IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 10.18

KILL THE TRAFFIC LIGHT!

Communicating wirelessly, cars will cooperatively navigate intersections—**and cut commute times by a third** P. 24

IEEE

IEEE SPECTRUM

On the cover and above **Photographs for *IEEE Spectrum* by Dan Saelinger**

IEEE SPECTRUM

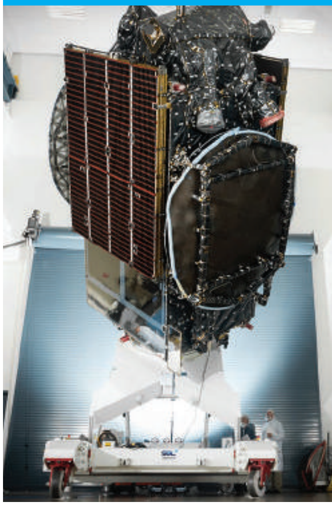

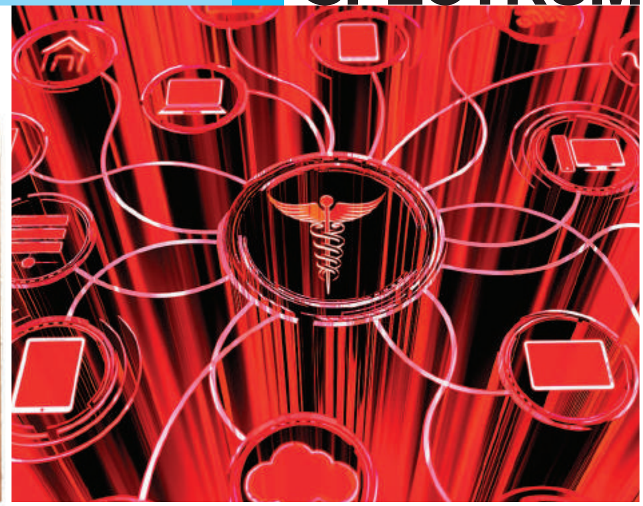

ADDITIONAL RESOURCES

Tech Insider / Webinars
Available at spectrum.ieee.org/webinars

- Magnetic Transmission Design Using Multiphysics Modeling

White Papers
Available at spectrum.ieee.org/whitepapers

- Increase Reliability and Efficiency in Power Converter Designs
- Five Steps to Medical Device Commercialization in the U.S.
- Shielding Effectiveness of Expanded Metal Foils

The Institute
Available at theinstitute.ieee.org

- **INNOVATIONS IN POWER** Learn about the latest advancements in the power field in such areas as batteries, electronics, and the smart grid.
- **SOCIETY ANNIVERSARY** The IEEE Power Electronics Society got its start 35 years ago.
- **VOLUNTEER TRAINING** The IEEE Center for Leadership Excellence is a training portal for volunteers that provides them with resources to perform their duties.

FROM LEFT: SSL; RANDI KLETT; ISTOCKPHOTO

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2018 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 55, issue No. 10, 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.

MATTHIEU WYART

WHEN CALCULATORS WERE REVOLUTIONARY

HISTORICAL RESEARCH CAN TAKE its practitioners to unexpected places. So it was for Ksenia Tatarchenko. In “The Great Soviet Calculator Hack,” in this issue, she delves into an ambitious 1980s program to make teenagers in the USSR computer literate.

Because most Soviet households and schools lacked actual computers, many students made do with programmable calculators. A popular science magazine spurred on their interest using the computer game *Lunar Lander* and a serialized sci-fi novel. In the game, you try to land on the moon by entering power levels for your rocket as you descend. In the novel, an engineer and a pilot take their tiny lander from the moon all the way back to Earth. “Many people played the game,” says Tatarchenko, who is a lecturer at the University of Geneva. “And so I thought, Well, it must be easy.”

Her mother, who lives in Novosibirsk, Russia, bought a 1980s calculator and sent it to Tatarchenko. Operating the calculator, she discovered, was far from easy. To play the game, you had to enter its 97 steps into the calculator. But rendering equations in reverse Polish notation, in which you punch in all the numbers and then the operators, was counterintuitive. And the assignments included with each of the novel’s installments required a familiarity with the physics of spaceflight. Some readers went well beyond those tasks, creating their own games. “I tried and tried, but I never reached the stage of inventing my own game.”

Still, she says, she enjoyed learning about this underappreciated device. In the West, Tatarchenko notes, the personal computer is what’s credited with sparking the digital age. But in the Soviet Union, the programmable calculator proved just as revolutionary. ■

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, Past Forward is in *IEEE Spectrum*, Vol. 55, no. 10 (INT), October 2018, p. 52, or in *IEEE Spectrum*, Vol. 55, no. 10 (NA), October 2018, p. 56.

CONTRIBUTORS_

Behnood Gholami

Gholami trained as an aerospace engineer and worked on controllers for airplanes. Now he's using that expertise to develop controllers for hospital systems, as described in "AI in the ICU" [p. 30]. He cofounded Autonomous Healthcare with Wassim M. Haddad, an aerospace engineering professor at Georgia Tech and an IEEE Fellow, in hopes of providing a smoother ride for patients. "The human body is just as dynamic as any airplane," Gholami says.

Nicholas Little

Little's illustrations for "Saving Software From Oblivion" [p. 36] make good use of a style he's been perfecting since earning a graduate degree from New York's School of Visual Arts, in 2015. The style mimics the graphics found in old computer games. "There's been a resurgence of that aesthetic," says Little, who's based in Brooklyn, N.Y. "I had a Nintendo, so I grew up right at the end of that era."

Dan Saelinger

Saelinger, who works out of Portland, Ore., is a conceptual photographer: He shoots artfully arranged objects to represent an idea. He created this month's cover–for a story on virtual traffic lights [p. 24]–using paper, paint, and model cars. "It's the antithesis of high tech," Saelinger says. "We take this system of complex and futuristic technology and then break it down into a beautiful and graphic visual created from simple items."

Mahadev Satyanarayanan

Satyanarayanan is a professor of computer science at Carnegie Mellon University, in Pittsburgh. Almost a decade ago, he began collaborating with IBM, developing software that allowed a virtual copy of one computer to be created on another computer. He later applied the same strategy to archive old software in executable form, which he discusses in this issue [p. 36]. "We're basically re-creating lost worlds–like the dinosaur movies," he says.

Ozan K. Tonguz

Tonguz is a professor of electrical and computer engineering at Carnegie Mellon University. He's also the founder and CEO of the CMU spin-off Virtual Traffic Lights, which uses vehicle-to-vehicle communication and distributed artificial intelligence to fight traffic jams, "thus giving back to people years of life now wasted during their daily commutes," he says. Tonguz describes the technology in "Red Light, Green Light–No Light," on p. 24.

Celebrating IEEE and Its Members

Marking the anniversary of the first technical meeting of EEs and the origins of IEEE

During October, IEEE members around the world celebrate IEEE Day, commemorating the anniversary of the first technical meeting of the American Institute of Electrical Engineers, the society that eventually merged with the Institute of Radio Engineers in 1963 to become IEEE. ¶ IEEE is renowned for its conferences, publications, and standards. But IEEE's *members*—each and every one of its members—are the key, the secret ingredient, to its continued success as a world-class global technology association. ¶ Why do people join IEEE? As the editor of *IEEE Spectrum*, I've had the pleasure of talking to many members about what they do and why they choose to stay close to IEEE. Sometimes these member stories come from unexpected sources. ¶ I once went to a new dentist and found, to my surprise, copies of *Spectrum* in his waiting room. Why? Before he turned to oral surgery, he'd taken an undergraduate EE degree at what was then the Polytechnic Institute of Brooklyn, where he became a student member of IEEE. A big fan of *Spectrum*, he'd kept his membership going long after he became an oral surgeon. He liked keeping up with new technical developments, and I think he also had some lingering "what ifs"—what if he'd pursued his interest in chips and computer programming instead of implant technology and imaging? ¶ Other professionals become and remain members for more directly practical reasons—to network with their colleagues or meet new ones, to stay current with rapidly changing developments in technology, to publish in IEEE's many important journals and transactions, to develop "soft skills" like people management and résumé writing.

For me, the networking component is very important. In an age when you can go online and idly scroll through your phone to research the most ridiculously complicated subjects, still nothing beats a conversation with a fellow carbon life-form to learn something new—while making a new friend or two along the way. Vint Cerf—one of the "fathers" of the Internet, Google evangelist, and IEEE Fellow—says, "It's the interaction among people, the side conversations, and the chatting in front of a whiteboard that makes IEEE so valuable."

But, as is the case with most things in life, you get back what you put into whatever group or community you belong to—your school, your company, your family and friends, your IEEE chapter, section, or region.

Are you an active IEEE member? Do you attend meetings, volunteer? If so, we'd like to hear from you. Why did you become a member? Do you come from a long line of members? How has being a member helped you? Be sure to include your name and member grade when you send me your story (s.hassler@ieee.org). Or post your comments online.

We are 374,778 members strong, in 160 countries, with 3,005 student branches and over 2,000 chapters uniting local members. Our members are engineers, scientists, and allied professionals whose technical interests are rooted in electrical and computer sciences, engineering, and related disciplines. Happy IEEE Day!

—SUSAN HASSLER, *IEEE Senior Member*

Correction: The article "Hydrogen Hybrids Debut as Zero-Emission Delivery Trucks" [September] had two incorrect specifications. The converted UPS van has a 48-kilowatt-hour battery pack, and each of its two hydrogen tanks weighs 5 kilograms.

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/ieeeday1018

NEWS

NEW SPOT BEAMS BOOST SATELLITES' BANDWIDTH

SSL's satellites constructed for Telesat also feature 3D-printed components

HI-DEF BROADCAST: Telstar 19V's spot beams will improve data rates across the Americas.

SSL

Over the past few months, the spacecraft manufacturing company SSL put the finishing touches on three massive communications satellites. It built two of them for Canadian operator Telesat. Then it initiated its plan to have them launched, one after the other, into distinct geosynchronous orbits 36,000 kilometers above Earth—perches that will keep each satellite hovering over a particular spot, even as the planet turns.

From their respective vantage points, the pair built for Telesat will deliver high-speed communications services, including broadband Internet, to Asia and the Americas for the next 15 years.

At more than 7,000 kilograms, Telstar 19 Vantage became the heaviest commercial communications satellite ever launched when it left Earth on 22 July. But the record-setting 19V and its sibling, Telstar 18 Vantage (which was launched on 10 September), are also noteworthy in other ways. »

What makes them so special? The 18V and 19V feature antennas that can transmit a type of beam that will vastly improve their data throughput. What's more, SSL used 3D-printing techniques to build and customize both satellites.

Geosynchronous communications satellites send and receive signals over an exceptionally wide coverage area. "With geo, you can cover about one-third of the Earth with one satellite," says Carolyn Belle, a senior analyst at Northern Sky Research, in Cambridge, Mass., a telecommunications consulting firm. From certain angles, these satellites can have almost half of the global population within range.

That's a lot of potential users, spread over a massive area, which is why communications satellites have historically operated with wide-beam coverage. But, as SSL's acting chief technology officer Rob Schwarz says, spot beams are becoming more popular. Spot beams are focused rays of electromagnetic energy that can transmit more data to smaller areas than wide beams can manage.

"An analogy is the human eye and the insect compound eye," says Schwarz. "The human eye is one wide view, while the compound eye is like spot beams." Like the components of the multifaceted eye of a dragonfly, the spot beams produced by antennas onboard both the 18V and 19V focus in slightly different directions, allowing the same spectrum bands to be used multiple times without interference.

Eventually, the 18V and 19V will look down on Earth with their spot beams and wide beams, as if through a compound contact lens fitted over an eye. "There will be a backdrop of broad-area beams, and in high-density areas, we're adding spot beams to aid coverage," says Schwarz.

Spot beams ultimately function similarly to cell towers. The satellite transmits along a specific band in the general area around the ground station, and the station, tuned into that band, will receive the signal. "It's like locating a cellphone in a specific cell," says Schwarz.

The 18V and 19V use the Ku (12 to 18 gigahertz) and Ka (26 to 40 GHz) bands for their spot beams and wide beams. The bands have risen in popularity in recent years because they deliver higher data rates than the 4-to-8-GHz C band commonly used for satellite communications. This shift came despite the fact that higher-frequency bands have a greater signal loss in the atmosphere, particularly in rainy climates.

Both of SSL's new communications satellites—and a third called Merah Putih (launched 7 August) that SSL built for the state-owned operator Telkom Indonesia—

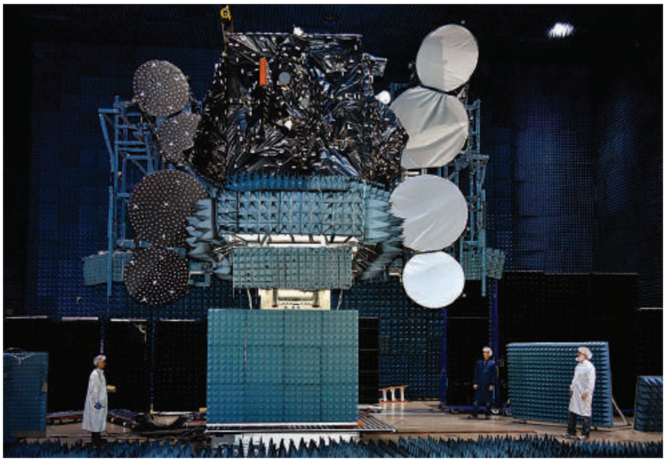

FAR-OUT VIEWS: Coverage from Telstar 18 Vantage [top] will span from India and Pakistan to Hawaii. Prior to launch, Telesat lined up agreements to provide satellite Internet from Telstar 19 Vantage [middle] to five countries in South America and 25 communities in Nunavut, Canada's northernmost territory. Merah Putih [bottom] will serve remote parts of Indonesia.

SSL (3)

are based on the popular SSL 1300 series framework. Additionally, the satellites required specialized antenna arrays, which in Telesat's case had to support spot beams.

The Telstar satellites launched by SSL, which is based in Palo Alto, Calif., have 3D-printed components in the antenna struts, which support antennas, tracking equipment, and satellite-control equipment. The 3D printing speeds up the design process and can create joints that are just as strong but have more complex shapes than those made with traditional assembly, without requiring difficult welding.

This technique also cuts down on costs, though geosynchronous satellites remain expensive. "Geosatellites are up in the billions of dollars per satellite" after factoring in the cost of launching, says analyst Belle. Although some startups are pursuing constellations of smaller, low-Earth-orbit satellites as an alternative, Belle says larger satellites remain viable thanks to their efficient power usage and very low cost per bit of data.

As for the trio of new communications satellites, each one must spend about 10 days under its own propulsion after launch to reach its final, geosynchronous orbit. Once there, they deploy their solar arrays and open their antennas. As SSL makes contact with each satellite, it runs checks on every system to ensure the equipment is in optimal condition. Then SSL shepherds the satellites for 30 to 40 days before handing them over to their operators.

If they do their jobs well for the next 15 years, then just before each satellite runs out of propellant, it will be boosted into a higher, graveyard orbit. "The graveyard orbit is for the good satellites who have completed their operations," says Schwarz.
—MICHAEL KOZIOL

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/ssl1018

MIKE ELIASON/SANTA BARBARA COUNTY FIRE DEPARTMENT/AP

UTILITIES ROLL OUT REAL-TIME GRID CONTROLS

Synchrophasor tech enables rapid response to broken power lines and other emergencies

SLOW BURN: A 2017 wildfire near power lines in Montecito, Calif., burned for almost six months before firefighters subdued it.

Amidst what could be California's worst wildfire season on record, San Diego Gas & Electric is counting on technology to reduce dangerous sparking from its power lines. Last month, the utility completed the initial rollout of a home-grown automated control technology that taps ultrafast synchrophasor sensors to detect and turn off broken power lines before they hit the ground.

Projects such as this mark a turning point for grid control. Synchrophasor sensors send out time-stamped measurements of power and its phase—the angular position of the alternating current and voltage waves—up to 60 times per second. That is at least 120 times as fast as most utilities' industrial control systems. And the GPS-synchronized time stamps allow data assembled from multiple sensors to create a precise wide-area view of power grids.

The grid's human operators have progressively attained a wider view since the synchrophasor device's invention 30 years ago. But only recently have they begun to exploit the speed of these phasor measurement units (PMUs) for real-time grid control.

San Diego's line-break-protection system works by spotting quick voltage changes. PMUs arrayed along a circuit report continuously via a high-speed wireless radio communications network to a controller in a substation. If the controller spots a sudden voltage spread between adjacent sensors, it orders the closest relays to isolate and de-energize the iffy segment. Generally, it's all over in less than half a second.

San Diego Gas & Electric and its parent company, Sempra Energy, started looking at synchrophasor sensors in 2010 and quickly identified dozens of potential uses. A broken-line-detection and control system became the utility's flagship project after engineer William O'Brien calculated that it could spot broken lines two to three times as fast as gravity could pull them down, allowing the controller to stop the flow of electricity before a line touched the ground, and thus greatly reduce the risk of fire. (O'Brien developed and patented the concept with Eric Udren, an executive adviser at Quanta Technology, a consultancy based in Raleigh, N.C.)

As of last month, the system had been installed in what the utility expects will be its final form on six circuits emanating from three substations in the fire-prone territory east of San Diego. The utility has 18 more substation build-outs planned and expects to ultimately deploy the system across its entire grid.

This system for detecting and disarming broken lines marks the first deployment of PMU-based automation on a distribution system. But a few utilities elsewhere have already integrated synchrophasor-based controls into their high-voltage transmission grids. One of the first installations, initially completed in Iceland in 2014 and substantially upgraded last year, tunes the island nation's 50-hertz AC frequency.

Iceland has a relatively small grid whose supply and demand can easily be thrown out of balance when power plants, transmission lines, or big factories unexpectedly go off line. For years, the resulting AC frequency fluctuations regularly caused the grid's eastern and western zones to split into electrical islands, which often led to power outages.

A wide-area PMU network and added controls, provided by GE's grid solutions business, enabled Iceland's grid operator, Landsnet, to rapidly locate power imbalances and automatically fix them by tweaking demand from aluminum smelters and other big consumers. The June 2017 updates have cut the magnitude of Iceland's AC frequency deviations roughly in half, according to GE senior power systems engineer Sean Norris. "Events that we previously would have expected to cause splits in the system have occurred, and the system has remained intact," says Norris.

Emerging frequency challenges for Great Britain's much-larger grid have prompted a three-year research effort directed by London-based National Grid. England and Scotland's many fossil-fueled power plants and the inertia in their heavy rotating generators currently hold the United Kingdom's AC frequency steady. But that frequency-stabilizing inertia is disappearing as coal and gas plants shutter.

Simulations conducted earlier this year at the University of Strathclyde, in Glasgow, showed that synchrophasor-driven controls, running on an expanded version of GE's technology, could keep the U.K. grid stable with fewer inertia-rich generators. By early next year, the research team hopes to begin testing its control platforms at National Grid substations.

Ultimately such real-time controls will take over grid operation, according to Patrick Lee, president of control developer PXiSE Energy Solutions (another Sempra Energy subsidiary). As renewable generation grows, industrial control systems aided by human operators watching PMU readings will no longer suffice. According to Lee, "As the system gets more renewable integration and becomes more dynamic, you have less time to respond. If you don't have this high-speed synchrophasor-based technology, you really will have no chance." –PETER FAIRLEY

Synchrophasors could sense a power line break before gravity pulls the loose ends to the ground

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/utilities1018

MACHINE VISION TO CURB PIG PUGNACITY

3D cameras can help predict when pigs are about to nip each other's tails

Pig farmers want human diners to bite into the delicious pork they produce, not for swine to bite each other. (Yes, it happens.) Now, using 3D cameras and machine-vision algorithms, scientists are developing a way to automatically detect when a pig might be about to chomp down on another pig.

Pigs have an unfortunate habit of biting one another's tails. Infections from these bites can render up to 30 percent of a pig farm's swine unfit for human consumption. Docking, or cutting, pig tails can reduce such biting but does not eliminate it, and the routine use of docking is banned in the European Union. There are a wide range of potential triggers for outbreaks of tail biting–among them genetics, diet, overcrowding, temperature variations, insufficient ventilation and lighting, disease, and even the season–so it's an unpredictable problem. "Tail biting is a very frustrating challenge," says John Deen, a veterinarian and epidemiologist at the University of Minnesota. "Controlling it has not always been that effective."

To predict and potentially prevent tail biting, researchers in Scotland monitored 667 undocked pigs on a farm using both time-of-flight and regular video cameras that recorded continuously for 52 days. The pigs were checked at least twice a day for evidence of biting.

TASTY TAILS: In extreme tail-biting outbreaks, up to 30 percent of pigs raised together may be affected by bites so severe that their carcasses are no longer fit for human consumption.

Each time-of-flight camera emitted pulses of infrared light from LEDs 25 times a second, and recorded the amount of time needed to detect reflected pulses. This data allowed scientists to track each pig's position and posture. Machine-vision algorithms from farm-technology company Innovent Technology, in Aberdeenshire, Scotland, then determined which activities might serve as possible early warning signs of tail biting.

The scientists found that before outbreaks of biting, pigs increasingly held their tails down against their bodies. Moreover, the software could detect when these changes in tail posture occurred with 73.9 percent accuracy. "It looks like good technology, and I'm very interested in how it could be applied on a farm," says Deen, who did not take part in this project.

If farmers think a biting outbreak is likely to happen in a pigpen, they could deploy distractions such as straw, knotted ropes, or shredded cardboard, which tap into the pigs' instincts to root and chew.

"Another thing that people try is to apply bad-tasting stuff such as Stockholm Tar to tails," says Richard D'Eath, an animal behavior scientist at Scotland's Rural College, in Edinburgh, who worked on this research. An early warning system could help farmers use such remedies only when needed, which would save money.

