific project by testing different materials and designs.

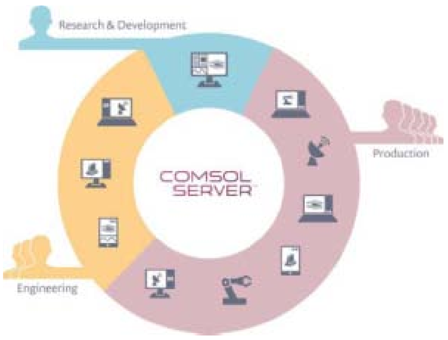

FIGURE 1: Apps are created by the R&D team using the Application Builder, and then used by other departments by connecting to the COMSOL Server™ environment.

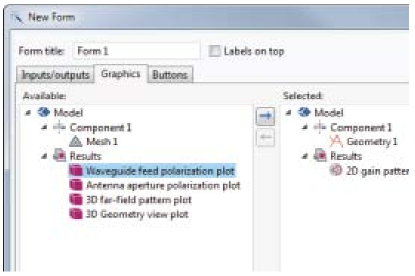

FIGURE 2: New form created with the Form Wizard for the corrugated circular horn antenna demo application, which is available in the COMSOL® software.

In short, "simulation apps offer a new line of communication between the professional modeling expertise of the model builder, and the rest of the production team," says Bjorn Sjodin.

» *BUILDING A SIMULATION APP WITH THE APPLICATION BUILDER*

ALL OF THE NECESSARY TOOLS for building a simulation app are accessible within the COMSOL Desktop® environment where the Model Builder and the Application Builder are fully integrated. As an example, consider the model of a corrugated circular horn antenna (available as a demo application in the Application Libraries). After switching from the Model Builder mode to the Application Builder mode, a simple app can be built in just a few clicks, and can then be further customized and enhanced.

To start, the Form Wizard suggests a list of elements from the model that are suitable for use as inputs and outputs, modeling operations, and graphics (see Figure 2).

The template app created using the Form Wizard can then be updated and customized using two tools: the Form Editor and the Method Editor. With the Form Editor, form objects such as input fields, buttons, graphics windows, and results tables can be added to the app using drag and drop operations (see Figure 3).

After adding these forms, the app can be further customized for additional functionality using the Method Editor, which is essentially a Java® programming environment that allows users to combine the COMSOL® interface with Java® programs and libraries.

In the corrugated horn antenna example (see Figure 4), features created using methods include input-within-bounds restrictions for frequency and wavelength fields, the ability to reset the input parameters back to the default values, and 2D and 3D results visualizations. Additionally, when the app has finished running, it generates a simulation report, which can be automatically emailed to a specified recipient when a computation is complete. The finalized app is shown in Figure 4.

Simulation apps are fully customizable based on the specific needs of the app user and can be created for diverse projects, including for sales representatives looking to modify a design to accommodate the needs of a customer or a design engineer looking to quickly acquire data about the feasibility of a new project. Apps created with the Application Builder and COMSOL Multiphysics make simulation accessible to a wider group of people working in product development, design, or production. The democratization of simulation is happening now! ◎

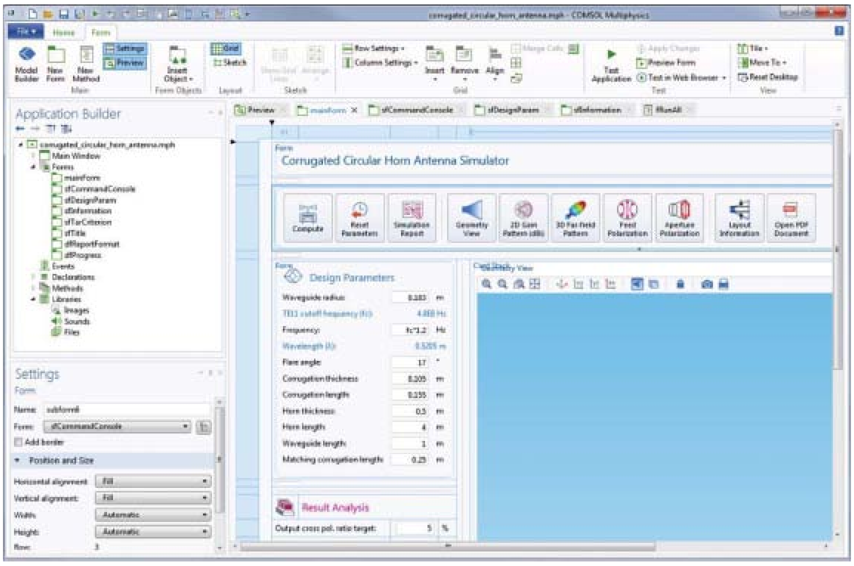

FIGURE 3: The finalized main form of the corrugated horn antenna app consist of forms, card stacks, text labels, input fields, units, images, etc. These objects are available from a drop-down menu in the Form Editor.

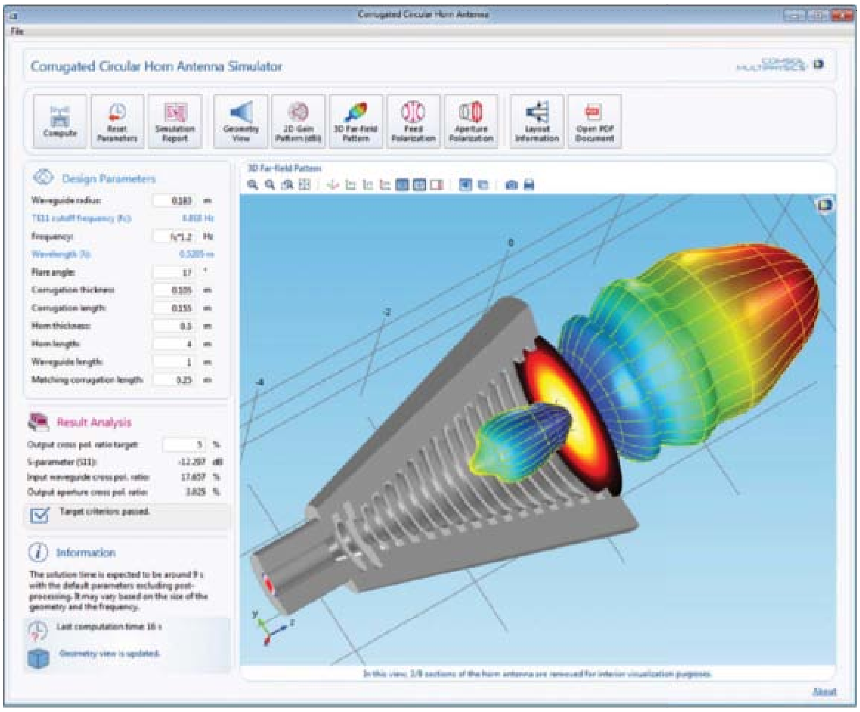

FIGURE 4: An app showing the far-field radiation of a corrugated circular horn antenna. The geometry parameters and operating frequencies can be changed to optimize the antenna's performance.

> *"Simulation apps offer a new line of communication between the professional modeling expertise of the model builder and the rest of the production team."*
>
> **—BJORN SJODIN, VP OF PRODUCT MANAGEMENT, COMSOL**

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

WITRICITY LEVERAGES MAGNETIC RESONANCE FOR FLEXIBLE WIRELESS CHARGING

Engineers at WiTricity have used multiphysics simulation in the innovative development of wireless power transfer technology that extends efficiency and charging ranges beyond the reach of existing methods.

By **LEXI CARVER**

IMAGINE COMING HOME and dropping your phone, laptop, and Bluetooth® headset on your kitchen table so that they all recharge, simultaneously. Or driving your electric car into a garage, parking above a mat, and knowing it will be charged in the morning. Or being told by your doctor that there is a new medical implant to replace the one you wear—and the new version does not include power cords or the need to replace batteries.

Wireless power transfer is making these scenarios and other applications a reality by delivering a cordless way to charge electronic devices. WiTricity, a Watertown, MA-based company that develops wireless charging technology based on magnetic resonance, has launched the most consumer-friendly method available for the wireless transfer of electrical power. Invented at the Massachusetts Institute of Technology (MIT) by Professor Marin Soljacic and a team of researchers, WiTricity's technology has the ability to charge multiple devices at once, over distances and through materials like wood, plastic, granite, and glass. Companies such as Toyota, Intel, and Thoratec have already licensed the technology for use in hybrid-electric vehicles, smartphones, wearable electronics, and heart pumps.

» MAGNETIC RESONANCE WIDENS THE TRANSMISSION GAP

OTHER OPTIONS FOR WIRELESS ENERGY TRANSFER require precise device positioning on a pad or holder, very close proximity to (often resting directly on) the charging source, and the source can only charge a single device with a single coil. Now, the engineers at WiTricity have leveraged the power of magnetic resonance to rethink these limitations.

Their system, dubbed "highly resonant wireless power transfer", relies on oscillating time-varying magnetic fields generated by alternating current passing through a coil that functions as a power source. A power amplifier connected to this source coil controls the power levels and operating frequency, driving the magnetic field levels.

A capture device, which acts as a receiver and captures the magnetic field, contains another coil tuned to the same frequency as the source (see Figure 1). The field converts the magnetic energy back to radio-frequency alternating current in the receiver, which can then be used as a new local power source after being rectified and regulated by power electronics.

The notable difference between WiTricity's technology and other approaches is the use of magnetic resonance. With

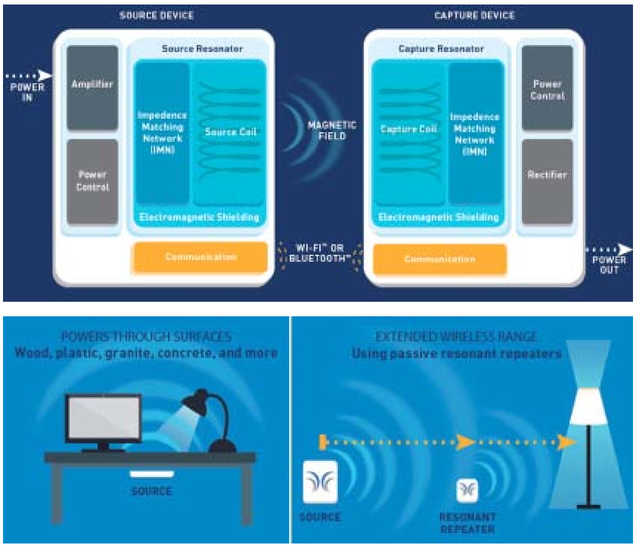

FIGURE 1: Top: Concept behind WiTricity's technology, including coils, electronics that control power output and regulation, and wireless communication between the source and capture devices to ensure the correct power output. Bottom: The technology supports charging through surfaces of different materials (left); the wireless range may be extended using resonant repeaters (right).

Bluetooth is a registered trademark of Bluetooth SIG, Inc.

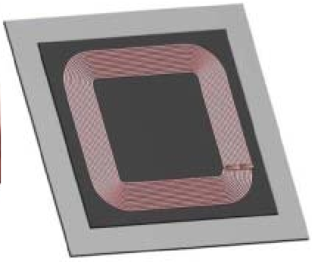


FIGURE 2: Left: A capture resonator, a resonant repeater, and a source resonator. Center: A WiTricity source resonator designed for consumer electronics applications. Right: An electric vehicle charging wirelessly, parked above a charging pad.

both coils tuned to the same resonant frequency, the receiving coil is able to capture maximum power through the magnetic field with very low losses, and power can be transmitted without the source and capture device sitting next to each other or being perfectly aligned.

"One major advantage is the flexibility of motion and positioning. The receiving coil doesn't have to be in direct contact with the device; for instance, while driving your car you could drop your phone into a cup holder positioned near the capture device, rather than arranging it on a charging pad," explained Andre Kurs, co-founder of WiTricity. "And you can charge everything together, including electronics that have different power requirements."

Extending the wireless range is also easy enough: resonant repeaters that each contain another circuit and coil may be placed between the source and receiver, allowing power to 'hop' over greater distances (see Figure 2).

Transfer occurs effectively even with barriers (such as people and concrete walls) between the power source and the receiver.

» MODELING THE ELECTROMAGNETICS OF WIRELESS POWER TRANSFER

IN DESIGNING FOR MAXIMUM EFFICIENCY using coils with the same resonant frequency, Kurs and his team had to account for variables such as number of coil turns, diameter, and necessary power input. From the early stages of development, they relied on computer simulation to test key details, verify designs, and optimize the system. Using a COMSOL Multiphysics® software model, Kurs analyzed the electromagnetic and thermal behavior of different coil configurations, and was able to quickly validate new designs.

One challenge lay in making the technology scalable for a wide range of devices: a car, for instance, needs a different charging configuration than a smartphone. "We have such a wide range of applications, prototyping and testing is time-consuming and expensive, and design mistakes can set you back a lot—in such a competitive market, we can't afford that," he said. "Design validation in COMSOL was cost-effective and time-saving and allowed us to virtually test our concepts before building the real device."

He created simulations with different setups for each application, and included electromagnetically relevant components such as coil windings, specially shaped ferrites and metal surfaces used to guide the electromagnetic field, plates for shielding sensitive electronics, and large objects that might perturb the field, such as a car chassis.

Then he ran a multiphysics study to analyze the resulting electromagnetic and thermal performance as a function of power drawn by the devices, coil displacements, and the effects of perturbing objects (see Figure 3, top).

He extracted circuit parameters from the results to guide the design of the electronics, as well as predictions of power dissipation and thermal loading on different components (see Figure 3, bottom). The team adjusted their designs accordingly, determining the viable range of coil displacements and power levels as a function of size, weight, and thermal constraints.

"The simulation allowed us to disentangle various

> *"Design validation in COMSOL was cost-effective and time-saving and allowed us to virtually test our concepts before building the real device."*
>
> **—ANDRE KURS, CO-FOUNDER, WITRICITY**

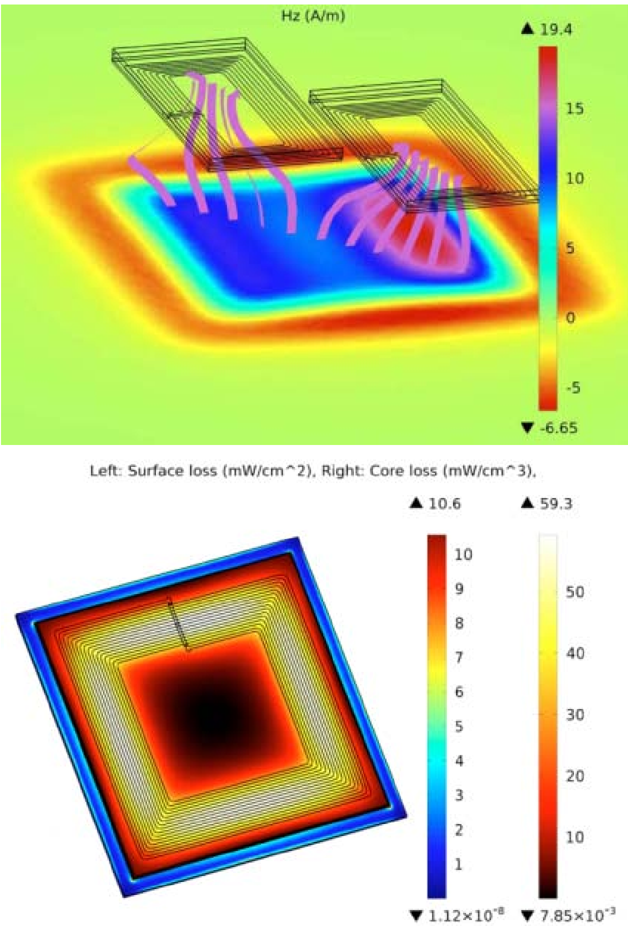

FIGURE 3: Simulation results showing the magnetic field levels (top) and power dissipated (bottom) in a source resonator for consumer electronics applications.

effects that we couldn't isolate just by testing, like power dissipation and heat transfer," Kurs remarked. "The flexibility of COMSOL was particularly useful; we built a suite of simulation apps around it that would allow any of our engineers—whether they knew how to use COMSOL or not—to quickly test and validate designs even if they didn't understand the entire model."

» *KEEPING A SAFE DISTANCE, RIGHT UP CLOSE*

SINCE SUCH DEVICES are near to or in contact with people's bodies, electronics manufacturers must adhere to safety limits on the electromagnetic fields emitted by their products. The magnetic fields needed for WiTricity's wireless transfer are usually fairly weak, but each new application needs be checked for compliance.

To make sure that the field levels and resulting body temperatures would meet regulations, the team ran several more COMSOL simulations to study different body tissues in close proximity to the device. Their models calculated the electric field based on the operating frequency of the charging system, and confirmed that the results were well within FCC safety guidelines (see Figure 4).

» *RETHINKING A GROWING INDUSTRY FOR WIRELESS CHARGING*

WITRICITY'S DESIGNS BASED on magnetic resonance are a major improvement over other wireless charging methods, allowing reliable wireless power transfer in a flexible, consumer-friendly product. Thanks to their simulation work in COMSOL Multiphysics, the WiTricity team optimized their designs for better efficiency and longer ranges before building costly prototypes.

In addition to being frontrunners in game-changing wireless power transfer technology, WiTricity is on the board of the Alliance for Wireless Power (A4WP), an organization dedicated to building a "global wireless ecosystem" and creating standards for wireless charging. Another board member, Intel, has licensed WiTricity's technology to develop a wireless desktop system. The A4WP is an innovative group comprising leading companies that are ushering in a new way of thinking about wireless power: they are imagining a future where everyday surfaces—desks, cup holders, and even your kitchen countertop—become zones for charging the electronic devices we depend on so much. ◎

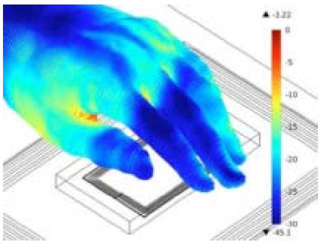

FIGURE 4: COMSOL simulation showing the specific absorption rate (SAR) in a hand above a charging cell phone. SAR is a measurement of electromagnetic energy absorbed and turned into heat. Results are in dB relative to the FCC limit (a value of zero represents the limit).

Andre Kurs, co-founder, WiTricity.

MEDTRONIC ADVANCES ABLATION TECHNOLOGY WITH MULTIPHYSICS SIMULATION

The new technology will enhance physicians' abilities to plan and implement ablation procedures, potentially leading to better patient outcomes.

By **GARY DAGASTINE**

ABLATION, OR THE USE OF HIGH-FREQUENCY electromagnetic (EM) energy to destroy soft-tissue tumors, has been in existence for a few decades, but in recent years its underlying technology has evolved.

The benchmark of minimally invasive tissue treatment has long been the application of electrical current to kill abnormal tissues. This is done by heating tissues until they break down, a process called thermal ablation. Energy is delivered at 500 kHz, within the radio frequency (RF) range of the EM spectrum, hence these systems are called RF ablation systems.

In recent years, microwave (MW) ablation technology has also become commercially available and increasingly popular. At MW frequencies, oscillating EM fields are utilized to perform thermal ablation. Medtronic, one of the world's premier medical technology and services companies, is a leader in both RF and microwave ablation technologies.

With both RF and MW systems, the energy for ablation is applied using one or more needle-like probes.

Medtronic's latest innovation, the Emprint™ ablation system with Thermosphere™ technology, offers more predictable and repeatable results than other techniques and devices (Figure 1). These advantages come from the fact that Thermosphere™ technology enables precise control of an EM field independent of the surrounding tissue environment.

» STRIVING FOR BETTER PREDICTABILITY

ACCORDING TO RESEARCH, physicians rate predictability as their number one concern with ablation performance. The higher the level of predictability, the easier it is for a physician to plan a treatment procedure that will be safer, more effective, and less time-consuming.

Because of its nature, it's challenging to be certain that RF ablation procedures will achieve the desired results. Given their different electrical conductivities, some tissues are less amenable to effective RF heating than others. Moreover, as the temperature in targeted tissue approaches 100°C, water in the tissue begins to vaporize and electrical conductivity rapidly decreases. This can make it difficult to generate temperatures high enough to cause cell breakdown.

MW ablation technology attempts to overcome these limitations by using an EM field radiated into the tissue (Figure 2). However, in practical application, tissue type and the vaporization of water during ablation cause the size and shape of the EM field to vary.

The Emprint™ ablation system with Thermosphere™ technology realizes the promise of predictability. It gives physicians the ability to easily control the thermal energy delivered by allowing precise control of the EM field across tissues and temperatures. This allows clinicians to accurately predict the boundaries and characteristics of the ablation zone.

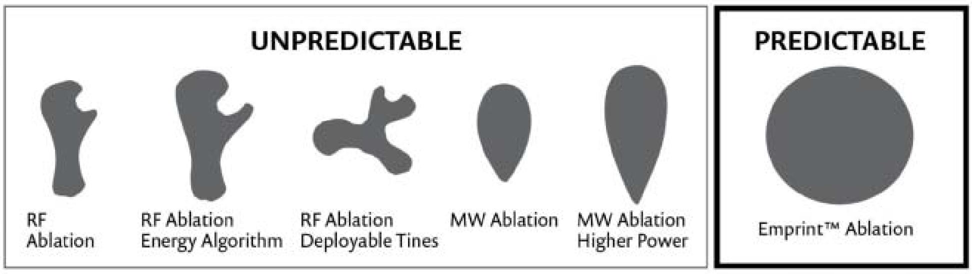

FIGURE 1: At left, shapes of tissue ablation zones that can result unpredictably from the use of various ablation technologies. At right, Medtronic's Emprint™ ablation system with Thermosphere™ technology yields predictable spherical ablations regardless of target location or tissue type.

Emprint and Thermosphere are trademarks of a Covidien company.

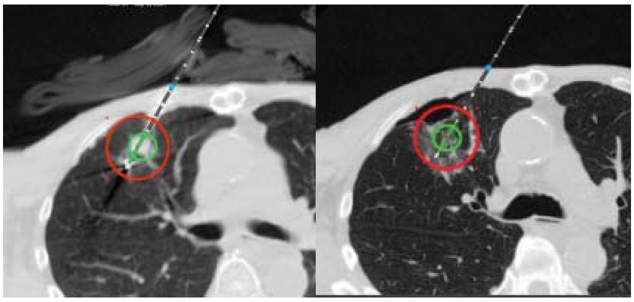

FIGURE 2: The photo at left illustrates placement of an ablation probe. The green circle delineates the target (where the lesion is located) and the red circle delineates the margin the ablation is meant to achieve. The image at right shows the site after ablation has taken place.

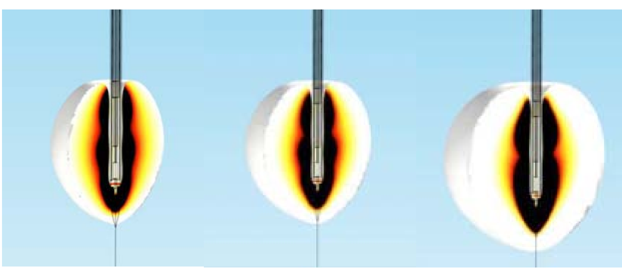

FIGURE 3: These results from a COMSOL® software simulation show the power dissipation density, or the extent of the ablation, as determined by the thermal damage calculation. The antenna and the surrounding tissue are initially well-matched, and the match (i.e., the antenna pattern) changes over time as tissue temperature increases during the procedure (left to right).

» REAL-TIME MONITORING OF ABLATIONS

"**THE CHALLENGE NOW** is to monitor the ablation performance in real-time," said Casey Ladtkow, principal engineer in the Early Technologies unit of Medtronic's Minimally Invasive Therapies Group (MITG). "At present, when performing ablations, physicians don't have continuous real-time feedback on the effectiveness of their procedure. If they could know exactly what is happening in real-time from start to finish, the effectiveness of ablation treatment would increase," he said.

With some 40 staff members focused on interventional oncology, the mission of his unit is to deliver procedural solutions that alleviate pain, restore health, and extend life. He and his team are using COMSOL Multiphysics® software to develop new ablation probes in order to achieve even higher levels of predictable performance and effectiveness.

One development-stage project is to optimize the design of these probes so they can both create a more precise ablation zone and also provide real-time feedback using radiometers.

Radiometers measure EM radiation and enable the characterization of the spatial distribution of an EM field. Ladtkow's team is incorporating radiometers into Medtronic probes in order to give clinicians real-time feedback about the ablation zone. This will enable a clinician to fine-tune the zone as needed during the procedure, and to make sure the radiation destroys the targeted tissues while minimizing effects on the surrounding healthy tissue.

The team uses COMSOL Multiphysics and its RF Module to help them model the probes and better understand and optimize their emitting/radiating and receiving/monitoring properties. "The performance and accuracy of MW ablation systems are affected by a number of dynamic factors that arise simultaneously in multiple physics domains. COMSOL® software gives us the ability to perform the relevant complex modeling quickly and easily, to help us understand these coupled effects and improve our design," Ladtkow said.

> *"Multiphysics simulation enabled the rapid development, evaluation, and optimization of our design, which would not have been possible otherwise."*
>
> —**CASEY LADTKOW, PRINCIPAL ENGINEER, MEDTRONIC**

» SIMULATION ENABLES FAST AND SAFE DESIGN, OPTIMIZATION, AND PROTOTYPING

FOR SUCH A COMPLEX device, the traditional approach of building and evaluating a series of physical prototypes is all but out of the question because of the complexity and relationships among the many physics-based factors that impact device performance.

The team used COMSOL to model the energy radiator and test designs that incorporated radiometric sensing in the same device. They simulated coupled thermal and electromagnetic effects around the radiative probe hardware to determine radiometric performance under different conditions (Figure 3).

Ladtkow analyzed heat transfer in living tissue using a bioheat equation, which included a perfusion term, to account for blood flow cessation once the tissue coagulated (Figure 4). This helped his team understand heat transport to cells around the tumor and predict the temperature distribution to ensure efficient and predictable energy delivery.

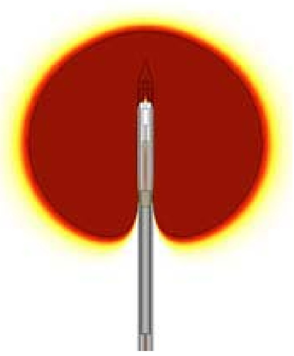

FIGURE 4: A COMSOL plot showing a cross section of the predicted ablation volume, or predicted tissue damage. This information is used to modify the bioheat equation and thereby to modify perfusion conditions in the tissue. Red areas represent coagulated tissue where no perfusion is present, and white areas represent areas of normal perfusion. This makes the model more accurate by creating a realistic on/off condition for the perfusion term in the bioheat equation.

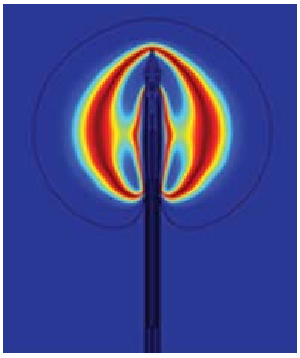

FIGURE 5: COMSOL results show the change in heat capacity in the tissue surrounding the probe, dominated by phase change of water in that tissue. Knowing where water is boiling is important because the MW radiation wavelength is dramatically different for liquid water than for vaporized water.

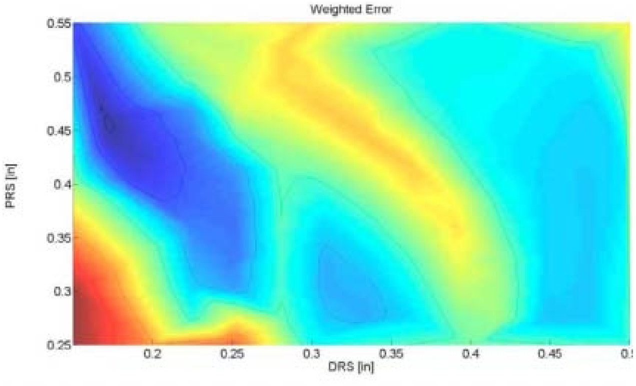

FIGURE 6: A COMSOL weighted error plot. Blue areas indicate regions where reflected power is low for delivered ablation energy and also where the receiver quality is good. They represent antenna configurations that are both good ablation devices and good radiometers.

He performed other studies as well: investigations of temperature dependence of reaction rates (to understand the size of the ablation zone); radiometry modeling to determine how much energy enters the tissue and how much is reflected back into the radiator; and liquid-to-gas phase-change dynamics (Figure 5). "The latter is critical to knowing what the wave pattern will look like, because knowing how much water is in the tissue is critical to knowing how a radiometer will behave, because of the change in wavelength," he said. "Implementing this model in COMSOL is straightforward."

Simulation showed that lengthening the proximal radiating section (PRS) and shortening the distal radiating section (DRS) of an antenna would produce an efficient ablation radiator and an efficient receiver. These studies (Figure 6) resulted in versions of a prototype ablation radiator with an integrated radiometer, along with results showing the performance of the integrated probe.

» FROM IMPOSSIBLE TO POSSIBLE

"**WITHOUT COMSOL** to help us perform these analyses, it simply would be impossible to do enough experiments to find an optimum solution that integrates an emitter and a receiver. COMSOL helps us see that certain architectures —which we'd never have investigated otherwise—might make an integrated device possible," Ladtkow continued.

His team uses COMSOL® software in conjunction with MATLAB® software, and he said that the combination gives him a powerful ability to optimize complex models with highly sophisticated algorithms quickly and easily. He also hopes to integrate the Application Builder available in COMSOL Multiphysics into their modeling workflow. This would enable the team to create simulation apps allowing partners to test and verify different designs, while protecting their proprietary models.

"Based on our simulations, we are now realizing the potential to introduce ablation devices that will allow clinicians to not only deliver a precise energy dose, but also monitor ablations in real time," Ladtkow said. "Multiphysics simulation enabled the rapid development, evaluation, and optimization of our design, which would not have been possible otherwise." ◎

From left: The Medtronic team consists of Morgan Hill, Casey Ladtkow, and Robert Behnke.

MATLAB is a registered trademark of The MathWorks, Inc.

GRAPHENE PAVES THE WAY FOR NEXT-GENERATION PLASMONICS

Simulation tools bring the complex physics of two-dimensional materials and plasmonics together in a way that could change the face of optoelectronic devices.

By **DEXTER JOHNSON**

EVER SINCE A SINGLE-ATOM-THICK FILM of graphite was first successfully synthesized back in 2004 and called graphene, it has been on a decade-long ride through applications ranging from photovoltaics and next-generation batteries to electronics.

While graphene's list of desirable properties—like its electrical and thermal conductivity—initially made it attractive for electronics, its equally attractive optoelectronic capabilities were initially overlooked. But it soon became clear that graphene has incredible potential as a transparent conducting electrode and could be an alternative to the commonly used indium tin oxide (ITO). Graphene offers comparable or better optoelectronic performance in addition to its mechanical strength and flexibility. Other potential uses are diverse and include applications such as transparent conductors used in touchscreens and photovoltaics (see Figure 1), lab-on-chip devices for the sensing of viruses or proteins, improved night vision, mid-IR imaging applications, and solar cells.

FIGURE 1: Bendable and lighter smartphone and laptop screens are just one of the many applications of graphene. Others include energy, computing, engineering, and health technologies and devices.

» GRAPHENE AND PLASMONICS MEET

IN ADDITION TO OPTOELECTRONICS, graphene's star has shone particularly bright in photonics when it is used in combination with the field of plasmonics, a subfield of photonics that grew out of the need to continually explore properties and applications of light on ever-smaller scales.

Traditionally, photonics has dealt with structures on the micrometer scale, but squeezing light into smaller dimensions is fundamentally challenging due to a property of light known as the diffraction limit. Plasmonics helps with addressing this challenge and enables light confinement even at the nanoscale.

This is achieved by coupling incident light into oscillations of electrons known as plasmons—hence the name plasmonics. Today, plasmonics is an important, actively developing branch of photonics that deals with the efficient excitation, control, and use of plasmons.

» GRAPHENE-ENABLED PLASMONICS IS LEADING TO PRACTICAL DEVICES

COMPUTATIONAL NANOPHOTONICS efforts at Birck Nanotechnology Center, Purdue University, led by Alexander V. Kildishev, associate professor of electrical and computer engineering, have been leading the way in combining graphene with plasmonics to bring it closer to practical optoelectronic applications.

The work of Kildishev and his colleagues deals with a fundamental problem in graphene research: it is currently difficult to fabricate high quality, large-area graphene films. Until graphene production improves, Kildishev and his team are leveraging simulation tools to perform design and optimization of devices made from graphene.

Through both simulation and experimental testing, Kildishev and his colleagues have been able to dem-

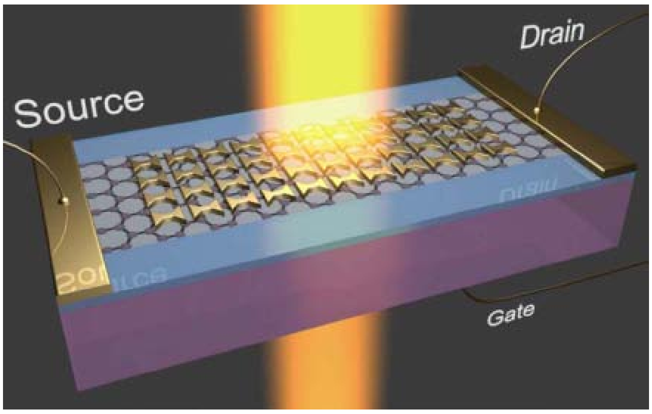

FIGURE 2: Design of Fano resonant plasmonic antennas on top of a single-layer graphene sheet optimized with COMSOL® software and its Wave Optics Module to achieve resonance at a 2 μm wavelength. The design tunability has been successfully validated in experiments using ion-gel top electrolyte gating[2].

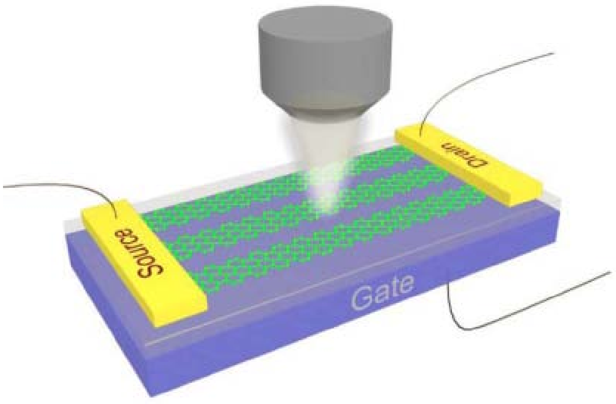

FIGURE 3: 3D artistic sketch of the experimental setup used for studying plasmon resonance in graphene nanoribbons (GNRs), simulated with COMSOL Multiphysics® software using the surface current approach. The lattice orientation of GNRs is for illustration only and dimensions are not to scale.

onstrate tunable graphene-assisted damping of plasmon resonances in nanoantenna arrays, which is important for designing tunable photonic devices in the mid-infrared range[1]. Since the mid-infrared is where fundamental vibrational resonances reside for a wide range of molecules, it is critical to have tunable plasmonic devices that work in that range for applications in sensing and imaging.

On the other hand, moving closer to even shorter infrared (IR) waves, e.g., the telecom range, is also of ultimate importance for telecommunications and optical processing. The group at Purdue has shown efficient dynamic control of Fano resonances in hybrid graphene-metal plasmonic structures at near-infrared wavelengths. Fano resonances are seen in the transmission of specifically coupled resonant optical systems. Researchers are currently leveraging the properties of Fano resonances for use in optical filtering, sensing, and modulators (see Figure 2).

Leveraging the predictive power of COMSOL Multiphysics® software models is a vital step for designing tunable elements for the next generation of plasmonic and hybrid nanophotonic on-chip devices such as sensors and photodetectors, according to Kildishev. The photodetectors could ultimately find use in the sensing of infrared electromagnetic radiation for multicolor night vision and thermal imaging. Another application may be in biosensing, where the resonant lines of plasmonic elements are tuned to match the resonances of the spectral optical responses of viruses or proteins.

In their work, the Purdue researchers combined the unique properties of graphene with plasmonic nanoantennas to modulate the antenna's optical properties. Having a tunable resonant element along an optical path is as critical to optoelectronics as having a transistor in an electric circuit.

"By using the nanopatterned graphene with an electrical gating (see Figure 3), it's possible to modulate light flow in space with unparalleled spatial resolution," said Dr. Naresh Emani, a former Ph.D. student advised by Kildishev, now with DSI, Singapore. "The reduced dimensionality and semimetallic behavior of graphene plasmonic elements gives us, along with its other properties, a very vital feature—electrical tunability. This critical functionality is not attainable with conventional metal plasmonics."

Plasmonic devices based on noble metals lack this level of control over electrical tunability. Noble metals possess a large number of electrons in the conduction band, and consequently the electrical conductivity of metals cannot be easily modulated. But since graphene is a tunable semimetal, it does not contain any electrons in the conduction band in its pristine state. Therefore, its electron concentration—and hence its electrical conductivity—can be tuned chemically, modulated electrically, or even modulated optically.

» THE ROLE OF SIMULATION AND MODELING

NUMERICAL MODELING has been a critical tool for the researchers, allowing them to optimize their designs without complications and the significant cost of nanofabrication processes.

"Compared to experimental work, mathematical modeling is low-cost, has the opportunity to validate its output through a reduced number of prototypes, has predictive power, and, finally, allows you to optimize for a desired functionality," explained Kildishev.

In a field where the quality of the graphene material can vary, it is critical that there always be a tight connection between numerical results and experiments in order to better understand the impact of all variables involved.

"In most cases, by fitting model

parameters to experiments, we can retrieve the actual physics of a given process," said Kildishev. "Having a validated mathematical model in hand always provides better understanding and interpretation. Once you understand the phenomena in terms of a mathematical model, you gain comprehensive knowledge of the whole mechanism that can be applied to other new ideas."

Of course, mathematical modeling has its own barriers. "Unfortunately, many problems do not have analytical solutions and we must revert to alternative options," he added.

This is where numerical techniques come in as powerful tools for circumnavigating these hurdles, according to Dr. Ludmila Prokopeva, a high-performance computing specialist on Kildishev's team. Properly designed simulation tools provide stability, accuracy, and speed. There is often a need for substantial high-performance computational machinery, especially for nanostructured devices that require full three-dimensional (3D) simulations.

"The multiphysical and multiscale essence of computational nanophotonics necessitates the use of powerful simulation tools," said Kildishev.

It is never one simulation tool that works in all situations. "We have a whole zoo of our own software and commercial software, and we are always looking for ways to incorporate new and interesting physics," said Kildishev. "COMSOL Multiphysics is one tool that we have relied on for about ten years, and its key advantage is its flexible operation within its unique equations-driven framework, which is unparalleled."

He added: "COMSOL allows users to couple several physics interfaces sharing the same mesh or even having separate meshes. We can also link the solvers to complex material

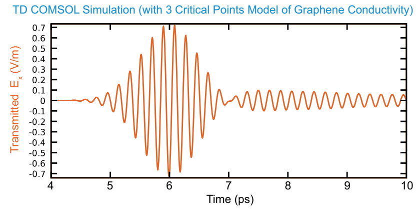

FIGURE 4: The time-dependent electric field of a Gaussian pulse transmitted through an array of graphene nanoribbons.

functions: for example, my team has implemented several complex dielectric models for graphene—written in MATLAB® software—which are seamlessly incorporated using COMSOL® software. Some of these dielectric functions are impossible to handle for a straightforward explicit input in terms of plain arithmetic or look-up tables. We are also able to introduce nonlinear effects, couple these to a heat transfer analysis, add quantum emitters—the list goes on."

"Another strength of COMSOL is its capability to model two-dimensional (2D) materials natively in terms of surface conductivity (i.e., surface current)," noted Prokopeva. "Because of its atomic thickness, graphene behaves like a 2D material, but many researchers use a thin artificial thickness and have to resort to a 3D model in their simulations just because of the inability to treat 2D materials naturally in their software. The 3D approach brings unphysical shifts, uncertainty in optimization procedures, and significant complications to the numerical calculations."

While waiting for manufacturing techniques to mature, the Purdue team used a theoretical model for graphene's optical conductivity and simulated the device response in COMSOL to numerically investigate the system properties (see Figures 2 and 4).

"We've been very fortunate to collaborate with our 'next-door' experiment-oriented teams of Profs. Yong Chen, Alexandra Boltasseva, Vlad Shalaev, Ashraf Alam, David Janes, and Gary Chen, here at the Birck Nanotechnology Center at Purdue. Collaboration with the Ted Norris and Vinod Menon groups within the C-PHOM NSF MRSEC center is also of critical importance. As the experimental studies are focused on very diverse facets of novel graphene applications, including IR sensors, hybrid photovoltaic electrodes, and even other 2D materials, they give us an excellent base for validating our new modeling approaches. They offer indispensable feedback from the fabrication and optical characterization of real-life graphene-based nanostructures."

» LOOKING AHEAD TO QUANTUM OPTICS, BETTER NIGHT VISION, AND FLEXIBLE TOUCHSCREENS

THE PURDUE TEAM is continuing their simulation work to understand and predict the behavior of graphene so that it may be put to use in devices such as photovoltaics, optical modulators, and—one day—flexible touchscreens. They are looking to make graphene nanoribbons so that they can begin fabricating a preliminary light modulation device[3,4].

"Generation and modification of short optical pulses is an important aspect of imaging and sensing," explained Kildishev. "Currently, the devices capable of achieving this functionality at mid-IR wavelengths are rather bulky and are not tunable. We envision a prototype device that can dynamically change the frequency content of an incoming optical pulse or light beam. This will enable higher sensitivity detection for night vision and mid-IR imaging applications."

They also have longer-range aspi-

MATLAB is a registered trademark of The MathWorks, Inc.

> “Another strength of COMSOL is its ability to model two-dimensional materials natively in terms of surface current.”
>
> —**LUDMILA PROKOPEVA, HIGH-PERFORMANCE COMPUTING SPECIALIST, BIRCK NANOTECHNOLOGY CENTER**

rations to explore the plasmonic properties of graphene in the quantum optics regime. Kildishev and his colleagues believe the quantum optics regime will be the next frontier for the science of light and has been relatively unexplored in the mid-IR wavelengths.

“Semiconductor quantum wells show some interesting quantum properties but are restricted to low temperatures so far,” said Kildishev. “If we successfully address some of the challenges in graphene research, it might end up outperforming semiconductor quantum wells. If we are able to do this, we could significantly reduce the size of many devices.” They continue to move forward on the cutting edge of research with many unknowns, toward a future that contains unbelievable possibilities. ◎

Alexander V. Kildishev, associate professor at Purdue’s Birck Nanotechnology Center.

References

[1] *N. K. Emani et al., Nano Lett. 12, 5202–5206 (2012).*

[2] *N. K. Emani et al., Nano Lett. 14, 78-82 (2014).*

[3] *D. Wang et al., “Plasmon Resonance in Single- and Double-layer CVD Graphene Nanoribbons,” in CLEO: 2015, p. FTu1E.3.*

[4] *L. Prokopeva and A. V. Kildishev, “Time Domain Modeling of Tunable Graphene-Based Pulse-Shaping Device (invited),” in ACES 2014.*

SIMULATING GRAPHENE

BY ANDREW STRIKWERDA

WHAT IS THE BEST WAY to simulate graphene? More specifically, should graphene actually be modeled as a 2D layer or rather as a 3D material that is extremely thin? Many researchers have used the latter approach because it is the only one supported in their numerical software. With COMSOL Multiphysics® software, you can use either method. As stated in the article, Professor Kildishev and his colleagues have found that simulating graphene as a 2D material yields better agreement with experimental results. Let’s take a closer look at how this is implemented in COMSOL® software.

Ohm’s law states that, in the frequency domain, the current density is simply the product of the conductivity and the electric field:

$$J = \sigma E$$

In COMSOL Multiphysics, this can be implemented in 2D using a Surface Current boundary condition where the induced current is expressed, according to Ohm’s law, as the product of the graphene conductivity (calculated, for example, from a Random Phase Approximation) and the tangential electric field.

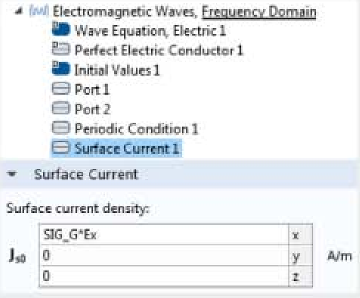

For time-domain simulations, the required surface current density can be a little more difficult to calculate, since Ohm’s law is now a convolution of the electric field and the conductivity:

$$J(t) = \int_{-\infty}^{t} \sigma(t-\tau)E(\tau)d\tau$$

To implement this in COMSOL (see Figure 4), Professor Kildishev’s group used a Padé approximation to represent the frequency-dependent optical conductivity of graphene. They then applied a Fourier transform of the terms in the Padé series to obtain second order partial differential equations in time, which can be solved in COMSOL.

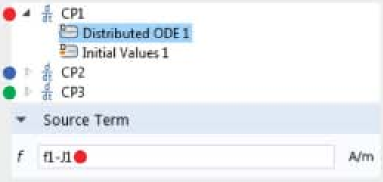

The solutions to these equations, representing contributions to the time-dependent surface current, can then be linked to the Surface Current boundary condition.

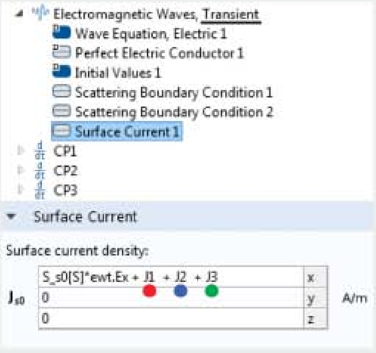

If you would like to learn more about how to simulate graphene, watch the webinar by Alexander Kildishev on comsol.com/webinars and download his COMSOL models available at comsol.com/community/exchange/361.

DEFYING CONVENTION TO ACHIEVE FASTER SIGNAL AND SIMULATION SPEEDS

Innovative optimization methods combining space mapping algorithms and electromagnetic simulation make it possible to improve signal speed and integrity in package interconnects. Faster simulations ensure the latest high-speed interconnect technology is available in less time.

By **JENNIFER SEGUI**

ESTABLISHED AS LEADERS in the electronics and computer hardware industry, it is easy to assume that the researchers and engineers at Intel rely on powerful computing clusters and servers to efficiently simulate and optimize their designs. While we would be correct in our assumptions, there's much more to their methods. Multiphysics simulation software and an unconventional approach to design optimization developed at the Intel Guadalajara Design Center may be behind the latest high-speed interconnects for electronics packaging.

Take printed circuit boards (PCBs), for example. A mainstay in electronics packaging, PCBs are found in almost every electronic device from handheld computers and cellphones to state-of-the-art satellite communication systems. PCBs have many integrated high-speed interconnects enabling the transfer of electronic signals between components that are attached at the surface. To demonstrate, a PCB research prototype is shown in Figure 1, where probes are in contact with surface traces at each end of a package interconnect.

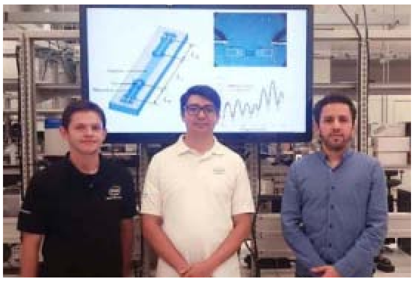

Engineers from the Intel Guadalajara Design Center, pictured from right to left, are: Juan C. Cervantes-Gonzalez, Carlos A. Lopez, and Isaac G. Farias-Camacho.

In making electronic devices smaller, the size and spacing of package interconnects is scaled down by necessity, which can make computational design optimization to improve signal speed and integrity more time consuming.

Interconnects operating at higher frequencies—or signal speeds—will also consume more power. The geometry and materials of package interconnects have to be redesigned to minimize power consumption and prevent signal loss for a given application, which applies particularly well in the case of PCBs given their range of uses.

"Using simulation to optimize the design of package interconnects is essential, requiring accurate models that capture the non-negligible couplings arising from complex 3D structures," explains Juan C. Cervantes-González, an engineer at Intel. "To make electromagnetic simulations of package interconnects even faster, we have exploited and validated a space mapping optimization algorithm. With this optimized approach to simulation, we can further decrease the length of the design cycle and time to market of the latest, high-speed interconnect technology."

» *MODELING HIGH-SPEED PACKAGE INTERCONNECTS*

FULL-WAVE ELECTROMAGNETIC SIMULATION is necessary to model signal propagation in package interconnects operating at higher frequencies. By solving the complete set of Maxwell's equations without any simplifying assumptions, simulations can accurately account for non-negligible electromagnetic couplings and impedance mismatch in complex 3D structures, important contributors that cause crosstalk and reflection compromising signal integrity.

Using COMSOL Multiphysics® software, the engineers at Intel developed a model of a single-ended interconnect line embedded in a PCB structure. A cross section of the model geometry is shown at left in Figure 2, and highlights the relevant design parameters that are optimized in their work.

Single-ended interconnects are

> *With this optimized approach to simulation, we can further decrease the length of the design cycle and time to market of the latest, high-speed interconnect technology."*
>
> **—JUAN C. CERVANTES-GONZALEZ, ENGINEER, INTEL**

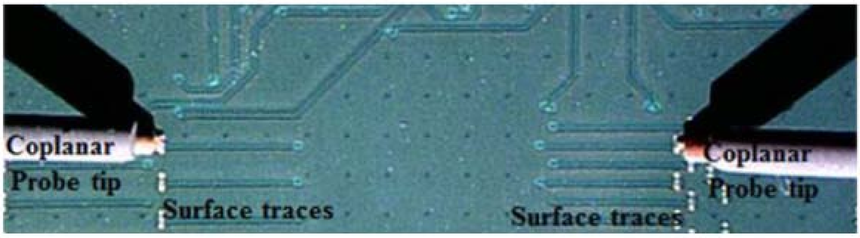

FIGURE 1: An Intel PCB research prototype featuring package interconnects.

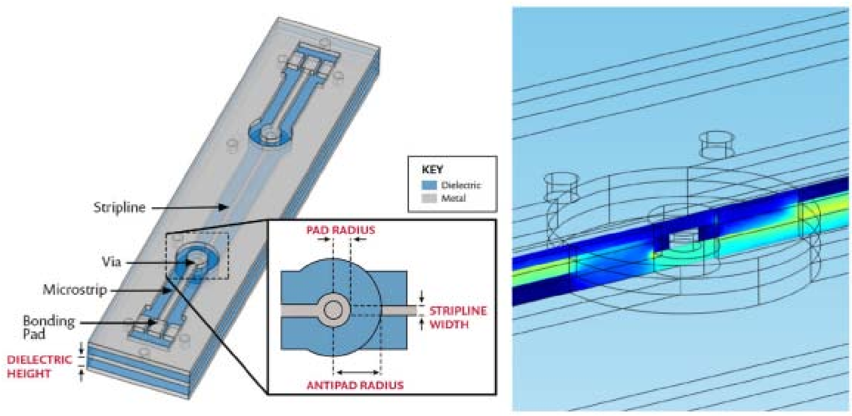

FIGURE 2: Left: Model geometry of a single-ended interconnect set up in COMSOL Multiphysics® software. Parameters highlighted in red are optimized using simulation. Right: Electric field distribution through a via in the modeled package interconnect.