This research was part of a £160 million push by the United Kingdom to support innovative farming technology through its Agri-Tech Catalyst program. Agriculture and food already help generate more than £108 billion annually and support 3.9 million employees. A recent industry-led review suggested that incorporating digital technologies such as robotics and autonomous systems into food manufacturing could add £58 billion to the U.K. economy over the next 13 years.

A three-year project called TailTech is now furthering the development of this early warning system with up to £676,000 in funding from Innovate UK, a government agency. The aim is to test a prototype system on more than 16,000 pigs at nine farms throughout Europe for roughly 18 months, with each time-of-flight camera capable of monitoring up to 300 pigs, D'Eath says.

TailTech will compare the system's efficacy for different types of pig farms, including ones with docked and undocked pigs, with and without straw on their floors, and with pigs of varying genetics, diet, and group sizes. The project will analyze what fraction of pigs hold their tails low, and for how long, before scientists are sure a tail-biting outbreak will occur. And it will assess how predictors vary between farms. The researchers also aim to improve the accuracy of their system. "In practice, though, 73.9 percent is good enough for the system to work well," D'Eath says.

The ultimate aim is to have an early warning system "that reads out continually on a screen and also sends alerts to the farmer's smartphone," he adds. "No technical expertise will be needed once the system is installed." The software will compute trends to give farmers a better idea of their herds' current level of risk, D'Eath says.

A farmer might not buy technology designed solely to detect tail biting, but D'Eath notes that this system is being developed as an add-on to a camera-based automatic pig-weighing system called Qscan that Innovent already produces to help farms meet contractual weight targets. Deen, the veterinarian and epidemiologist, thinks this strategy could make all the difference. As he says, "If Qscan is already in place, I think farmers can quite easily justify adding this system at little cost." —CHARLES Q. CHOI

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/pigtails1018

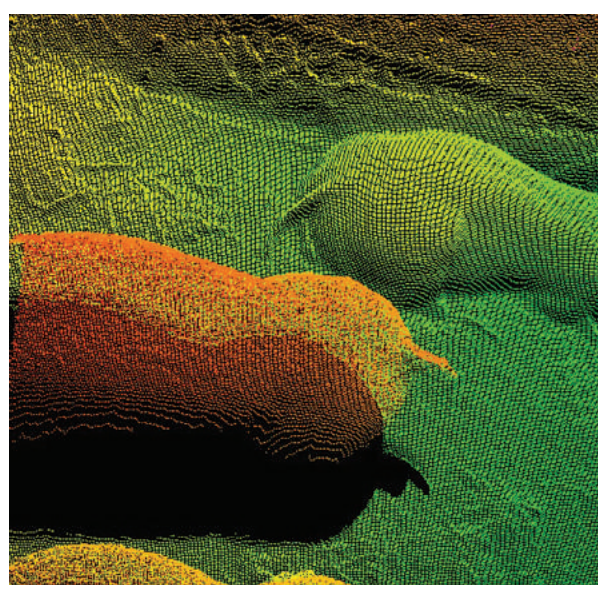

PIG VISION: The TailTech system monitors pigs with cameras that produce 3D images. Each pixel is created by measuring the reflection of infrared pulses as an indicator of the distance from objects. Software then represents the distances with colors.

TOP: ARCTIC IMAGES/GETTY IMAGES; BOTTOM: MARIANNE FARISH/SRUC

NEWS

ICELAND'S CONSUMERS TRY DRONE DELIVERY

The startup Aha takes on Amazon with basic drones bearing burgers

An Icelandic startup called Aha is using a Chinese-made drone and an Israeli logistics system to deliver hot food, groceries, and electronics to households in Iceland's capital city of Reykjavik.

It's a world's first, and one that flouts the standard aviation safety mantras. These drones don't sense and avoid obstacles–in fact, they don't even have cameras, radar, or any other imaging systems. They fly according to GPS coordinates, along routes certified free of trees, buildings, and other impediments. And with some 500 deliveries completed in the past five months, no injuries have been reported.

It works like this: You punch your order into an app on your smartphone ("Two hamburgers, hold the onions") and Aha's cook loads the food onto the drone. Then you track the delivery, go outside to welcome it, and if all's well at the drop-off point, you agree to accept it. Then the drone lowers your burgers on a line and buzzes home.

Delivery is around US $7. That's enough to cover operational costs, says Maron Kristófersson, the chief executive officer of Aha. "The electricity comes to 25 cents," he says.

A delivery can be completed in as little as 4 minutes, versus 25 minutes when delivering by road, under heavy traffic. That's perfect for hot food (burgers are the most popular item), but customers also seem to value it for groceries (where bananas rule) and hardware (mostly electronics).

Kristófersson began to look into autonomous delivery schemes in 2014 as a way to deal with skyrocketing labor costs. But when he turned to drone companies, he found that nearly all were lukewarm about Iceland's tiny market.

Then, in 2015, he contacted Flytrex, a Tel Aviv startup that was selling GPS trackers to drone companies. Flytrex developed a logistics system based on those trackers rather

AHA

AIR DROP: Aha's drone delivery service operates until 7 p.m. in Reykjavik on days that are not too windy, too snowy, or too rainy.

than making a drone of its own. “This was much more pragmatic, in my view,” Kristófersson says.

“FedEx doesn’t produce trucks; they procure them,” says Yariv Bash, chief executive officer of Flytrex. “We do the same with drones. We have the know-how to choose the right drones and modify them to our needs, but our core competence is in logistics and cloud capability.”

Kristófersson told the Icelandic air authorities that they had a choice: Iceland could remain at the bottom of the list of countries that would develop drone delivery or it could be at the top. And to be at the top, Iceland had to put safety regulations in place.

After more than a year of back-and-forth talks, Aha got the go-ahead. The first experimental deliveries came earlier this year.

Aha is using the DJI Matrice 600 drone, which carries 3 kilograms of cargo and weighs 15 kg (33 pounds) fully loaded. It can fly 8 kilometers—4 km out and 4 km back—enough to cover the outer ring of Reykjavik from Aha’s operations center in the middle of town. The routes are surveyed and reviewed constantly to avoid new construction. To further decrease the risk, the drones are directed over water and unpopulated industrial areas as much as possible.

Initially, the drones arrived at a dozen or so set drop-off points on the outskirts of the city. A company representative would have to wait to receive the drone and pull the package out of a cargo compartment. These deliveries gave experience to Aha and peace of mind to the safety regulators.

Then, in August, the company got permission to lower packages on a line to select homes. “It can drop a package almost on a carpet in your backyard,” says Bash. All you have to do is get your neighbors to agree to the plan.

That’s rarely a problem, says Kristófersson. When you mow your lawn, you cause much more disturbance, he insists.

Flying blind has been a contentious issue in other jurisdictions around the world. In the United States, the Federal Aviation Administration is sticking to its sense-and-avoid requirement. But the agency has recently allowed test flights, including one by Wing, a subsidiary of Alphabet (Google’s parent company), to go farther than the operator can see with unaided sight.

Bash says that two countries—he won’t give their names or even their continents—are about to allow projects like the one Flytrex and Aha have started in Iceland. And why not, he argues, when drone delivery poses less risk than some other practices that we allow in our lives.

“Which is safer,” he asks, “a 16-year-old driving a 2-ton vehicle to your house every time you want [a] hamburger, or sending a 30-pound drone that didn’t miss any sleep the night before?”

—PHILIP E. ROSS

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/drones1018

PHOTOGRAPH BY Annie France Noël

A LAMENT FOR SHORTWAVE

IN 2017, ARTIST Amanda Dawn Christie commemorated Radio Canada International's New Brunswick shortwave site, which was torn down in 2014 after decades of operation. For the three-person performance, dubbed *Requiem for Radio: Full Quiet Flutter*, model radio masts were rigged with capacitive sensors. Touching the sensors played tones derived from sounds that the RCI masts made in the wind. For another part of the performance, Christie [far left] commissioned a five-channel electroacoustic work, based on a requiem mass, from composer Lukas Pearse. The channels were simulcast by shortwave broadcasters and incorporated Morse code and voices calling to the silent RCI site. The simulcast created bemusement on shortwave-listener message boards, says Christie, with a laugh: "There was this long thread of people going, 'Hey, I'm hearing this weird thing on 11580. I'm hearing this on 5140. Oh, they're not quite the same, but similar.' And then someone says, 'I've decoded the Morse code. It's in Latin!' "

For more about Christie's shortwave-inspired art, see "The Radio Elegist," in this issue.

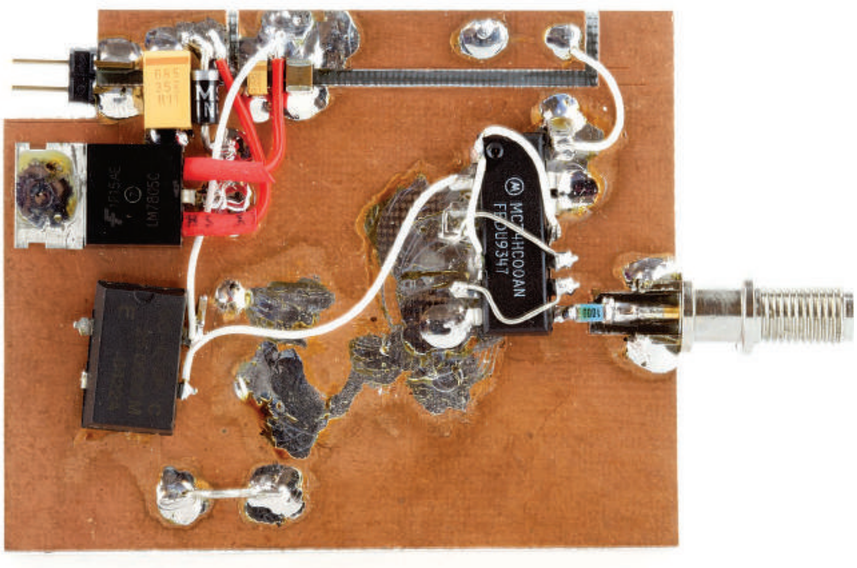

RANDI KLETT

RESURRECTING THE "DEAD BUG" METHOD

AN UNCONVENTIONAL APPROACH TO PROTOTYPING HIGH-FREQUENCY CIRCUITS

DON'T GET ME WRONG—I LOVE PRINTED CIRCUIT BOARDS. PCBs are, of course, essential in mass-produced products. Even for hobbyists, a small run assures almost perfectly repeatable circuits. And PCBs with a good ground plane are essential for high-frequency circuits operating at more than a few megahertz. A ground plane is a large area of copper that's used as a low-inductance electrical return path from components to a circuit's power supply. It prevents parasitic capacitance from smearing high-frequency signals into noise, and the absence of a ground plane is why you can't build a high-frequency circuit using a breadboard and expect it to work well, or at all. ● But rapid prototyping with PCBs has drawbacks compared with the speed and ease of building a circuit on a breadboard. You can quickly make your own PCBs—as long as you don't mind the mess and some stained clothing and are willing to drill your own through holes. Or you can send your PCB layout to be made by a commercial service, but this takes several days at least and is more expensive. ● So I began thinking about practical alternatives for high-frequency circuits that can provide maker-friendly prototypes that are fast to build, and easy to probe and alter. In this article, I'll be presenting one key idea; some follow-on strategies will appear on the *IEEE Spectrum* website in the coming weeks. I should say that I make no claims of originality: Indeed I employ some oft-forgotten, decades-old techniques, but they turn out to be surprisingly useful in an age of surface-mount components operating at gigahertz frequencies.

The general approach is to start with a standard raw board, typically one made using FR4-grade resin, with its copper layer untouched. Instead of using etched traces, you interconnect components with lead wires while leaving a large ground. To demonstrate, I built a so-called comb generator circuit.

A comb generator produces a set of sharply defined harmonics across a wide range of frequencies, in this case up to 1 gigahertz, and it is a useful building block in microwave systems. The heart of my generator is a 74HC00 integrated circuit, which houses four NAND logic gates. A signal from a 25-megahertz surface-mount generator feeds two of the NAND gates in series so as to produce two square-wave signals that are slightly delayed. These signals go to a final NAND gate that generates narrow "sliver" pulses, which form the harmonic spectrum.

To create a circuit, I divided the copper layer into two lands. In this case, I wanted one small area along the top to serve as a 5-volt supply rail. Everything else forms the ground plane.

To isolate the lands from each other, I stripped off three thin rectangles of copper to form the boundary of the supply rail. I did this by marking parallel lines with a scriber. Then I held a steel ruler very firmly against the scribe marks and used a hobby knife to cut all the way through the copper along the length of the rule (this takes a fair amount of force, and often several passes). Then, using my soldering iron to heat the copper between the lines, I peeled each strip away using tweezers.

So how do you mount an integrated circuit on a board that's mostly a single ground plane with no through holes? You bend the IC's ground pins back so that they touch the surface, and solder them to the ground plane, holding the IC in place. You bend the

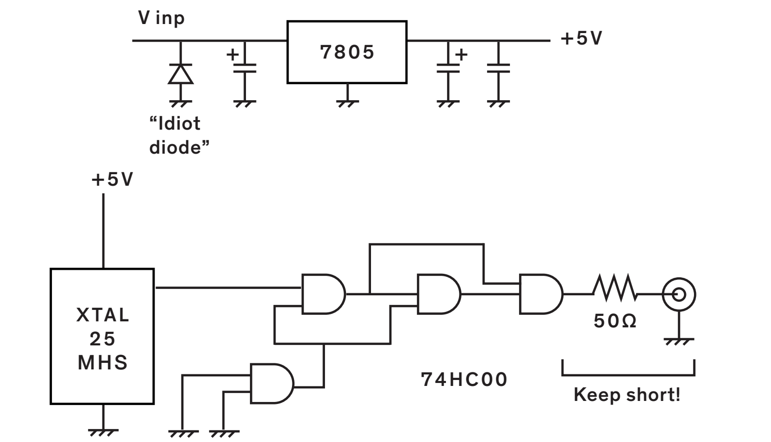

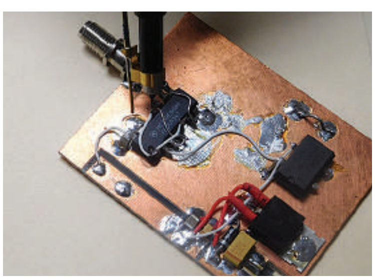

FIGURE AND GROUND: The comb generator uses NAND gates to create high-frequency pulses [top]. To implement the circuit, I remove copper from the board [middle] to create a power rail and ground plane [bottom].

other pins parallel to the board and solder connecting wires directly to them. Sometimes this is referred to as the "dead bug" method because of the way ICs look with their legs sticking out. As a bonus, the dead-bug method makes soldering surface-mount components easier than with a conventional PCB, as the contacts are more accessible. The ground plane also provided a convenient place to attach the heat sink of my comb generator's power regulator.

With a bit of practice cutting and peeling strips from the copper layer, it is possible to form an isolated land in the middle of a board to act as a tie point for both surface-mounted and through-hole components. Such an isolated land has very little capacitance to ground.

Another advantage of this construction technique is that you can easily check if your high-frequency circuit is in fact working as designed. Using a spectrum analyzer with a 500-ohm resistive probe (such as the Tektronix P6056) works well for such circuits, as long as the probe's shield is connected to ground close to the circuit node being probed. By attaching a spring-mounted Pogo pin to the probe's ground shield that extends down to the board, I'm able to make a ground connection right beside whatever pin I'm touching the probe to. (If you can't find a P6056 or similar probe, you can make one yourself using a 450-ohm resistor in series with a 50-ohm length of coaxial cable, but remember to use a 50-ohm terminator at the analyzer end.)

The results of these methods are not always the prettiest-looking boards, but I've had good success in using these techniques to prototype microwave circuits. I hope you'll follow me online to see some further elaborations of the ground-plane method, along with some other tips and tricks that I and some of my fellow circuit builders have been using for high-frequency circuits. —**ROBERT MELVILLE**

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/deadbug1018

TRITON SUBMARINES' DIVE INTO LUXURY

THESE MANEUVERABLE SUBS ARE PRIZED BY SCIENTISTS AND YACHTERS ALIKE

If you're looking for the perfect add-on to your megayacht, how about a personal submarine? Triton Submarines can set you up. The company, based in Vero Beach, Fla., specializes in high-end submersibles that can dive as far as 1,000 meters deep. Now, Triton has partnered with luxury carmaker Aston Martin, based in England, to build a limited-edition model. Due in early 2019, it combines ultimate style with hydrodynamic performance.

Over the past 12 years, Triton's subs have earned a reputation for safety, maneuverability, and comfort. But back in 2008, when the company was founded, the idea of a personal submersible was a tough sell. Too many potential buyers had seen too many Hollywood action movies featuring doomed submarines, recalls Triton's president, Patrick Lahey, who cofounded the company with CEO L. Bruce Jones.

"People thought [submarines] must be massively complicated and dangerous," Lahey says. "I'll forever be grateful to our first customer. Putting our sub on his vessel and having it displayed at boat shows really got the conversation start-

SUBMERGE IN STYLE: Triton's Project Neptune will bring luxury to dives beneath the seas.

ed." Today, Triton's preorders and word-of-mouth recommendations continue to propel the firm's growth.

Triton subs all feature spherical transparent cabins, which provide the widest possible window on ocean flora, fauna, and landforms while resisting the deep's crushing pressures. Figuring out how to build the cabins took some doing. In 2011, production of Triton's most popular model—the US $3.8 million three-person 3300/3—hit a wall when suppliers were unable to cast the 2.1-meter-diameter, 2.2-metric-ton acrylic bulb. "It actually threatened to take us out of business because we had a couple of orders that we couldn't fill," says Lahey. Triton turned to German acrylics pioneer Evonik Industries, which developed a more uniform thermal-forming process.

To shrink the profile of its subs, Triton engineers shifted from lead-acid batteries to fire-resistant lithium iron magnesium phosphate batteries. The lithium iron batteries pack 90 watt-hours per kilogram, more than double the energy density of lead batteries. The slimmed-down subs can be tucked inside a yacht's hangar, rather than sitting on deck. The batteries' fire-safe chemistry eased acceptance from international ship certification firms such as Norway's DNV GL. Triton subjects all of its subs to rigorous independent certification, comparable to what commercial aircraft go through. "We don't build experimental subs," says Lahey.

Codeveloping a sub with Aston Martin is about maxing out styling, creature comforts, and performance. The new sub features a more powerful quartet of thrusters and streamlined hydrodynamics, which will propel the vessel at a relatively snappy 6 knots (11 kilometers per hour). The thrusters also offer greater control when navigating through coral reefs and other tricky terrain, and for holding steady in strong currents.

"Jacques Cousteau had a great saying: Speed is the enemy of observation," says Lahey. "You don't pull up to the Louvre and put on your running shoes and sprint through the place. You stop and take the time to drink it in."

Bioluminescence expert Edith Widder says the maneuverability of Triton subs is already top-notch, likening it to flying a helicopter. In 2012 Widder, who's CEO of the Florida-based Ocean Research & Conservation Association, was on a dive campaign in Japan that used a Triton 3300/3 and scored the first-ever sighting of a giant squid in its habitat.

Widder says privately owned subs fill a need when government funding isn't forthcoming. "We're back to the time of the Medicis, where scientists get access through wealthy people," she says. These days, governments prefer to fund cheaper, deeper-diving unmanned underwater vehicles, but Widder says crewed submersibles remain unbeatable for venturing into the unknown. "When you know what you want to do, you can build a robot to do it. But when you're exploring, there's nothing more adaptable than a human."

Soon Triton may be taking yachters and scientists deeper. Lahey says a new design will push well beyond Triton's current 1,000-meter limit, with details to be revealed as early as mid-October.

—PETER FAIRLEY

Triton Submarines **Founded:** 2008 **Headquarters:** Vero Beach, Fla. **Employees:** 36 **Founders:** Patrick Lahey, L. Bruce Jones **Capital raised to date:** Organic growth through sales

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/triton1018

THE RADIO ELEGIST
AN ARTIST COMMEMORATES THE FADING GLORY OF SHORTWAVE

Some artists work in oils, say, or marble. Amanda Dawn Christie works in radio. Not radio in the sense of performing on air. But radio in the sense of the giant cultural and technological phenomenon that is broadcasting, and specifically shortwave broadcasting.

For decades, shortwave was the only way to reach a global audience in real time. Broadcasters such as the BBC World Service and Voice of America used it to project "soft power." But as the Internet grew, interest in shortwave diminished.

Christie's art draws from shortwave's history, representing it in sculpture, performance, photography, and film. Her focus is the life of the Radio Canada International (RCI) transmitter complex, located in Sackville, New Brunswick, near Christie's hometown. The transmitter was in operation from the 1940s until 2012. "Those towers were always just a part of the landscape that I grew up around," says Christie. It took a radio-building workshop to spark her interest: "I built a radio out of a toilet-paper tube.... I thought I did a *great* job because I picked up Italian radio. It turned out I did not—I was just really close to this international shortwave site."

She began talking to locals about the complex. "Some people would hear the radio in their sink, or their fridge.... I was jealous because my sink didn't play the radio," she says. Research led her to the rusty bolt effect, in which corroded metal acts as a radio diode, something that engineers normally strive to prevent. "I thought, 'My gosh, if there's instructions on how to stop this, I should be able to reverse engineer and create it.'"

This led to *The Marshland Radio Plumbing Project*, in which Christie attached a sink to a long loop of copper pipe, with a wonky

TRITON SUBMARINES

solder joint as the diode. She took it to the salt marsh surrounding the RCI transmitters and got to know the engineers on-site: "We looked at topographical maps to see where I could take the sink and at what time it would be likely to pick up a signal." Christie's sink radio turned out to be about 40 meters too short in the pipe department, but soon she was hearing all kinds of stories about RCI from neighbors. For example, "in the 1950s, there were two dairy farms, [and anytime there was a broadcast to Africa] their lights would come on, not all the way, but just sort of glow," she says. Christie began recording these stories, which led to her next major radio-inspired work, a 2016 film called *Spectres of Shortwave*. The film is made up mostly of the recordings played over footage of the RCI site's 13 masts.