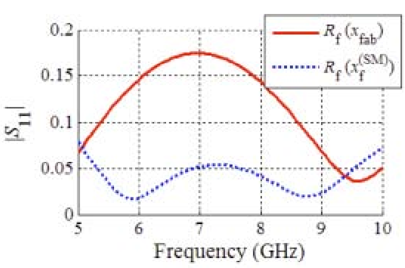

FIGURE 3: Magnitude of the reflected signal (|S11|) obtained by solving the fine model. Significant reflection is observed for the design parameters of the fabricated interconnect prototype (red curve) compared with the space-mapped optimized solution (blue curve).

known for their high signal speeds and for taking up less space in electronics packaging, enabling their use in densely packed layouts. Signals propagate laterally through interconnects along metallic microstrips and striplines with ground planes separated by dielectric material. Signals travel vertically through vias, often traversing more than 10 layers of dielectric and metal in a typical PCB, making them the primary culprits contributing to impedance mismatch.

In COMSOL® software, full-wave electromagnetic simulation is used to optimize the geometric parameters of a fabricated prototype in order to minimize impedance mismatch and the resulting signal reflection or return loss. Simulation results showing the electromagnetic field distribution in a single-ended interconnect are presented at right in Figure 2.

» OPTIMIZING INTERCONNECTS FOR LOW RETURN LOSS

TO OPTIMIZE THE DESIGN of high-speed package interconnects and minimize the magnitude of the reflected signal over a range of frequencies, simulations are run on a high-performance server with dual Intel® Xeon® X5670 CPU at 2.93 GHz and 160 GB of RAM. To make simulations run even faster, a Broyden-based input space mapping optimization algorithm was implemented.

The space mapping approach to electromagnetic simulation involves solving two separate interconnect models in COMSOL. The first is a "coarse" model and is a 2D simplification of the model geometry in Figure 2, which neglects electromagnetic losses and is discretized using a very coarse mesh designed to provide a fast result. The second is a "fine" model whose topology is identical to the first, but represents the full 3D geometry shown in Figure 2 and is solved using a much finer mesh, providing greater accuracy at the cost of speed.

The signal response for the optimal 2D model design was determined first using conventional optimization methods, and is an important first step that reduces the overall computation time. The objective of the space mapping algorithm, implemented in MATLAB® software, is to find the 3D model design parameters that make its resulting signal response close to the optimal 2D model response. Using this method, the interconnect design parameters are optimized within just four iterations. The results in Figure 3 were obtained by solving the full 3D model, and show a significant reduction in reflected signal for the optimized design compared with the original fabricated interconnect prototype.

"By using full-wave electromagnetic simulation along with space mapping optimization, a much better interconnect design is achieved with lower return loss, and in far less time than it would take to make and test many different prototypes," says Cervantes.

Although their initial model only solved for electromagnetic wave propagation in order to validate the space mapping optimization method, heat transfer and solid mechanics can also be included in fully coupled multiphysics models providing innovative, if not unconventional, design capabilities. ◎

Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.
MATLAB is a registered trademark of The MathWorks, Inc.

INCREASING LIFESPANS OF HIGH-POWER ELECTRICAL SYSTEMS

Using a combination of experimental testing and multiphysics simulation, researchers at ABB Semiconductors have redesigned the insulated-gate bipolar transistor modules (IGBT modules) used in high-power electrical components to increase device lifetime.

By **DEXTER JOHNSON**

THE HIGH-POWER electrical systems found in locomotives must be able to withstand the enormous amounts of stress brought on by the high currents and voltages surging through them. At the heart of these electrical systems are insulated-gate bipolar transistors (IGBTs), electronic switches that are used because of their high efficiency and fast switching to deliver power to locomotive systems. As a train travels from one station to the next, IGBT power modules are exposed to repeated electrical, thermal, and mechanical fatigue, which can degrade the module and cause failure.

"Typically, traction motors for driving locomotives are designed so the devices can withstand the harsh load profile for 30 years," explained Samuel Hartmann, principal R&D engineer at ABB Semiconductors in Lenzburg, Switzerland. If the IGBT modules wear out during the lifetime of the traction motor, they must be replaced. In order to meet the requirements of the traction motor's long lifetime and boost the reliability of these systems, Hartmann and his colleagues are leveraging computer simulation to better understand how the power cycling performance of IGBT modules can be increased.

FIGURE 1: ABB HiPak power module with rated blocking voltage of 6500 V and nominal current of 750 A.

"Our team is looking into ways to improve the performance of ABB's HiPak power modules," said Hartmann (see Figure 1). "The modules are composed of many paralleled IGBT chips, which, in their 'on' state, can conduct high levels of current, and in their 'off' state can resist very high voltages." The modules are also used for applications such as industrial drives and renewable energy.

"During use in a locomotive, IGBT power modules are exposed to high temperatures, and as a result, the joints between different components can degrade due to thermo-mechanical stress," described Hartmann. "After the weakest bond fails and the wire pulls away from the emitter, electrical contact is lost and the remaining wire bonds interconnecting the semiconductor device and its packag-

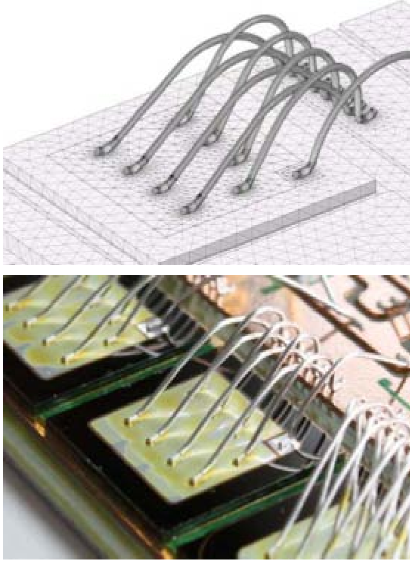

FIGURE 2: Top: The meshed COMSOL® software model of the reference wire bond layout. Bottom: Photo of the reference emitter.

ing must conduct higher currents. This eventually results in a cascading failure as the thermo-mechanical stress in the remaining bonds increases. If we can strengthen the weakest joint, then we can increase the overall lifetime of the device." By increasing the HiPak power module's usable lifetime, ABB can reduce the number of modules needed to reach the 30-year lifetime typically required of these

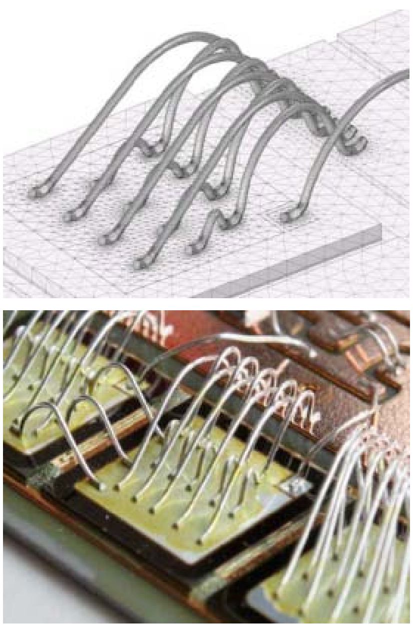

FIGURE 3: Top: The mesh of the stitch-bonded wire bond layout. Bottom: Photo of the stitch-bonded layout now used in some of the HiPak power modules.

motors, thereby saving resources and reducing the time needed for repairs.

» SIMULATIONS BRING CLARITY TO EXPERIMENTAL RESULTS

ABB'S HIPAK POWER modules typically consist of a baseplate, circuit boards, IGBT and diode chips, wire bonds, and conductor leads. In order to increase the lifetime of the power module, Hartmann explored a few different ways to increase the durability of the wire bond connections from the circuit board to the emitter bond pad.

"We explored two different methods for improving the design," described Hartmann. "In one case, we looked at different ways the wires were connected to the emitter to see if stitched bonding techniques could prevent component degradation and extend device lifetime." The meshed models and photos of the device for the commonly used reference wire bond layout and the stitch-bonded layout are shown in Figure 2 and Figure 3, respectively.

"For the second case, we used new joining techniques to bond a stress buffer between the emitter's silicon chip and the aluminum wire bonds," Hartmann continued. "The coefficient of thermal expansion (CTE) of the stress buffer is between the CTE of silicon and aluminum, and thus results in reduced

> *"Experimentally, we have assessed several variants of these wire bond connections and used multiphysics simulation to understand why one variant is better than the other."*
>
> **—SAMUEL HARTMANN, PRINCIPAL R&D ENGINEER, ABB SEMICONDUCTORS**

thermal and mechanical loading."

The ABB team leveraged multiphysics simulation to gain a better understanding of the underlying mechanisms at play in the deterioration of the IGBT chips, such as the electro-thermal and thermo-mechanical response of different designs when exposed to repeated power cycling tests. "The higher the power cycling capability, the more durable and reliable the design," explained Hartmann.

"Experimentally, we have assessed several variants of these wire bond connections and used multiphysics simulation to understand why one variant is better than the other." Hartmann believes that ABB's use of the COMSOL Multiphysics® simulation software was key to the success of their design.

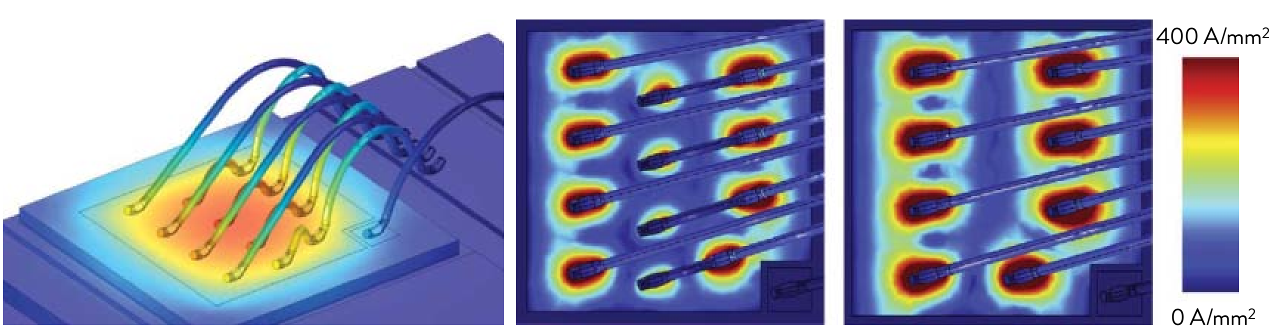

FIGURE 4: Left: COMSOL results showing the temperature distribution obtained for the stitch-bond layout. Right: Current density in the stitch-bonded and reference layout showing the reduction in current around the wire's feet for the new design.

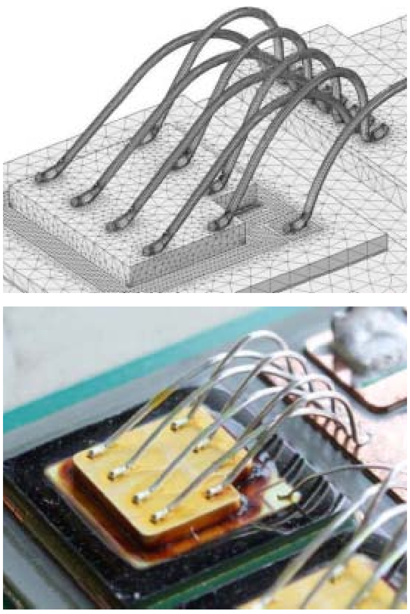

FIGURE 5: Top: The mesh of the reinforced emitter contact. Bottom: Photo of the reinforced emitter contact.

» *EXPLORING DIFFERENT IGBT MODULE DESIGNS WITH SIMULATION*

IN A FIRST EXPERIMENT, the ABB team tested two different bonding techniques: the reference wire bond layout (see Figure 2), and the stitch-bonded layout, where the wire is bonded to the surface of the chip more than once (see Figure 3).

Using simulation and experimental testing, Hartmann compared three different stitch-bonded layouts to the reference layout. "As expected, we found that, with more wire bonds on a single chip, the current density within the wires, especially at their feet, was also reduced," described Hartmann. "And thanks to simulation, we gained an unexpected insight: the stitch-bonded layout did not reduce temperature gradients or mechanical stress: the improved performances are due to the current density reduction resulting from a lower current density in the chip's metallization around the bond feet."

The additional bonds provide more locations for current to pass through the wires, therefore decreasing the amount of current dissipated by each wire (see Figure 4). "The new bond layout resulted in an IGBT design that has a power cycling capability that is four times higher than the reference layout. This new design is now being used in some of our HiPak power modules."

For a second experiment, Hartmann and his colleagues compared wires bonded directly on the chips with wires bonded to a metal plate attached to the chip that serves as reinforcement for the emitter (see Figure 5).

Using simulation, Hartmann found that for the reinforced emitter contact, the current density, the temperature variation, and the mechanical stress experienced by the wires at the bond interface was much lower than in the reference module (see Figure 6). This resulted in wire bonds that were

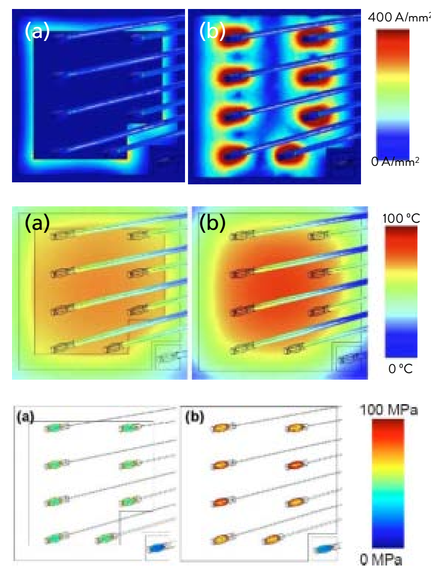

FIGURE 6: From top to bottom: Results from ABB's COMSOL model showing the simulated current density, temperature variation, and von Mises stresses at the bond-to-chip interface for the (a) reinforced and (b) reference emitter contacts on the front metallization and wire bonds.

Emre Oezkol (left) and Samuel Hartmann (right), holding the circuit board and the HiPak power module, respectively.

"In the IGBT modules that were reinforced, we saw that the wire bonds showed a cycling performance that was ten times that of the standard modules. With simulation we were able to verify that the mechanical stress was reduced, and this explained the dramatically increased durability."

—SAMUEL HARTMANN

less likely to detach from the emitter.

"In the IGBT modules that were reinforced, we saw that the wire bonds showed a cycling performance that was ten times that of the standard modules," said Hartmann. "With simulation we were able to verify that the mechanical stress was reduced, and this explained the dramatically increased durability."

» INCREASED LIFETIME FOR IGBT POWER MODULES

THE RESULT OF THE WIRE bond simulations and the new joining techniques has resulted in the lengthening of the lifetimes of ABB's power modules by a factor of 4 for the stitch-bonded layout and a factor of 10 for the reinforced emitter contact. This improved lifetime translates into higher power output per device, which in turn leads to lowering costs for ABB's power module customers.

"If the power cycling capability is increased, as in the case of the improved wire bond layout now available in our new modules," explained Hartmann, "then a lower number of power modules are needed to reach the 30-year lifetime of a traction motor, which is the standard in the traction industry. This directly reduces the cost of a locomotive and enhances the competitiveness of our power modules." ◎

JOULE HEATING AND THERMAL EXPANSION

BY HENRIK SÖNNERLIND

AN ELECTRIC CURRENT WILL generate heat through resistive losses, an effect called Joule heating. Since the resistivity usually has a strong dependence on the temperature, the heat transfer problem and the electrical problem must be solved simultaneously in order to accurately find the temperature and current distribution.

As an effect of the heating, thermal expansion will induce deformations. Large strains and stresses may then occur for several reasons. Deformations in materials with different coefficients of thermal expansion will not be compatible with each other; and there may also be large temperature gradients within a single material.

There are also certain cases where heat distribution is affected by structural deformations. For example, when objects come into contact with each other or large deformations cause changes in the electrical or thermal boundary conditions, a dramatic shape change occurs. If the heating cycle is repeated, the corresponding stress and strain cycles will be repeated as well. This may ultimately lead to a fatigue failure of the material.

In COMSOL Multiphysics® software, you can directly combine all these effects by selecting *Joule Heating and Thermal Expansion* in the list of available *Structural Mechanics* physics interfaces.

When doing so, the three contributing physics interfaces (*Solid Mechanics, Heat Transfer in Solids,* and *Electric Currents*) are added to the application, along with the necessary multiphysics couplings added through the *Multiphysics* node.

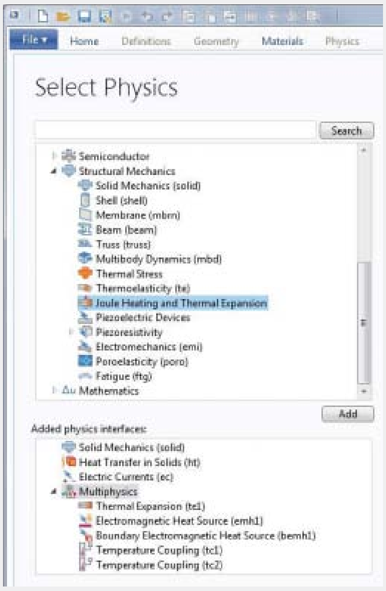

You can then choose settings for how to solve for the three physics interfaces. One approach would be to solve for the electric currents and temperatures together in a time-dependent study, and then solve the structural mechanics problem as stationary. Since the highest stresses could appear at any time during the thermal cycle, it is necessary to check stress values at several time steps.

For a qualitative comparison, it is sufficient to look at the computed stresses, but adding a *Fatigue* interface would make it possible to also make lifetime predictions.

BRINGING GLUCOSE MONITORING TO NEW LEVELS THROUGH INTEGRATED SENSOR DESIGN

Researchers and designers at Roche Diagnostics are developing glucose sensors with greater measurement accuracy for diabetic care, aided by multiphysics simulation.

By **LEXI CARVER**

CLOSE METABOLIC CONTROL through glucose monitoring is a well-known way for persons with diabetes to maintain good health and avoid medical complications. The current generation of glucose monitors relies on electrochemical methods to facilitate unprecedented measurement accuracy, and has given diabetics a reliable way to control their diet and insulin intake.

However, the chemical reactions that take place on the sensing strips used in glucose monitors are sensitive to environmental conditions and chemical interferences. Sensors are shipped worldwide, stored under uncertain conditions, and needed by users with different levels of knowledge and experience. Robust design is crucial for enabling sensors to survive these environments, deliver accurate results, and detect conditions that would cause errors. Now multiphysics simulation is used alongside experiments and calculations, enabling scientists to understand the chemical, electrical, and biological phenomena interacting in these systems so they can optimize their design and measurement methods.

» GAINING GROUND WITH A NEW KIND OF SENSOR

ENGINEERS AT ROCHE DIABETES CARE, a worldwide leader in diabetes diagnostic products and services, are currently pursuing a better understanding of the electrochemistry in their existing devices and are designing new sensing methods to provide more accurate monitoring. Like other amperometric biosensors, their glucometers (an example is shown in Figure 1) measure the electric current that results when a voltage is applied to an electrode system. The resulting current is proportional to the glucose levels in an electrolyte solution (such as a blood sample combined with a chemical reagent).

A set of gold traces lie on each glucose test strip, running from the electrode system in the strip to electrical contacts that insert into the glucose meter (see Figure 2). The reagent, which consists of a glucose-reactive enzyme and a very stable chemical referred to as a proto-mediator, is deposited on these electrodes during manufacturing and then dried. A capillary channel constructed over the electrode system receives a blood sample that rehydrates the reagent, causing it to react with glucose in the blood. "The initial reaction of glucose with the enzyme converts the proto-mediator to a reactive, low-potential mediator, which carries out the rest of the reaction," explains Harvey Buck, principal scientist at Roche Diagnostics Operations, Inc.

FIGURE 1: Photograph of an ACCU-CHEK Aviva® and ACCU-CHEK Nano® created at Roche Diagnostics.

» SIMULATION UNVEILS CHEMICAL AND ELECTRICAL MYSTERIES

THE CURRENT RESPONSE to a DC voltage applied at the electrodes during the reaction predicts glucose concentration in a blood sample, providing crucial information that tells a patient what action to take to correct their blood sugar levels. But configuration and manufacturing of the test strip affect this response accuracy. Using two COMSOL Multiphysics® software simulations, the Roche team was able to study a new test strip design—one of several they are investigating—and isolate the chemical reactions from the electrical, mechanical, and temperature conditions so that they could analyze the voltage response.

The isolated system contains many parameters and coupled variables, such as concentrations of different chemical species. The reagent system has so many complex interactions between the chemicals and their reactions that it was difficult to predict the response to different measurement methods or interfering substances. So the team made the simplifying assumption that mass transport of chemicals only occurs in a very thin layer above the electrode, thin

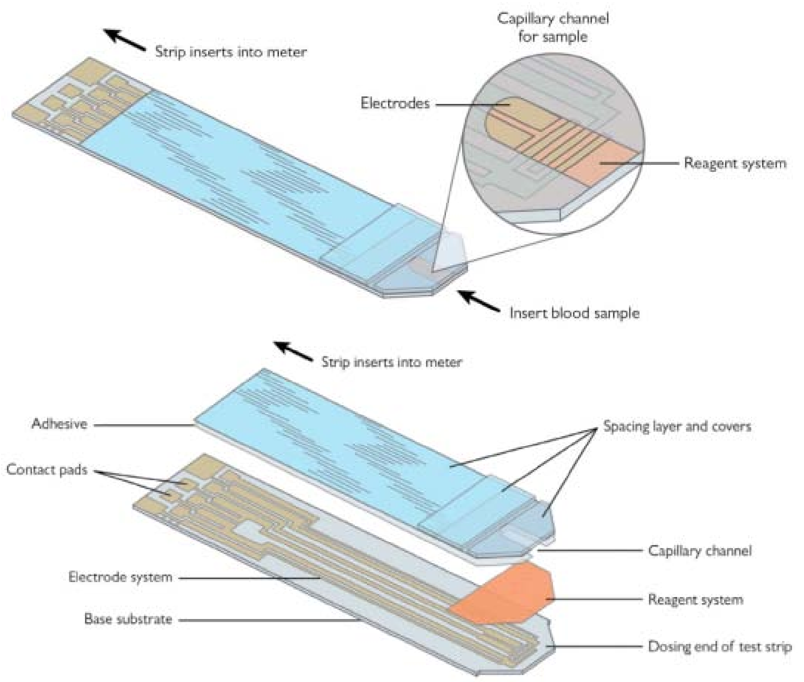

FIGURE 2: Schematic of test strip components. The chemical reaction occurs right on top of the electrodes. Adhesives and spacing layers form the curve of the capillary channel and bind together the electrodes, reagent system, and top and bottom covers.

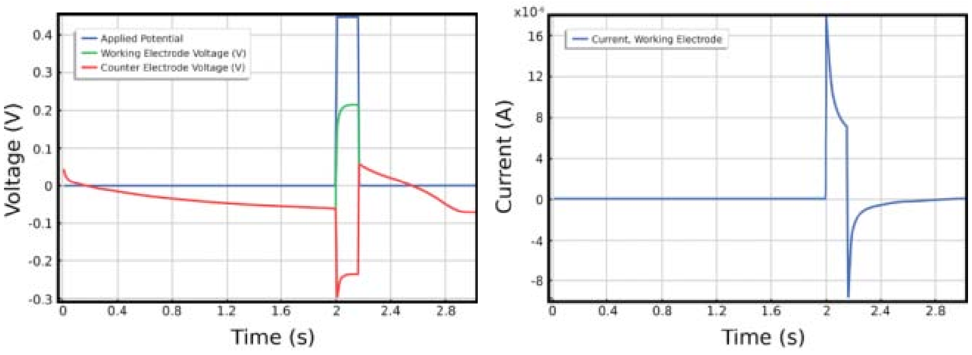

FIGURE 3: Simulation results showing the applied potential difference and the working and counter electrode potentials in the Roche sensor (left), as well as the current response to a potential difference step (right). The current response is proportional to the glucose concentration in the sample. The working and counter electrode potentials (green and red, respectively) are not measurable and are only available through the simulation.

enough for the reactivity to be considered uniform in the direction perpendicular to the surface. "We built a one-dimensional model that lets us understand and predict the responses, which required a combination of Michaelis Menten enzyme kinetics and mixed Butler-Volmer electrode kinetics," Buck comments.

Having established rates for the different reactions, the relevant equations were then easy to implement in the software. By restricting the model to one dimension, it was possible to predict the sensor response to different DC potential profiles with reasonable solution times (see Figure 3).

But the DC current is also affected by temperature and red blood cell fraction in the sample (called hematocrit), so prior to the DC measurement, an AC signal is applied to obtain impedance information used to compensate for these effects (see Figure 4). These are combined with the DC measurements in a mathematical algorithm, giving the sensor the information needed to make a truly accurate glucose prediction.

The capabilities of COMSOL® software proved particularly valuable for interpreting these complex measurements. "We quickly found during our modeling process that when you try to apply a large potential step to create diffusion-limited flux at an electrode, you risk causing unrealistically high potential calculations," says Buck. "In COMSOL it's very easy to use a log transform of the concentration variables, which really simplified the analysis process."

"The impedance measurements are very sensitive to the sample and not very sensitive to the reagent," Buck continues. "The electrode arrangement to enable impedance measurement is an integral part of the sensor design, and has a great influence on the measurement sensitivity." Buck's team built a second model of the cell to solve the electrical problem, this time in 3D. "The sample conductivity in the cell serves as a proxy for hematocrit variation. We're able to investigate different electrode configurations and materials, and predict the sensitivity of the

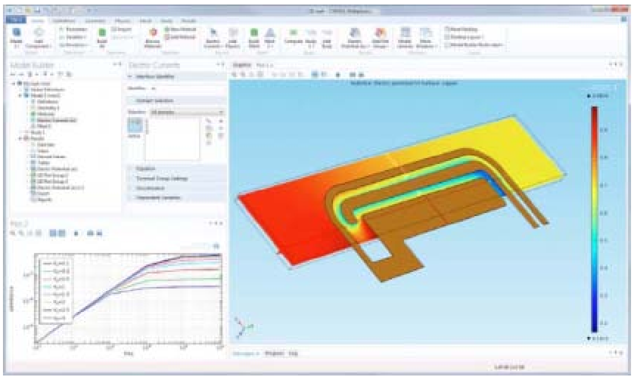

FIGURE 4: Buck's 3D COMSOL simulation showing the admittance response for different conductivities, plotted with a log scale (lower left) and a plot of the electric potential in the sensor measurement zone (right). The gold electrodes contact the electrolyte at a surface impedance interface.

> *"We're able to investigate different electrode configurations and materials, and predict the sensitivity of the impedance measurements to hematocrit as well as to other mechanical properties of the sensor, such as capillary height and spacer placement."*
>
> —**HARVEY BUCK, PRINCIPAL SCIENTIST, ROCHE DIAGNOSTICS**

impedance measurements to hematocrit as well as to other mechanical properties of the sensor, such as capillary height and spacer placement." (See Figure 5.)

The electrodes are fabricated from sputtered metal films whose resistance significantly affects the impedance measurements and potential distribution. It's impossible to measure the potential drop across the electrodes or within the electrolyte in the measurement cell without physically disturbing the system, but it is relatively easy to simulate it.

Relying on the COMSOL results for guidance, Buck adjusted the shape, length, and spacing of the working and counter electrodes until he had optimized the electrode design for impedance measurements. Ultimately, he was able to maximize the electrode sensitivity to hematocrit levels while minimizing sensitivity to manufacturing tolerances—thereby ensuring an accurate impedance measurement for the DC signal compensation. This paved the way for the new configuration to move toward production.

» APPROACHING NEW HORIZONS FOR GLUCOSE MONITORING

THROUGH THE CHEMICAL and electrical response correction modeled in COMSOL, the researchers at Roche have gained greater insight into their new sensor design and are delivering glucose monitors that correct the DC signal for more accurate measurements. Their innovative system, including its built-in sensing capabilities, sets a new standard for biosensing devices. Simulation allowed them to investigate parameters that were impossible to measure experimentally, make informed design decisions, and optimize their electrode configuration. Their continued research and modeling work is leading to the production of these new sensors and, ultimately, better care for persons with diabetes. ◎

Harvey Buck, principal scientist, Roche Diagnostics.

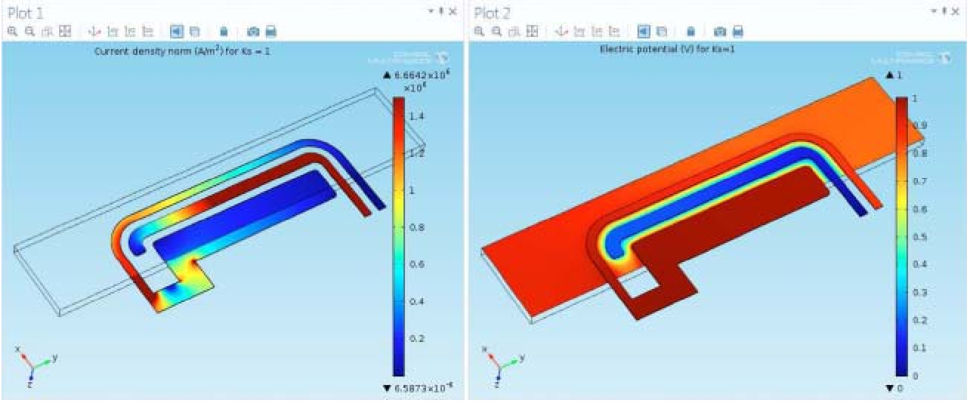

FIGURE 5: COMSOL simulation results showing the current distribution (left) and electric potential (right) in the electrodes and electrolyte.

SIMULATING PRINTHEAD UNIMORPH ACTUATORS AT FUJIFILM DIMATIX

Engineers at FUJIFILM Dimatix have used multiphysics simulation to gather compliance data for improving industrial printhead actuator performance.

By **LEXI CARVER**

THE REACH OF INDUSTRIAL INKJET PRINTERS is truly incredible—from commercial packaging and wide-format graphics to signage, textiles, and even electronic applications, inkjet printing enables the information sharing and communication that surrounds our everyday activities. FUJIFILM Dimatix, a premier producer of commercial inkjet printheads, is now using multiphysics simulation in the development of the MEMS actuators driving their newest ink deposition products.

» PRINTING WITH MICRON-SCALE PIEZOELECTRIC ACTUATION

A PRINCIPAL SCIENTIST on the research team at FUJIFILM Dimatix, Chris Menzel, is studying printhead actuation in order to design FUJIFILM's newest unimorph diaphragm actuators. These actuators are created in a MEMS fabrication process using a high-performance thin-film piezoelectric layer. This layer is a high-quality proprietary sputtered version of lead zirconium titanate (PZT), an electroceramic that changes shape under an applied electric field and is used in many transducers. The PZT is bonded to a silicon membrane and the actuators are then arrayed across the surface of a wafer, with each one corresponding to a tiny jet consisting of flow channels and a nozzle (see Figures 1 and 2). Thousands of these systems are packed tightly together in the printhead.

The components of each jet (the fluidic channels and the actuator) combine to form a resonant fluidic device. Upon electrical stimulation of the PZT by pulses tuned to stimulate the jet's resonance, the actuator deflects and generates acoustic waves within the closely coupled flow channels. The jet design effectively converts the pressure wave into an oscillating flow, which has to overcome the surface tension at the nozzle in order to throw an ink drop. When the resulting fluid momentum is large enough, the droplet is propelled outward and onto a substrate.

The goal of Menzel's design work was to define an actuator and jet flow channels that combine to generate a droplet meeting a target mass at a given

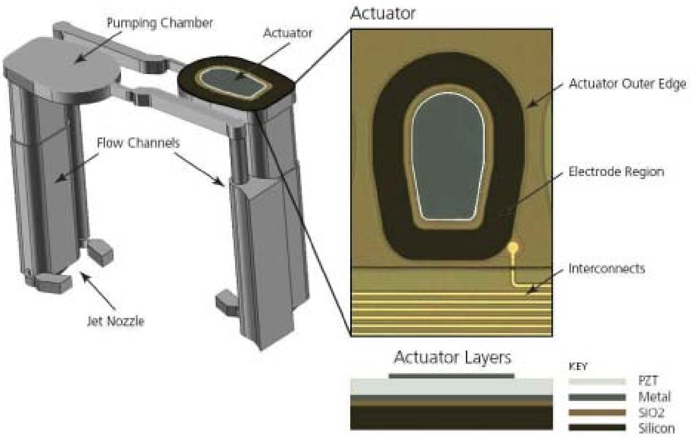

FIGURE 1: The printhead geometry developed by FUJIFILM. Each actuator sits on top of a pumping chamber containing a reservoir full of ink. Below the chamber are flow channels that carry ink to the nozzle.

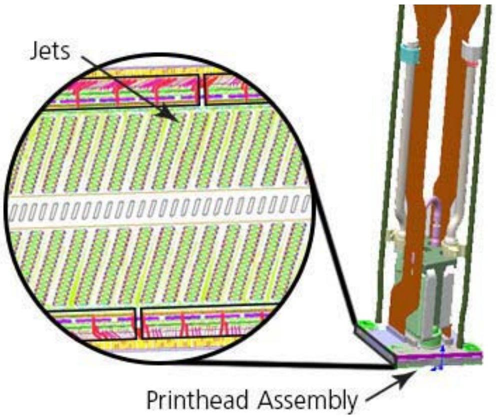

FIGURE 2: Magnification of jets on the wafer and their location in the printhead assembly.

velocity, with a target maximum firing frequency for the available voltage. Implicit in this design process is the need for miniaturization and the associated lower cost. With this in mind, the primary concerns in actuator design are maximizing deflection, minimizing size, and matching the actuator's impedance to the flow channels and the nozzle.

» SIMULATION REVEALS ACTUATOR COMPLIANCE AND OUTPUT

A TWO-STAGE MODELING approach was needed because the actuator performs its function within a jet system. In the first stage, Menzel determined functional parameters for various actuator geometries. He then used these parameters in a complete jet model to determine how the whole system would respond.

"We set up a COMSOL Multiphysics® software simulation to determine the actuator functionality," Menzel said. "Simulations offer an understanding of the relationships between functional parameters and the many layer thicknesses, boundary conditions, and sizes our process can generate. The software's ability to efficiently sweep through a large set of these variables and deliver easy-to-interpret results is of great value. It allows us to easily optimize our total system response, and hence, our product."

He modeled half of the actuator geometry along its central axis and included different layers for the silicon, metals, insulators, and PZT (see Figure 3, top). He also included a section of the ink-filled pumping chamber below the actuator and a section of a neighboring flow channel, then performed a simulation to extract the actuator's deflection under a pressure load (known as compliance) and the deflection under a voltage load (known as output) (see Figure 3, bottom). Menzel ran the study over a wide range of actuator geometries. The resulting values were applied to a larger-scale model used for system-level design optimization.

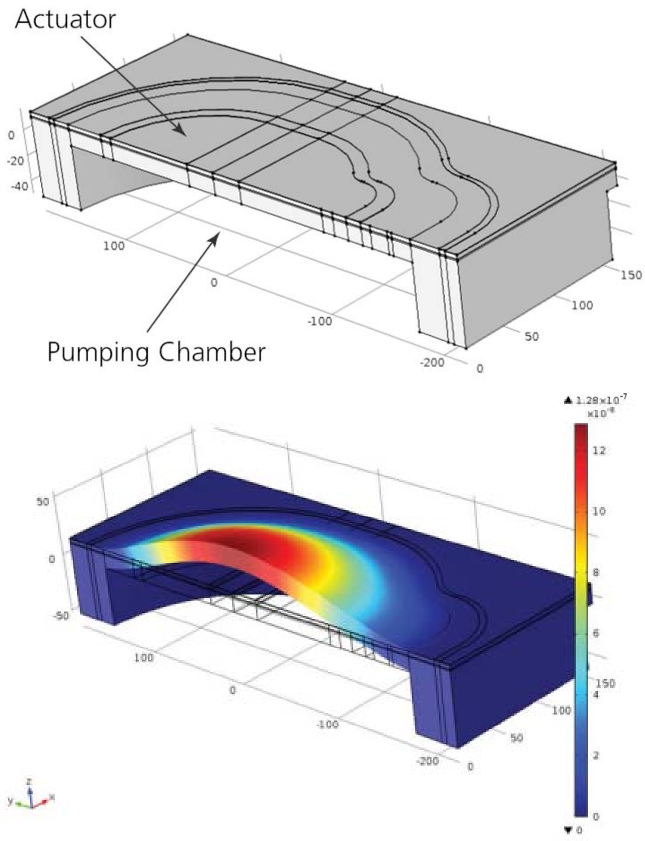

FIGURE 3: Top, Menzel's COMSOL® software model showing half of the actuator geometry with metal, silicon, PZT, electrodes, and pressurized ink chamber. Bottom, simulation results showing the deflection of the actuator.

» LOOKING AHEAD TO FASTER, SMALLER PRINTHEADS

THE COMSOL RESULTS led Menzel to an updated design by giving him the information needed to fit a new device to tight specifications and smaller actuator geometries. The multiphysics model revealed valuable information that allowed the engineering team to better understand the ins and outs of their actuator and jet. Modeling remains the starting point for evaluating actuator concepts and product feasibility; the associated reduction in design time is critical to effective and efficient product release. Even higher quality printing will soon be on the market as FUJIFILM Dimatix continues to lead the industry in printhead design, supported by simulation.

> *"The software's ability to efficiently sweep through a large set of these variables and deliver easy-to-interpret results is of great value. It allows us to easily optimize our total system response, and hence, our product."*
>
> —**CHRIS MENZEL, PRINCIPAL SCIENTIST, FUJIFILM DIMATIX**

SIMULATING THE PIEZOELECTRIC EFFECT

BY YESWANTH RAO

PIEZOELECTRIC MATERIALS are a family of solids, some natural and some man made, that become electrically polarized as a result of mechanical strain, a phenomenon known as the direct piezoelectric effect. They also exhibit an inverse piezoelectric effect where a mechanical strain results from an applied electric field. Piezoelectric materials are natural transducers that are used in many kinds of sensors and actuators.

The COMSOL Multiphysics® software offers a predefined *Piezoelectric Devices* interface that couples electrostatics and structural mechanics (Figure 1), which are essential for modeling these phenomena.

For accurate modeling, material properties and orientation must be carefully described. The *Piezoelectric Devices* interface allows the user to specify material properties in Stress-Charge or Strain-Charge form (Figure 2), with options for defining material orientation using Euler angles.

Piezoelectric materials are usually one of many components in a device. To capture the true behavior of the device as a whole, it is necessary to model the interactions between the piezoelectric devices and the surrounding materials. The multiphysics modeling capabilities in the COMSOL® software allow the *Piezoelectric Devices* interface to be readily coupled with physics such as pressure acoustics, fluid flow, and structural vibrations (Figure 3). It is also important to describe damping mechanisms that may affect device performance. COMSOL allows users to include mechanical damping, dielectric losses, conduction losses, and piezoelectric coupling losses.

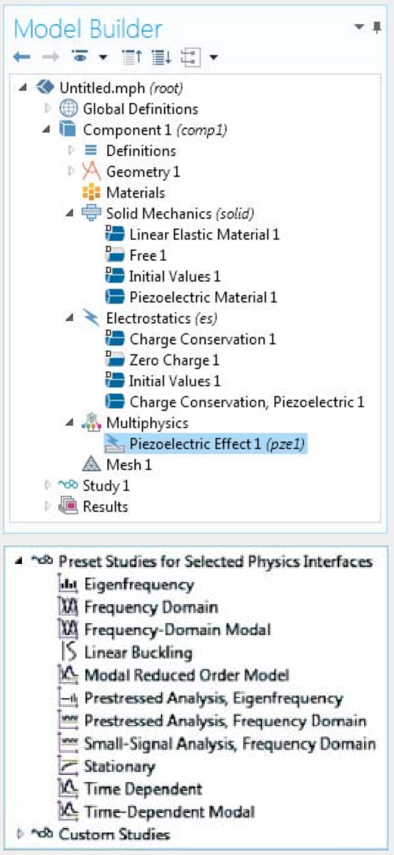

FIGURE 1: COMSOL® software Model Builder showing the setup for simulating the piezoelectric effect (top) and studies available to the user to simulate a piezoelectric application (bottom).

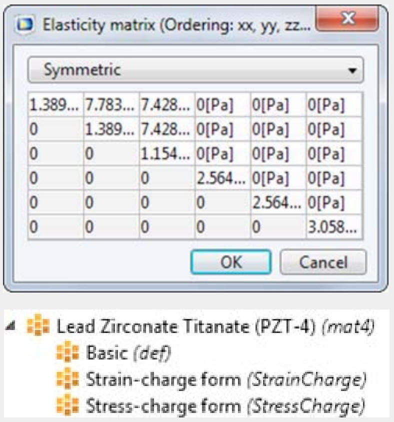

FIGURE 2: Strain-charge form settings showing the elasticity matrix, accessible from the Materials node in the Model Builder.

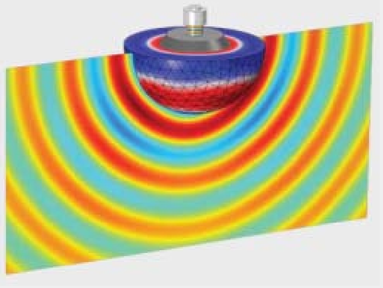

FIGURE 3: Simulation of a piezoelectric tonpilz transducer with results showing the acoustic pressure levels, including the far-field and voltage distribution in the piezoceramic rings. These transducers are used for low-frequency, high-power sound emission.

HPC-ENABLED SIMULATION AIDS IN THE DESIGN OF CUSTOMIZED HIGH-POWER ELECTRICAL DEVICES

COMSOL Multiphysics and its HPC capabilities get the best designs to customers more quickly than ever before.

By **DEXTER JOHNSON**

BLOCK TRANSFORMATOREN-ELEKTRONIK is a leading manufacturer in the field of coiled products that are used in a wide variety of industries, especially for electronics applications.

BLOCK engineers design custom transformers, power supplies, EMC filters, and reactors (see Figure 1), which usually have to meet precise specifications concerning working frequencies, product sizes and weights, electrical power losses, electrical insulation, as well as varying environmental conditions, including dirt, temperature changes, or moisture. Additionally, such equipment must often have product lifetimes of 30 years.

"Depending on the customer's application, there are restrictions to the materials that may be used," said Marek Siatkowski, who is responsible for all of BLOCK's simulation activities. "For example, in railway applications, the materials must meet strict requirements like flammability standards, smoke toxicity in case of fire, etc. We don't just open our catalog and they pick a device. The customer specifies a size and their requirements and each time we must do a new set of calculations."

Under all these circumstances, BLOCK found it increasingly more difficult to design inductors and transformers with aging simulation software. To save costs and in order to provide improved services to their customers, the company needed to find a way to reduce the number of prototypes it created before finalizing a design.

With this in mind, the company turned to the COMSOL Multiphysics® software for its ease-of-use, flexibility, and HPC (high-performance computing) capabilities.

FIGURE 1: Layout of a line reactor used to filter out spikes of current and reduce injection of harmonic currents into the power supply.

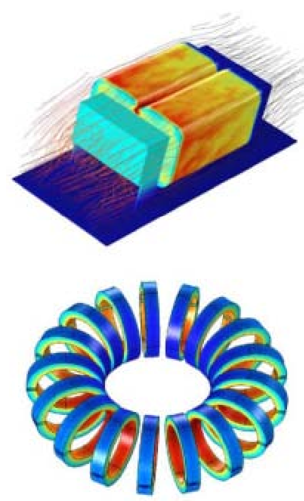

FIGURE 2: Simulation of an air cooled DC choke where temperature distribution and velocity streamlines are shown (top). Magnetic flux density in a toroidal choke (bottom). Its inductance is numerically determined as a function of inner and outer radius and wire thickness.

"We can model new devices and find critical areas, where, for example, electromagnetic losses are high or the temperature of the device reaches some threshold," said Siatkowski. "With COMSOL Multiphysics, we can identify these areas and simulate the relevant physics effects so that we can quickly and accurately find ways to improve the design."

The research department is analyzing magnetic characteristics and hysteresis losses in several soft magnetic materials in the BLOCK testing laboratory. One of the main reasons the company uses COMSOL® software is that it allows them to easily insert their own formulas developed over years for all these characteristics and to use them for their simulations.

» HPC LEADS TO GREATER THROUGHPUT

IN ADDITION to using multiphysics simulation, BLOCK is benefiting from the HPC capabilities that COMSOL Multiphysics offers: they can run their simulations on a multicore workstation with no limit to the number of cores and on a cluster with no limit to the number of compute nodes. This offered them improved efficiency regardless of whether a simulation is run on a workstation or a cluster; their R&D team can now quickly deliver the best products to customers.

Siatkowski uses COMSOL to set up models for many of BLOCK's devices, which are often difficult to calculate analytically, but have a geometry that can be based on a few parameters and specific customer's needs. One example of a model that

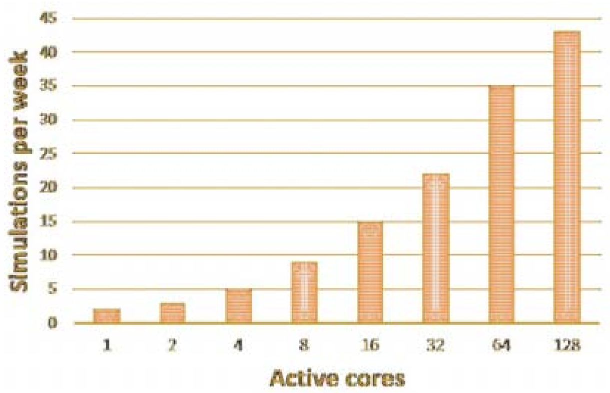

FIGURE 3: Moving from a single workstation with eight cores to a modestly sized cluster can lead to a significant performance increase.

Siatkowski built was for a DC choke (see Figure 2).

"With COMSOL Multiphysics, I can run a simulation that has parameters like width, height, thickness of the wires, etc. and explore the entire design space defined by our teams and customers. Our product developers and sales teams can now work more efficiently and easily find the best solution," explained Siatkowski.

» IT'S ALL IN THE ARCHITECTURE

"**FOR SMALLER MODELS,** I can build a model on my workstation and run the computation there," explained Siatkowski. "But for the larger models, my workstation is not fast enough and does not have enough memory."

This is when the flexible nature of COMSOL came into play and BLOCK fully benefited from the available HPC capabilities supported by the software architecture and generous licensing. Siatkowski instead runs his models on several computers with multiple cores.

"I'm currently using a cluster with 22 cores and 272 GB of RAM and I can easily run my simulations remotely on it," said Siatkowski. "COMSOL supports distributed memory computing where each node of a cluster can also benefit from local shared memory parallelism; this means that I'm getting the most out of the hardware available." The speedup obtained in terms of simulations per week for a large electrical study is shown in Figure 3.

After executing the simulation on the high-performance computer, Siatkowski reviews the result on his workstation, where he can then also perform postprocessing. "The benefit of this is that during the simulation itself, my workstation is free and I can continue with other work and even do pre- or postprocessing on other models. The architecture that the COMSOL software has allows us to be more productive and service our customers better." ◎

From left to right: C. Kliesch (Bachelor Student), Dr. M. Siatkowski (Advanced R&D), M. Owzareck (Advanced R&D), A. Bimidi (Student Apprentice), Y. Kumar (Master Student), Dr. D. Kampen (Head of Advanced R&D)

CONTINUED FROM PAGE 32

These tools allow our engineers to tailor our designs to meet customers' needs with minimal time and input. In the past, these recurring analyses took hours and required an employee specializing in simulation; with a COMSOL application, employees at all levels of our organization can run simulations nearly effortlessly.

All told, multiphysics simulation and application design through COMSOL allows our designers to make better, more competitive products. Efficiency is core to our company philosophy—doing more, using less. This is not limited to the efficiency of our products, but also in the way we conduct business, generate ideas, and create new designs. The Application Builder is now a vital element in helping APEI build the best wide band gap solutions possible. ◎

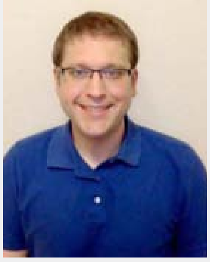

BRICE MCPHERSON *is a senior engineer at APEI, with 11 years of experience in high performance, extreme environment wide band gap power semiconductor packages. He specializes in the parametric CAD design and analysis of APEI's power modules and conversion systems.*

SIMULATION APPS STREAMLINE THE DESIGN OF POWER ELECTRONICS

By **BRICE MCPHERSON**

POWER ELECTRONICS ARE ESSENTIAL in nearly every application that uses electricity, from cell phone chargers to industrial scale power distribution. Different applications require converting power from one form to another. For example, driving the motor on an electric vehicle requires power switches, drivers, filters, sensors, and control circuitry. These conversion systems need to process power as efficiently, safely, and cost-effectively as possible.

At APEI we are pushing the limits of power electronic systems, developing advanced solutions utilizing wide band gap semiconductors that can block higher voltages, carry larger amounts of current, switch on and off more quickly, and more effectively dissipate waste heat than traditional semiconductors. These benefits are key to processing large amounts of power in increasingly smaller volumes and at higher efficiencies.

The COMSOL Multiphysics® software and the Application Builder are fundamentally changing the way that we design, support, and market our products. In the real world, most problems are not constrained to singular, isolated physical phenomena. For electronic systems, thermal, electrical, and mechanical behavior are closely intertwined; their effects and interactions must be studied simultaneously in order to see the full picture of factors driving performance. COMSOL Multiphysics simulations have been essential tools for our engineers to extract a more detailed understanding of our products, virtually assess real-world performance, and reduce the amount of prototyping needed.

A COMSOL® software simulation application created at APEI for predicting the performance of different power module designs.

When the Application Builder was released, we were eager to try it out. It was surprising to see how easy it was to build our first application—a tool to analyze the fusing current and impedance of the tiny bond wires used to interconnect semiconductor devices. It took little time to transition an existing COMSOL® software simulation to an application designed for ease of use, while still based on a powerful multiphysics model. A drag-and-drop graphical interface, straightforward controls and entry fields, and full integration within the COMSOL environment narrowed the learning curve considerably. In short, if you can build a model, you can easily build an app from it.

We now have multiple apps ranging from simple design tools to comprehensive analyzers that extract all relevant performance and design metrics for custom configurations of our power modules.

CONTINUED ON PAGE 31

The story behind Airbnb's pricing algorithm

How Much Is Your Spare Room Worth?

by **DAN HILL**

Photography by **THE VOORHES**

How much should you charge someone to live in your house? Or how much would you pay to live in someone else's house? Would you pay more or less for a planned vacation or for a spur-of-the-moment getaway? • Answering these questions isn't easy. And the struggle to do so, my colleagues and I discovered, was preventing potential rentals from getting listed on our site—Airbnb, the company that matches available rooms, apartments, and houses with people who want to book them. • In focus groups, we watched people go through the process of listing their properties on our site—and get stumped when they came to the price field. Many would take a look at what their neighbors were charging and pick a comparable price; this involved opening a lot of tabs in their browsers and figuring out which listings were similar to theirs. Some people had a goal in mind before they signed up, maybe to make a little extra money to help pay the mortgage or defray the costs of a vacation. So they set a price that would help them meet that goal without considering the real market value of their listing. And some people, unfortunately, just gave up. • Clearly, Airbnb needed to offer people a better way—an automated source of pricing information to help hosts come to a decision. That's why we started building pricing tools in 2012 and have been working to make them better ever since. This June, we released our latest improvements. We started doing dynamic pricing—that is, offering new price tips daily based on changing market conditions. We tweaked our general pricing algorithms to

consider some unusual, even surprising characteristics of listings. And we've added what we think is a unique approach to machine learning that lets our system not only learn from its own experience but also take advantage of a little human intuition when necessary.