Christie also began using contact microphones to record the sounds of the masts vibrating in the wind. These droning tones were filtered to create a musical scale. The masts were finally demolished in 2014, and at the end of *Spectres of Shortwave*, Christie uses these drones as an affecting chorus that gradually stills as each mast falls.

The unusual nature of *Spectres of Shortwave* has meant that Christie has had difficulty showing it. "It's too experimental for mainstream festivals, and it's too narrative for experimental festivals!" she says ruefully. (Readers in Montreal between 15 November 2018 and 26 January 2019 can catch screenings at the Dazibao art center.)

The tower sounds led to performances. For *Requiem for Radio: Pulse Decay*, Christie plays a theremin rigged to sound the different drones and trigger images of the masts. This solo piece was also incorporated into a complex, three-person 2017 performance called *Requiem for Radio: Full Quiet Flutter*. (See The Big Picture, in this issue, for details.)

Although the RCI site is gone, Christie's art is still focused on radio waves: "I just built a cello that has a loop antenna and AM transmitter instead of a resonating body. I'm looking forward to refining it and composing music for it."

—**STEPHEN CASS**

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/radio1018

ARTFUL AIRWAVES: Amanda Dawn Christie, pictured with her radio cello [top], is also the creator of a sink designed to pick up shortwave [right].

TOP: VINCENT LAFRANCE; BOTTOM: AMANDA DAWN CHRISTIE

ENGINEERS SAY "NO THANKS" TO SILICON VALLEY

SOME ARE TURNING DOWN COMPANIES, CITING ETHICAL CONCERNS

Anna Geiduschek usually has no time to respond to the recruitment emails that arrive in her in-box each week. But Geiduschek, a software engineer at Dropbox, recently made a point of turning down an Amazon Web Services recruiter by citing her personal opposition to Amazon's role in hosting another tech company's service used by U.S. government agents to target immigrants living in the country illegally for detention and deportation.

"I'm sure you're working on some very exciting technical problems over there at AWS, however, I would never consider working for Amazon until you drop your AWS contract with Palantir," Geiduschek wrote in her email response, which she shared on Twitter.

Tech companies such as Amazon, Facebook, Google, and Microsoft have faced growing internal unrest from employees who raise ethical concerns about how the companies deploy their high-tech services and products. Tech workers have signed open letters opposing Google's Project Maven contract with the U.S. military, Microsoft's contract with U.S. Immigration and Customs Enforcement (ICE), and Amazon's sale of facial-recognition technology to law enforcement.

That chorus of dissent is now growing louder as even outside engineers voice their concerns to recruiters working for those tech companies.

The protests of employees have proven persuasive because Silicon Valley firms compete fiercely to recruit and retain engineering talent. For example, Google's leadership sought to reassure employees by declaring it would not renew its Pentagon contract and issuing a set of ethical principles for future uses of Google-developed technologies.

By the same logic, engineers who are approached by tech recruiters also have leverage. "I might be a one-off example, but it could be different if Amazon gets a lot of people emailing them saying, 'Hey, I won't work for you because of this,' " Geiduschek says.

Jackie Luo, a software engineer at Square, took a similar stance with a tech recruiter who sought to interest her in a career with Google. The recruiter happened to contact Luo when she was reading about Google's plans to re-enter the Chinese market with a censored version of the company's Internet search engine.

"I won't be considering a job at Google now or in the future unless it seriously rethinks the way it does business by putting human rights before profit," Luo told the recruiter in a response email that she shared on Twitter. She described the China search engine plans as a "huge deal-breaker for me" and expressed other concerns about Google's Pentagon contract and the company's work environment for "women, underrepresented minorities, trans people, etc."

ATILGAN OZDIL/ANADOLU AGENCY/GETTY IMAGES

TARGETING TECH: In July, protestors in New York City demanded that Microsoft cancel contracts with U.S. Immigration and Customs Enforcement (ICE).

Individual engineers such as Luo and Geiduschek seem to be responding spontaneously to tech recruiters. But some tech employees have joined organized efforts, such as the #TechWontBuildIt movement spearheaded by the Tech Workers Coalition labor advocacy group.

It's likely no accident that underrepresented groups in the tech industry are among the more prominent voices speaking out through recruitment channels, Luo says. And their opinions may carry even more weight because of Silicon Valley's eagerness to recruit more diverse workforces, Geiduschek says.

Raising ethical concerns to recruiters isn't a new tactic for tech employees. "Many, many women pushed back at Uber at the peak of its transgressions," Luo says. She pointed to the example of Kelly Ellis, a software engineer at MailChimp, who described having "rejected Uber so hard in so many recruiting emails" in a tweet sent out in August 2017.

Crucially, Luo and Geiduschek both took firm but polite stances when responding to the tech recruiters who reached out to them.

"The recruiter has little to no control over what government contracts their company is accepting or how their company handles sexual harassment," Luo says. "But, still, they're the ones who can communicate to the company when they're struggling to find candidates...and then the company can then reassess if they need to make a change."

—JEREMY HSU

A version of this article appears in our View From the Valley blog.

POST YOUR COMMENTS at https://spectrum.ieee.org/nothanks1018

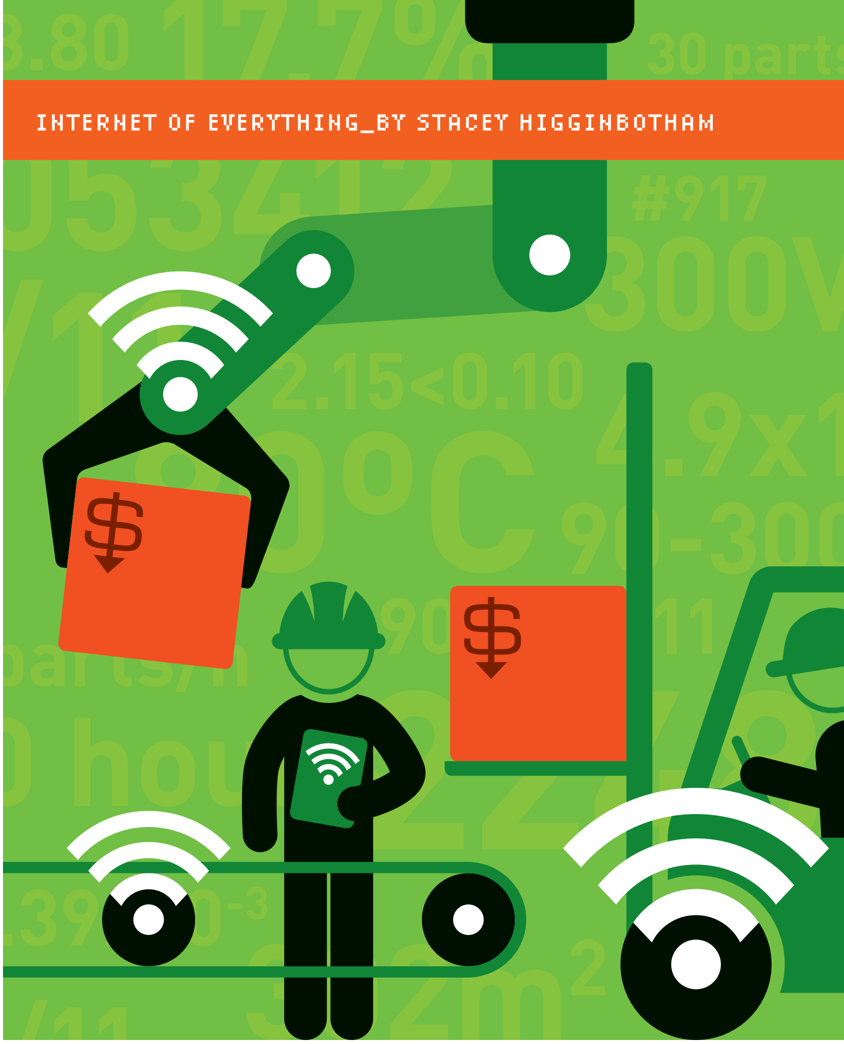

THE FRUGAL FACTORY

THE NEXT BIG EFFORT to reduce carbon emissions and hold the line on climate change will be enabled by the Internet of Things. Companies can rethink their costs of operations, taking into account the energy used, with a combination of more granular data from cheap sensors and faster, more in-depth analytics from cheap computing. ● At Schneider Electric's factory in Lexington, Ky., workers make electric components, including load centers and switches. The plant is four years into a company-mandated five-year goal to reduce energy consumption by 5 percent each year. The first two years, it achieved that goal. Then, management decided to deploy sensors and richer analytics to take a closer look at the mix of products made in the plant and the order in which those products were manufactured. They realized that by changing the production mix and order, they could save a lot of energy. ● How much? After tweaking the production mix, the plant reduced consumption by 12 percent in year three and 10 percent in year four. "The entire group had been focused on the processing side, and now every process decision is dictated by energy savings," says Andy Bennett, former senior vice president of Schneider Electric's EcoStruxure platform, which drove the factory innovations. "What has changed in the last five years is the technology and the drive and need to have a sustainable message." ● Bennett says that despite the United States pulling out of the Paris Agreement, many U.S. business leaders are still focused on reducing carbon emissions. Reduced energy consumption improves the bottom line, but it's something that manufacturers hadn't focused on because they didn't have the tools or impetus.

A similar shift in thinking played out in the 2000s in data centers, as companies such as Amazon, Facebook, and Google realized that power was a significant aspect of their costs of doing business. To address this, they prioritized the metric of performance per watt and forced Intel and Advanced Micro Devices, their suppliers, to focus on that metric.

The results were impressive. In 2016, a report by Lawrence Berkeley National Laboratory showed that energy consumption in those companies' data centers had remained flat, despite the growth in computing power, and had saved them roughly US $60 billion in energy costs each year.

Now it's the manufacturing world's turn. Schneider Electric isn't the only company using sensor data and artificial intelligence to optimize its energy conservation. This year at the Bosch ConnectedWorld IoT conference, in Berlin, the German industrial giant exhibited software that tracks the energy consumption of industrial processes and calculates how much energy they require.

Ikea, another huge European company, also tracks its energy use, and has gone a step further, calculating the energy used in production to determine whether it should make the product at all. The company, famous for its conservation efforts, looks at energy consumption used in manufacturing and the expected lifetime of a product. If something requires a lot of energy or can't be recycled effectively, it doesn't get made, according to Lena Pripp-Kovac, Ikea's sustainability manager.

Ikea may have an entire business unit dedicated to sustainability, but that's not feasible for many manufacturing companies. Fortunately for them, the Internet of Things will allow them to re-create similar analyses with much less fuss. And that will be good for all of us. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/frugalfactory1018

ILLUSTRATION BY Patrick George

OCTOBER 1958: FIRST BOEING 707 TO PARIS

DATING THE DAWN OF THE JET AGE is hard because there were so many different "firsts." The first experimental takeoff of a jet-powered airplane was that of the German Heinkel He 178, in August 1939. The first flight of the first commercial design, the British de Havilland DH 106 Comet, was in July 1949, and its first BOAC commercial flight was in 1952. The Comet, redesigned after catastrophic crashes, made the first transatlantic flight on 4 October 1958. Meanwhile, the Soviet Tupolev Tu-104 entered service in September 1956. ● But you can make a strong argument that the jet age began on 26 October 1958, when a Pan Am Boeing 707 took off from Idlewild Airport (now JFK International Airport) to Paris on the first of its daily scheduled flights. ● Several reasons justify that choice. The redesigned Comet was too small and unprofitable to begin a design dynasty, and there were no successor models. Tupolev's aircraft were used only by the countries of the Soviet bloc. ● But the Boeing 707 inaugurated the industry's most successful design family, one that progressed relentlessly by adding another 10 models to its varied lineup. The three-engine Boeing 727 was the first follow-on, in 1963; the four-engine 747, introduced in 1969, was perhaps the most revolutionary design in modern aviation; and the latest addition, the 787 Dreamliner series, introduced in 2011, is made mostly of carbon fiber composites and is now able to fly on routes longer than 17 hours. ● The 707 had a military pedigree: The plane started as a prototype of an in-air refueling tanker, and further development led to the KC-135A Stratotanker and finally to a four-engine passenger plane powered by small-diameter Pratt & Whitney turbojet engines, each with about 50 kilonewtons of thrust. By comparison, each of the two General Electric GEnx-1B high-bypass turbofan engines powering today's 787 delivers more than 300 kN at takeoff. ● The first scheduled flight of the 707 *Clipper America* on 26 October 1958 was preceded by a welcoming ceremony, a speech by Juan Trippe, Pan Am's president, and a performance by a U.S. Army band. The 111 passengers and 12 crew members had to make an unscheduled stop at Gander International Airport, in Newfoundland, Canada, but even so they were able to land at Paris-Le Bourget Airport 8 hours and 41 minutes after leaving New York. By December the plane was flying the New York-Miami route, and in January it began to make the first transcontinental flights, from New York to Los Angeles.

Before the introduction of the wide-bodies—first the Boeing 747 and then the McDonnell Douglas DC-10 and the Lockheed L-1011 in 1970—Boeing 707s were the dominant long-distance jetliners. One of them brought me from Europe to the United States in 1969.

Gradual improvements in the Boeing family resulted in a vastly superior plane. In a standard two-class (business and economy) configuration, the first Dreamliner could seat about 100 more people than the 707-120, with a maximum take-off weight nearly twice as great and a maximum range almost twice as long. Yet the Dreamliner consumes 70 percent less fuel per passenger-kilometer. And because it is built from carbon composites, the 787 can be pressurized to simulate a lower altitude than an aluminum fuselage will allow, resulting in greater comfort for passengers.

Eventually Boeing made just over 1,000 707s. When Pan Am brought the plane out of retirement for a 25th-anniversary commemorative flight, in 1983, it flew most of the original crew as passengers to Paris. But that was not the end of 707 service. A number of non-U.S. airlines flew different models until the 1990s, and Iran's Saha Airlines did so as late as 2013.

Although today the 707 can be found only in jet junkyards, the plane's place in history remains secure. It represents the first effective and rewarding step in the evolution of commercial jet flight. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/boeing707.1018