In the online world, a number of companies use algorithms to set or suggest prices. eBay, for example, tells you what similar products have sold for and lets you choose a price based on that information. eBay's pricing problem was relatively simple to solve: It didn't matter where the sellers and buyers were or whether you're selling the product today or next week. Meanwhile, over at Uber and Lyft, the ride-sharing companies, geography and time do matter–but these two companies simply set prices by decree; there is no user choice or need for transparency in how the prices are determined.

At Airbnb, we faced an unusually complex problem. Every one of the million-plus listings on our site is unique, having its own address, size, and decor. Our hosts also vary in their willingness to play concierge, cook, or tour guide. And events–some regular, like seasonal weather changes; others unusual, like large local events–muddy the waters even further.

Three years ago we started building a tool to provide price tips to potential hosts based on the most important characteristics about a listing, like the number of rooms and beds, the neighboring properties, and certain amenities, like a parking space or even a pool. We rolled it out in 2013, and it did well, for the most part. But it had limitations. For one, the way its price-setting algorithms worked didn't change. If we set them to consider that the Pearl District in Portland, Ore., say, had a certain boundary, or that rooms on a river were always worth a certain amount more than rooms a block from that river, the algorithm would apply those metrics forever, unless we went in manually to change them. And our pricing tools weren't dynamic–price tips didn't adjust based on when you were booking a room or how many other people seemed to be booking rooms at the same time.

Since mid-2014, we've been trying to change that. We wanted to build a tool that learns from its mistakes and improves by interacting with users. We also wanted the tool to adjust to demand and, when necessary, drop price tips to fill rooms that would otherwise stay empty or raise them in response to demand. We've started to figure that out, and we began to let our hosts use this new tool in June. We'll tell you about how these tools evolved and how they work today. We'll also tell you why we think our latest tool, Aerosolve, will eventually do a lot more than just price home rentals. That's why we're releasing it into the open-source community.

To get an idea of the problem we faced, consider three different situations.

Imagine you had lived in Brazil during the last football (soccer) World Cup. Your hometown will see a huge influx of travelers from all over the world, all united by the greatest football tournament on the planet. You have a spare room in your house, and you want to meet other football lovers and make some extra cash.

For our tool to help you figure out a price, there were a few factors to consider. First, this was a once-in-a-generation event in that country, so we at Airbnb have absolutely no historical data to look at. Second, every hotel was sold out, so clearly there was a massive imbalance between supply and demand. Third, the people coming to visit already had paid immense sums for their tickets and international travel, so they'd probably be prepared to pay a lot for a room. All of that had to be considered in addition to the obvious parameters of size, number of rooms, and location.

Or imagine you've inherited a castle in the Highlands of Scotland and, in order to pay the costs of cleaning the moats, operating the distillery, and feeding the falcons, you decide to turn the turret into a bed-and-breakfast. Unlike the World Cup situation, you'd have some comparative data, based on nearby castles. Some of that data would likely span many years, providing information about the seasonality of tourism and travel. And you'd know, because there are a number of other accommodation options in the area, that the supply and demand for tourist rooms is pretty balanced right now. Yet this particular castle is the only one in Scotland with a uniquely designed double moat. How can a system calculate what this rare and unique feature would be worth?

As a final example, imagine you own a typical two-bedroom apartment in Paris. You're taking a few weeks of your August vacation and heading south to Montpellier. There are lots of comparative properties, so it's relatively easy to price. But say, after you receive a wave of interest based on the first listing, you decide to start increasing the price gradually to try to maximize the amount earned. That's a tricky proposition–what happens if you go too high, or cut it too close to the booking date and lose the chance to make any money? Or perhaps the opposite occurs: You take the first inquiry at a lower price and spend the next few months wishing you'd been brave enough to take a little more risk. How do we help hosts get better information to prevent this kind of uncertainty and regret?

These are the kinds of questions we faced. We wanted to build an easy-to-use tool to feed hosts information that is helpful as they decide what to charge for their spaces, while making the reasons for its pricing tips clear.

AROUND THE GLOBE: Airbnb's pricing tools handle a variety of accommodations in many different countries, including [from top] a yurt in London, a castle in Ireland, and a tree house in Hawaii. At the time of a quick spot check in July, prices set by hosts were US $147 for a room for two in the castle, $159 for two beds sleeping four people in the yurt, and $275 for the tree house, which sleeps two.

AIRBNB (3)

The overall architecture of our tool was surprisingly simple to figure out: When a new host begins adding a space to our site, our system extracts what we call the key attributes of that listing, looks at other listings with the same or similar attributes in the area, finds those that are being successfully booked, factors in demand and seasonality, and bases a price tip from the median there.

The tricky part began when we tried to figure out what, exactly, the key attributes of a listing are. No two listings are the same in design or layout, there are listings in every corner of a city, and many aren't just apartments or houses but castles and igloos. We decided that our tool would use three major types of data in setting prices: similarity, recency, and location.

For data on similarity, we started with all the quantifiable attributes we know about a listing and then looked to see which were most highly correlated with the price a guest would pay for that listing. We arrived at how many people the space sleeps, whether it's an entire property or a private room, the type of property (apartment, castle, yurt), and the number of reviews.

Perhaps the most surprising attribute here is the number of reviews. It turns out that people are willing to pay a premium for places with many reviews. While Amazon, eBay, and others do rely on reviews to help users make selections about what to buy or whom to buy it from, it's not clear that the number of reviews makes a big difference in price. For us, having even a single review rather than no reviews makes a huge difference to a listing.

We considered recency, because markets change frequently, especially in travel. On top of that, travel is a highly seasonal business, so it's important to look for the market rate either as it is today, or as it was this time last year; last month may not be relevant.

In highly developed markets like London or Paris, obtaining this market data is easy enough—there are thousands of listings being booked on our site to compare with. For new and emerging markets, we classified them into groups of similar size, level of tourism, and stage of growth for Airbnb. This way, we are able to compare listings not only in the actual city a space is in but also in other markets with similar characteristics. So if a Japanese host is one of the first Airbnb users to list an apartment in Kyoto, we might look at data from Tokyo or Okayama, back when Airbnb was similarly new in those cities, or from Amsterdam, a more mature market for Airbnb but one with a similar size and level of tourism.

Finally, we needed to consider location, a rather different problem for us than for hotels. Hotels are typically grouped in just a few main locations; we have listings in almost every corner of a city.

Early versions of our pricing algorithms plotted an expanding circle around a listing, considering similar properties at varying radii from the listing location. This seemed to work well for a while, but we eventually discovered a crucial flaw. Imagine our apartment in Paris for a minute. If the listing is centrally located, say, right by the Pont Neuf just down from the Louvre and Jardin des Tuileries, then our expanding circle quickly begins to encompass very different neighborhoods on opposite sides of the river. In Paris, though both sides of the Seine are safe, people will pay quite different amounts to stay in locations just a hundred meters apart. In other cities there's an even sharper divide. In London, for instance, prices in the desirable Greenwich area can be more than twice as much as those near the London docks right across the Thames.

We therefore got a cartographer to map the boundaries of every neighborhood in our top cities all over the world. This information created extremely accurate and relevant geospatial definitions we could use to accurately cluster listings around major geographical and structural features like rivers, highways, and transportation connections.

So now, for example, for the first weekend of October, the price tip for a basic private room for two in London on the Greenwich side of the Thames comes up at US $130 a night; a room with similar attributes across the river comes up at $60 a night.

We were pretty happy with our algorithms—that is, after we fixed a bug that had caused our system to give a price tip of $99 on a large number of new listings, no matter what their particular characteristics. It didn't happen for long, and not in every region, but we recognize that when this happens it may cause people to question whether our pricing tools are working.

We improved our algorithms over time until they were able to consider thousands of different factors and understand geographic location on a very detailed level. But the tool still had two weaknesses. The tips it gave were static—it understood local events and peak tourist seasons, and so would suggest different prices for the same property for different dates of the year. It didn't, however, change those prices as the date approached, as airlines do, dropping prices when bookings were slow and raising them when the market heated up.

And the tool itself was static. Its tips did improve somewhat as it tapped into ever more historical data, but the algorithms themselves didn't get better.

Last summer, we started a project to address both of these problems. On the dynamic pricing side, our goal was to give each host a new pricing tip every day for each date in the future the property is available for booking. Dynamic pricing isn't new. Airlines began applying it several decades ago, adjusting prices—often in real time—to try and ensure maximum occupancy and maximum revenue per seat. Hotels followed suit as consolidation made the large chains larger, bringing them an ever-increasing amount of data about their business, and hotel marketing moved online, allowing the chains to change prices multiple times a day.

So investing in dynamic pricing—once we had several years of historical data about a large number of properties to tap—made a lot of sense for us despite the fact that it requires more computing resources.

Making the algorithms improve themselves over time was harder, particularly because we wanted our system to allow humans to easily interpret, and in some cases influence, the computer's "thought process" as it did so. Machine-learning systems of the size and complexity required to handle our needs often work in mysterious ways. The Google Brain that learned to find cat videos on the Web, for example, has layers upon layers of algorithms that classify data, and the way it gets to its conclusions—cat video or not—is virtually impossible for a human to replicate.

We selected a machine-learning model called a classifier. It uses all of the attributes of a listing and prevailing market demand and then attempts to classify whether it will get booked or not. Our system calculates price tips based on hundreds of attributes, such as whether breakfast is included and whether the guest gets a private bath. Then we began training the system by having it check price tips against outcomes. Considering whether or not a listing gets booked at a particu-

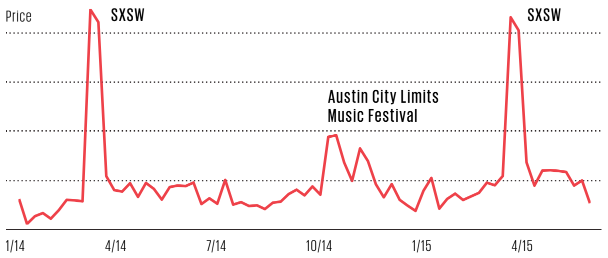

UPS AND DOWNS: Seasonal demand and local events can cause prices to differ dramatically. In Austin, Texas, as shown here, prices jump during the South by Southwest (SXSW) and Austin City Limits festivals.

AIRBNB

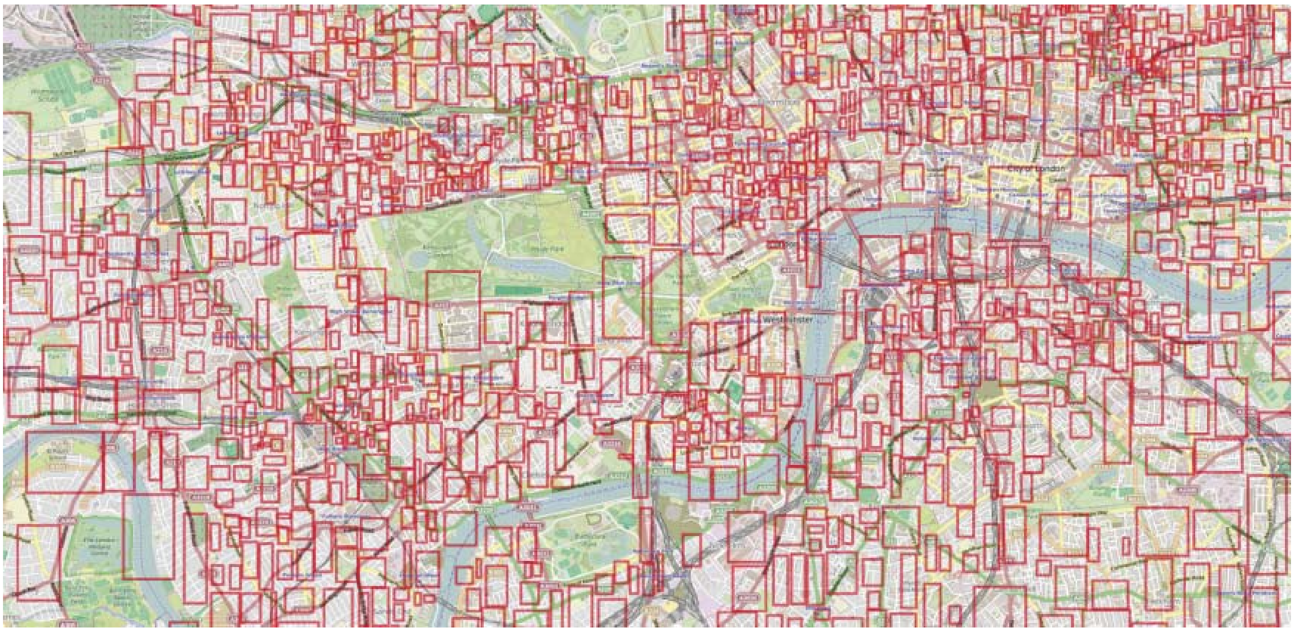

WON'T YOU BE MY NEIGHBOR? Algorithms use historical pricing data to group properties into detailed microneighborhoods, as shown on this map of London.

AIRBNB

lar price helps the system hone its price tips and estimate the probability of a price being accepted. Our hosts, of course, can choose to go higher or lower than the price tip, and then our system adjusts its estimate of likelihood accordingly. It later checks back on the fate of the listing and uses this information to adjust future tips.

Here's where the learning comes in. With knowledge about the success of its tips, our system began adjusting the weights it gives to the different characteristics about a listing–the "signals" it is getting about a particular property. We started out with some assumptions, such as that geographic location is hugely important but that usually the presence of a hot tub is less so. We've retained certain attributes of a listing considered by our previous pricing system, but we've added new ones. Some of the new signals, like "number of lead days before booking day," are related to our dynamic pricing capability. We added other signals simply because our analysis of historical data indicated that they matter.

For instance, certain photos are more likely to lead to bookings. The general trend might surprise you–the photos of stylish, brightly lit living rooms that tend to be preferred by professional photographers don't attract nearly as many potential guests as photos of cozy bedrooms decorated in warm colors.

As time goes on, we expect constant automatic refinements of the weights of these signals to improve our price tips.

We can also go in and influence the weighting if we believe we know something that the model has yet to figure out. Our system can produce a list of factors and weights considered for each price tip, which we have our people looking at. If we think something isn't well represented, we will add another signal manually to the model.

For example, we know that a listing in Seattle without Wi-Fi access to a broadband Internet connection is extremely unlikely to get booked at any price, so we don't have to wait for our system to figure it out. We can adjust that metric ourselves.

Our system also constantly adjusts our maps to reflect changes in neighborhood boundaries. So instead of relying on local maps to tell us, say, where Portland's Sunnyside neighborhood ends and Richmond begins, we are relying on the data on bookings and price differentials within a city to draw those kinds of lines. This approach also lets us spot microneighborhoods that we were not previously aware of. Such areas may have a large number of popular listings that don't necessarily map to standard neighborhood boundaries, or there may be some local feature that makes a small section of a larger traditional neighborhood more desirable.

These tools are generating price tips for Airbnb properties globally today. But we think it can do a lot more than just better inform potential hosts as they choose prices for their online rentals. That's why we've released the machine-learning platform on which it's based, Aerosolve, as an open-source tool. It will give people in industries that have yet to embrace machine learning an easy entry point. By clarifying what the system is doing, it will remove the fear factor and increase the adoption of these kinds of tools. So far, we've used it to build a system that produces paintings in a pointillist style. We're eager to see what happens with this tool as creative engineers outside of our industry start using it. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/airbnb0915

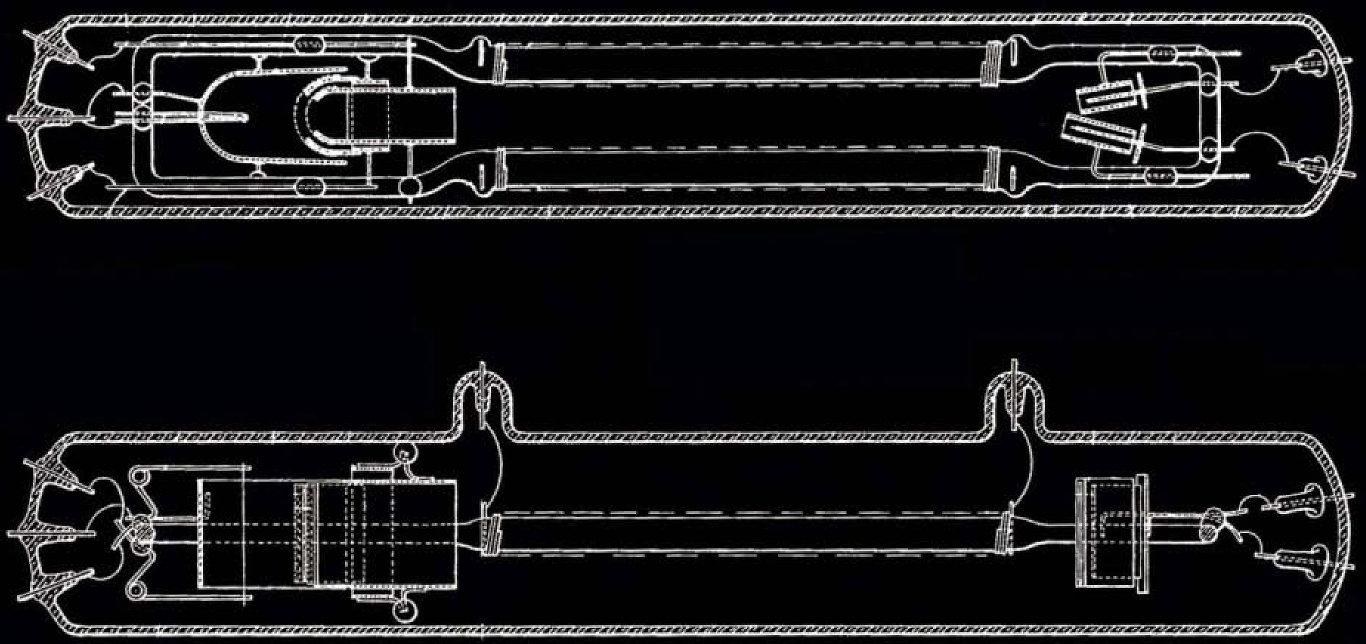

PATENT PENDING RECOGNITION: The patent Haeff filed in 1933 for a primitive type of traveling-wave tube has been largely ignored.

LEFT: U.S. PATENT AND TRADEMARK OFFICE; RIGHT: THE ANDREI V. HAEFF PAPERS

An intense-looking young man stepped out of his engineering lab at Caltech to watch what was happening. In the university's nearby high-voltage laboratory, gigantic bolts of electricity were leaping eerily from outlandish equipment. It was 1931, and a Hollywood crew was filming the spark-filled special effects for the creation scene in Boris Karloff's first Frankenstein movie. The serious-minded young engineer loved cinema, but as he walked back to his bench in the Kellogg Radiation Laboratory, he probably had no idea that a new kind of vacuum tube he was working on would in time revolutionize the movie business, enabling TV broadcasters to bounce *Frankenstein* and countless other films off satellites straight into people's homes.

The young man was Andrei "Andy" Haeff (father of coauthor Andre Haeff), and the device he was working on was an early form of what was later called the traveling-wave tube. This rather exotic type of vacuum tube was a key component in early telecommunications, radar, and television-broadcast systems. In 1952, for example, the British Broadcasting Corp. used the technology to forge a chain of microwave

THE TRUE HISTORY OF THE

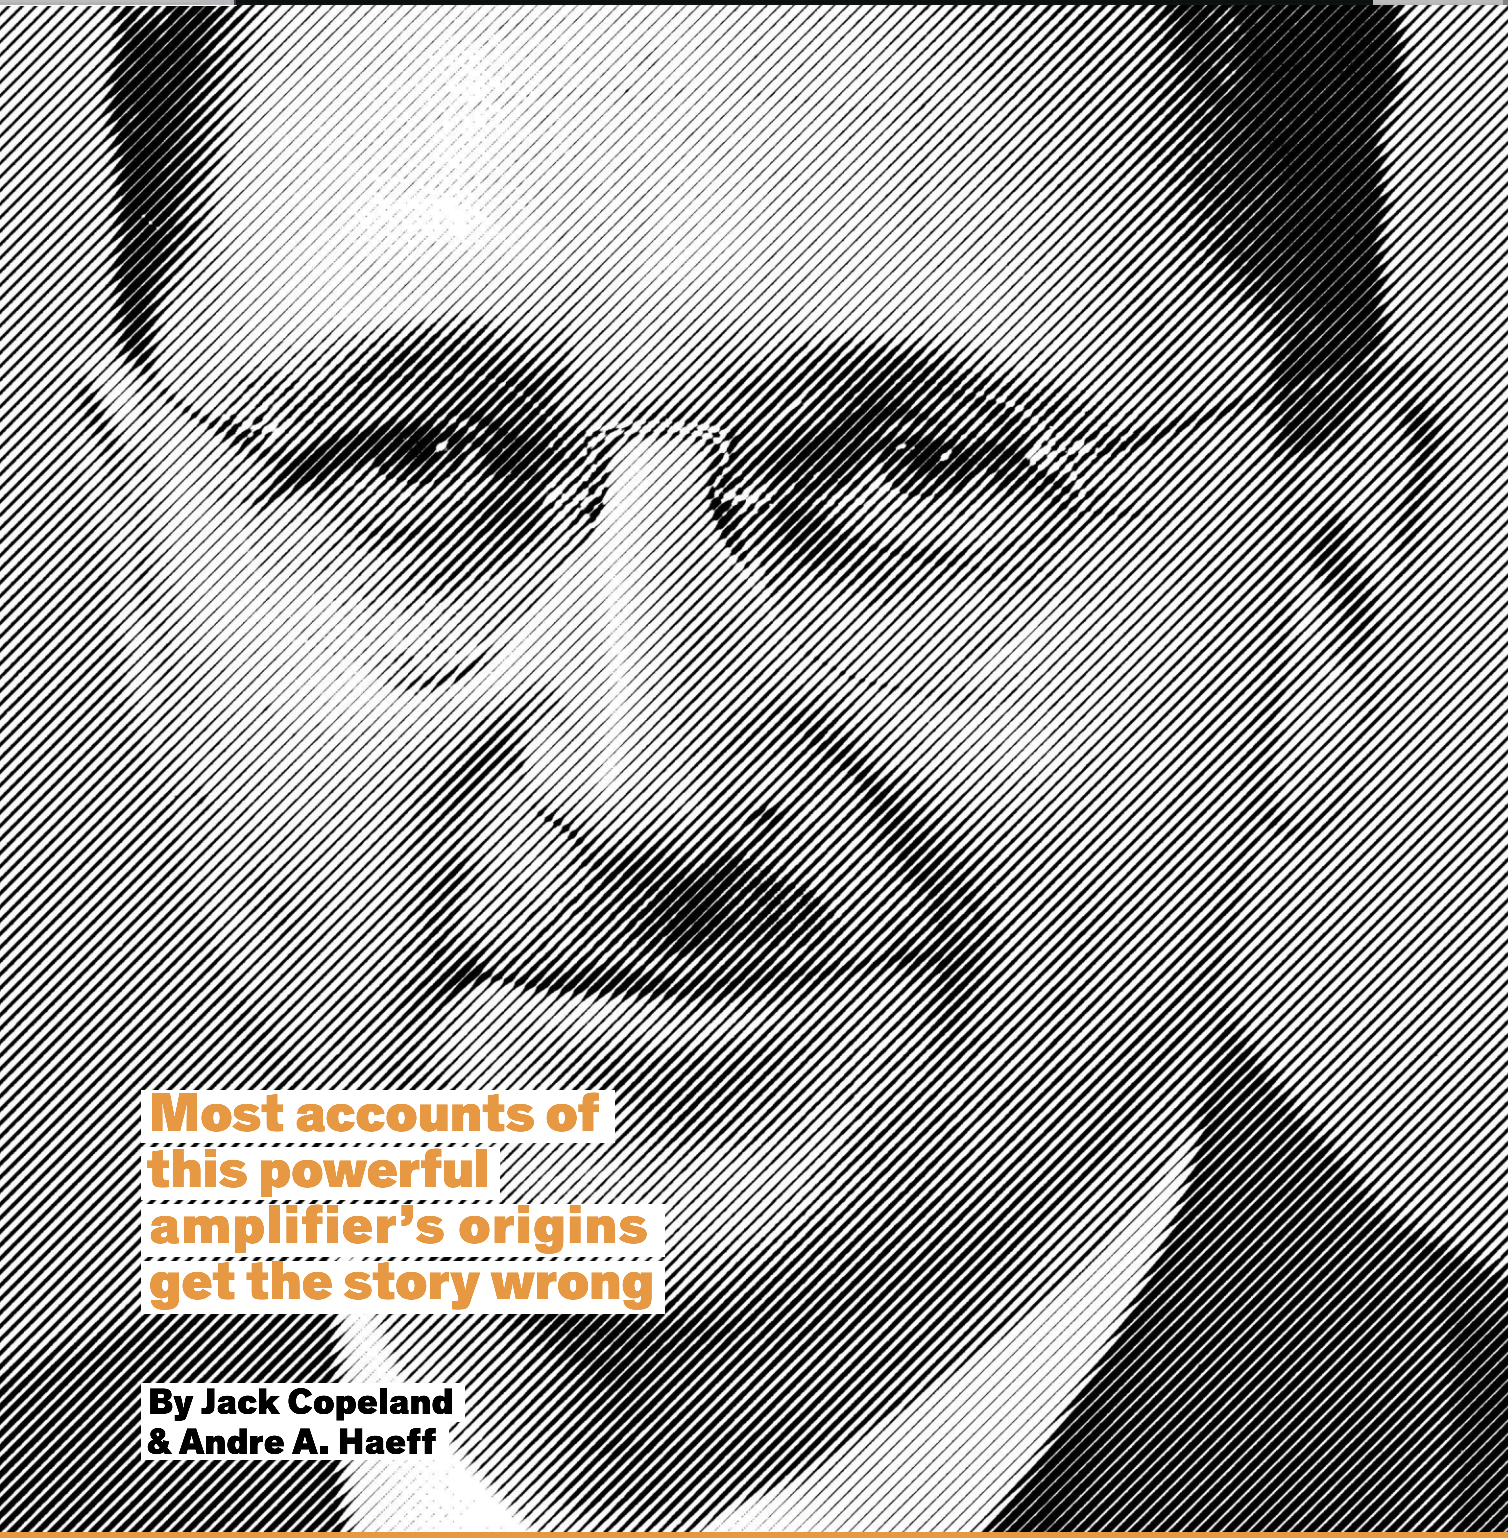

Most accounts of this powerful amplifier's origins get the story wrong

By Jack Copeland & Andre A. Haeff

TRAVELING–WAVE TUBE

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

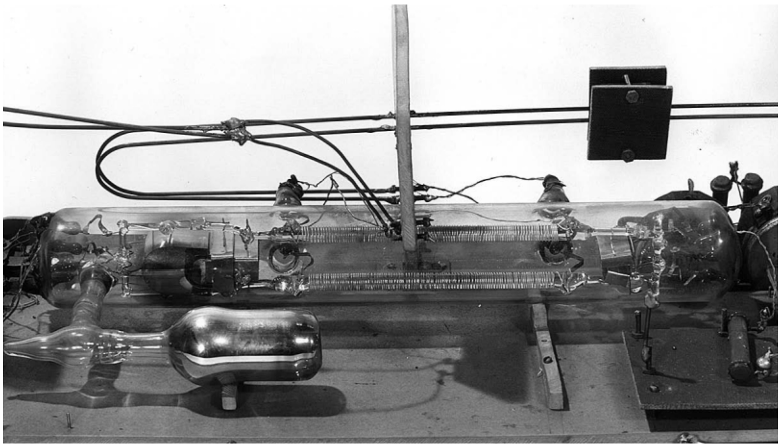

CALIFORNIA DREAMIN': A young Andy Haeff poses in Caltech's Kellogg Radiation Laboratory in 1932 [top]. It was here that he constructed a primitive form of traveling-wave tube [bottom], which he developed for working with high-frequency radio signals.

THE ANDREI V. HAEFF PAPERS (2)

links between Manchester in England and a broadcasting station midway between Glasgow and Edinburgh, allowing the company to distribute television programming to Scotland without stringing lengthy cables. But it was the advent of communications satellites that allowed the low-weight, power-efficient traveling-wave tube to really shine. Nowadays everybody's favorite signals–broadband, phone, TV–can reach almost anywhere on the globe, thanks to satellite-based traveling-wave tubes.

Despite the importance of this device, the origins of the traveling-wave tube are little known. Most authorities credit an Austrian architect named Rudolf Kompfner with inventing it in England during the Second World War. Indeed, textbooks that cover the topic tend to acknowledge only Kompfner.

In fact, though, the traveling-wave tube has much earlier roots. Its originator was not Kompfner but the shy and intense Andy Haeff. And so it's worth revisiting the fundamental contributions that Haeff made to this important technology and to explore how it was that he came to be written out of this slice of technological history.

The story begins in California in 1931. Haeff had arrived in the United States only a few years before from China, where he had moved with his parents and brother and sisters in 1920, not long after the Bolshevik revolution in his native Russia.

Haeff had recently obtained a master's degree at Caltech and was doing research that he hoped would soon earn him a doctorate. His thesis subject was microwave electronics, which in that era represented a high-tech frontier. Nobody knew how to effectively amplify microwave signals, which were too high in frequency for standard vacuum tubes to handle.

The young engineer began working on a tube in which a high-frequency radio wave corkscrews around, following a helical copper electrode, while electrons projected in a beam flow parallel to the axis of the helix. Moving around the helix reduces the wave's progress along the axis of the tube enough to make it match the speed of the slower-moving electrons, enabling a strong interaction that draws energy from the beam and amplifies the wave. Haeff's prototype used two parallel helical electrodes with an electron beam running in between them. When he filed his first patent on the tube in October 1933, he used the term "traveling wave" to refer to the high-frequency wave moving around the helix. His was the first example of what later came to be called a helix traveling-wave tube.

Haeff got the idea for this novel tube after watching surfers on Santa Monica Beach and realizing that the velocities of the board and the wave had to match for a surfer to use the wave's energy effectively.

Haeff obtained his Ph.D. in 1932, but he continued working on his revolutionary new tube as a research fellow in Caltech's electrical engineering department. The few accounts of the origin of the traveling-wave tube that mention Haeff at all typically state—quite incorrectly—that he overlooked the use of the tube as an amplifier. In fact, when Haeff applied for a patent on his design, he emphasized three ways of using the tube, including as a microwave signal detector and to "amplify…extremely high radio frequencies." So there's no doubt that he recognized the ability of his tube to amplify frequencies that lay beyond the upper limits of conventional vacuum tubes.

What's more, Haeff's new tube was not just a paper design: Shortly before filing his patent application, he used a traveling-wave tube to build a portable radio transmitter and receiver operating at 750 megahertz—much higher than other radio equipment of the time.

In March 1934, Haeff left Caltech to join RCA's research and engineering department in Harrison, N.J. The head of Caltech, the physicist Robert A. Millikan, acting on behalf of the university's governing body, granted him all rights to his invention. Haeff soon sold RCA the patent rights to the traveling-wave tube, together with a working prototype, for the then tidy sum of US $12,000—approximately $200,000 in today's money. Later, in 1936, RCA applied for a second patent on Haeff's design. The company did not, however, permit Haeff to develop the invention further, wanting him to concentrate on developing miniature tubes and circuits for use in television receivers.

The traveling-wave tube was subsequently rediscovered not once but twice. The first time was in 1940, when Nils Lindenblad, an antenna specialist at RCA, filed for a patent on an improved traveling-wave tube. This was six years after Haeff sold RCA the rights to his 1933 patent. It's hard to imagine that Lindenblad was unaware of Haeff's earlier work: The same RCA patent attorney, Harry G. Grover, filed both Lindenblad's patent and RCA's 1936 patent on Haeff's traveling-wave tube.

How Lindenblad, who was not a tube designer, came to be investigating helical electrodes is unclear. Although both he and Haeff were employed by RCA at the time, their labs were about 100 kilometers (roughly 60 miles) apart, and we have no evidence that they ever met.

One strong possibility is that Fred Kroger, an associate of Haeff's from Caltech, may have steered Lindenblad to Haeff's patents. In 1933 Kroger had helped Haeff build the microwave transmitter and receiver he used to demonstrate the capabilities of traveling-wave tubes. And when Haeff moved from California to New Jersey, it was Kroger who transported the prototype traveling-wave tube there in the trunk of his car. He also had a hand in the lucrative sale of Haeff's patent rights to RCA. Kroger, like Lindenblad, worked at RCA Communications in Rocky Point, N.Y., and they were both key players in RCA's early forays into television broadcasting. Indeed, it was for this purpose that Lindenblad revived the traveling-wave tube.

Within a few years the traveling-wave tube was discovered a third time—at the University of Birmingham, in England. In 1940, physicists at a special laboratory set up by the British admiralty had developed a powerful tube for generating microwaves: the multicavity magnetron. The magnetron is said to have had more influence on the course of the Second World War than any other single invention, because it was key to constructing small, powerful, accurate radar sets, ones compact and light enough to be fitted into airplanes.

When Kompfner joined this laboratory in 1941, he worked by day on magnetrons and klystrons—another kind of tube for amplifying microwaves—and in the evenings on various designs of his own, including a traveling-wave amplifier. Kompfner built his first traveling-wave

HAEFF'S OTHER TUBES

This prolific engineer kept inventing for decades

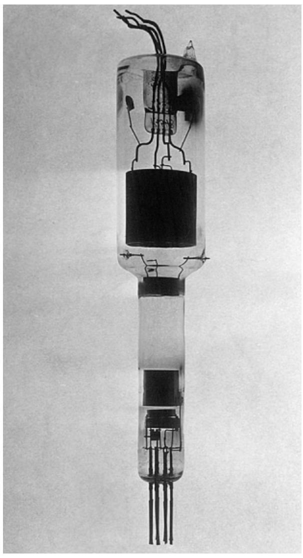

In addition to his traveling-wave tube, and his extensive contributions to military radar and electronic warfare, Andrei Haeff invented an array of important vacuum tubes. These included two ingenious amplifiers, his electron-wave and resistive-wall tubes. Another of his breakthroughs, also an amplifier, was the inductive-output tube, or IOT [a prototype is shown at left], an invention that made RCA's pioneering 1939 TV broadcasts possible, with a transmitter atop the Empire State Building in New York City and a series of relay stations spanning the length of Long Island, N.Y.

Haeff's contributions to computing included his bistable storage technology for computer monitors and his very early work on the display of graphics and text. He was the first to discover a way for the content of the screen to be retained so that it displayed for more than a second or two.

Haeff also anticipated aspects of virtual reality, as long ago as 1964 inventing a laser device for scanning works of art, buildings, museum treasures, and other 3-D objects so that technicians could re-create them virtually. Yet the traveling-wave tube remains the greatest of all his inventions. *–J.C. & A.A.H.*

THE ANDREI V. HAEFF PAPERS

tubes in the Birmingham lab, winding the long copper helices himself on a lathe. The results were good, and the British Admiralty lodged a patent application on his behalf in June 1944.

Kompfner described the form of traveling-wave tube that became the prototype for future commercial versions in a 1946 article in *Wireless World*. His tube, though, did not differ essentially from Lindenblad's less well-known design, for which a U.S. patent had in fact already been granted in October 1942–ironically, about the same time Kompfner was rediscovering the traveling-wave tube.

While the tube Kompfner described in *Wireless World* made use of the same basic principles as Haeff's, it was a distinct improvement on Haeff's design, where the electron beam ran close to the circumference of the helix. Like Lindenblad in his earlier design, Kompfner used a precision electron gun, unavailable at the time of Haeff's pioneering work, to direct the beam along the central axis right inside the helix. That method resulted in much higher amplification because it brought different physical principles into play–velocity modulation and electron bunching.

Those phenomena were discovered and investigated during the second half of the 1930s–initially by the husband-and-wife team of Oskar Heil and Agnes Arsenjewa-Heil–making Lindenblad's and Kompfner's designs possible. And indeed, Haeff was one of the first to harness these principles, in his 1939 inductive-output tube [see "Haeff's Other Tubes"].

Kompfner recounted his early work on the traveling-wave tube in a 30-page pamphlet published in 1964, *The Invention of the Traveling Wave Tube*. In it, Kompfner revealed that the crucial idea of using a helical electrode was not his own. While attempting to design a high-frequency oscilloscope, he had realized–like Haeff before him–that if an electric field could be slowed down from the speed of light to the speed of an electron beam, the field and the beam could interact. But Kompfner knew of no way of achieving this reduction in velocity. In September 1942, he discussed his high-frequency oscilloscope with other tube experts at Birmingham, and "the suggestion was made to use a helix," Kompfner stated.

Kompfner first experimented with traveling-wave tubes that had the electron beam running outside the helix, as Haeff had done a decade earlier, the so-called deflection type of traveling-wave tube. It was not until April of 1943 that Kompfner landed on the same idea as Lindenblad–of using a precision gun to shoot a stream of electrons down the center of the helix. In November of that year, Kompfner demonstrated that this approach boosted amplification.

Kompfner wrote a great deal about the circumstances of his "discovery" of the traveling-wave tube, most notably in his 1964 pamphlet. But in these extensive writings he never so much as mentioned Haeff's earlier work–yet the U.S. patent on Kompfner's traveling-wave tube design, granted in 1953, did cite Haeff's prior patent.

BELL LABS TEAM: John Pierce [left] and Rudolf Kompfner—to whom invention of the traveling-wave tube is often exclusively credited—worked together at Bell Labs, starting in 1951.

Although he was silent about Haeff's early traveling-wave tube, Kompfner emphasized the importance of Haeff's work on velocity modulation and electron bunching. Remarking in 1949 that the early attempts of Heil and Arsenjewa-Heil were "not really successful," Kompfner continued: "Then, in the space of a few months (1938-39), there appeared a series of papers by American workers, notably Haeff, [William] Hahn and [George] Metcalf, [Russell] Varian and [Sigurd] Varian, [William Webster] Hansen and [David] Webster, which transformed the situation."

Haeff was by nature modest and reluctant to blow his own horn. Although far from pleased that his early contributions were being overlooked, he didn't object publicly. But in a summary of his life's

BEAMS AND HELICES: A drawing from the U.S. patent Haeff filed in 1933 shows how the beamed electrons in his device [dashed lines] traveled outside the helical electrode [left, from "Device for and Method of Controlling High Frequency Currents," U.S. patent number 2,064,469]. A drawing from the U.K. patent that Rudolf Kompfner and the British Admiralty filed in 1944 shows a more advantageous configuration, with a narrow electron beam directed down the center of the helical electrode [right, from "Improvements in or relating to Electron Discharge Devices," U.K. patent number 623,537].

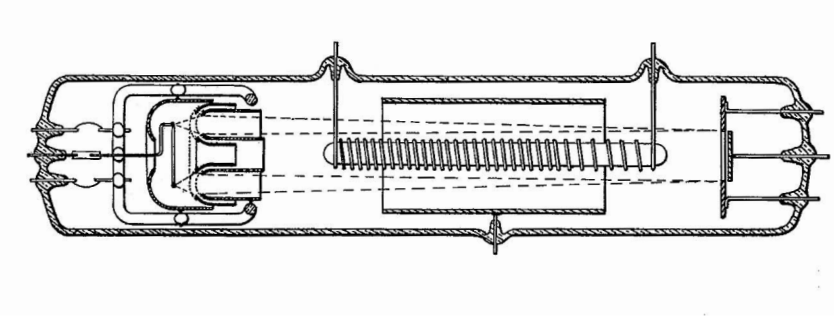

TOP: YALE JOEL/THE LIFE PICTURE COLLECTION/GETTY IMAGES work, composed in about 1950, he wrote: "It is of interest to note that the operation of the recently announced 'revolutionary' device known as the 'Travelling Wave Tube,' is based on the principle first disclosed by me in my patents No. 2,064,469 and 2,233,126 which describe the fundamental ideas of the travelling wave tube." But he did little beyond that to associate himself with this remarkable device, which was to become immensely important after John R. Pierce brought the concept of the traveling-wave tube to Bell Labs, archrival of RCA, soon after meeting Kompfner in Britain in 1944. In 1951, Kompfner himself joined Bell Labs as Pierce's protégé.

Pierce, an engineer who is perhaps most famous now for coining the word *transistor* in 1948, rose to become an executive director at Bell Labs. And his relationship with Haeff was cordial enough for him to have dined at the Haeff home on one occasion in the early 1950s. Nevertheless, Pierce relentlessly pressed Kompfner's–and, indirectly, Bell Labs'–claim to priority. He wrote in 1956: "About the origin of the traveling-wave tube, I will say only that it was invented in England during the war by an Austrian architect, Rudolf Kompfner, who had always wanted to be a physicist."

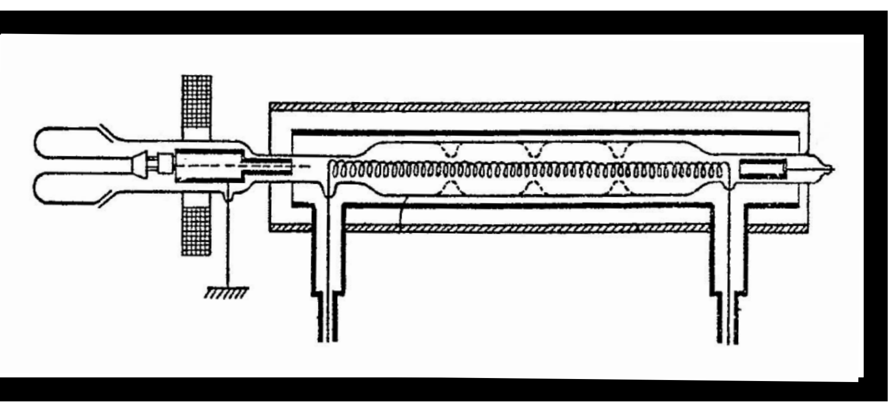

BOTTOM, LEFT: U.S. PATENT AND TRADEMARK OFFICE; RIGHT: UK PATENT OFFICE

Pierce's definitive 1950 textbook *Traveling-Wave Tubes* did not so much as mention Haeff's contribution. In Europe, on the other hand, Haeff's early work on the traveling-wave tube featured in a French textbook, published in 1951. Ten years later, Pierce told Haeff in a letter of his plans to write a historical article about microwave-tube development that would refer to Haeff's "use of traveling waves with electron beams back around 1935." But Pierce, who died in 2002, never did author an article setting the historical record straight. And only a few specialists realized the importance of Haeff's early work, like Victor Granatstein at the University of Maryland, who wrote in 2000 that Haeff "made an indispensible contribution to the development of the helix TWT" (traveling-wave tube), noting that the "key" idea of using a helix to slow a high-frequency wave to the same speed as an electron beam was "revealed for the first time in Andrew Haeff's patent."

In 1950 Haeff, who had left RCA in 1941 to work on radar at the U.S. Naval Research Laboratory, joined the research and development laboratories of the rapidly expanding Hughes Aircraft Co. in Culver City, Calif. There he set up and directed the Hughes Electron Tube Laboratory, eventually becoming a vice president of Hughes and director of research. Haeff's lab quickly began developing advanced versions of the traveling-wave tube.

The launch of the Russian Sputnik satellite in October 1957 gave him even more impetus for doing so, as communication via satellite became the new frontier for radio and TV engineers. Researchers at Bell Labs proposed launching an orbiting relay station and using a network of ground-based stations to track it continuously as it moved across the sky. In a 1945 issue of *Wireless World*, the science fiction writer Arthur C. Clarke had described his idea of placing communications satellites in stationary orbits, which would avoid the tracking problems associated with satellites that move relative to Earth's surface. When Hughes engineer Harold Rosen suggested that the company develop such a stationary satellite, many of his colleagues were skeptical, but Haeff convinced the general manager, Lawrence "Pat" Hyland, that a stationary satellite was preferable to the Bell Labs scheme.

Haeff established a special task force to study the problems of commercial satellite communication. In October 1959 the task force recommended that Hughes proceed at lightning pace with a "major space program...under tight security." The aim was to place a simple broadband repeater satellite into an approximately stationary orbit above the Atlantic.

The task force further advocated developing a special lightweight, high-efficiency traveling-wave tube, weighing around 1 pound–about one-twentieth the weight of Hughes's standard version. This tube would form "the heart of the proposed satellite electronic system," amplifying the signals the satellite picked up before relaying them back to Earth at a different frequency.

Hughes's Syncom became the world's first geosynchronous communications satellite, successfully reaching orbit in 1963. This launch was just a NASA-sponsored test, but it soon gave rise to the first commercial communications satellite, placed in orbit in 1965, the Syncom-based Early Bird, also known as Intelsat I. Early Bird's success ultimately allowed Hughes to dominate the manufacture of communications satellites for many years to come.

It must have been an occasion for pride–and some mixed feelings–for Haeff when the device he pioneered in his twenties, one rarely associated with his own name, successfully carried a multitude of human voices across space. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/travelingwavetube0915