PHOTO-ILLUSTRATION BY Stuart Bradford

RED LIGHT, GREEN LIGHT—NO LIGHT

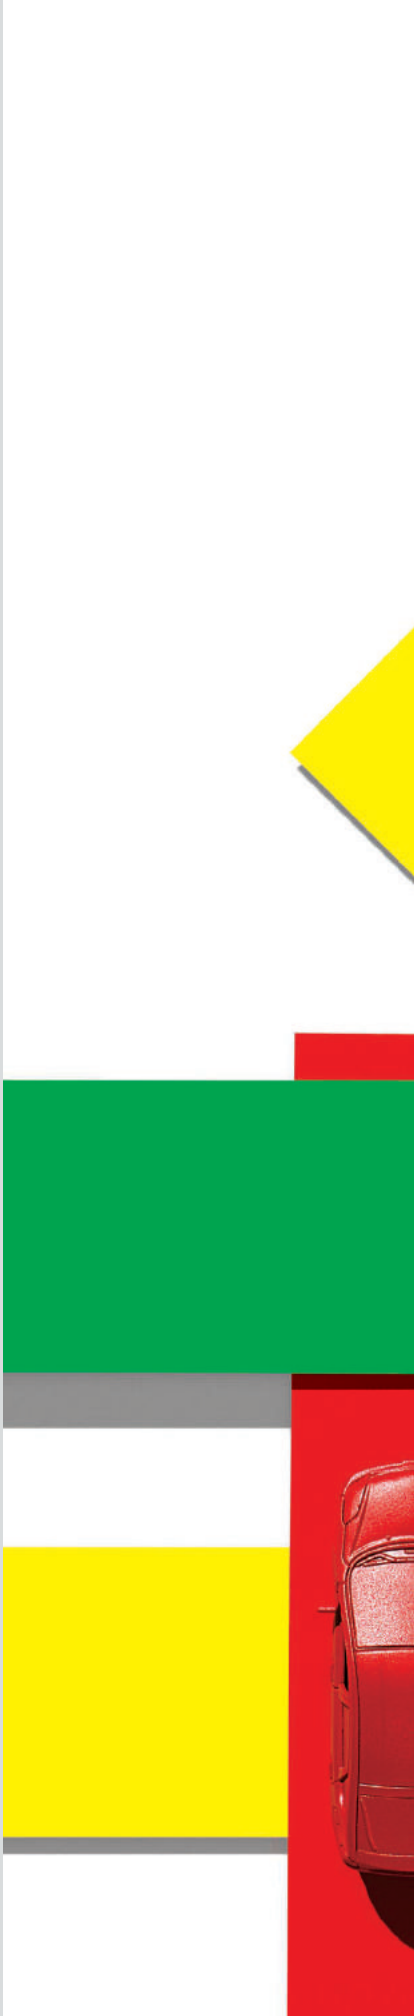

Tomorrow’s communicative cars could take turns at intersections

By Ozan K. Tonguz • • • Photography by Dan Saelinger

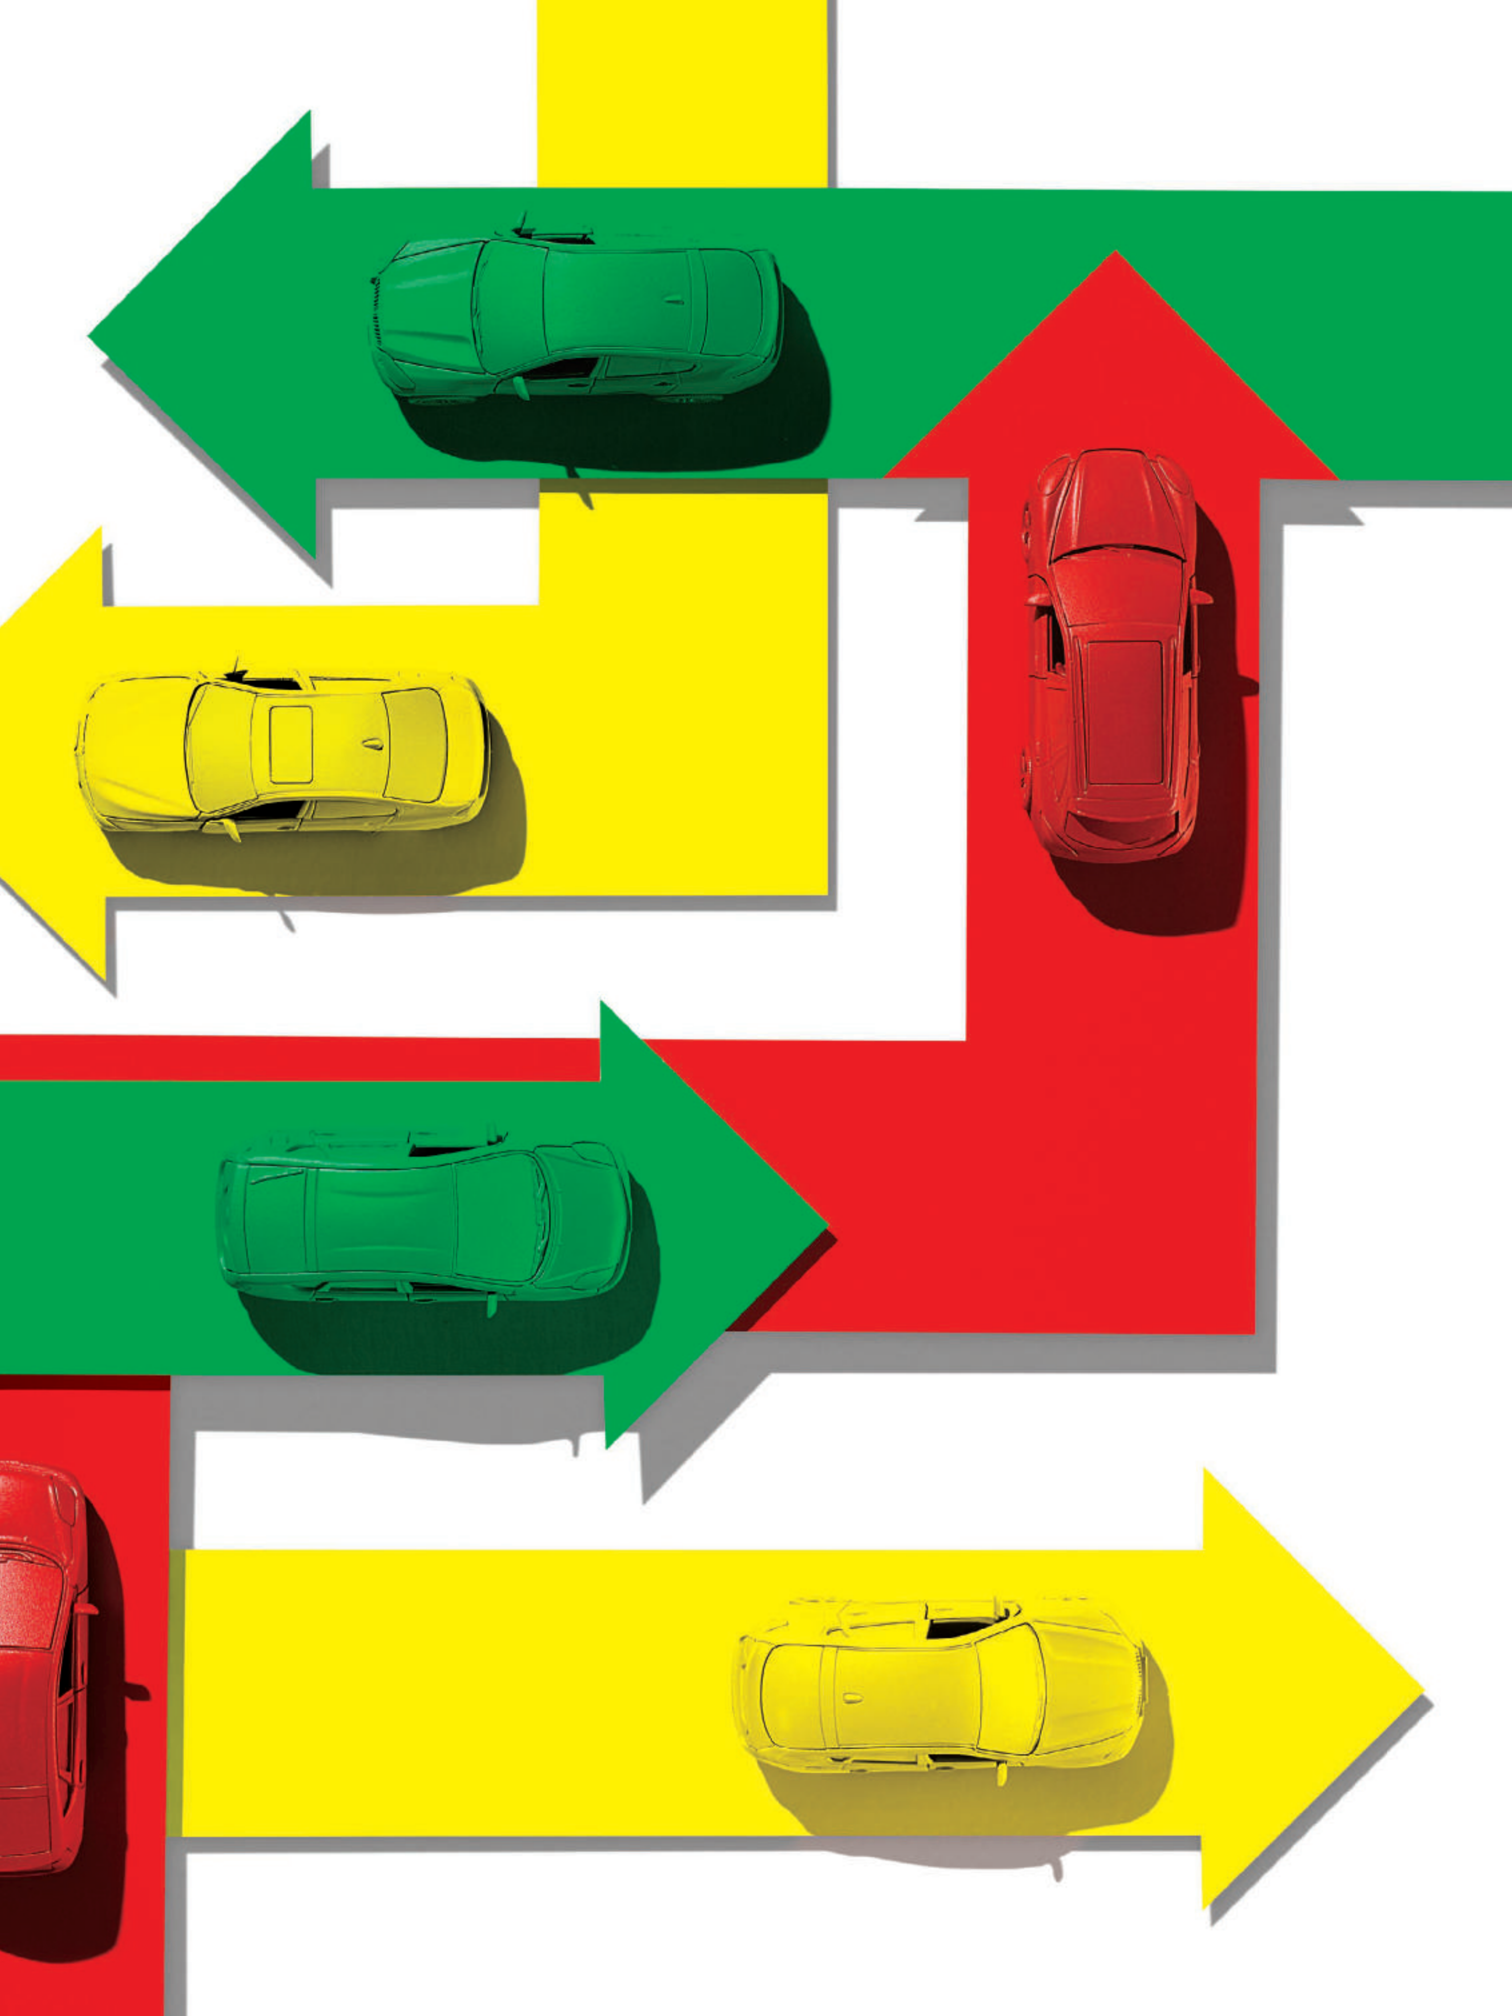

Life is short, and it seems shorter still when you're in a traffic jam. Or sitting at a red light when there's no cross traffic at all. In Mexico City, São Paolo, Rome, Moscow, Beijing, Cairo, and Nairobi, the morning commute can, for many exurbanites, exceed 2 hours. Include the evening commute and it is not unusual to spend 3 or 4 hours on the road every day. Now suppose we could develop a system that would reduce a two-way daily commute time by a third, say, from 3 to 2 hours a day. That's enough to save 22 hours a month, which over a 35-year career comes to more than 3 years.

Take heart, beleaguered commuters, because such a system has already been designed, based on several emerging technologies. One of them is the wireless linking of vehicles. It's often called vehicle-to-vehicle (V2V) technology, although this linking can also include road signals and other infrastructure. Another emerging technology is that of the autonomous vehicle, which by its nature should minimize commuting time (while making that time more productive into the bargain). Then there's the Internet of Things, which promises to connect not merely the world's 7 billion people but also another 30 billion sensors and gadgets.

All of these technologies can be made to work together with an algorithm my colleagues and I have developed at Carnegie Mellon University, in Pittsburgh. The algorithm allows cars to collaborate, using their onboard communications capabilities, to keep traffic flowing smoothly and safely without the use of any traffic lights whatsoever. We've spun the project out as a company, called Virtual Traffic Lights (VTL), and we've tested it extensively in simulations and, since May 2017, in a private project on roads near the Carnegie Mellon campus. In July, we demonstrated VTL technology in public for the first time, in Saudi Arabia, before an audience of about 100 scientists, government officials, and representatives of private companies.

The results of that trial confirmed what we had already strongly suspected: It is time to ditch the traffic light. We have nothing to lose except countless hours sitting in our cars while going nowhere.

THE PRINCIPLE BEHIND THE TRAFFIC LIGHT has hardly changed since the device was invented in 1912 and deployed in Salt Lake City, and two years later, in Cleveland. It works on a timer-based approach, which is why you sometimes find

VTL Algorithm: Letting Cars Control Their Own Traffic

Transceivers (using IEEE Standard 802.11p) send out a basic safety message every tenth of a second. The message tells recipients where the transmitting vehicle is by latitude, longitude, and heading.

Cars "Elect" a Leader—

1. Each vehicle computes its own distance to the intersection, the distance of the vehicles approaching the intersection from other directions, and each vehicle's speed, acceleration, and trajectory. Together they elect one vehicle to serve as the leader for a certain amount of time.

2. The leader vehicle decides which direction has the right-of-way (the equivalent of a green light) and which direction has the red light.

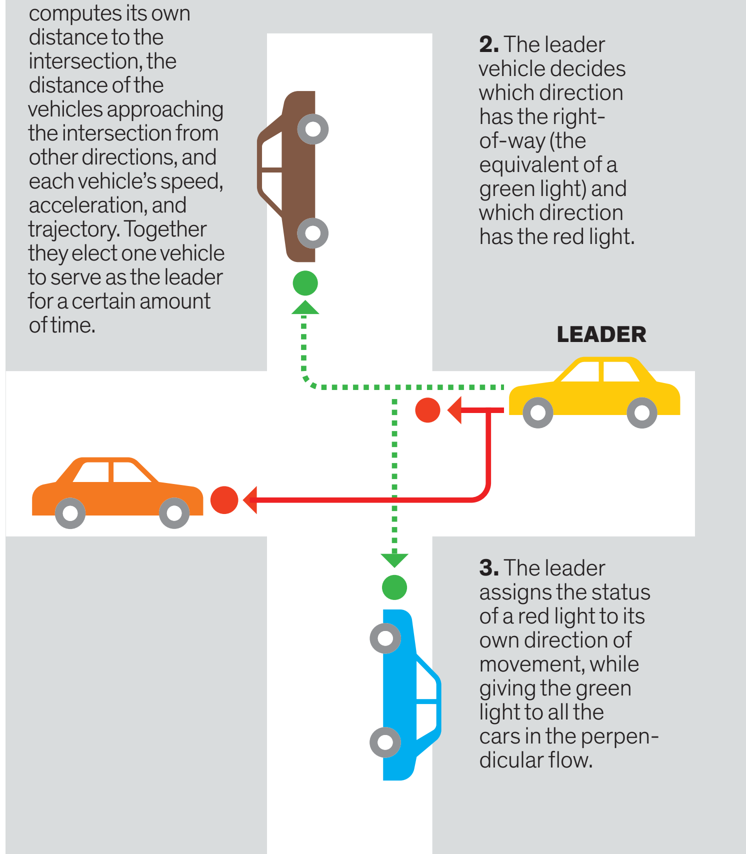

3. The leader assigns the status of a red light to its own direction of movement, while giving the green light to all the cars in the perpendicular flow.

ILLUSTRATIONS BY Anders Wenngren

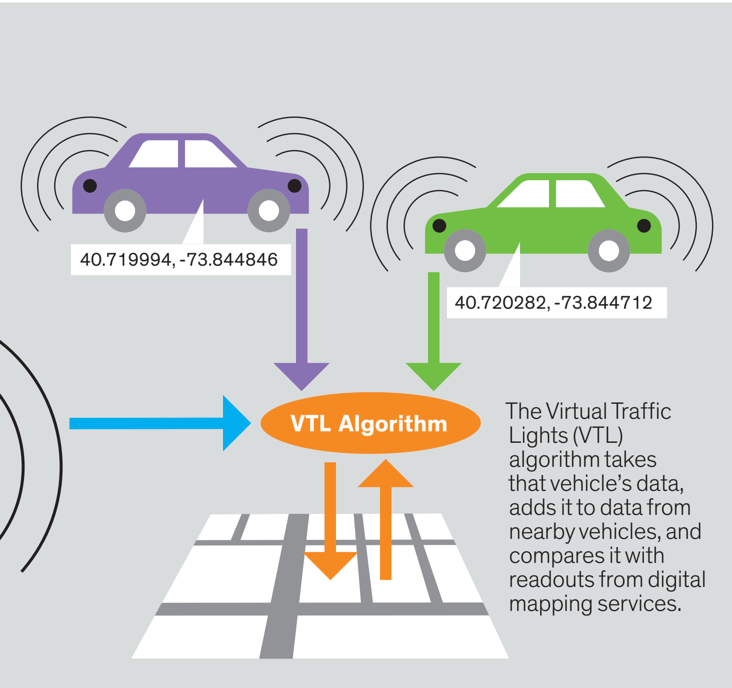

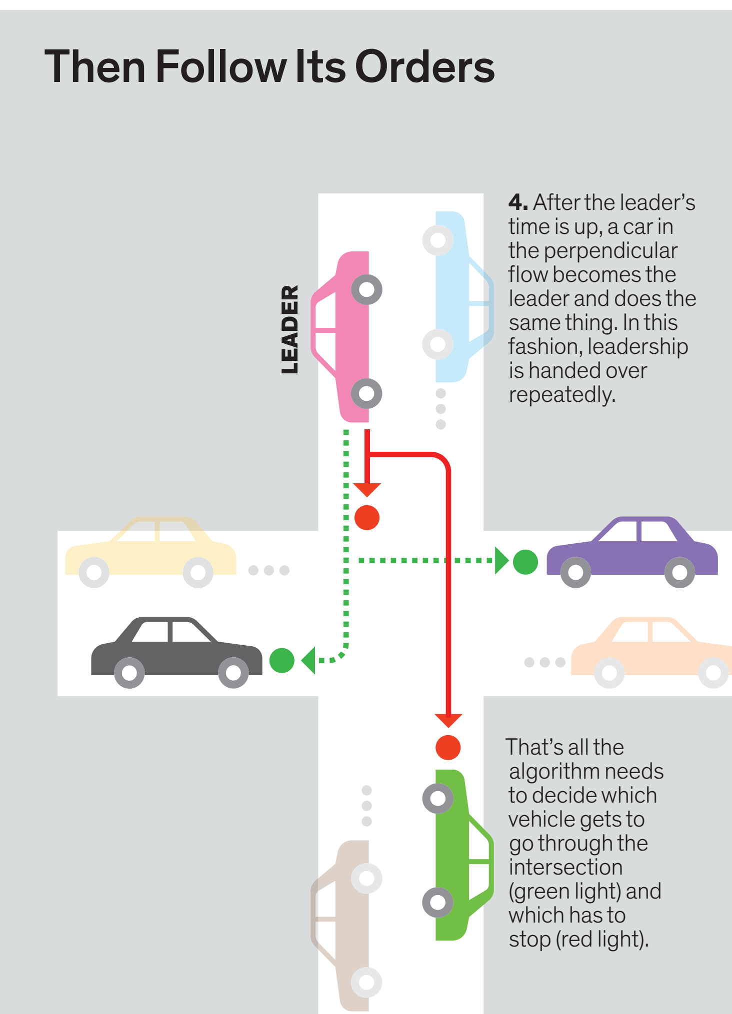

yourself sitting behind a red light at an intersection when there are no other cars in sight. The timing can be adjusted to match traffic patterns at different points in the commuting cycle, but that is about all the fine-tuning you can do, and it's not much. As a result, a lot of people waste a lot of time. Every day.

Instead, imagine a number of cars approaching an intersection and communicating among themselves with V2V technology. Together they vote, as it were, and then elect one vehicle to serve as the leader for a certain period, during which it decides which direction is to be yielded the right-of-way–the equivalent of a green light–and which direction has the red light.

So who has the right-of-way? It's very simple, and deferential. The leader assigns the status of a red light to its own direction of movement while giving the green light to all the cars in the perpendicular flow. After, say, 30 seconds, another car–in the perpendicular flow–becomes the leader and does the same thing. Thus, leadership is handed over repeatedly, in a round-robin fashion, to fairly share the responsibility and burden–because being the leader does involve sacrificing immediate self-interest for the common good.

With this approach, there is no need at all for traffic lights. The work of regulating traffic melts invisibly into the wireless infrastructure. You would never find yourself sitting at a red light when there was no cross traffic to contend with.

Our company's VTL algorithm elects leaders by consulting such parameters as the distance of the front vehicle in each approach from the center of the intersection, the vehicles' speed, the number of vehicles in each approach, and so on. When all else is equal, the algorithm elects the vehicle that's farthest from the intersection, so it will have ample time to decelerate. This policy ensures that the vehicle that's closest to the intersection gets the right-of-way–that is, the virtual green light.

It's important to note that VTL technology needs no camera, radar, or lidar. It gets all the orientation it needs from a wireless system called dedicated short-range communications. DSRC refers to radio schemes, including dedicated bandwidth, that were developed in the United States, Europe, and Japan between 1999 and 2008 to let nearby cars communicate wirelessly. DSRC developers envisioned various uses, including electronic toll collection and cooperative adaptive cruise control–and also precisely the function we are using it for, intersection collision avoidance.

Very few production cars are now equipped with DSRC transceivers (and it's possible that emerging 5G wireless technology will supersede DSRC). But such transceivers are readily available, and they provide all the functionality we need. These transceivers, designed to make use of IEEE Standard 802.11p, must each send out a basic safety message every tenth of a second. The message tells recipients where the transmitting vehicle is by latitude, longitude, and heading. Running on a processor in a vehicle, our VTL algorithm takes the data from that vehicle, throws in whatever it is receiving from neighboring vehicles, and overlays the result onto readouts from such digital mapping services as Google Maps, Apple Maps, or OpenStreetMap. In this way, each vehicle can compute its own distance to the intersection as well as the distance of the vehicles approaching the intersection from the other directions. It can also compute each vehicle's speed, acceleration, and trajectory. That's all the algorithm needs to decide who gets to go through the intersection (green light) and who has to stop (red light). And once the decision has been made, a head-up display in each car displays the light to the driver from a normal viewing position. Of course, the VTL algorithm solves only the problem of managing traffic at intersections, stop signs, and yield signs. It doesn't drive the car. But when functioning within its proper domain, VTL can do everything at a much lower cost than autonomous vehicle technology can. Self-driving cars require far more computing capability just to make sense of the individual data feeds coming from their lidar, radar, cameras, and other sensors, and more still to fuse those feeds into a single view of the surroundings.

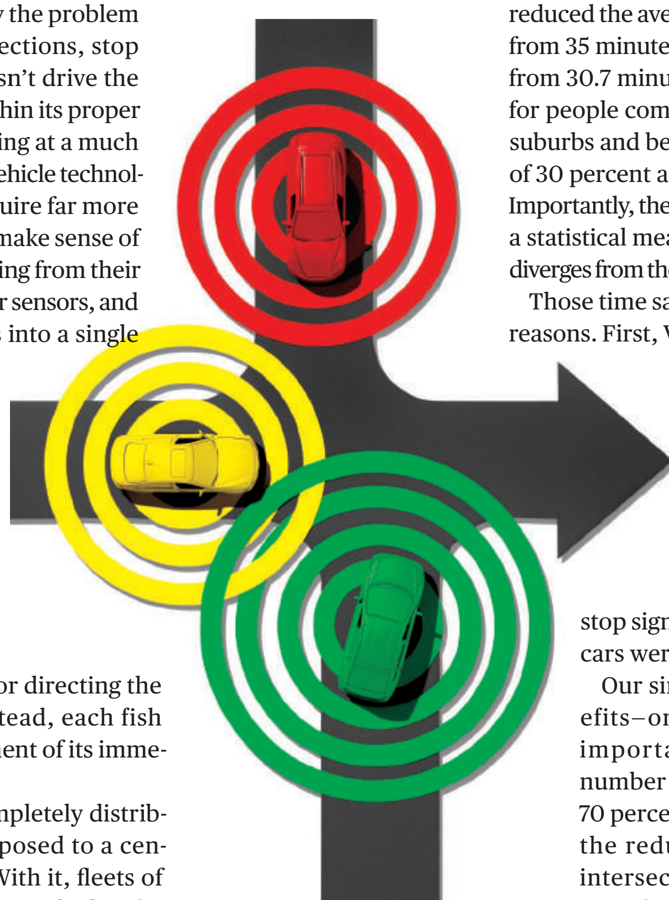

Think of our method as the substitution of a rule of thumb for true intelligence. The VTL algorithm lets the cars control their own traffic much as colonies of insects and schools of fish do. A school of fish shifts direction all at once, without any master conductor directing the members of the school; instead, each fish takes its cue from the movement of its immediate neighbors.

This is an example of a completely distributed system behavior as opposed to a centralized network behavior. With it, fleets of vehicles in a city can manage traffic flow by themselves without a centralized control mechanism and without any human intervention—no police, no traffic lights, no stop signs, and no yield signs.

• • •

We didn't invent the concept of intelligent intersections, which dates back decades. One early idea was to place a magnetic coil under the asphalt surface of a road to detect the approach of vehicles along a single route to an intersection and then adjust the duration of the green and red phases accordingly. Similarly, cameras placed at intersections can be used to count the vehicles in each approach and compute how best to time the lights at an intersection. But both technologies are expensive to install and maintain and therefore only a few intersections have been fitted out with them.

We started by running our VTL algorithm on a virtual model for two cities: Pittsburgh and Porto, Portugal. We took traffic data from the U.S. Census Bureau and the corresponding Portuguese agency, added map data from Google Maps, and fed it all into SUMO, the Simulation of Urban Mobility, an open-source software package developed by the German Aerospace Center.

SUMO simulated the rush-hour commuting time under two scenarios, one using the existing traffic lights, the other using our VTL algorithm. It found that VTL reduced the average commute to 21.3 minutes from 35 minutes in Porto and to 18.3 minutes from 30.7 minutes in Pittsburgh. Reductions for people commuting into the city from the suburbs and beyond were cut by a minimum of 30 percent and a maximum of 60 percent. Importantly, the variance of the commute time—a statistical measure of how much a quantity diverges from the mean value—was also reduced.

Those time savings came primarily for two reasons. First, VTL eliminated the time spent waiting at a red light when there were no cars crossing at right angles. Second, VTL introduced traffic control to every intersection, not just those that have active signals. So it was not necessary for cars to stop at a stop sign, for example, when no other cars were around.

Our simulations showed other benefits—ones that are arguably more important than saving time. The number of accidents was reduced by 70 percent, and—no surprise—most of the reduction was centered at the intersections, stop signs, and other interchanges. Also, by minimizing the time spent dawdling at intersections and accelerating and decelerating, VTL measurably reduces the average car's carbon footprint.

So, what would it take to get VTL technology out of the lab and into the world? To begin with, we'd have to get DSRC into production cars. In 2014, the U.S. National Highway Traffic Safety Administration proposed the adoption of the technology, but the Trump administration hasn't yet implemented the regulation, and it's not clear what the final decision will be. So U.S. manufacturers may now be reluctant to install DSRC transceivers, given that they'd add cost to a car and they'd be useful only if other cars carry them, too—the familiar chicken-and-egg problem. And until enough cars begin to carry the devices, the scale of manufacturing will remain low and the unit cost high. In the United States, only General Motors has begun to put DSRC radios into cars, all of them high-end Cadillacs. However, in Europe and Japan the outlook is a lot more favorable. A number of European automakers have committed to putting the transceivers in their

PHOTOGRAPH BY Dan Saelinger

cars, and earlier this year in Japan, where the government strongly supports the technology, auto giant Toyota reiterated its commitment.

And even if DSRC fails entirely, our VTL algorithm can be implemented with other wireless technologies, such as 5G or Wi-Fi.