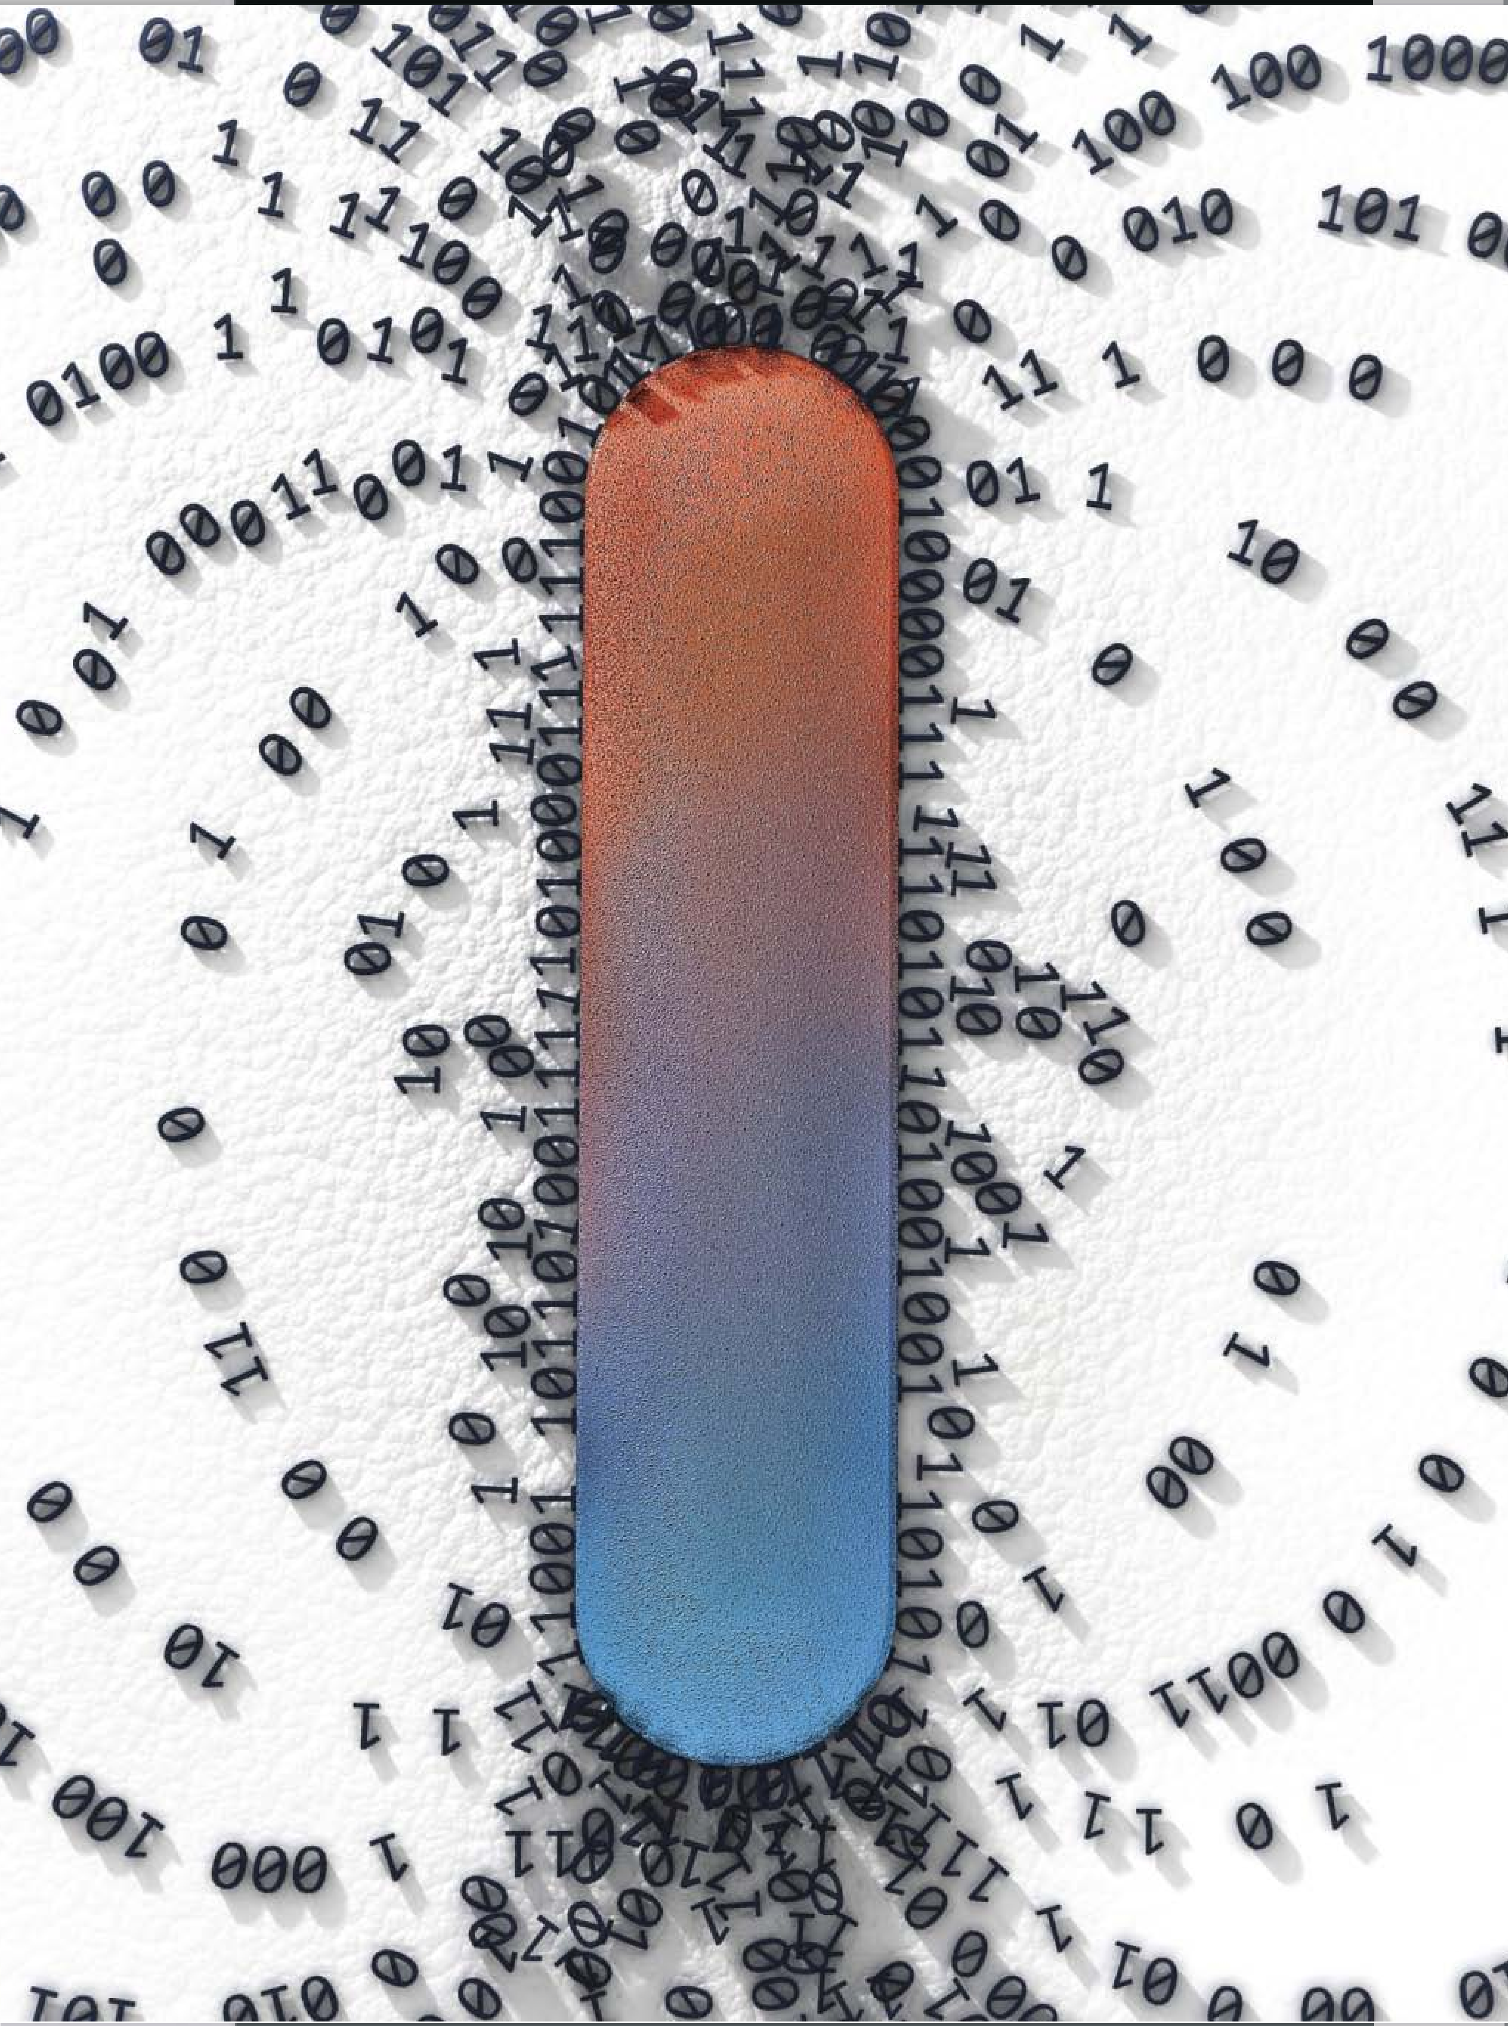

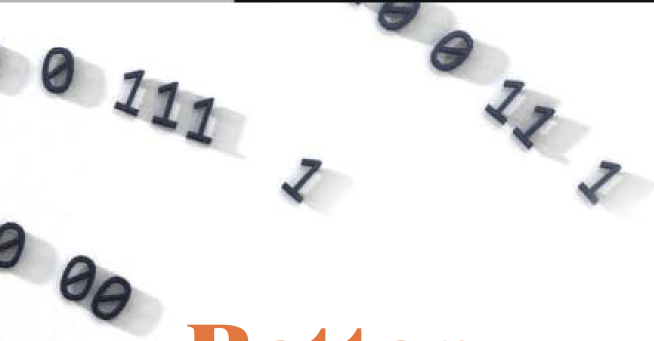

Better Computing With Magnets

The simple bar magnet, shrunk down to the nanoscale, could be a powerful logic device

By Wolfgang Porod & Michael Niemier

When it comes to computing, we can very easily divide the technologies we use into two distinct categories: speedy electronics and stable magnetics. Electrons, which move fast and interact strongly with one another, are ideal for performing computation. Magnets, by contrast, aren't known for their speed, but they're hard to perturb, making them the perfect medium for data storage.

But this division may soon disappear. Thanks to modern fabrication technology, we can now create nanometer-scale magnetic devices that can perform computations. These devices are not as speedy as state-of-the-art transistors, but they require far less energy to switch.

We'll need devices like these because modern chips are consuming too much power. Today, the amount of electricity needed simply to maintain data in a circuit—called standby power—is fast approaching the amount that's consumed when an actual computation is performed. Magnet-based devices, which require no power to save their state, could drastically cut down on this constant power drain, which is one of the main obstacles to continued progress in chips.

There are many ways to make logic devices using magnets. Our group at the University of Notre Dame and several others are working on one of the most straightforward approaches: building logic gates and wires out of small patches of magnetic material. These "nanomagnets" act just like tiny bar magnets. In circuits made with them, information isn't carried from one place to another electronically. Instead it's transmitted directly through the magnetic attractions and repulsions, flipping the polarity of north and south poles as it moves from one magnet to the next, in much the same way that a NOT gate reverses the logic state of a bit.

We have already demonstrated that we can build simple circuits, such as adders, with these nanomagnets. Now, with a new fabrication technique, we're starting to contemplate what can be done to build fully integrated logic chips. Though relatively modest, these achievements make it clear that this technology could someday be used to make ultralow-power chips. In some cases we expect that nanomagnet circuits could use a tenth or maybe even a hundredth of the power with no sacrifice in performance. Such capabilities could be ideal for sensors

Illustration by Emily Cooper

and display electronics, hardware accelerators on multicore chips, and computationally demanding applications such as machine-vision systems in autonomous vehicles.

THE IDEA OF COMPUTING with magnetism isn't new. Some early computers actually contained iron-based cores, which were used both to store information and perform computations. But that magnetic logic technology, which was based on stringing wires between ring-shaped magnets, was too bulky to compete with increasingly compact semiconductor technology.

Today's magnetic logic is fabricated at much smaller scales, and the semiconductor industry, which funds university-based research through the Semiconductor Research Corp., has been supporting a variety of approaches.

Many of those efforts are part of a field called spintronics, in which researchers exploit the spin of electrons, rather than their charge, to transmit and manipulate information. Spin is a quantum-mechanical property associated with magnetism. One such approach, pursued by a group led by Kang Wang at the University of California, Los Angeles, transmits information in the form of waves of spins propagating through a magnetic layer of material (see "The Computer Chip That Never Forgets," *IEEE Spectrum*, July 2015). Another scheme, being developed by Supriyo Datta and colleagues at Purdue University, uses similar pulse-carrying wires to connect magnetic dots.

Our strategy, which we expect to be more energy efficient than either of the above, is based on a simple premise: Make very small magnetic regions and allow them to interact with one another just as bar magnets would. These interactions can be harnessed to make wires that transmit information and to construct logic gates that can perform computations.

The key to this approach is miniaturization. Take a look at an ordinary piece of magnetic material under the microscope and you'll find that it is naturally divided into a number of tiny patches. Within each of these patches, called domains, the spins of the electrons all point in the same direction. That alignment is what causes magnetization. However, a bulk sample of the material will not be magnetic if all of these domains are aligned at random—that is, if the north poles of the domains point in different directions. A patchwork of randomly oriented domains helps minimize the amount of energy that's locked up in the magnetization of the material as a whole. But a lot of energy is required to pull all the domains into alignment to create a permanent magnet, and even more energy is needed to change the orientation of the poles of that magnet.

Make a magnet smaller than the material's domain size—typically around 100 nanometers across—and you can make a tiny version of a permanent magnet. As with a bar magnet or a compass needle, the magnetic field of a nanomagnet naturally prefers to align itself on the longer axis, with a north pole on one end and a south pole on the other. But by fine-tuning the aspect ratio—the ratio of height to width—we can make it so that it requires relatively little energy to flip the direction of the nanomagnet's north and south poles. In fact, if you make a nanomagnet too round—or too small, for that matter—little kicks from the thermal energy in the material can cause that flip to happen on its own at random.

Today we use similar tiny, flippable nanomagnets to store information in ordinary hard disk drives. They can also be found in magnetic RAM technologies now under development. In both of these cases, we don't want the magnetic bits to interact with one another; such interactions would risk corrupting the data that you're trying to store. But when building nanomagnetic logic, those interactions are exactly what you need in order to connect the devices and make them perform computations.

OUR WORK ON NANOMAGNET LOGIC grew out of some previous research we'd done in the 1990s on quantum dots. Our original idea was to try to make computers by using arrays of these dots, which are patches of semiconductor in which

Of Magnets and Wires

The magnetic field of an elliptical nanomagnet naturally prefers to align along the longer, or "easy," axis [top illustrations, near right]. Wires can be constructed in two ways. One is by arranging the magnets side by side [center], where their magnetizations prefer to alternate in direction. The other is end to end [far right], where all magnetizations prefer to point in the same direction. All arrows point north.

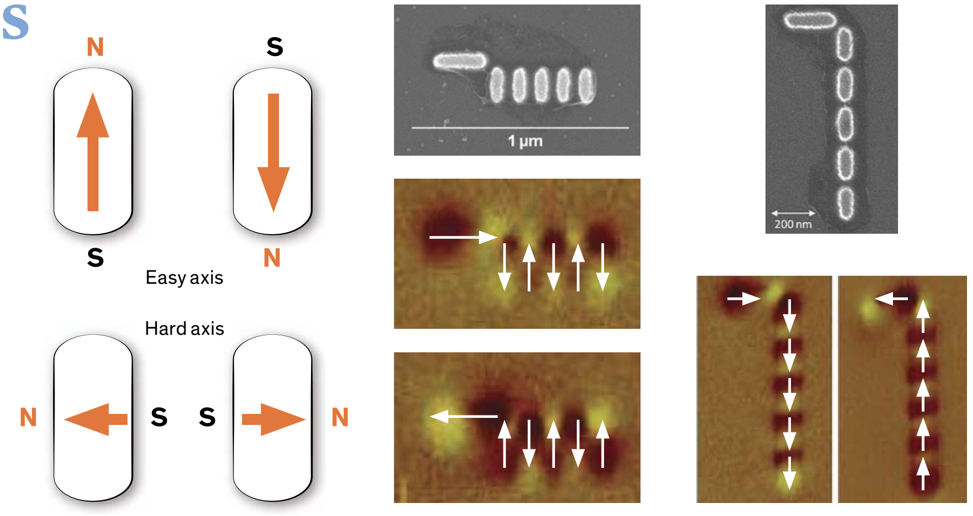

EDIT VARGA/UNIVERSITY OF NOTRE DAME (5)

electrons are confined to a space so small that they exhibit the same quantum behavior that they would in an atom. Rather than using wires, neighboring dots would transmit information by the Coulomb forces, which attract opposite charges and repel like charges. We showed that these physical interactions could be used in an appropriately structured array to perform logic operations, but we ran into technical limitations, in part because it was hard to control size variations when constructing quantum dots.

We realized early on that using magnetic dots–which are much more stable, easier to fabricate, and capable of operating at room temperature–might be a good alternative. And in the early 2000s, following some initial work done by Russell Cowburn, now at the University of Cambridge, in England, we began simulating and then experimenting with the devices.

We started out by making magnetic dots from a mix of nickel and iron. We gave them an elliptical shape, stretching them out in one direction to give each magnet a preferred north-south axis. But we limited the elongation, thus minimizing the energy required to flip the orientation to represent the 0s and 1s in binary logic.

Like a bar magnet, each nanomagnet has "fringing fields"–magnetic fields that extend a long distance from the magnet. If we place the nanomagnets close enough, we can use the fields from one magnet to influence the state of the ones near it, and so build up logic circuits that can pass magnetic information from one magnet to the next.

To get a sense of how such nanomagnetic logic would work, it's helpful to remember the most basic rules of magnetism: Like poles repel and unlike poles attract. So, just as bar magnets would, two nanomagnets placed end to end will naturally prefer to have their north poles point in the same direction, so the north pole of one is closest to the south pole of the other. If two nanomagnets are placed side by side, so their long axes are parallel, their north poles will point in opposite directions, so that each pole is closest to one with the opposite polarity. This second configuration, which is called antiferromagnetic coupling, can actually be considered an elementary circuit element: the inverter. One nanomagnet would be the input and another the output, flipping the state of the information.

The inverter is not a bad start, but to perform useful logic operations you also need AND and OR gates. And in 2006, we reported a proof-of-concept demonstration of a second logic element called a three-input majority-logic gate, which can be used to build both of those other gates. The majority gate takes on the state of the majority of its inputs. One way to build it is with a cross-shaped arrangement of five nanomagnets–one center dot surrounded by four others. Three of those dots act as inputs, and the center dot "calculates" the majority by naturally aligning its magnetization with the majority of the other spins. The fifth dot carries the result out of the device.

As for the wires that you'd need to connect logic gates, these can be constructed by simply lining up nanomagnets either end to end or side by side. Information can travel in either of these configurations, flipping spins as it goes–a little like falling dominoes. (In the side-by-side configuration, of course, you need to make sure you have an odd number of nanomagnets, so the information retains the same state at the end of the wire that it did at the beginning.)

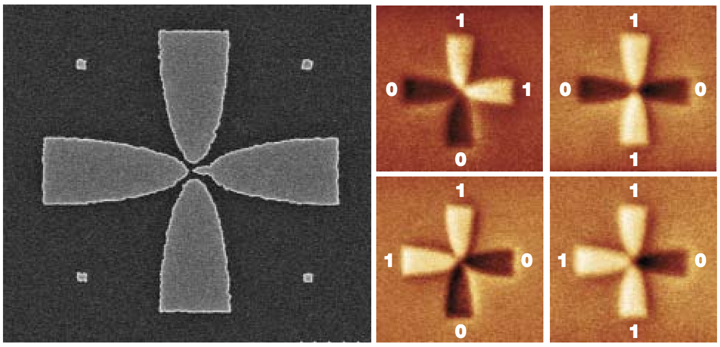

VOTING WITH MAGNETS: Four or five nanomagnets can be arranged to form a key circuit component: the majority logic gate. This gate outputs the state that the majority of the inputs are in—or in this case, the inverse of that state. Three 1s, for example, will yield a 0. Note that this gate is built from nanomagnets with an out-of-plane magnetization—more on this design in the final section of the article.

The fundamentals are quite straightforward. But you might have noticed that there's a catch. Unlike a line of dominoes, which all fall in the same direction after the first domino is tipped, there is nothing intrinsic in a line of nanomagnets that determines which way the information will flow. A nanomagnet in the middle of a line of nanomagnets will be equally affected by the nanomagnet to its left and the one to its right. If we can't reliably control the direction of computation, our computer will have a high error rate. To make up for this intrinsic symmetry in computational direction, we need to add some clocking circuitry–more on that in a moment.

NANOMAGNETS HAVE A FEW NICE PROPERTIES. They're inherently insensitive to radiation, they can switch pretty much indefinitely without degradation, and they're nonvolatile, requiring no energy to retain data when they're not switching.

At the same time, they're very slow by modern transistor standards, maxing out at about one-hundredth the speed of a traditional transistor. That means nanomagnet logic will likely never reach gigahertz speeds. But the potential energy savings still make it an attractive alternative for the many applications that don't require such speeds.

Much of the energy advantage comes in at the circuit level. Because of the way nanomagnets interact to perform logic operations, it can take as few as five magnets to add two 1-bit numbers together. For comparison, it can take 20 to 30 transistors to construct a similar adder in silicon.

STEPHAN BREITKREUTZ-V. GAMM/TECHNICAL UNIVERSITY OF MUNICH (5)

Moving Information

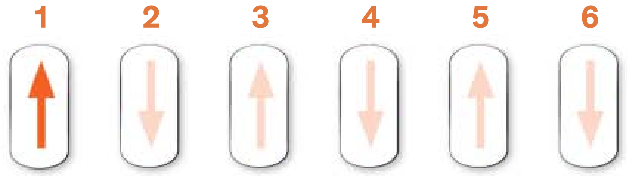

To start a chain of spin flips, first use a strong magnetic field to flip magnet 1.

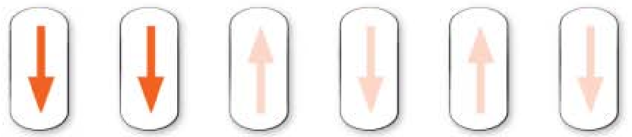

Magnet 2, originally in a down configuration, must flip up to propagate the information to the right. The magnetic field of magnet 1 isn't strong enough to do this on its own. So, the magnetic field from a wire, or clock, is first used to draw magnet 2's magnetization into the metastable "hard" axis.

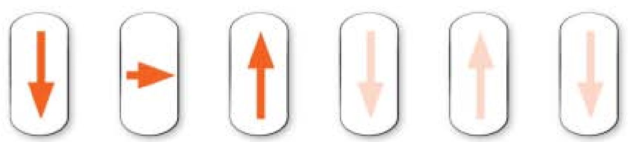

Magnet 2 can now relax into a new stable state. But it's not clear which way it will fall; the field from magnet 1 would have it point up, and the field from magnet 3 would have it return to a down state. To make sure that magnet 2 feels a stronger influence from magnet 1, magnet 3 can be clocked as well.

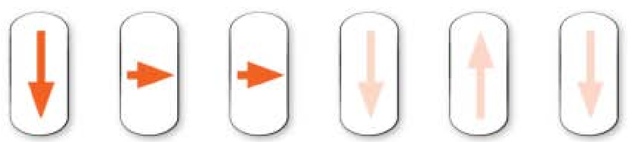

When the clock signal is now removed from magnet 2, it will relax into an up state, advancing information to the right. Magnet 4 can now be clocked to keep the flow of information going.

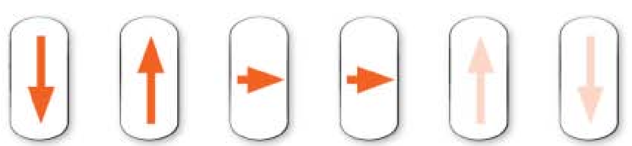

Controlling every nanomagnet would require clocking circuitry for each device, which would be a miniaturization challenge. In practice, a line of nanomagnets might be clocked in groups. A wire, for example, might control half a dozen nanomagnets at once. The result will be a little more error prone but easier to manufacture.

In 2011, we demonstrated that we could combine majority logic gates, inverters, and nanomagnet-based wires to create the first full nanomagnet-based circuit: a functional 1-bit full adder. And thanks in part to work done in collaboration with IBM and with funding from the U.S. Defense Advanced Research Projects Agency, we're confident that we can use magnetic RAM technology to connect nanomagnet circuits to the outside world. That's because these nanomagnets, other than being shaped differently, aren't much different from the magnetic bits used on such memory chips, which have already found their way into commercial production.

But even early on, we realized that any nanomagnet computer we build will be only as good as its clock. As we mentioned before, we need clocking circuitry to keep computations moving in the right direction. We also need a clock to get the nanomagnets to switch reliably. Because magnets are so stable—the very reason they're commonly used for data storage—switching them is often the tricky part. On its own, the fringing field surrounding one dot isn't strong enough to reliably induce a 180-degree switch in a neighboring magnet.

A clock signal can help a nanomagnet switch. It could be created with something as simple as a nearby wire, which would generate an extra magnetic field near a nanomagnet when it's carrying a current. To see how such a clock would work, imagine an elliptical nanomagnet. As noted before, its natural state is to have its magnetization along its longer axis. If we use the added wire to apply an additional magnetic field, called a switching field, we can rotate the nanomagnet's magnetization by 90 degrees into what's called the "hard" axis, the shorter axis of the two. This is an unstable state for the nanomagnet, and when the switching field is removed, the magnetization will begin to snap back into either one of two directions along its longer, "easy" axis. When it starts to do so, the fringing fields from the neighboring magnets will determine which direction it falls into.

To make sure that the nanomagnet is influenced only by the correct neighbor when it makes the transition, the clocking circuitry can also be used to pull nearby nanomagnets that are not supposed to influence the magnet—the ones that are downstream of the flow of information—into the same 90-degree state. In that orientation, they can't influence which direction a neighboring nanomagnet will switch to.

In 2012, we built such a clock system out of copper wires clad with ferromagnetic material on the sides and bottom to help concentrate the magnetic field. We demonstrated that nanomagnets can indeed be switched using the magnetic fields created when current was run down those wires.

The downside of this scheme is that wires are a relatively energy-intensive approach to clocking. They dissipate quite a bit of heat and emit unfocused magnetic fields, even with cladding. But it turns out that a single clock line could control many parallel ensembles. If a nanomagnetic circuit has enough devices—say 100,000

ERIK VRIELINK

or so—the clock energy needed per device will be acceptably low. As always, though, there will be a trade-off. The more devices you have that are governed by the same clock, the less control you'll have over each individual device, which will translate to a higher error rate.

OUR INITIAL WORK at Notre Dame has been based on dots patterned in a thin film of magnetic material. Each nanomagnet is a two-dimensional ellipsoid, in the plane of the chip. But we encountered a tricky circuit-design limitation with this approach. To take full advantage of the space on a chip, you must find a way to build circuits that take up the full two-dimensional surface and so can pass information in both the *x* and *y* directions.

The most straightforward way to do this is to make all the magnets on the chip the same, so they all have their magnetization pointing along the same axis. For the sake of argument, let's pick the *y*-axis. All the nanomagnets will be elongated in that direction. But the magnets that pass information along the *y*-axis will have to be strung end to end, while the ones that pass information in the *x* direction will have to be arranged side by side.

This dual arrangement makes circuit design quite complicated. The switching behavior of these two configurations of nanomagnets is different. As a result, signals will travel at different speeds depending on whether they're traveling in the *x* or the *y* direction, making synchronization of signals much more complicated than it is in traditional silicon circuits.

To get around this limitation, we began collaborating in 2009 with groups led by Doris Schmitt-Landsiedel and Paolo Lugli at the Technical University of Munich. Schmitt-Landsiedel's team has pioneered the construction of nanomagnet logic devices that have all their easy axes aligned out of the plane of the chip. Instead of building the magnets from small patches of material, the team constructs them out of multiple, alternating layers of elements such as cobalt and platinum. The interface between these layers is magnetized vertically, creating what is essentially a bar magnet standing up on one end, perpendicular to the plane of the chip.

The process begins by depositing thin films and then patterning them with a focused-ion-beam instrument. The ion beam destroys the clean interface between materials where it hits, so it can be used to pattern the smooth film into individual islands small enough to exhibit single-domain behavior. The Munich groups recently built the same sorts of logic gates and adders that our group demonstrated with in-plane nanomagnets.

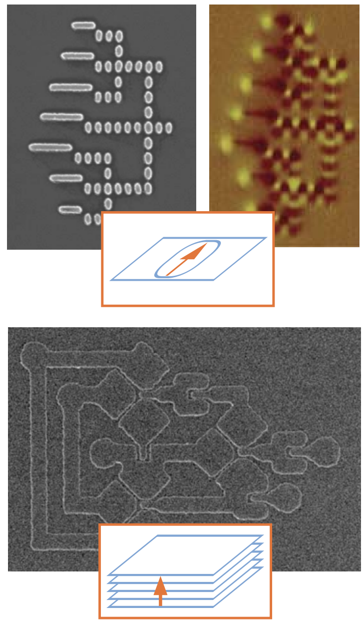

TALE OF TWO ADDERS: Nanomagnet circuits, such as the adders shown here, can be made in two ways. One method [top] builds nanomagnets in two dimensions, with their magnetization in the plane of the substrate. Another approach [bottom], which might prove easier to implement in large circuit designs, builds nanomagnets from many thinly deposited layers. In this case, the magnetization points out of the plane of the substrate.

TOP: EDIT VARGA/UNIVERSITY OF NOTRE DAME/IEEE (2); BOTTOM: STEPHAN BREITKREUTZ-V. GAMM/TECHNICAL UNIVERSITY OF MUNICH

Encouragingly, the out-of-plane design offers the opportunity for new ways of clocking. In late 2013, a team led by Sayeef Salahuddin at the University of California, Berkeley, showed that out-of-plane nanomagnets could be clocked using a sheet of magnetic material placed beneath them. Owing to a principle called the spin Hall effect, electrons with a particular spin can be made to concentrate underneath a nanomagnet, creating a magnetic field that can alter its orientation. At the circuit level, this approach could be hundreds or even thousands of times more energy efficient than clocking by using current-carrying wires.

And that's very good news indeed for the prospect of a highly efficient computer. One application where we expect nanomagnets to be particularly useful is in data-intensive, high-throughput applications, such as filtering, polynomial evaluation, and discrete Fourier transforms. These sorts of computations, which constitute the backbone of image and signal processing, can be greatly speeded up by being executed in a pipeline. This is a strategy where computations run in parallel as much as possible, and pieces of data that have already been computed are saved until the entire computation is ready to move on to the next step. In ordinary computers, this approach requires adding extra circuits to actively hold onto that data, adding to the energy consumed by the chip. But nanomagnets, which naturally retain states until they are changed, would passively hold the data until it's ready to be used.

When will we see this technology in chips? Those of us pursuing novel computing devices are in an awkward position now. Any such technology has to compete with silicon systems that have been endlessly optimized for decades. The success of nanomagnet logic—or any future logic, for that matter—will depend on all kinds of factors, not just technological ones.

At some point though, it's quite likely that the need for more-energy-efficient circuitry will trump the convenience of silicon. When that happens, magnetic logic will finally have its moment. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/nanomagnets0915

Phone to Fridge: SHUT UP!

OUR HYPERCONNECTED WORLD NEEDS BETTER PROTECTION AGAINST ELECTRONIC NOISE

By Mark A. McHenry, Dennis Roberson & Robert J. Matheson

WHEN ONE OF US (Roberson) used to live in Wheaton, Ill., his car's FM radio would blare static every time he drove near a pole-mounted electrical transformer. Now, when he's near a particular intersection in Chicago and an elevated train passes by, his mobile phone call gets dropped. The same thing happened to him in a rapid transit station in Washington, D.C., during a conference call with the other two authors of this article. One of them (Matheson) has had to train himself to wait until the commercials begin before turning on his electric toothbrush, because it always breaks up the picture and sound of the TV set in his bedroom.

Radio-frequency noise pollution is everywhere. You can't see, hear, taste, or smell this noise, of course. Nor can you summon it and study it at your leisure, because it comes and goes along with the movements of its sources or its victims. Start with the fact that any significant digital appliance has a high-speed clock and a digital bus, and both leak radiation profusely. Electric motors and generators generate RF noise with every small spark that jumps between their brushes and spinning commutators. Automobile engines sputter when spark plugs fire. Computers snap and pop during the sharp transitions between ones and zeroes. The high-voltage ballasts of neon signs and fluorescent lights blare a broad mix of frequencies. Industrial machinery, elevators, welders, relays, switching power supplies, even light dimmer switches add to the din. (Of course, natural sources of noise abound as well, including lightning and solar flares, but we will not deal with those problems here.)

The problems of RF pollution fall into four categories. First, it increases the cost of deploying wireless systems while it reduces the battery life of handsets. Second, it creates various levels of interference across a range of frequencies. Third, interference does not–but should–figure in policies on how best to share spectrum, given that the more interference you have, the more spectrum you'll need in order to transfer a given amount of infor-

ILLUSTRATION BY Lou Brooks

mation. So in practice, wireless channels do not always achieve the data rates they were designed to achieve. And fourth, it is expensive to trace RF pollution to a source and, when you do, it is often challenging to get offenders to stop offending.

The coming Internet of Things is going to make things worse. Much worse. It will do so by adding complex RF-control chips to countless common devices, like door locks, light switches, appliances of every type, our cars, and maybe even our bodies, which will enable them to connect to the Internet. Each of these chips is a potential source of noise. Plenty of technological fixes are available, of course, but the huge number of chips means that manufacturers will be more reluctant to add costly shielding and other noise-muffling features to their products. Silence is golden: It costs money to get it.

The time to start a public dialogue on this question is now. And, having studied the problem for decades, we're just the guys to do it.

BELIEVE IT OR NOT, there hasn't been a systematic study of radio-frequency noise in the United States since the mid-1970s, when the Institute for Telecommunication Sciences (ITS), a part of the National Telecommunications and Information Administration, last monitored federal use of the radio spectrum. By then, man-made electronic noise had been a problem for nearly half a century, having started with the inauguration of commercial radio transmissions. In 1934, the International Special Committee on Radio Interference was founded in Paris, and that same year a committee with the same mandate was established by the Institute of Radio Engineers, a forerunner of the IEEE. The noise problem grew with the use of electricity and the number of wireless systems, although many details have changed over time.

Only in 1927, with the invention of a method of suppressing RF noise, was it possible to listen to a car radio while the engine was running. By the end of the 1930s, 20 percent of all cars in the United States had built-in radios—which means that 80 percent were presumably not noise suppressed. When the ITS made extensive measurements of RF traffic noise in the mid-1960s to early 1970s, it discovered that many—perhaps most—cars had little noise suppression. Some made so much electrical noise that the men taking the measurements could hear the static from a couple of blocks away.

Today every car has a radio, and the noisy spark-plug problem is gone. The quieting process has continued: In the 1960s, quiet alternators began replacing noisy generators, and electronic ignition started replacing noisy distributors. In the meantime, electronic switches were replacing noisy relays. Electric-drive cars, which would otherwise produce much more interference than standard cars, are sufficiently noise suppressed for cellphones and car radios to work. That same improvement came to the workplace and the home. Early lamp dimmer switches were often electronically noisy; today most are much quieter. Personal computers, too, were redesigned for silence, if only for the sake of their own internal and external wireless data connections.

Nevertheless, the RF noise problem is increasing. Although most devices pollute less than their predecessors, we have far more of those devices. Other sources, such as the power grid, are expanding as wind farms and solar households connect to it. Such devices need to switch large amounts of DC power at a 60-hertz or an even faster rate whenever they feed excess generated power back into the grid. If not done properly, this could also feed large amounts of noise into the power grid. This risk is magnified when the solar and wind systems operate without expert maintenance inside millions of ordinary homes.

At the same time, today's machines are more sensitive to noise than ever before. Many new wireless systems, including smartphones, are designed to operate with the lowest possible power while still providing their intended function. This means that just a little more noise interference can decrease the coverage area.

Unwanted transmissions can be inherent in the design of a device, such as a microwave oven, whose RF cooking energy also contains a large amount of RF noise. Imperfect shielding can allow this energy to leak out and interfere with other RF devices. Noise can also result from a partial failure in a device—for instance, a tiny break in the ground shield around an insulator in a high-voltage power transmission system. Such a mishap creates an inadvertent transmitter, broadcasting at unpredictable frequencies, in unknown locations and at unexpected times. Complex though these patterns are, we can still often classify the causes of RF pollution by looking at a graph of noise power versus time.

Take the many sources related to arcing over a failed power-line insulator. In the United States, where power frequencies are at 60 Hz, such RF noise typically repeats at a 60-Hz rate (every 16.7 milliseconds) or at a 120-Hz rate (every 8.3 ms). This is because some noise-generating arcing occurs with either a positive or a negative peak line voltage—but not both—and therefore occurs 60 times a second; other arcing processes occur equally well with positive and negative peak voltages.

ILLUSTRATION BY Lou Brooks

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

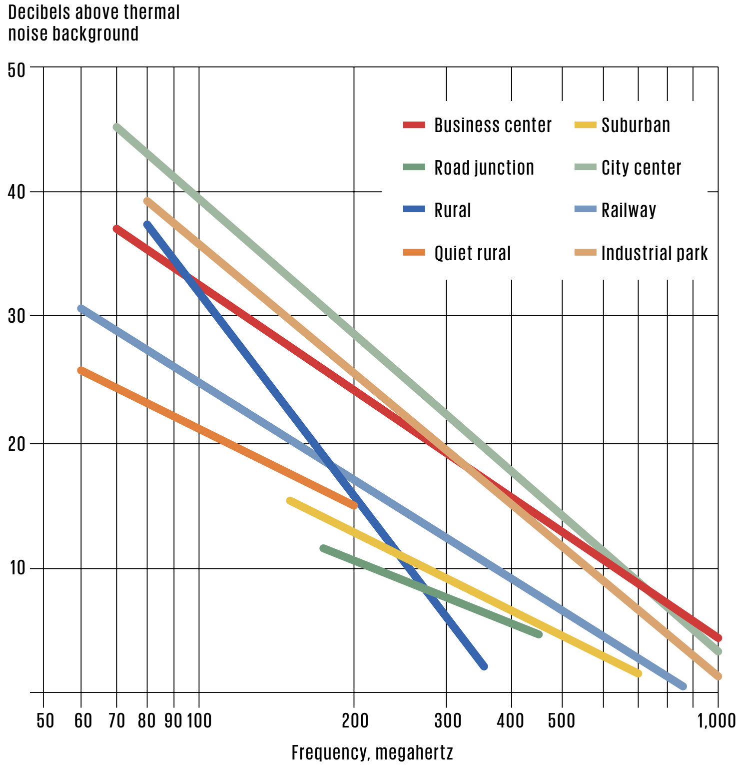

Source: Mass Consultants Limited (2003)

TOWN AND COUNTRY: Man-made noise is higher in cities than in suburbs and rural areas because cities have more electrical devices. Noise drops as frequency rises because arcing and other noisemaking processes generate less power at high frequencies and high frequencies are more readily absorbed in propagation.

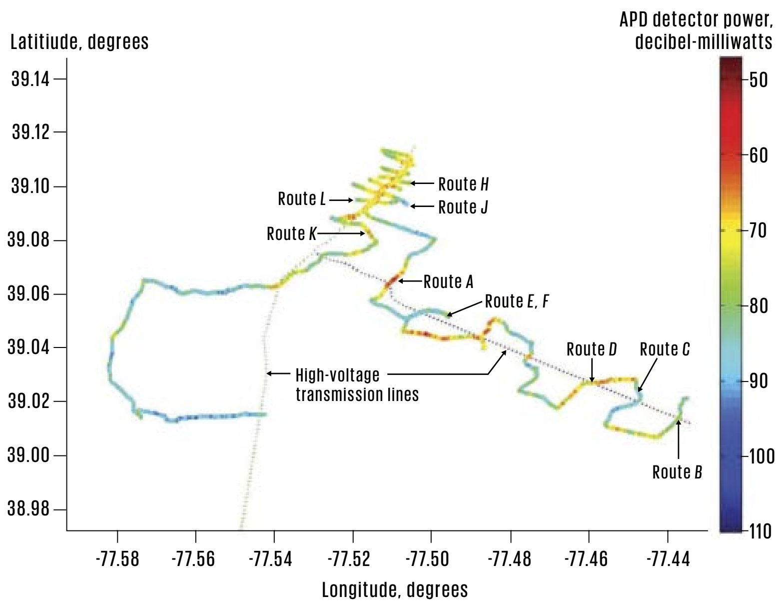

CACOPHONOUS CAR TRIP: The authors' crew drove a car with a roof-mounted antenna and an amplitude probability distribution (APD) detector, following a route that crossed high-voltage lines. Their measurements showed that the noise power usually (but not always) increased when the car was near a high-voltage transmission line.

These signatures point the electronic detective to either the power line itself or to a motor, power supply, or other appliance that is connected to it. Indeed, most outdoor man-made noise below 1 gigahertz stems in some fashion from the power grid.

Microwave ovens are another example of what's generally viewed as a noise generator with a characteristic pattern. In this case, residential microwave ovens are on for half of a 60-Hz cycle and off for the other half, emitting noise into the 2.4-gigahertz spectral band. This is an unlicensed but popular piece of radio spectrum heavily used by low-power nonlicensed devices, cordless phones, and Wi-Fi. Many other electronic appliances, including chargers for cellphones, tablets, and laptops, have switching power supplies that generate noise at the harmonics of the power supply switching frequency, in the hundreds of kilohertz.

Man-made noise levels are much higher in cities than in suburbs or out in the country because cities have a higher concentration of all forms of electrical equipment, computers and radios, and home appliances and systems, as well as industrial equipment. And in all these places the noise level, on average, is in inverse relation to the frequency, because the impulsive source mechanisms generate additional power at lower frequencies.

TO MAP THE EXTENT AND VARIATION in RF noise in a typical U.S. suburb, we outfitted a car with a roof-mounted antenna, a GPS receiver, and a small, well-shielded, custom-made spectrum measurement system. We drove the car through northern Virginia, on a route that crossed both high- and low-voltage transmission lines. We covered an area about 10 kilometers square.

We set the receiver at 10 frequencies, from 100 to 1,500 megahertz, that coincided with no known intended signal. In each case we measured the signals within a bandwidth of 0.7 MHz for a measurement duration of 55 ms. After measuring the highest frequency, we began again with the lowest frequency. To detect noise from the switching of a power supply, we looked at variations in the noise amplitude of around 100 kilohertz; to detect noise from power lines, we looked at variations in the noise amplitude at 60 or 120 Hz.

We could drive along a high-voltage transmission line for quite some time without detecting much RF noise, but the hot spots we found turned up at random. This implies that there was no systemic cause of the noise, no inherent problem with the design. A defective part—perhaps a loose nut or a cracked insulator—might be arcing. If the power companies did the maintenance, such noise wouldn't be a problem, but generally they repair such defects only if someone complains.

In our sample, the probability of detecting significant noise outdoors near a high-voltage line was a bit more than 10 percent at frequencies above 500 MHz. About 100 meters from such a line we measured surprisingly high noise levels, in some cases translating to almost 60 decibels (or a million times) greater than the normal background level of thermal noise (a theoretical minimum noise for everyday radio systems).

Our modest experiments offer just a basic sense of how RF noise is distributed in a suburban setting. Surprisingly, and unfortunately, nobody has made a complete and statistically reliable measurement of RF noise since the ITS did so more than 30 years ago. Noise can change significantly over time, locations, frequencies, and operating circumstances. For example, at frequencies above 100 MHz–where older models suggest no noise should be found–our limited measurements showed rather a lot of it.

The bottom line is, no one really knows whether outdoor noise has increased or decreased in recent years–and the same goes for measurements made indoors, where many wireless devices are found today. In fact, the question of noise levels inside buildings is particularly bothersome. The average middle-class home is now full of noise-generating electronic devices, like that electric toothbrush we mentioned. This includes laptop power supplies, power tools, LED light controls, and defective electronic equipment.

LET'S SAY THAT RF NOISE is wreaking havoc with your cellphone reception. How would you know it? That's just the problem. You probably wouldn't.

Modern emitters of radio signals, with their exotic digital modulation schemes, don't usually leave a clue as to exactly why they're malfunctioning. Older systems were much more transparent: "That sounds like a car engine interfering with my signal." And automatic error correction or retransmission schemes may completely hide the interference from the user. The only hint of its existence may be a reduction in performance or that the battery in your cellphone goes dead faster than it should. Therefore, even if we had a rich database of consumer complaints–and we don't–it might not help us get to the bottom of things.

To begin to solve the problem, we need to amass statistics on a broad scale. We have to determine where and when RF noise usually appears and at what frequencies, in each case tracing it to its source. We need samples in many regions to account for differences in equipment and practice, building codes, weather, and terrain.

We know a man who built a new house with LED lights in his kitchen ceiling that were connected to a dimmer switch. Several of his neighbors in a town near Boulder, Colo., began to notice that their garage door openers sometimes weren't working very well.

After several weeks, the man established a clear correlation between the neighbors' garage doors and his kitchen lights. After trying and failing at several attempts to fix the problem, he just replaced the dimmer switch with a simple on-off switch, and the noise apparently disappeared. Nobody made any measurements, and the local Federal Communications Commission office did not want to know about it. Was this a faulty dimmer switch or just a poor installation? If the switch itself was to blame, then are millions of these dimmers causing interference in the United States and throughout the world? Who knows?

If we had the relevant noise database in hand, it would be theoretically possible to revise today's regulatory standards and even to extrapolate, from the data, how those standards will have to change a decade from now. The hard part here will be in establishing a reasonable balance between the costs and benefits of such standards. However, the needed data is almost totally missing because we have only a very small database to work from—ours, basically—and we can't prove whether or not the problem is widespread.

One change that's clearly needed is in the way the regulators in most countries divvy up the electromagnetic spectrum. Today regulators tend to worry only about interference from an identified transmitter—say, a radio station—while ignoring noise. Yet noise can be far worse than the identified interference.

Regulators should decide this question, at least in part, by considering how much the communications of today's users would be degraded if new users were allowed to share in the spectrum. This will help to achieve the best trade-off between having some RF pollution that causes minimal interference due to spectrum sharing versus the expectations of empty spectrum.

Even if you've never lost a television signal to an electric toothbrush or a cellphone call to a malfunctioning transformer, you still shoulder some of the cost of such noise pollution. That's because you depend on wireless systems—and those systems must be engineered to handle ever-increasing levels of noise. The problem only grows with the proliferation of electronic devices and the ever-finer sharing of smaller and smaller pieces of the spectrum. No masterstroke can quell noise. Because the problem has to do with both design and maintenance, it incorporates an element of chance, and it is difficult to isolate from other causes of interference.

Our lack of the most basic data on the extent of the problem is unacceptable. We need that data and also a better understanding of the role that noise plays in wireless communications before we can strike the right balance between the costs and benefits of noise suppression. With the exponential increase in the number of noise-generating devices and of wireless systems, RF pollution will become a very expensive problem—unless we act now. ■

POST YOUR COMMENTS at http://spectrum.ieee.org/rfnoise0915

COLLEGE OF ENGINEERING
ELECTRICAL & COMPUTER ENGINEERING
UNIVERSITY OF MICHIGAN

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. The division invites candidates across all research areas relevant to ECE to apply.

The highly ranked ECE Division (**www.ece.umich.edu**) 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 7, 2015.**

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.

Georgia Tech | School of Electrical and Computer Engineering

The School of Electrical and Computer Engineering at the Georgia Institute of Technology invites applications for tenure-track faculty at the Assistant and Associate Professor level. Applicants should have an earned Ph.D. or equivalent. The School seeks individuals with outstanding potential for research achievement, and a strong aptitude and interest in undergraduate and graduate teaching. Related industry experience is desirable.

Candidates are sought with interdisciplinary and multidisciplinary backgrounds and interests in the following areas:

Computer Engineering with focus on software and hardware, with application to many-core systems, embedded/mobile platforms, high performance computing, resilient systems, non-conventional computing, hardware security, and cyber-physical systems

Energy systems with emphasis on control, protection, automation, cyber-physical security, integration of renewables, reliability, and power electronics applications.

Novel electronic materials by design, nano systems and technologies that address major social problems.

Exceptional candidates in other areas will also be considered. Diversity candidates are strongly encouraged to apply. For more information about the School of ECE at Georgia Tech, please visit **http://www.ece.gatech.edu**

Interested candidates should submit an application letter, curriculum vita, research and teaching statements, and names of three references electronically at: **http://facjobs.ece.gatech.edu**.

Review of applications will begin immediately. To receive the most serious consideration, interested applicants should submit their materials before **November 1, 2015**. Georgia Tech is an equal opportunity, affirmative action employer.

CLEMSON UNIVERSITY

College of Engineering and Science
Faculty Search

The Holcombe Department of Electrical and Computer Engineering at Clemson University is seeking applicants for multiple faculty positions at the rank of tenure-track assistant professor. The Department seeks applicants in technical areas associated with electronics and photonics (e.g., optoelectronics, lasers, semiconductor optical devices, integrated optics, fiber optics, microwave circuits, terahertz devices, lab-on-chip, bioelectronics, microsystems, nanoscale devices, flexible electronics) and computer engineering (high-performance computing, computer vision, networking, cyber security, big data and data-enabled science, embedded systems, and intelligent systems). Outstanding candidates will be considered for the Warren Owens Assistant Professorship.

The Holcombe Department of Electrical and Computer Engineering is one of the largest and most active at Clemson, with 30 tenured/tenure-track faculty members, 570 undergraduates and 190 graduate students. Many members of the faculty are known internationally; they include seven IEEE Fellows, three endowed chairs, and eight named professors. Externally funded research expenditures exceeded $5.7 million in 2014. Clemson University is a land-grant institution committed to academic excellence, world-class research and a high quality of life. Six interdisciplinary colleges and schools house strong programs in architecture, engineering, science, agriculture, natural resources, business, social sciences, arts and humanities, health care, and education. A faculty of 1,500 and staff of 3,000 support over 80 undergraduate degree offerings, and more than 70 master's and 40 Ph.D. programs. An annual operating budget of approximately $1 billion and an endowment of more than $500 million fund programs and operations. Major new research, graduate education, and economic development activities are enhanced by public-private partnerships at three innovation campuses and six research and education centers located throughout South Carolina. Today, Clemson University is ranked 20th among national public universities by U.S. News & World Report, and remains true to its roots as a science- and engineering-oriented research university with a strong commitment to teaching and student success. Clemson University is described by students and faculty as an inclusive, student-centered community characterized by high academic standards, a culture of collaboration, school spirit, and a competitive drive to excel.

Applicants must have an earned doctorate in electrical engineering, computer engineering, or a closely related field. Applicants should submit a current curriculum vitae and a minimum of five references with full contact information. Electronic submissions (PDF files) to ecefacsearch@clemson.edu are preferred, but applications and nominations can also be mailed to ECE Faculty Search, 105 Riggs Hall, Clemson University, Clemson, SC 29634, USA. Application material must be received by **November 1, 2015** to receive guaranteed consideration, though the search will remain open until the position is filled.

Clemson University is an AA/EEO employer and does not discriminate against any person or group on the basis of age, color, disability, gender, pregnancy, national origin, race, religion, sexual orientation, veteran status or genetic information. Clemson University is building a culturally diverse faculty committed to working in a multicultural environment and encourages applications from minorities and women.

Two Tenure Track Professor Positions and one Professor of Practice

Aalto University School of Electrical Engineering seeks experts in the following fields:

- Tenure Track position in Automatic control (open at all levels from assistant to full tenured professor)
- Tenure Track position in Networking technology and software (open at assistant professor level)
- Professor of Practice in Smart building technologies and services

Applicants must have a doctoral degree in a relevant field. For assistant and associate professor levels evaluation focuses on merits and potential for excellence, for full professors we look for demonstrated excellence in research and teaching. We especially encourage applications from young professionals. Relevant industrial experience is appreciated.

Application deadline is September 30, 2015. For further information and application details, please visit at: **www.elec.aalto.fi/tenuretrack**

Aalto University is a new university with over a century of experience. Created from a high-profile merger between three leading universities in Finland – the Helsinki School of Economics, Helsinki University of Technology and the University of Art and Design Helsinki – Aalto University opens up new possibilities for strong multidisciplinary education and research. The university has 20 000 students and a staff of 5 000 including 350 professors.

Aalto University
School of Electrical Engineering

aalto.fi

Faculty Positions in Robotics and Mechatronics

NAZARBAYEV UNIVERSITY

Nazarbayev University is seeking highly qualified full-time faculty members at the Assistant, Associate and Full Professor levels to join its rapidly growing program in Robotics and Mechatronics in the School of Science and Technology. Successful candidates must have a Ph.D. degree from a reputed research university, a demonstrated ability for research, excellent English-language communication skills 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. The university is committed to be a world-class research university and has made significant investments in research infrastructure, equipment and personnel.

Applications are particularly encouraged from candidates with expertise and research interests in the areas of medical robotics, unmanned aerial vehicles, sensor fusion, and mechanical systems design and control. Exceptional candidates with research interests in the broader field of Robotics and Mechatronics are also encouraged to apply.

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

Applicants should send a detailed CV, teaching and research statements, 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 September 1st, 2015. Successful appointments are expected to begin on January 4th, 2016. For more information please visit **http://sst.nu.edu.kz**.

上海科技大学
ShanghaiTech University

ShanghaiTech Faculty Search

ShanghaiTech University invites highly qualified candidates to fill multiple tenure-track/tenured faculty positions in the School of Information Science and Technology. Candidates should have exceptional academic records or demonstrate strong potential in cutting-edge research areas of information science and technology. English fluency is required and overseas academic connection or background is highly desired. ShanghaiTech is built as a world-class research university for training future generations of scientists, entrepreneurs, and technological leaders. Besides establishing and maintaining a world-class research profile, faculty candidates are also expected to contribute substantially to graduate and undergraduate education.

Academic Disciplines:
We seek candidates in all cutting-edge areas of information science and technology. Our recruitment focus includes, but is not limited to: computer architecture and technologies, nano-scale electronics, high-speed and RF circuits, intelligent/integrated signal processing systems, computational foundations, big data, data mining, visualization, computer vision, bio-computing, smart energy/power devices/systems, next-generation networking, as well as inter-disciplinary areas involving information science and technology.

Compensation and Benefits:
Salary and startup funds are highly competitive, commensurate with experience and academic accomplishment. We also offer a comprehensive benefit package to employees and eligible dependents, including housing benefits. All regular ShanghaiTech faculty members will be within its new tenure-track system with international practice for performance evaluation and promotion.

Qualifications:
- A detailed research plan and demonstrated record/potentials;
- Ph.D. (Electrical Engineering, Computer Engineering/Science, or related field);
- A minimum relevant research experience of 4 years.

Applications:
Submit (in English, PDF) a cover letter, a 2-page research plan, a CV plus copies of 3 representative publications, and names of three referees to sist@shanghaitech.edu.cn by October 31, 2015. For more information, visit **http://www.shanghaitech.edu.cn**.

REED RELAY REDUX

This simple device, once crucial to telephone exchanges, lives on in portable defibrillators

Invented in 1922 by a professor at Leningrad Electrotechnical University, the reed switch contains two slivers of wire that make contact in the presence of a magnet. Bell Telephone Laboratories evolved this device into the reed relay, in which the switch is controlled by an electromagnet. By the 1960s, reed relays were a fundamental component of telephone exchanges. Their ultralow power consumption also made them popular for space applications, including the Apollo missions.

Although solid-state components eventually eclipsed reed relays for telephone switching, they're still useful when you need to stop very low currents from leaking away (as in photomultiplier detectors) or are handling high voltages (as in cardiac defibrillators). Along the way, the price of reed switches and relays has come down considerably. A reed switch optimized for aerospace once cost about US $200. Modern reed switches go for well under a dollar, and vintage reed relays like the ones in this ad from *IEEE Spectrum*'s February 1965 issue can still be found on eBay.

—ROBERT COLBURN, IEEE HISTORY CENTER

PHOTOGRAPH BY Randi Klett