The concept of incomplete penetration of DSRC transceivers brings up one of the biggest potential obstacles to adoption of our VTL technology. Could it still work even if only a certain percentage of vehicles is equipped with DSRC? The answer is yes, provided that governments equip existing traffic signals with DSRC technology.

Governments may well be willing to do that, if only because they would rather not do away with hundreds of billions of dollars' worth of existing signal infrastructure. To address this problem, we've fitted out our Virtual Traffic Lights technology with a short-term solution: We can upgrade existing traffic lights so that they can detect the presence of DSRC-equipped vehicles in each approach and decide the green-red phases accordingly. The beauty of this scheme is that all vehicles could make use of the same roads and intersections, whether or not they are equipped with DSRC. This approach may not reduce commute time as much as the ideal VTL solution, but even so it is at least 23 percent better than the current traffic control systems, according to both our simulations and to field trials in Pittsburgh.

VTL REDUCED THE AVERAGE COMMUTE TO 21.3 MINUTES FROM 35 MINUTES IN PORTO AND TO 18.3 MINUTES FROM 30.7 MINUTES IN PITTSBURGH

Yet another challenge is how to handle pedestrians and bicyclists. Even in a regime mandating DSRC transceivers for all cars and trucks, we couldn't reasonably expect cyclists to install the devices or pedestrians to carry them. That might make it hard for those people to cross busy intersections safely.

Our solution for the short term, while physical traffic signals still coexist with the VTL system, is to provide pedestrians a way to give themselves the right-of-way. Ever since January of this year, our pilot program in Pittsburgh has provided a button to push that actuates a red light—real for the pedestrians, and virtual for the cars—at all four approaches to the intersection. It has worked every time.

In the longer term, the bicyclist and pedestrian challenge might be solved with Internet of Things technology. As the IoT expands, the day will finally come when everyone carries a DSRC-capable device at all times.

Meanwhile, under ideal conditions, with no physical signals at all, we have demonstrated that the vehicles voting on how to assign right-of-way can allot a portion of the signaling cycle to pedestrians. During these interludes, a virtual red light shines in all the vehicles at all four approaches, lasting long enough for any pedestrians there to cross safely. This preliminary solution wouldn't be optimal for traffic flow, and so we are also working on a method using cheap dashboard-mounted cameras to spot pedestrians and give them the right-of-way.

• • •

ULTIMATELY, WHAT MAKES VIRTUAL TRAFFIC SIGNALS so promising is the advent of self-driving vehicles. As envisioned today, such vehicles would do everything human drivers now do—stopping at traffic lights, yielding at yield signs, and so forth. But why automate transportation halfway? It would be far better to make such vehicles fully autonomous, managing traffic without any conventional signs or signals. The key in achieving this goal is V2V and vehicle-to-infrastructure communications.

This matters because today's self-driving cars are often unable to negotiate their way into and out of busy intersections. This is one of the hardest technical problems, and it continues to challenge even industry leader Waymo (a subsidiary of Google's parent company, Alphabet).

In our simulations and field trials, we have found that autonomous vehicles equipped with VTL can manage intersections without traffic lights or signs. Not needing to identify such objects greatly simplifies the computer-vision algorithms that today's experimental self-driving cars rely on as well as the computational hardware that runs those algorithms. These elements, together with the sensors (especially lidar), constitute the single costliest part of the package.

Because VTL has a largely modular software architecture, it would be easy to integrate it into the rest of an autonomous car's software. Furthermore, VTL can solve most, if not all, of the hard problems related to computer vision—say, when the sun shines straight into a camera, or when rain, snow, sandstorms, or a curving road obscure the view. To be clear, VTL is not really competing with the technology of self-driving cars; it is complementing it. And that alone would help to speed up the robocar rollout.

Well before then, we hope to have our system up and running for human-driven cars. Just this past July we staged our first public demonstration, in Riyadh, Saudi Arabia, in heat topping 43 °C (100 °F), with devices installed in the test vehicles. Representatives from government, academia, and corporations—including Uber—boarded a Mercedes-Benz bus and drove through the campus of the King Abdulaziz City for Science and Technology, crossing three intersections, two of which had no traffic lights. The bus, together with a GMC truck, Hyundai SUV, and a Citroën car, engaged the intersections in every possible way, and the VTL system worked every time. When one driver deliberately disobeyed the virtual red light and attempted to cross, our safety feature kicked in right away, setting off a flashing red light for all four approaches, heading off an accident.

I hope and believe that this was a turning point in transportation. Traffic lights have had their day. Indeed, they lasted over a century. Now it's time to move on. ■

↗ POST YOUR COMMENTS at https://spectrum.ieee.org/v2v1018


ST • I 0.3
ST • II 0.6
ST • III 0.2
ST • AVR -0.4
ST • AVL 0.0
10:41AM
10:42AM
120
50
93
SPO2
100
90
65
RR
30
8
23
PULSE 127
TEMP
42.2
42.6
(126)
160
90
128/76
97/73
(79)
RR
30
8
40

AI IN THE ICU

IN THE INTENSIVE CARE UNIT, ARTIFICIAL INTELLIGENCE CAN KEEP WATCH

BY BEHNOOD GHOLAMI, WASSIM M. HADDAD & JAMES M. BAILEY

ILLUSTRATIONS BY MCKIBILLO

IN A HOSPITAL'S INTENSIVE CARE UNIT (ICU), the sickest patients receive round-the-clock care as they lie in beds with their bodies connected to a bevy of surrounding machines. This advanced medical equipment is designed to keep an ailing person alive. Intravenous fluids drip into the bloodstream, while mechanical ventilators push air into the lungs. Sensors attached to the body track heart rate, blood pressure, and other vital signs, while bedside monitors graph the data in undulating lines. When the machines record measurements that are outside of normal parameters, beeps and alarms ring out to alert the medical staff to potential problems. ● While this scene is laden with high tech, the technology isn't being used to best advantage. Each machine is monitoring a discrete part of the body, but the machines aren't working in concert. The rich streams of data aren't being captured or analyzed. And it's impossible for the ICU team—critical-care physicians, nurses, respiratory therapists, pharmacists, and other specialists—to keep watch at every patient's bedside. ● The ICU of the future will make far better use of its machines and the continuous streams of data they generate. Monitors won't work in isolation, but instead will pool their information to present a comprehensive picture of the patient's health to doctors. And that information will also flow to artificial intelligence (AI) systems, which will autonomously adjust equipment settings to keep the patient in optimal condition.

At our company, Autonomous Healthcare, based in Hoboken, N.J., we're designing and building some of the first AI systems for the ICU. These technologies are intended to provide vigilant and nuanced care, as if an expert were at the patient's bedside every second, carefully calibrating treatment. Such systems could relieve the burden on the overtaxed staff in critical-care units. What's more, if the technology helps patients get out of the ICU sooner, it could bring down the skyrocketing costs of health care. We're focusing initially on hospitals in the United States, but our technology could be useful all around the world as populations age and the prevalence of chronic diseases grows.

The benefits could be huge. In the United States, ICUs are among the most expensive components of the health care system. About 55,000 patients are cared for in an ICU every day, with the typical daily cost ranging from US \$3,000 to \$10,000. The cumulative cost is more than \$80 billion per year.

As baby boomers reach old age, ICUs are becoming increasingly important. Today, more than half of ICU patients in the United States are over the age of 65—a demographic group that's expected to grow from 46 million in 2014 to 74 million by 2030. Similar trends in Europe and Asia make this a worldwide problem. To meet the growing demand for acute clinical care, ICUs will need to increase their capacity as well as their capabilities. Training more critical-care specialists is part of the solution—but so is automation. Far from replacing humans, AI systems could become part of the medical team, allowing doctors and nurses to deploy their skills when and where they're needed most.

IN ICUs TODAY, the data from the raft of bedside monitors is usually lost as the monitor screens refresh every few seconds. While some advanced ICUs are now trying to archive these measurements, they still struggle to mine the data for clinical insights.

A human doctor typically has neither the time nor the tools to make sense of the rapidly accumulating data. But an AI system does. It could also take actions based on the data, such as adjusting the machines involved in crucial ICU tasks. At Autonomous Healthcare, we're focusing first on AI systems that could manage a patient's ventilation and fluids. Mechanical ventilators come into play when a patient is sedated or suffers lung failure, a common ICU condition. And careful fluid management maintains the proper volume of blood flowing through a patient's circulatory system, therefore ensuring that all the tissues and organs get enough oxygen.

Our methodologies spring from an unlikely source: the aerospace industry. Two of us, Haddad and Gholami, are aerospace control engineers. We met at Georgia Tech's school of aerospace engineering, where Haddad is a professor of dynamical systems and control and Gholami formerly worked as a doctoral researcher. Bailey joined the collaboration in the early 2000s when he was an associate professor of anesthesiology at the Emory University school of medicine. Haddad and Bailey first worked on control methods to automate anesthesia dosing and delivery in the operating room, which we tested in clinical studies at Emory University Hospital, in Atlanta, and Northeast Georgia Medical Center, in Gainesville, Ga. We then set our sights on more complex and broader control problems for the ICU. In 2013, Haddad and Gholami founded Autonomous Healthcare to commercialize our AI systems. Gholami is the company's CEO, Haddad is chief science advisor, and Bailey is chief medical officer.

How is aerospace similar to medicine? Both fields involve vast amounts of data that must be processed quickly to make decisions while lives hang in the balance, and both require that many tasks be performed simultaneously to keep things running smoothly. In particular, we see a role for feedback control technology in critical-care medicine. These technologies use algorithms and feedback to modify the behavior of engineered systems through sensing, computation, and actuation. They have become ubiquitous in the safety-critical systems of flight control and air traffic control.

However, there's a key difference between an airplane and a hospital patient. An airplane's design and control is based on well-established theories of mechanics and aerodynamics, whereas the human body involves highly complex biological systems that function and interact in ways we don't yet entirely understand.

Consider the management of mechanical ventilation. ICU patients may require this support because of direct trauma, lung infection, heart failure, or an inflammatory syndrome such as sepsis. The ventilator alternates between forcing air into the lungs and allowing the lungs to passively deflate. The device can be dialed up or down to do all of the work or to assist the patient's own efforts.

The interaction between human and machine is a subtle thing to manage. The human body has its own automatic mechanism to govern breathing, in which the nervous system triggers the diaphragm muscle to contract and pull downward on the lungs, thus initiating the intake of air. The ventilator must work with this innate drive; it should be synchronized with the patient's natural transitions between inhaling and exhaling, and it should match the natural air volume of the patient's breathing.

Unfortunately, mismatches between the patient's demand and the machine's delivery are all too common, which can cause a patient to "fight the ventilator." For example, a patient may naturally need more time to inhale, but the ventilator transitions to the exhalation prematurely. This and other synchronization problems with mechanical ventilation are associated with longer stints on the ventilator, longer stays in the ICU, and increased risk of death. Experts don't yet know why asynchrony has these detrimental effects, but patients clearly experience discomfort when trying to breathe out while the machine is

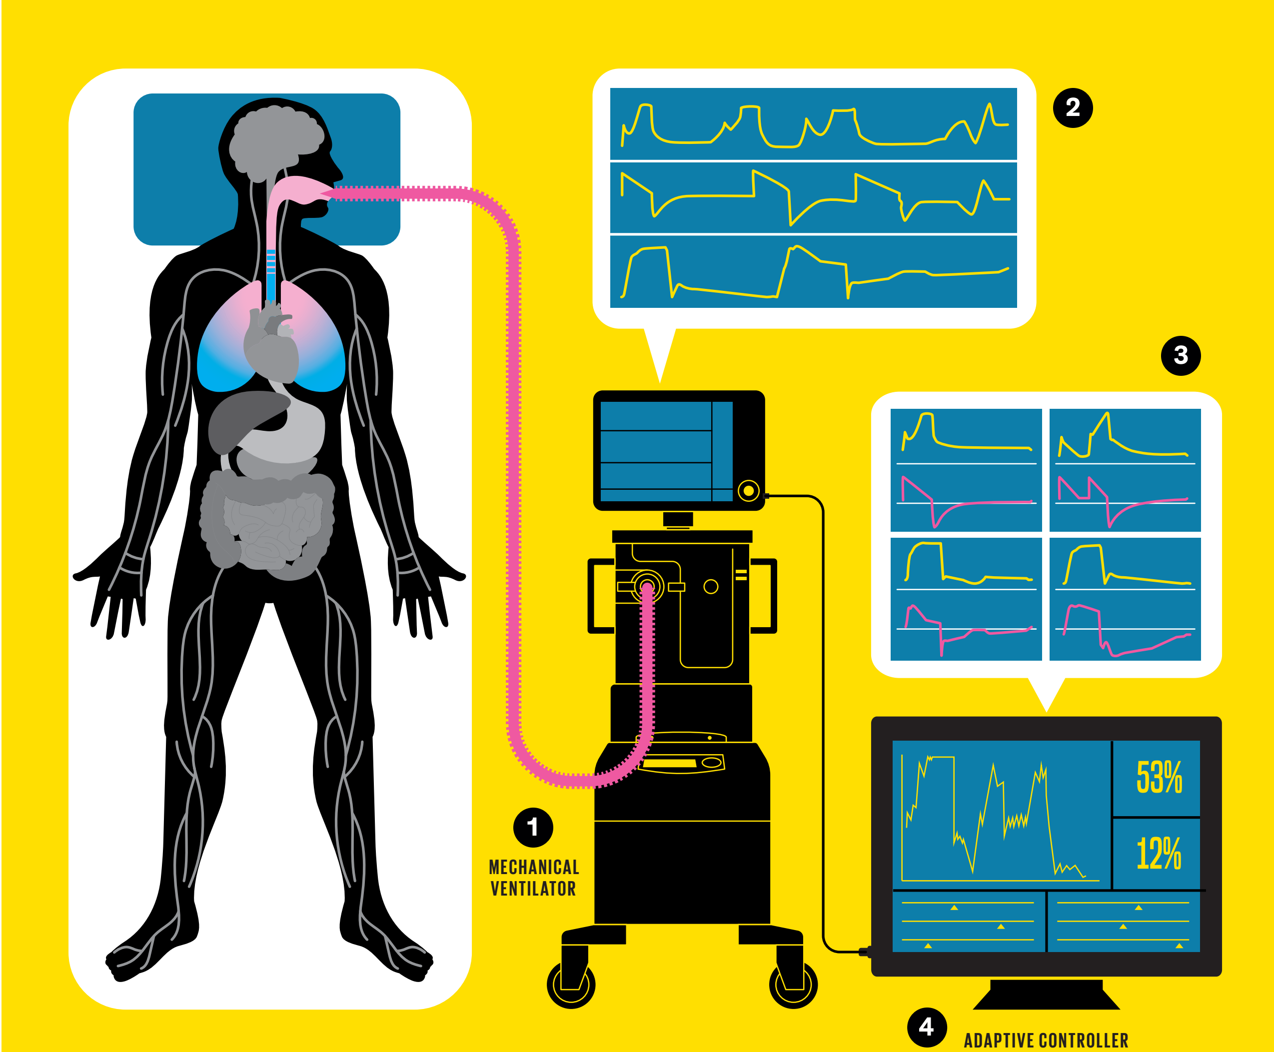

BREATHING EASIER

CRITICALLY ILL PATIENTS who need help breathing are put on mechanical ventilators [1]. These machines push air into the lungs, but the rhythm can get out of sync with natural breathing patterns, causing patients to "fight the ventilator." A smart control system could read airflow measurements [2] and identify different types of ventilator asynchrony [3] in real time via a machine-learning algorithm. In a fully autonomous system, an adaptive controller [4] would constantly adjust the ventilator's airflow to keep it in sync with the patient. As a step toward the goal of full autonomy, a similar system could be used as a decision-support tool in the ICU, providing recommendations that respiratory therapists could use to make adjustments.

pushing air into their lungs, and their laboring muscles experience an additional workload. In ICUs in the United States, the share of patients on ventilators who experience severe asynchrony has been estimated to be between 12 and 43 percent.

The first step in addressing this problem is to detect it. Experienced respiratory therapists can identify different types of asynchrony if they continuously monitor the waveforms on a ventilator's display screen indicating the pressure and flow. But in an ICU, one respiratory therapist typically oversees 10 or more patients and can't possibly monitor all of them constantly.

At our company, we've designed a machine-learning framework that replicates that human expertise in detecting different types of asynchrony. To train our system, we used a data set of waveforms from patients on ventilators, in which each waveform had been evaluated by a panel of clinical experts. Our algorithm learned the signatures of different asynchrony types—such as a particular dip in the flow signal at a specific point in time. In our first assessments of the algorithm's performance, we focused on what's called cycling asynchrony, which is the most challenging type to detect.

ILLUSTRATION BY MCKIBILLO

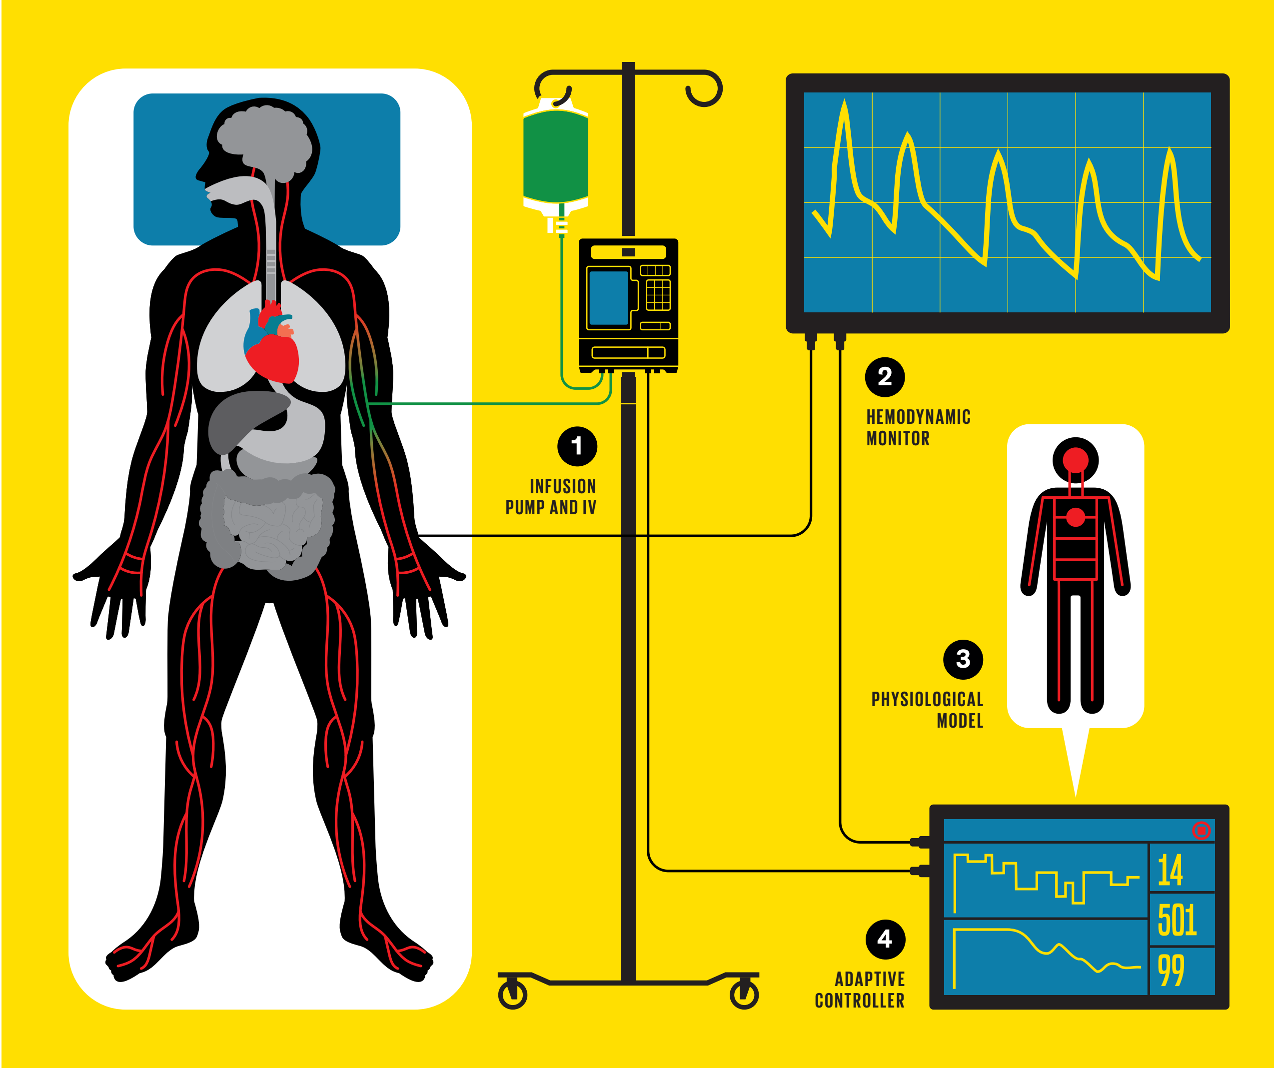

FLUID MOVEMENTS **MOST ICU PATIENTS** require infusion pumps and IVs [1] to drip fluid into their veins. Getting the fluid volume right is crucial: If levels are either too low or too high in the circulatory system, serious complications can arise. A smart control system could track real-time measurements [2] such as arterial blood pressure and the amount of blood pumped by the heart; the system could then feed the data into a physiological model [3] that represents how fluids move through the body's blood vessels and tissues. In a fully autonomous system, an adaptive controller [4] could continuously adjust fluid inputs to keep the patient stable. Initially, ICU physicians could use the technology as a decision-support system that provides recommendations.

Here the ventilator's initiation of the exhale doesn't match the patient's own exhalation. The accuracy of our algorithm in detecting cycling asynchrony in a new data set matched that of human experts.

We're now testing the algorithm at Northeast Georgia Medical Center's ICU to detect respiratory asynchrony in real patients and in real time. The technology has been incorporated into a clinical-decision support system, which is designed to help respiratory therapists assess a patient's needs. This framework can also provide researchers with a tool to better understand the underlying causes of asynchrony and its impact on patients. Our long-term goal is to design mechanical ventilators that can automatically adjust their own settings in response to each patient's needs.

WHEN YOU PICTURE AN ICU, your mental image probably includes patients with plastic bags hanging from stands by their bedsides, fluids continually dripping into their veins through IVs. About 75 percent of patients require such fluid management at some point during their stay in the ICU.

However, calibrating the correct amount of fluid is far from an exact science. Just tracking a patient's fluid levels is a hard

ILLUSTRATION BY MCKIBILLO

task: No existing medical sensors can directly monitor fluid volume, so doctors rely on indirect indicators like blood pressure and urine volume. The amount of fluids that patients need depends on their illness and medications, among other things.

Getting the fluids right is particularly important for patients with sepsis, a life-threatening syndrome characterized by inflammation throughout the body. In these patients, the blood vessels dilate, thus reducing blood pressure, and fluid leaks from the tiniest vessels, the capillaries. As a result, less oxygen-carrying blood reaches the organs, which can cause organs to fail and patients to die. Doctors combat sepsis by dispensing drugs to boost blood pressure and pumping extra fluids into patients' circulatory systems.

It's important to add enough fluid, but not too much—an excess can cause complications such as pulmonary edema, a buildup of fluid in the lungs that can interfere with breathing. Studies have shown that fluid overload is associated with more days on mechanical ventilators, longer stays in the hospital, and higher rates of mortality.

Doctors therefore aim to maintain their patients' fluids at certain levels, which are based on models for an average patient. When the doctors come through the ICU on their rounds, they try to determine whether the patient is holding steady at the goal level by checking the mix of gases in the blood and monitoring blood pressure and urine output. Deciding when to add fluids and how much to add is highly subjective, and there's considerable debate about the best practices.

An AI system could do better. Rather than basing its decisions on goals established for an average patient, it could analyze a wide variety of physiological indicators for an individual patient in real time, and continuously dispense fluids according to that patient's specific needs.

At Autonomous Healthcare, we've developed a fully automated system that looks at indirect measurements of a patient's fluid levels (such as blood pressure and variation in the volume of blood pumped out by each heartbeat) and then feeds the data into a sophisticated physiological model. Our system uses these measurements to assess how fluids are moving between the body's blood vessels and tissues, constantly adjusting the parameters as new measurements come in. Our proprietary adaptive controller then adjusts the fluid-infusion settings accordingly.

One advantage of our technology is its attention to what control engineers call closed-loop system stability, which means that any perturbations to a normal state lead to only small and fleeting variations. Many engineering applications use control systems that guarantee closed-loop stability—when an airplane runs into powerful turbulence, for instance, the autopilot system compensates to keep the shaking to a minimum. However, most control systems for medical devices have no such guarantee. If doctors judge that a sepsis patient's fluid levels have dramatically dropped, they might push a large volume of fluid into the bloodstream, perhaps overcompensating.

We've already tested our automated fluid-management system in collaboration with William Muir, a veterinary anesthesiologist and cardiovascular physiologist. Working with dogs that were experiencing hemorrhages, we used our system to regulate their fluid infusions. Our system successfully kept the dogs in stable condition as measured by the volume of blood pumped with every heartbeat.

We need to do more testing in order to win regulatory approval for a fully automated fluid-management system for humans. As with our work on ventilator management, we can start by building a decision support system for the ICU. This "human in the loop" system will present information and recommendations to the clinician, who will then adjust the settings of the infusion pump accordingly.

LOOKING BEYOND VENTILATION and fluid management, other key aspects of patient care that could be automated include pain management and sedation. In the ICU of the future, we envision many such clinical operations being monitored, coordinated, and controlled by AI systems that assess each patient's physiological state and adjust equipment settings in real time.

To make this vision a reality, though, it won't be enough for engineers to produce reliable technology. We must also find our way through many regulatory barriers and institutional requirements at hospitals.

Clearly, regulators need to scrutinize any new autonomous medical system. We suggest that regulatory agencies make use of two testing frameworks commonly used in the automotive and aerospace industries. The first is in silico trials, which test an algorithm through computer simulations. These tests are useful only if the simulations are based on high-fidelity physiological models, but in certain applications this is already possible. For example, the U.S. Food and Drug Administration recently approved the use of in silico testing as a replacement for animal testing in efforts to develop an artificial pancreas for diabetics.

The second useful framework is hardware-in-the-loop testing, where hardware stands in for the object of interest, whether it's a jet engine or the human circulatory system. You can then test a device—an autonomous fluid pump, say—on the hardware platform, which will generate the same type of data you'd see on an actual patient's bedside monitor. These hardware-in-the-loop trials can show that the device performs well in real time and in the real world. Once these technologies have been proven with stand-ins for critically ill humans, testing can begin with real patients.

To bring these technologies into hospitals, the final step is to win the trust of the medical community. Medicine is a generally conservative environment—and for good reason. No one wants to make changes that might threaten the health of patients. Our approach is to prove our technologies in stages: We'll first commercialize decision-support systems to demonstrate their efficacy and benefits, and then move to truly autonomous systems. With the addition of AI, we believe ICUs can be smarter, safer, and healthier places. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/smarticu1018

SAVING SOFTWARE FROM OBLIVION

A PROTOTYPE ARCHIVING SYSTEM LETS VINTAGE CODE RUN ON TODAY'S COMPUTERS

In early 2010, Harvard economists Carmen Reinhart and Kenneth Rogoff published an analysis of economic data from many countries and concluded that when debt levels exceed 90 percent of gross national product, a nation's economic growth is threatened. With debt that high, expect growth to become negative, they argued. ■ This analysis was done shortly after the 2008 recession, so it had enormous relevance to policymakers, many of whom were promoting high levels of debt spending in the interest of stimulating their nations' economies. At the same time, conservative politicians, such as Olli Rehn, then an EU commissioner, and U.S. congressman Paul Ryan, used Reinhart and Rogoff's findings to argue for fiscal austerity. ■ Three years later, Thomas Herndon, a graduate student at the University of Massachusetts, discovered an error in the Excel spreadsheet that Reinhart and Rogoff had used to make their calculations. The significance of the blunder was enormous: When the analysis was done properly, Herndon showed, debt levels in excess of 90 percent were associated with average growth of positive 2.2 percent, not the negative 0.1 percent that Reinhart and Rogoff had found. ■ Herndon could easily test the Harvard economists' conclusions because the software that they had used to calculate their results—Microsoft Excel—was readily available. But what about much older findings for which the software originally used is hard to come by? ■ You might think that the solution—preserving the relevant software for future researchers to use—should be no big deal. After all, software is nothing more than a bunch of files, and those files are easy enough to store on a hard drive or on tape in digital format. For some software at least, the all-important source code could even be duplicated on paper, avoiding the possibility that whatever digital medium it's written to could become obsolete. ■ Saving old programs in this way is done routinely, even for decades-old software. You can find online, for example, a full program listing for the Apollo Guidance Computer—code that took astronauts to the moon during the 1960s. It was transcribed from a paper copy and uploaded to GitHub in 2016.

By Mahadev Satyanarayanan
Illustrations by Nicholas Little

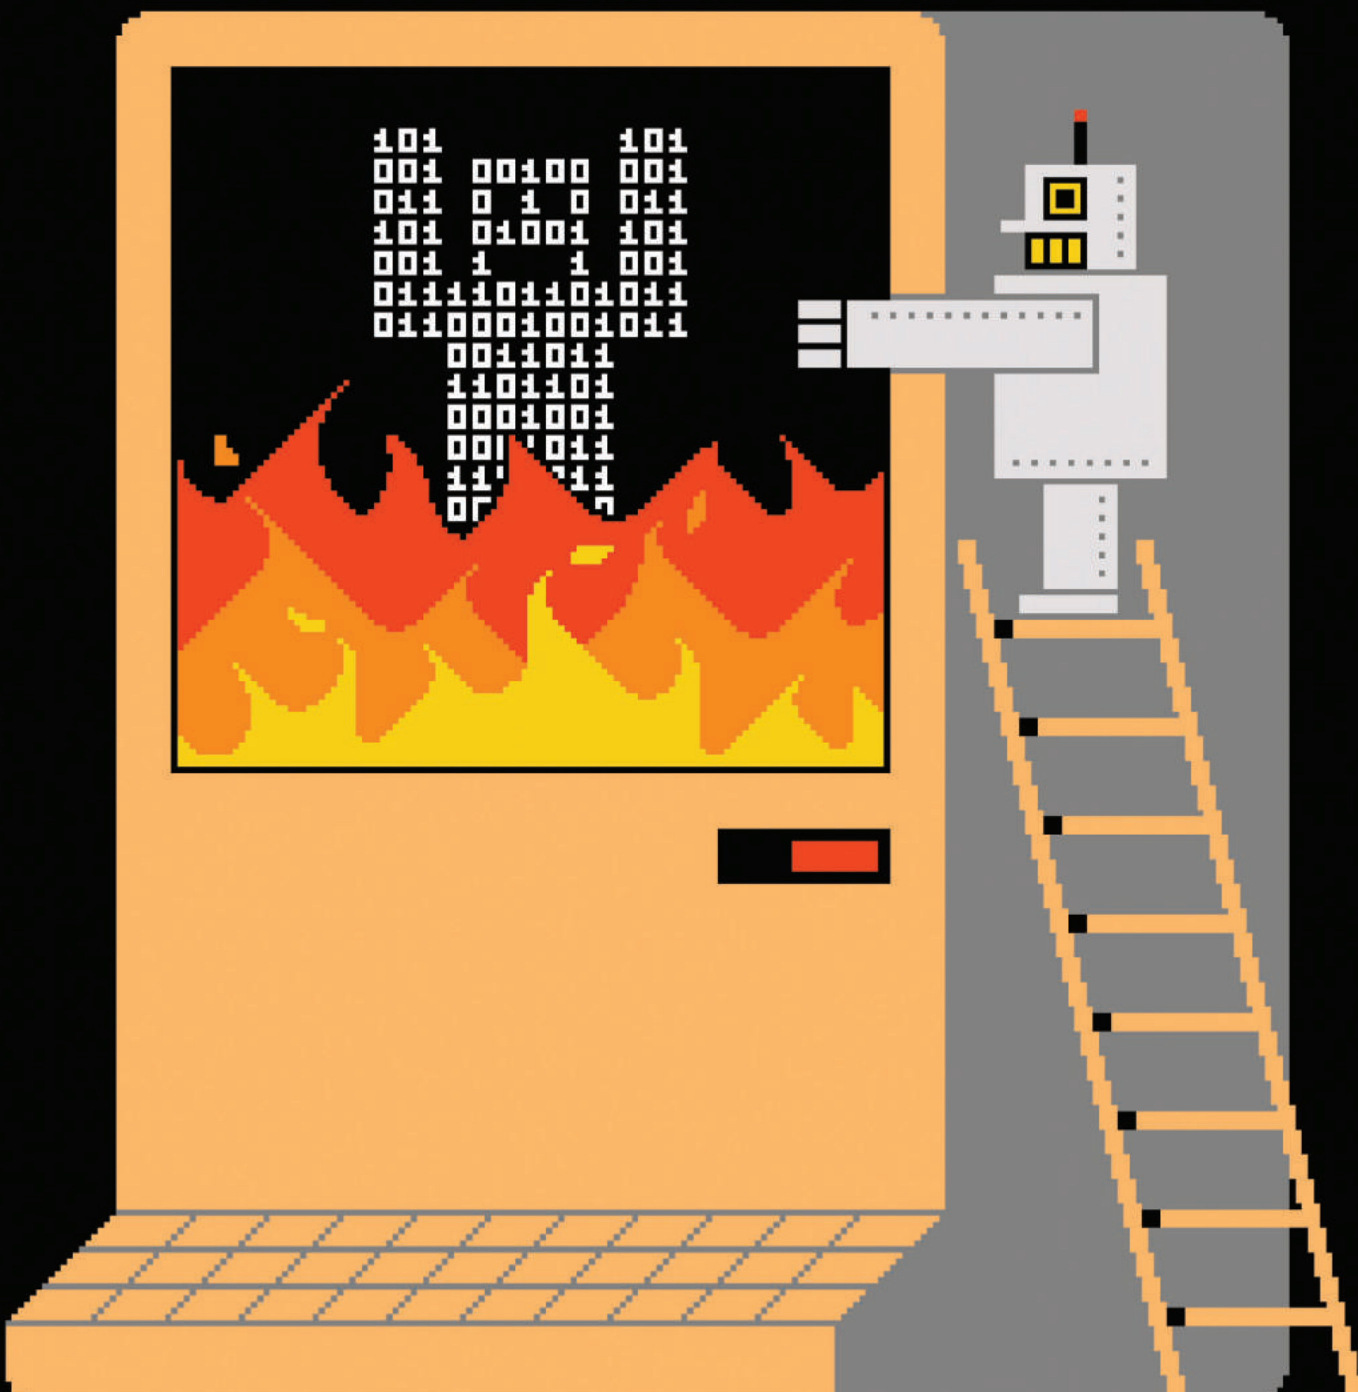

While perusing such vintage source code might delight hard-core programmers, most people aren't interested in such things. What they want to do is *use* the software. But keeping software in ready-to-run form over long periods of time is enormously difficult, because to be able to run most old code, you need both an old computer and an old operating system.

You might have faced this challenge yourself, perhaps while trying to play a computer game from your youth. But being unable to run an old program can have much more serious repercussions, particularly for scientific and technical research.

Along with economists, many other researchers, including physicists, chemists, biologists, and engineers, routinely use software to slice and dice their data and visualize the results of their analyses. They simulate phenomena with computer models that are written in a variety of programming languages and that use a wide range of supporting software libraries and reference data sets. Such investigations and the software on which they are based are central to the discovery and reporting of new research results.

Imagine that you're an investigator and want to check calculations done by another researcher 25 years ago. Would the relevant software still be around? The company that made it may have disappeared. Even if a contemporary version of the software exists, will it still accept the format of the original data? Will the calculations be identical in every respect—for example, in the handling of rounding errors—to those obtained using a computer of a generation ago? Probably not.

Researchers' growing dependence on computers and the difficulty they encounter when attempting to run old software are hampering their ability to check published results. The problem of obsolescent software is thus eroding the very premise of reproducibility—which is, after all, the bedrock of science.

The issue also affects matters that could be subject to litigation. Suppose, for example, that an engineer's calculations show that a building design is robust, but the roof of that building nevertheless collapses. Did the engineer make a mistake, or was the software used for the calculations faulty? It would be hard to know years later if the software could no longer be run.

That's why my colleagues and I at Carnegie Mellon University, in Pittsburgh, have been developing ways to archive programs in forms that can be run easily today and into the future. My fellow computer scientists Benjamin Gilbert and Jan Harkes did most of the required coding. But the collaboration has also involved software archivist Daniel Ryan and librarians Gloriana St. Clair, Erika Linke, and Keith Webster, who naturally have a keen interest in properly preserving this slice of modern culture.

Because this project is more one of archival preservation than mainstream computer science, we garnered

BRINGING BACK YESTERDAY'S SOFTWARE

The Olive system has been used to create 17 different virtual machines that run a variety of old software, some serious, some just for fun. Here are several views from those archived applications.

1. NCSA Mosaic 1.0, a pioneering Web browser for the Macintosh from 1993 2. Chaste (Cancer, Heart and Soft Tissue Environment) 3.1 for Linux from 2013 3. *The Oregon Trail* 1.1, a game for the Macintosh from 1990 4. *Wanderer*, a game for MS-DOS from 1988 5. *Mystery House*, a game for the Apple II from 1982 6. The Great American History Machine, an educational interactive atlas for Windows 3.1 from 1991 7. Microsoft Office 4.3 for Windows 3.1 from 1994 8. ChemCollective, educational chemistry software for Linux from 2013

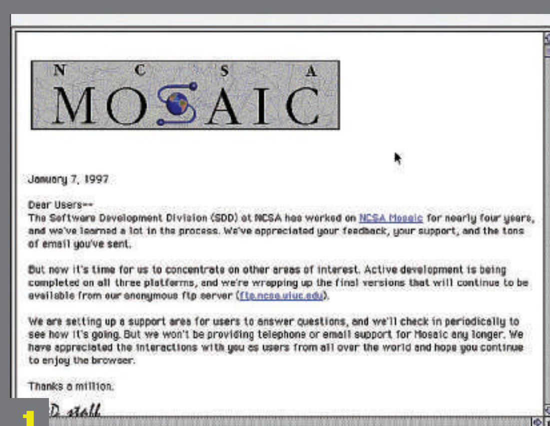
N C S A MOSAIC

January 7, 1997

Dear Users--
The Software Development Division (SDD) at NCSA has worked on NCSA Mosaic for nearly four years, and we've learned a lot in the process. We've appreciated your feedback, your support, and the tons of email you've sent.

But now it's time for us to concentrate on other areas of interest. Active development is being completed on all three platforms, and we're wrapping up the final versions that will continue to be available from our anonymous ftp server (ftp.ncsa.uiuc.edu).

We are setting up a support area for users to answer questions, and we'll check in periodically to see how it's going. But we won't be providing telephone or email support for Mosaic any longer. We have appreciated the interactions with you as users from all over the world and hope you continue to enjoy the browser.

Thanks a million.

D. staff

1

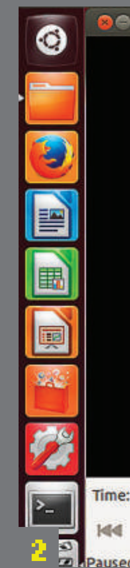

2

4

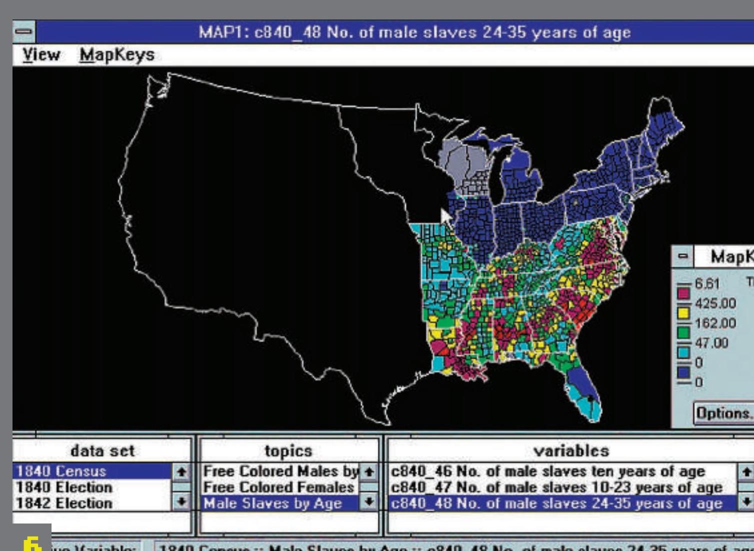

6

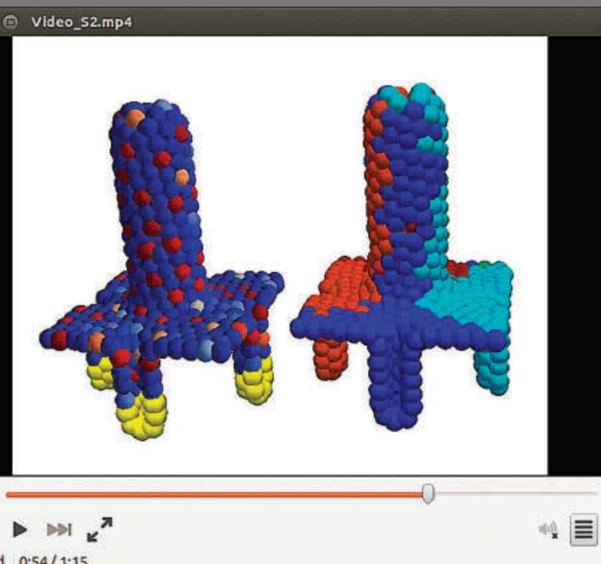
Video_S2.mp4
0:54 / 1:15


Oregon Trail
Map
Guide
Status
Rations
Buy
The Oregon Trail
Mountains
River
Trail
Landmark
Fort
0 Miles 300
• April 1, 1848 •
You started down the trail with 10 oxen, 20 sets of clothing, 1,000 bullets, 3 wagon wheels, 3 wagon axles, 3 wagon tongues, 1,000 pounds of food, and $810.00.
You decided to continue.
• April 6, 1848 •
Zeke was bitten by a snake.
Conditions
April 6, 1848
Sunny
Hot
Warm
Cold
Distance
To Landmark: 4 mi.
Traveled: 98 mi.
Wagon
Pace: Steady
Rations: Filling
Food Left: 925 lbs.
Health: Good
Wagon: Moving
Time Out
Trade
Talk
Rest
Pace
Hunt

3

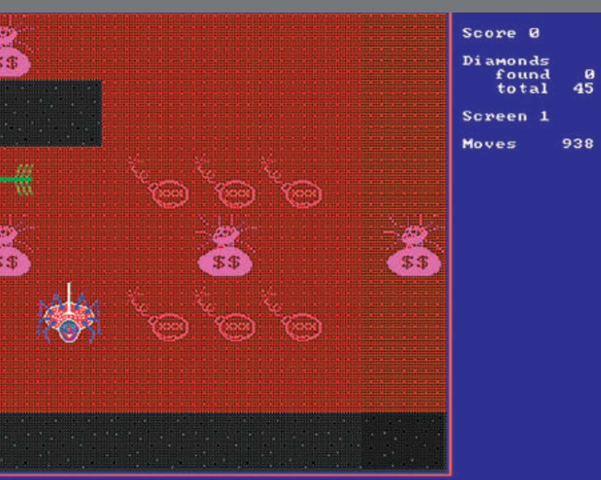
Score 0
Diamonds found 0
total 45
Screen 1
Moves 938

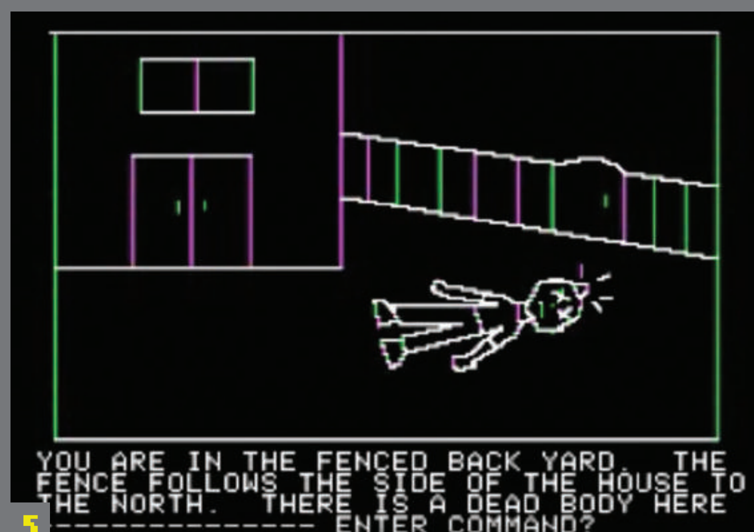
YOU ARE IN THE FENCED BACK YARD. THE FENCE FOLLOWS THE SIDE OF THE HOUSE TO THE NORTH. THERE IS A DEAD BODY HERE
------------- ENTER COMMAND?

5

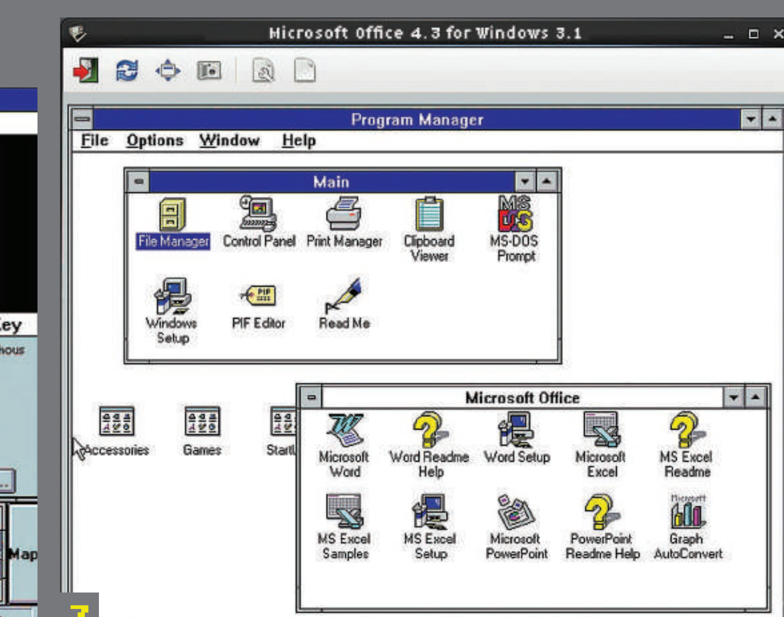
Microsoft Office 4.3 for Windows 3.1
Program Manager
File Options Window Help
Main
File Manager
Control Panel
Print Manager
Clipboard Viewer
MS-DOS Prompt
Windows Setup
PIF Editor
Read Me
Accessories
Games
Microsoft Office
Microsoft Word
Word Readme Help
Word Setup
Microsoft Excel
MS Excel Readme
MS Excel Samples
MS Excel Setup
Microsoft PowerPoint
PowerPoint Readme Help
Graph AutoConvert

7

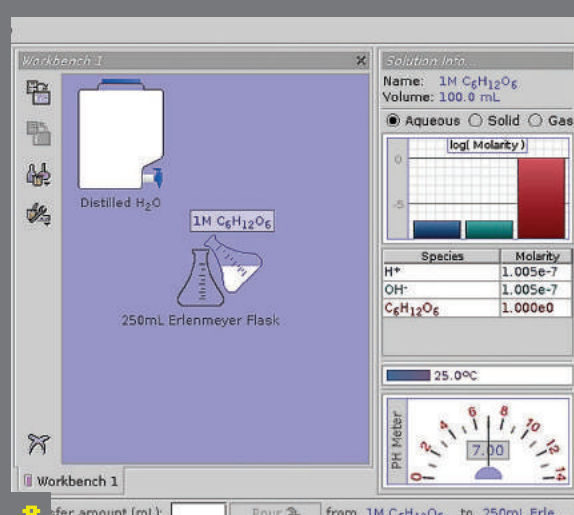
Workbench 1
Solution Info
Name: 1M C6H12O6
Volume: 100.0 mL
Aqueous Solid Gas
log(Molarity)
Species Molarity
H+ 1.005e-7
OH- 1.005e-7
C6H12O6 1.000e0
25.00C
PH Meter
7.00
Distilled H2O
1M C6H12O6
250mL Erlenmeyer Flask
Workbench 1
Pour
from 1M C6H12O6 to 250mL Erle...

8

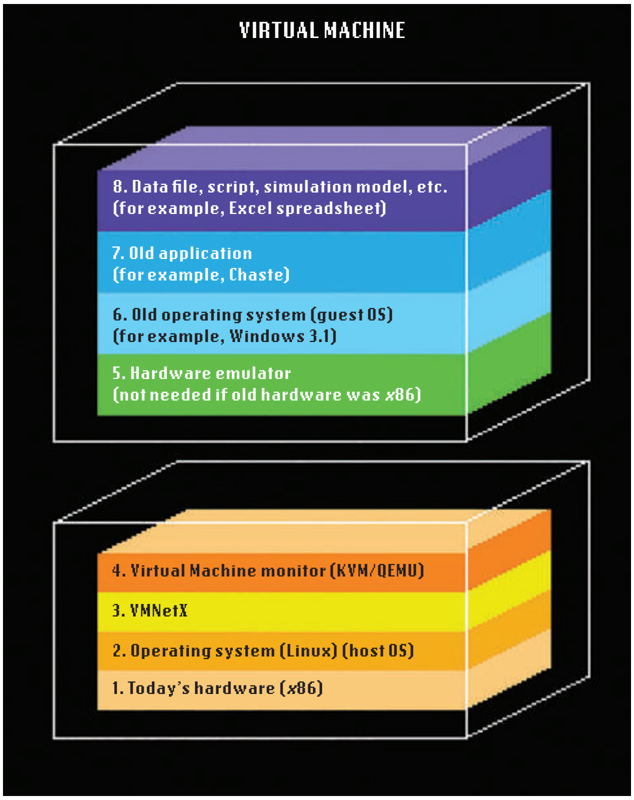

LAYERS OF ABSTRACTION: Olive requires many layers of software abstraction to create a suitable virtual machine. That virtual machine then runs the old operating system and application.

financial support for it not from the usual government funding agencies for computer science but from the Alfred P. Sloan Foundation and the Institute for Museum and Library Services. With that support, we showed how to reconstitute long-gone computing environments and make them available online so that any computer user can, in essence, go back in time with just a click of the mouse.

We created a system called Olive—an acronym for Open Library of Images for Virtualized Execution. Olive delivers over the Internet an experience that in every way matches what you would have obtained by running an application, operating system, and computer from the past. So once you install Olive, you can interact with some very old software as if it were brand new. Think if it as a Wayback Machine for executable content.

TO UNDERSTAND HOW OLIVE can bring old computing environments back to life, you have to dig through quite a few layers of software abstraction. At the very bottom is the common base of much of today's computer technology: a standard desktop or laptop endowed with one or more *x*86 microprocessors. On that computer, we run the Linux operating system, which forms the second layer in Olive's stack of technology.

Sitting immediately above the operating system is software written in my lab called VMNetX, for Virtual Machine Network Execution. A virtual machine is a computing environment that mimics one kind of computer using software running on a different kind of computer. VMNetX is special in that it allows virtual machines to be stored on a central server and then executed on demand by a remote system. The advantage of this arrangement is that your computer doesn't need to download the virtual machine's entire disk and memory state from the server before running that virtual machine. Instead, the information stored on disk and in memory is retrieved in chunks as needed by the next layer up: the virtual-machine monitor (also called a hypervisor), which can keep several virtual machines going at once.

Each one of those virtual machines runs a hardware emulator, which is the next layer in the Olive stack. That emulator presents the illusion of being a now-obsolete computer—for example, an old Macintosh Quadra with its 1990s-era Motorola 68040 CPU. (The emulation layer can be omitted if the archived software you want to explore runs on an *x*86-based computer.)

The next layer up is the old operating system needed for the archived software to work. That operating system has access to a virtual disk, which mimics actual disk storage, providing what looks like the usual file system to still-higher components in this great layer cake of software abstraction.

Above the old operating system is the archived program itself. This may represent the very top of the heap, or there could be an additional layer, consisting of data that must be fed to the archived application to get it to do what you want.

The upper layers of Olive are specific to particular archived applications and are stored on a central server. The lower layers are installed on the user's own computer in the form of the Olive client software package. When you launch an archived application, the Olive client fetches parts of the relevant upper layers as needed from the central server.

THAT'S WHAT YOU'LL FIND under the hood. But what can Olive do? Today, Olive consists of 17 different virtual machines that can run a variety of operating systems and applications. The choice of what to include in that set was driven by a mix of curiosity, availability, and personal interests. For example, one member of our team fondly remembered playing *The Oregon Trail* when he was in school in the early 1990s. That led us to acquire an old Mac version of the game and to get it running again through Olive. Once word of that accomplishment got out, many people started approaching us to see if we could resurrect their favorite software from the past.

The oldest application we've revived is *Mystery House*, a graphics-enabled game from the early 1980s for the Apple II computer. Another program is NCSA Mosaic, which people of a certain age might remember as the browser that introduced them to the wonders of the World Wide Web.

Olive provides a version of Mosaic that was written in 1993 for Apple's Macintosh System 7.5 operating system. That operating system runs on an emulation of the Motorola 68040 CPU, which in turn is created by software running on an actual *x*86-based computer that runs Linux. In spite of all this virtualization, performance is pretty good, because modern computers are so much faster than the original Apple hardware.

Pointing Olive's reconstituted Mosaic browser at today's Web is instructive: Because Mosaic predates Web technologies such as JavaScript, HTTP 1.1, Cascading Style Sheets, and HTML 5, it is unable to render most sites. But you can have some fun tracking down websites composed so long ago that they still look just fine.

What else can Olive do? Maybe you're wondering what tools businesses were using shortly after Intel introduced the Pentium processor. Olive can help with that, too. Just fire up Microsoft Office 4.3 from 1994 (which thankfully predates the annoying automated office assistant "Clippy").

Perhaps you just want to spend a nostalgic evening playing *Doom* for DOS—or trying to understand what made such

first-person shooter games so popular in the early 1990s. Or maybe you need to redo your 1997 taxes and can't find the disk for that year's version of TurboTax in your attic. Have no fear: Olive has you covered.

On the more serious side, Olive includes Chaste 3.1. The name of this software is short for Cancer, Heart and Soft Tissue Environment. It's a simulation package developed at the University of Oxford for computationally demanding problems in biology and physiology. Version 3.1 of Chaste was tied to a research paper published in March 2013. Within two years of publication, though, the source code for Chaste 3.1 no longer compiled on new Linux releases. That's emblematic of the challenge to scientific reproducibility Olive was designed to address.

To keep Chaste 3.1 working, Olive provides a Linux environment that's frozen in time. Olive's re-creation of Chaste also contains the example data that was published with the 2013 paper. Running the data through Chaste produces visualizations of certain muscle functions. Future physiology researchers who wish to explore those visualizations or make modifications to the published software will be able to use Olive to edit the code on the virtual machine and then run it.

For now, though, Olive is available only to a limited group of users. Because of software-licensing restrictions, Olive's collection of vintage software is currently accessible only to people who have been collaborating on the project. The relevant companies will need to give permissions to present Olive's re-creations to broader audiences.

We are not alone in our quest to keep old software alive. For example, the Internet Archive is preserving thousands of old programs using an emulation of MS-DOS that runs in the user's browser. And a project being mounted at Yale, called EaaSI (Emulation as a Service Infrastructure), hopes to make available thousands of emulated software environments from the past. The scholars and librarians involved with the Software Preservation Network have been coordinating this and similar efforts. They are also working to address the copyright issues that arise when old software is kept running in this way.

OLIVE HAS COME A LONG WAY, but it is still far from being a fully developed system. In addition to the problem of restrictive software licensing, various technical roadblocks remain.

One challenge is how to import new data to be processed by an old application. Right now, such data has to be entered manually, which is both laborious and error prone. Doing so also limits the amount of data that can be analyzed. Even if we were to add a mechanism to import data, the amount that could be saved would be limited to the size of the virtual machine's virtual disk. That may not seem like a problem, but you have to remember that the file systems on older computers sometimes had what now seem like quaint limits on the amount of data they could store.

Another hurdle is how to emulate graphics processing units (GPUs). For a long while now, the scientific community has been leveraging the parallel-processing power of GPUs to speed up many sorts of calculations. To archive executable versions of software that takes advantage of GPUs, Olive would need to re-create virtual versions of those chips, a thorny task. That's because GPU interfaces—what gets input to them and what they output—are not standardized.

Clearly there's quite a bit of work to do before we can declare that we have solved the problem of archiving executable content. But Olive represents a good start at creating the kinds of systems that will be required to ensure that software from the past can live on to be explored, tested, and used long into the future. ■

↗ POST YOUR COMMENTS at https://spectrum.ieee.org/olive1018

The

GREAT SOVIET CALCULATOR HACK

Programmable calculators and a sci-fi story brought Soviet teens into the digital age ***By KSENIA TATARCHENKO***

ILLUSTRATION BY Chris Malbon

DESPITE THE UBIQUITY of computers in modern society, the vast majority of today's students never study computer science or computer programming. Those who are exposed to these subjects typically learn low-level skills rather than undertaking any deeper exploration of computational concepts or theory. In earlier decades, a few countries did promote computer education at the national level. In the 1980s, for example, the British government launched a popular and quite successful initiative that brought thousands of BBC Micros into classrooms.

But the most ambitious computer literacy program ever conceived is one you've probably never heard of, and it originated in a very unlikely place: the Soviet Union.

Perhaps you're smiling to yourself, recalling the old trope about how the USSR invented *Tetris* and yet lost the Cold War. Implicit in this dismissal is the failure of the Soviets to fully appreciate the awesome power of the digital age. It's true that the Soviet government never embraced a national computer network or provided its citizens with affordable personal computers. But if you subscribe to this narrative of technological stumbles and political failure, then you're missing an important, not to mention fascinating, part of the story of global computerization–one in which Soviet teenagers latched onto a popular sci-fi novel of adventure and self-discovery and taught themselves and each other how to program using the only means available to them: the programmable calculator.

IN SEPTEMBER 1985, ninth graders all across the USSR began studying a new subject: Basics of Informatics and Computing Technology. The rollout of the compulsory course, which aimed to make programming a universal skill, was to be accompanied by new textbooks in 15 national languages, training for some 100,000 teachers, and a million computers for the 60,000 or so middle schools across the Soviet republics.

None of this went smoothly. The state didn't supply schools with equipment, efforts to print and distribute course literature were uneven, and many teachers never received the requisite training.

Meanwhile, the move sparked an international debate among computer experts over the very definition of "computer literacy." The U.S. computer scientist and entrepreneur Edward Fredkin argued that his country's experience should inform the Soviets:

> We now understand that computer literacy is not knowing how to program. It is not understanding how [a] computer works. It is not knowing about bits and bytes and flip-flops and gates.... We now know that true computer literacy means having the skills to use the advanced application programs, such as word processing and spreadsheet systems.

In response, the computer scientist Andrei Ershov quipped that coding and typing were not mutually exclusive. Ershov was head of the Akademgorodok Computer Center in the Siberian science city of Akademgorodok, and he had emerged as the computer literacy campaign's key promoter. In stark contrast to Fredkin, he viewed computer literacy as nurturing a set of intellectual habits, which he called "algorithmic thinking."

That idea grew in part out of Ershov's time as a student of Aleksei Liapunov, a towering figure of Soviet cybernetics. From Liapunov, Ershov learned to think in terms of cybernetic metaphors and to draw connections between technology and society. He conceived of algorithms as a form of communication between humans and machines.

CALCULATOR RULES: **At a time when few Soviet households had personal computers, programmable calculators like the Elektronika B3-34 took on many roles.**

SERGEI FROLOV/SOVIET DIGITAL ELECTRONICS MUSEUM

Ershov also drew on ideas from the West. In September 1958, he was among an elite group of Soviet computer experts to meet with their American counterparts. His exchange with the computer pioneer Alan Perlis, who would later become the first recipient of the Turing Award, proved particularly fruitful. Perlis shared with Ershov his enthusiasm for developing a universal algorithmic language, called Algol, which aimed to make software portable and international. Ershov embraced Algol's agenda, and he went on to develop one of the most ambitious compilers for the language in the early 1960s. The universalist aspirations of the Algol community would inform his views on computer education.

Ershov's educational agenda was also inspired by a visit to MIT in the early 1970s, where he met Seymour Papert and learned of his computer education experiments with Logo, a programming language designed for use by children.

And yet, while Ershov closely followed developments in Western computer science, he believed the Soviet Union should forge its own path to the information age, one imbued with socialist values, less dependent on computers as black-boxed commodities and more focused on building citizens' skills and habits of mind. By learning to program, he argued, students would develop abstract reasoning and a goal-oriented, problem-solving mind-set. By the end of the 1970s, Ershov and his team in Akademgorodok had formulated their literacy program, developing their curriculum with the help of Siberian students and testing it in local schools.

Of course, Ershov knew he needed much broader support to implement such a curriculum nationally. He began tirelessly promoting his idea of programming as a "second literacy" to Soviet authorities, computer experts, educators, parents, and children, as well as to the international community.

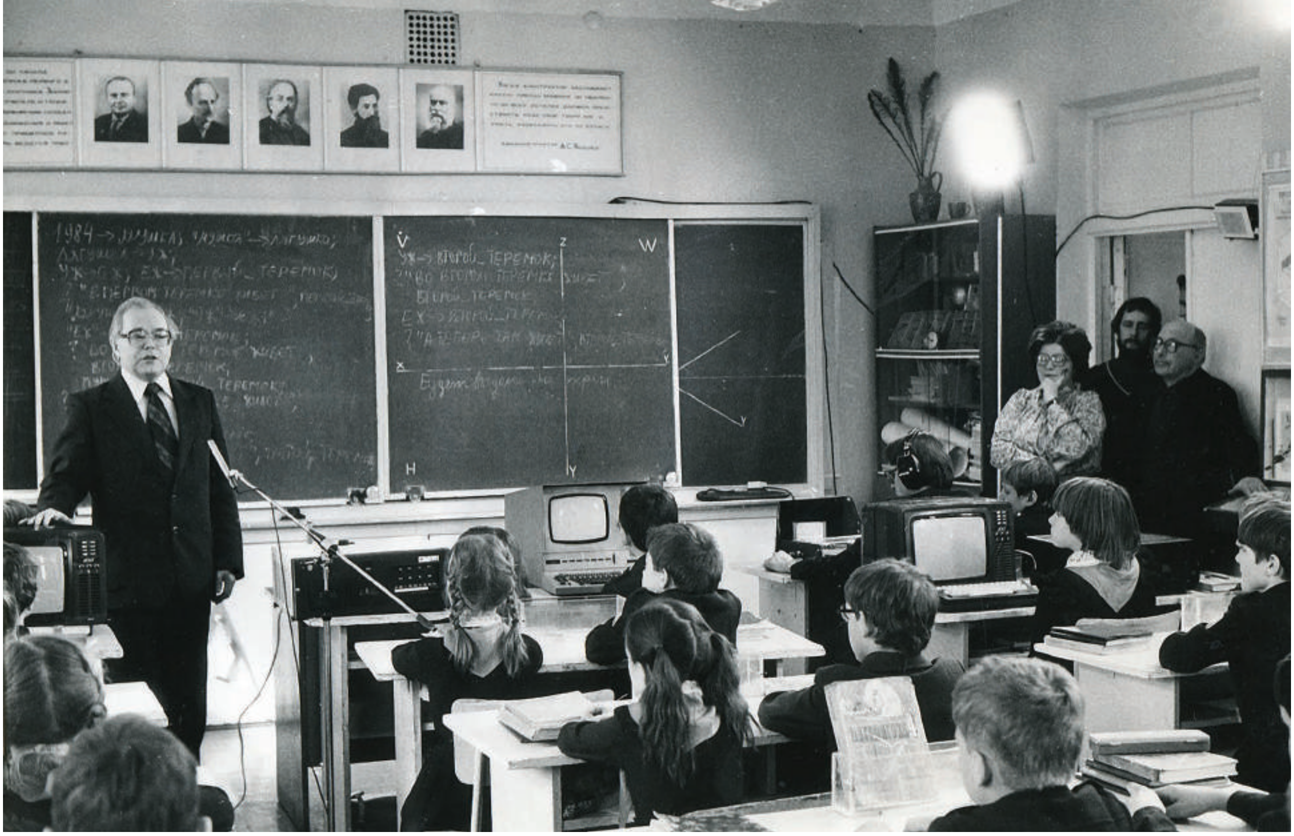

ALGORITHMIC THOUGHT: **Computer scientist Andrei Ershov [standing, left] championed computer literacy for all citizens of the Soviet Union. Beginning in 1985, ninth graders took a compulsory course entitled Basics of Informatics and Computing Technology.**

PHOTOGRAPHIC ARCHIVE/SIBERIAN BRANCH, RUSSIAN ACADEMY OF SCIENCES

Finally, in 1985, in a wave of transformative policies adopted with Mikhail Gorbachev's ascendance to power, the Akademgorodok informatics curriculum was officially adopted.

The inefficiencies of Soviet planning and economics meant that most ninth graders studied the curriculum without computers on which to test their new skills. This wasn't seen as an obstacle by the reformers. Instead, the teaching materials encouraged writing out programs on paper and engaging in imaginative exercises. Students, for instance, acted out the role of a robot named Dezhurik (from the Russian word *dezhurnyi*, the person responsible for maintaining the classroom), who was programmed to "close window" or "clean blackboard." When students from the remote city of Khabarovsk complained about the lack of classroom computers, Ershov commended them for taking the initiative to write and emphasized that the youths still had the chance to "catch up to the train to the future."

But he refused to commiserate with them. What they were learning—how to devise an algorithm and write a program for it—was the essential part, he said, whether or not they ever got to run the program on an actual computer. Ershov's letter to the students concluded: "If the teacher may have pity on you and give you a satisfactory grade, the computer will not forgive you any errors. It will stay there, an impenetrable piece of metal, up to the end of the school year. Without an algorithm, without a program, without a plan, there is no point in sitting in front of the computer."

SOVIET CITIZENS MAY have lacked access to PCs, but many millions of them did have access to computational devices, in the form of scientific programmable calculators. These handheld devices could store instructions and numbers in memory for later execution. Popular in the West following the 1974 introduction of the HP-65 by Hewlett-Packard, programmable calculators still have their fans and their uses.

In the Soviet Union, from the mid-1970s on, the microelectronics industry produced electronic calculators by the millions, primarily for use by what was then the world's largest population of engineers. As in the West, Soviet calculator users were instrumental in shaping the development of programs and applications for the devices. Unlike in the West, few Soviets had home computers, and so the calculator took on many more roles—including as a makeshift computing platform for computer education and a thriving game culture.

These two roles converged in a popular science magazine called *Tekhnika Molodezhi* (Technology for Youth), which was published by the Communist youth organization Komsomol. The publication was aimed at teenagers and had a subscriber base of 1.5 million. In January 1985, the magazine took up Ershov's computer literacy campaign and began devoting a section to programming with the most popular Soviet calculator, the Elektronika B3-34, which sold for 85 rubles. Reader response to the column was disappointing, however.

Then in August 1985, *TM* began serializing the space-travel novel *Kon-Tiki: A Path to the Earth*. In this tale of an epic quest, an engineer and a pilot attempt, against all odds, to fly a lunar lander from the moon back to Earth. The premise of the novel was the popular U.S. computer game *Lunar Lander*, in which players controlled thrusters and calculated trajectories to safely guide their landers to the lunar surface. The Soviet version was called *Lunalet*. Each installment of the novel invited readers to take up their calculators, transforming themselves into pilots and their devices into spaceships.

The *Kon-Tiki* serial was an instant hit, and the magazine soon became one of the most prominent forums for younger users of programmable calculators. The futuristic narrative of each chapter was combined with puzzles on the physical laws of space travel and tricks for programming the B3-34. But what kept readers reading was the dramatic plot and the novel's focus on overcoming human and technological limits.

A reference to Thor Heyerdahl's 1947 journey by raft across the Pacific Ocean, *Kon-Tiki* was also the name of the tiny vessel chosen by the novel's protagonists for their earthbound voyage. The story line evolved far beyond the game's original goal of landing the spacecraft. The "path to the Earth" became a journey of self-discovery. At one point, the pilot, called Moon Hawk, reflects on his own fallibility: "I am not a computer; I am a human, and it is typical for me to make mistakes. Because of that I can't choose a path that does not allow for mistakes. Of course, in cases when I have a choice, I will prefer the way that gives me the right to make a mistake and simultaneously an opportunity to correct it."

The heroes stumble even at the novel's conclusion: They arrive back on Earth only to touch down in the ocean and are forced to send out an SOS. "After all, I am a cosmonaut, not a sea captain," admits the pilot, as they await rescue.

Credit for the novel's clever intertwining of programming and storytelling goes to its author, Mikhail Pukhov, who was also editor of *TM*'s sci-fi section. The son of a prominent mathematician, Pukhov graduated from the country's most prestigious engineering school, the Moscow Institute of Physics and Technology. Abandoning a promising career at the Central Scientific-Research Radio Engineering Institute, he had turned to writing and editing.

Before starting *Kon-Tiki*, Pukhov thoroughly explored the calculator's functions as well as its malfunctions. Calculator users in the West and in the East alike were quick to discover and exploit the devices' undocumented features, pushing them to do things their designers never intended. Such exploration became known as errorology, from the "ЕГГОГ" message that would frequently appear on the small display at the execution of an undocumented feature. Pukhov's novel glorified errorology with poetic descriptions of "fishing" for unusual combinations of symbols.

And readers responded, writing to *TM* about their own calculator exploits. "I inform you that I obtained an easy way for creating any combination from the numbers and symbols 'E,' 'Г,' 'C,' 'L,' '-,' which do not begin from zero on the display of the B3-34," boasted one reader. To have their programs and names printed in *TM* was the highest aspiration of many readers.

Thus did *TM* and its sci-fi editor help cultivate a generation of hackers and computer enthusiasts. If you find it odd that a major state-sanctioned Soviet magazine promoted hacking practices, consider how U.S. hacker culture emerged—as a form of hands-on technological investigation. In his 1984 book *Hackers: Heroes of the Computer Revolution*, Steven Levy traced the origin of MIT's hackers to a club of railroad buffs. Similarly, in the Soviet Union, a combination of state interests and grassroots initiatives had nurtured a hands-on culture among radio amateurs. For a radio engineer like Pukhov, as well as for the educators who used the novel in the classroom, subverting the calculator's design specifications was a way to encourage technical skills.

GAMIFICATION: **In August 1985, the Soviet science magazine *Tekhnika Molodezhi* began publishing a serialized novel about a pair of explorers trying to fly a lunar lander from the moon to Earth. Each installment included tasks to be worked out on a programmable calculator.**

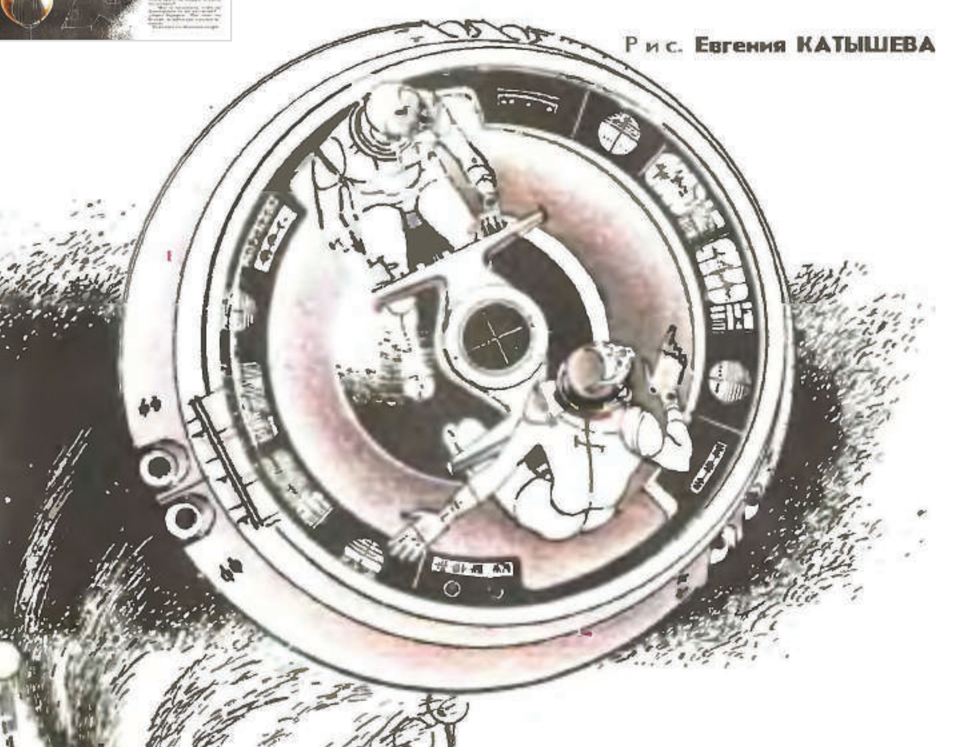

TEKHNIKA MOLODEZHI (4)

The community of readers and players that formed around *Kon-Tiki* unwittingly embraced the goals of programming literacy as conceived by Ershov. In their letters to *TM*, many requested more games as well as flowcharts for rewriting the programs for other kinds of calculators. One reader wrote that he aspired "to see the program as a conscious pattern of actions, and not as a thoughtless row of symbols. To be able, with the help of your magazine, not only to execute the available programs but to create [new programs] myself." In this sense, the novel and its community of calculator users contributed to the spread of Ershov's vision of computer literacy.

But was Ershov's curriculum a success? The results of any educational initiative are, of course, hard to gauge. Soviet statisticians no doubt monitored the reform effort in some fashion, but that data would hardly capture real-world experiences in the classroom and beyond.

I posted on several Russian calculator-user forums hoping to hear from *Kon-Tiki* readers. The responses I received were tinged with nostalgia. Some wrote that their fascination with the novel spurred them to get a calculator. "For half a year, like a vacuum cleaner, I was absorbing all information available on programming and calculators in particular," wrote one forum member, explaining how he learned the principles of programming before acquiring a device of his own. For others, the calculator was but a steppingstone; eventually they earned enough money to buy computer kits (available in perestroika-era street markets) and assemble their own machines. Meanwhile, copies of *TM* continued to circulate via secondhand shops, where new groups of readers discovered the novel long after its original publication. Today you can easily find electronic versions of the magazine online, along with calculator emulators.

The extent to which such school-age experiences influenced people's professional lives is less clear. Unless you lived through it, you may not fully appreciate the enormous disruption brought on by the economic crisis following the collapse of the Soviet Union—in Russia, those years are known as "the wild 1990s." *Kon-Tiki*'s readers came of age as citizens of the emerging sovereign states. Few of them had full control over their career options, and for many, coding became a calling, a gig, and a gateway. These days in Russia, the "universality" of programming skills is no longer associated with creating a computer-literate society. Instead, it raises the prospect of migration, as skilled programmers choose to leave the country to pursue their careers.

PAPER PROGRAMMING: Andrei Ershov works with students at a summer school for young programmers. Most Soviet schools didn't have computers, though, and students were encouraged to write out their programs on paper. "Without an algorithm, without a program, without a plan, there is no point in sitting in front of the computer," Ershov wrote.

PHOTOGRAPHIC ARCHIVE/SIBERIAN BRANCH, RUSSIAN ACADEMY OF SCIENCES

SOVIET-ERA EFFORTS to foster computer literacy cast Western assumptions about the information age in a different light. Unlike events in the West, the Soviet digital revolution was not one of geeks and geniuses but of state-sponsored academics, writers, and educators, who worked with government officials, industrialists, and programmers toward a shared goal. It was not predicated on the personal computer but instead made do with calculators, pencil and paper, and students' own imaginations.

Despite the passion of digital enthusiasts like Ershov and Pukhov, the campaign's ideal of universality was hard to attain. The reform seemed to work best where you would expect it to, in the elite schools of the capitals and in a few remote schools blessed with wealthy patrons, such as those supported by the oil and gas industry. *TM* transcended some geographical and economic barriers and provided a motivation, an entry point, and a community to students who lacked inspiring teachers or computers of their own. But the magazine failed to bridge another familiar divide—that of gender. Unlike the population of Soviet professional programmers, and unlike the compulsory and gender-neutral informatics classes, the readers who wrote to *TM* about their calculator exploits were predominantly male.

And so, the digital socialist society that Ershov and others strove for was imperfectly realized. Neither Ershov, who died in 1988, nor the country itself survived long enough for the experiment to run its course. And yet, we shouldn't be so quick to dismiss a vision of computer literacy that considered all students capable of thinking algorithmically. The Soviets did not prefigure the many challenges of the information age. But what we choose to remember from our computing pasts can help determine how we solve our present-day conundrums. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/sovietcalculator1018

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

Edward S. Rogers Sr. Department of Electrical and Computer Engineering

The Edward S. Rogers Sr. Department of Electrical and Computer Engineering (ECE) at the University of Toronto invites applications for up to four full-time tenure-stream faculty appointments at the rank of Assistant Professor. The appointments will commence on **July 1, 2019**.

Within the general field of electrical and computer engineering, we seek applications from candidates with expertise in one or more of the following strategic research areas: 1. Computer Systems and Software; 2. Electrical Power Systems; 3. Systems Control, including but not limited to autonomous and robotic systems.

Applicants are expected to have a Ph.D. in Electrical and Computer Engineering, or a related field, at the time of appointment or soon after.

Successful candidates will be expected to initiate and lead an outstanding, innovative, independent, competitive, and externally funded research program of international calibre, and to teach at both the undergraduate and graduate levels. Candidates should have demonstrated excellence in research and teaching. Excellence in research is evidenced primarily by publications or forthcoming publications in leading journals or conferences in the field, presentations at significant conferences, awards and accolades, and strong endorsements by referees of high international standing. Evidence of excellence in teaching will be demonstrated by strong communication skills; a compelling statement of teaching submitted as part of the application highlighting areas of interest, awards and accomplishments, and teaching philosophy; sample course syllabi and materials; and teaching evaluations, as well as strong letters of recommendation.

Eligibility and willingness to register as a Professional Engineer in Ontario is highly desirable.

Salary will be commensurate with qualifications and experience.

The Edward S. Rogers Sr. Department of Electrical and Computer Engineering at the University of Toronto ranks among the best in North America. It attracts outstanding students, has excellent facilities, and is ideally located in the middle of a vibrant, artistic, diverse and cosmopolitan city. Additional information may be found at **http://www.ece.utoronto.ca**.

Review of applications will begin after **September 1, 2018**, however, the position will remain open until **November 29, 2018**.

As part of your online application, please include a cover letter, a curriculum vitae, a summary of your previous research and future research plans, as well as a teaching dossier including a statement of teaching experience and interests, your teaching philosophy and accomplishments, and teaching evaluations. Applicants must arrange for three letters of reference to be sent directly by the referees (on letterhead, signed and scanned), by email to the ECE department at **search2018@ece.utoronto.ca.** Applications without any reference letters will not be considered; it is your responsibility to make sure your referees send us the letters while the position remains open.

You must submit your application online while the position is open, by following the submission guidelines given at **http://uoft.me/how-to-apply**. Applications submitted in any other way will not be considered. We recommend combining attached documents into one or two files in PDF/MS Word format. If you have any questions about this position, please contact the ECE department at **search2018@ece.utoronto.ca**.

The University of Toronto is strongly committed to diversity within its community and especially welcomes applications from racialized persons / persons of colour, women, Indigenous / Aboriginal People of North America, persons with disabilities, LGBTQ persons, and others who may contribute to the further diversification of ideas.

As part of your application, you will be asked to complete a brief Diversity Survey. This survey is voluntary. Any information directly related to you is confidential and cannot be accessed by search committees or human resources staff. Results will be aggregated for institutional planning purposes. For more information, please see **http://uoft.me/UP**.

All qualified candidates are encouraged to apply; however, Canadians and permanent residents will be given priority.

上海科技大学
ShanghaiTech University

TENURE-TRACK AND TENURED POSITIONS

ShanghaiTech University invites highly qualified candidates to fill multiple tenure-track/tenured faculty positions as its core founding team in the School of Information Science and Technology (SIST). We seek candidates with exceptional academic records or demonstrated strong potentials in all cutting-edge research areas of information science and technology. They must be fluent in English. English-based overseas academic training or background is highly desired. ShanghaiTech is founded as a world-class research university for training future generations of scientists, entrepreneurs, and technical leaders. Boasting a new modern campus in Zhangjiang Hightech Park of cosmopolitan Shanghai, ShanghaiTech shall trail-blaze a new education system in China. Besides establishing and maintaining a world-class research profile, faculty candidates are also expected to contribute substantially to both graduate and undergraduate educations.

Academic Disciplines: Candidates in all areas of information science and technology shall be considered. Our recruitment focus includes, but is not limited to: computer architecture, software engineering, database, computer security, VLSI, solid state and nano electronics, RF electronics, information and signal processing, networking, security, computational foundations, big data analytics, data mining, visualization, computer vision, bio-inspired computing systems, power electronics, power systems, machine and motor drive, power management IC 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 on-campus housing. All regular ShanghaiTech faculty members will join its new tenure-track system in accordance with international practice for progress evaluation and promotion.

Qualifications:
-Strong research productivity and demonstrated potentials;
-Ph.D. (Electrical Engineering, Computer Engineering, Computer Science, Statistics, Applied Math, or related field);
-A minimum relevant (including PhD) research experience of 4 years.

Applications: Submit (in English, PDF version) a cover letter, a 2-page research plan, a CV plus copies of 3 most significant publications, and names of three referees to: sist@shanghaitech.edu.cn. ***For more information, visit*** http://sist.shanghaitech.edu.cn/NewsDetail.asp?id=373

Deadline: The positions will be open until they are filled by appropriate candidates.

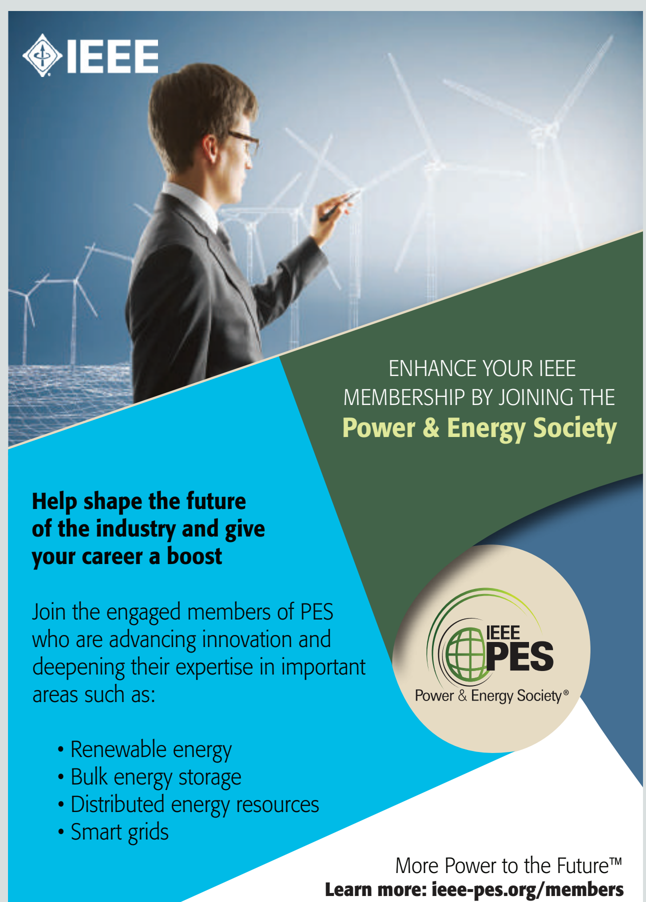

BETTMANN/GETTY IMAGES

ELEKTRO THE MOTO-MAN

Convinced that every home would eventually have a robot—doing chores, minding the kids—engineers at Westinghouse pioneered a voice-controlled robot in the 1930s. They unveiled Elektro the Moto-Man at the 1939 New York World's Fair. The 2.1-meter-tall, 118-kilogram robot "walked" by rolling along a track with its left leg bent, and it "talked" by way of prerecorded snippets played back from 33⅓-rpm records. It could also smoke a cigarette. Voice commands from a human assistant were turned into electrical pulses, which were then transmitted through telephone relays to the robot's control unit, located offstage. One of Elektro's pet lines was, "My brain is bigger than yours." At 25 kg, it certainly was. ■

↗ For more on Elektro, see https://spectrum.ieee.org/pastforward1018

IEEE

Endless Possibilities...

and countless risks

IEEE Member Group Term Life Insurance can help you wrap your hands around your family's financial security

You work hard so your family can enjoy the endless possibilities your future holds. *But what if something happened and you could no longer be there for your family? Would they have enough life insurance to replace your income and maintain their current lifestyle—let alone pay future expenses?*

You come across countless risks every day. An accident on the way to work or a terminal disease can happen regardless of age. Yet many families know they don't have enough life insurance. In fact, nearly 50% of households have a $200,000 shortfall in life insurance benefits.[1]

IEEE understands the importance of protecting your family and wants to make sure you have enough. That's why through the IEEE Member Group Term Life Insurance Plan,* IEEE members can use the flexible benefit amounts as their sole source of family protection or as a supplement to employer-provided or other life benefits.

NEW QuickDecision℠ for Group Term Life Insurance available to IEEE members! Apply for up to $500,000 in term life coverage.†

- **No medical exam required—just answer some health questions and other information.**
- **Everything is handled conveniently and securely online.**
- **You'll find out in as little as one day if you're approved.**

Take advantage of this valuable income protection today!**

1-800-493-IEEE (4333) | IEEEinsurance.com/Possibilities

1 LIMRA, Facts About Life 2016.

*The Group Term Life Insurance Plan is available only for residents of the U.S. (except territories), Puerto Rico and Canada (except Quebec). This plan is underwritten by New York Life InsuranceCompany, 51 Madison Ave., New York, NY 10010 on Policy Form GMR. This plan is administered by Mercer Health & Benefits Administration LLC. This coverage is available to residents of Canada (except Quebec). Mercer (Canada) Limited, represented by its employees Nicole Swift and Suzanne Dominico, acts as broker with respect to residents of Canada.

**For information on plan features, costs, eligibility, renewability, limitations and exclusions visit IEEEinsurance.com/Possibilites

†QuickDecision℠ is not available in Canada or Puerto Rico. May not available in all states. If you are not approved for coverage under the QuickDecision℠ tool, you may still apply online, but will follow the fully underwritten process.

Program Administered by Mercer Health & Benefits Administration LLC

In CA d/b/a Mercer Health & Benefits Insurance Services LLC

AR Insurance License #100102691

CA Insurance License #0G39709


DLR