FOR THE TECHNOLOGY INSIDER

MARCH 2023

IEEE Spectrum

Extending the Human Body With Extra Robotic Limbs

IEEE

VOLUME 60 / ISSUE 3 MARCH 2023

Contents

ON THE COVER: Illustration by John MacNeill

Illustration by Eddie Guy

Contents

Photography by Gabriela Hasbun

CONTRIBUTORS

DAVID C. BROCK

Brock is a historian of technology, director of curatorial affairs at the Computer History Museum, and director of the museum's Software History Center. Brock is occasionally lucky enough to use the museum's restored Alto, a computer whose origins and influences he details in "The Machine that Transformed Computing" on page 40.

CLAIRE BRUST

Brust, an AI research engineer at Duolingo, writes with Klinton Bicknell and Burr Settles about personalized language learning in this issue [p. 28]. Bicknell is the head of AI at Duolingo. Settles was formerly the research director and head of AI at Duolingo. Before joining the company, Brust was preparing to visit China, but Duolingo was no help because it lacked a Chinese course at the time. So when applying to work there, Brust says, "I made sure to mention that I was expecting a Chinese course soon!" And it swiftly appeared.

DARIO FARINA

Farina is an IEEE Fellow who holds the chair in neurorehabilitation engineering at Imperial College London. With Etienne Burdet, also at Imperial College London, Carsten Mehring of the University of Freiburg, and Jaime Ibáñez of the University of Zaragoza, Farina writes about a noninvasive system that may let humans control extra robotic limbs [p. 22]. He says his work exploits "a natural biological amplifier—the muscle—which lets us look at the activity of individual neurons" without piercing the skin.

MATT SVRCEK

Svrcek cofounded Mainspring Energy and leads research and development as its chief technology officer. While completing his doctorate at Stanford University, he received funding from the Global Climate and Energy Project to investigate the efficiency and emissions impact of reacting various fuels under high compression. There he helped design and develop the early prototype that served as the foundation for Mainspring's linear generator technology, described on page 34.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2023 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 60, Issue No. 3. 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, ℅ 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 3401 N. Heartland Dr, Liberty, MO 64068, U.S.A. Printed at Spenta Multimedia Pvt. Ltd., Plot 15, 16 & 21/1, Village Chikhloli, Morivali, MIDC, Ambernath (West), Dist. Thane. 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 https://www.ieee.org/about/corporate/governance/p9-26.html.

News

ENERGY

Big Push for Small Reactors >

Can nuclear power be reborn 100 megawatts at a time?

BY RAHUL RAO

On the island of Hainan, in southern China, one possible future is taking shape within a compact labyrinth of concrete and metal. Last July, a crane hoisted the upper half of a steel containment shell into place. Slowly, steadily, workers are piecing together a miniature nuclear reactor.

This is Linglong One, whose diminutive size is a drastic shift from the gigawatt-scale megaprojects that dominate nuclear energy today. But if one persistent cadre of nuclear optimists are right, then Linglong One could be a model for fission's future in an age of clean energy.

Small reactors won't save the day quite yet; depending on the country, there's still plenty of regulatory and logistical issues to hammer out. But, experts say, the 2020s could help set the foundations for a nuclear blossoming in decades to come.

"It's a really exciting time for the nuclear industry," says Victor Ibarra Jr., a nuclear engineer at the Nuclear Innovation Alliance think tank.

When it comes on line in 2026, Linglong One will have a capacity of 125 megawatts of electricity—equivalent to around 40 onshore wind turbines. Next to a large reactor (some 1,000 MW of electric power), 125 MW may seem insignificant. Why, after all, would an ambitious nuclear reactor designer want to go small? In part,

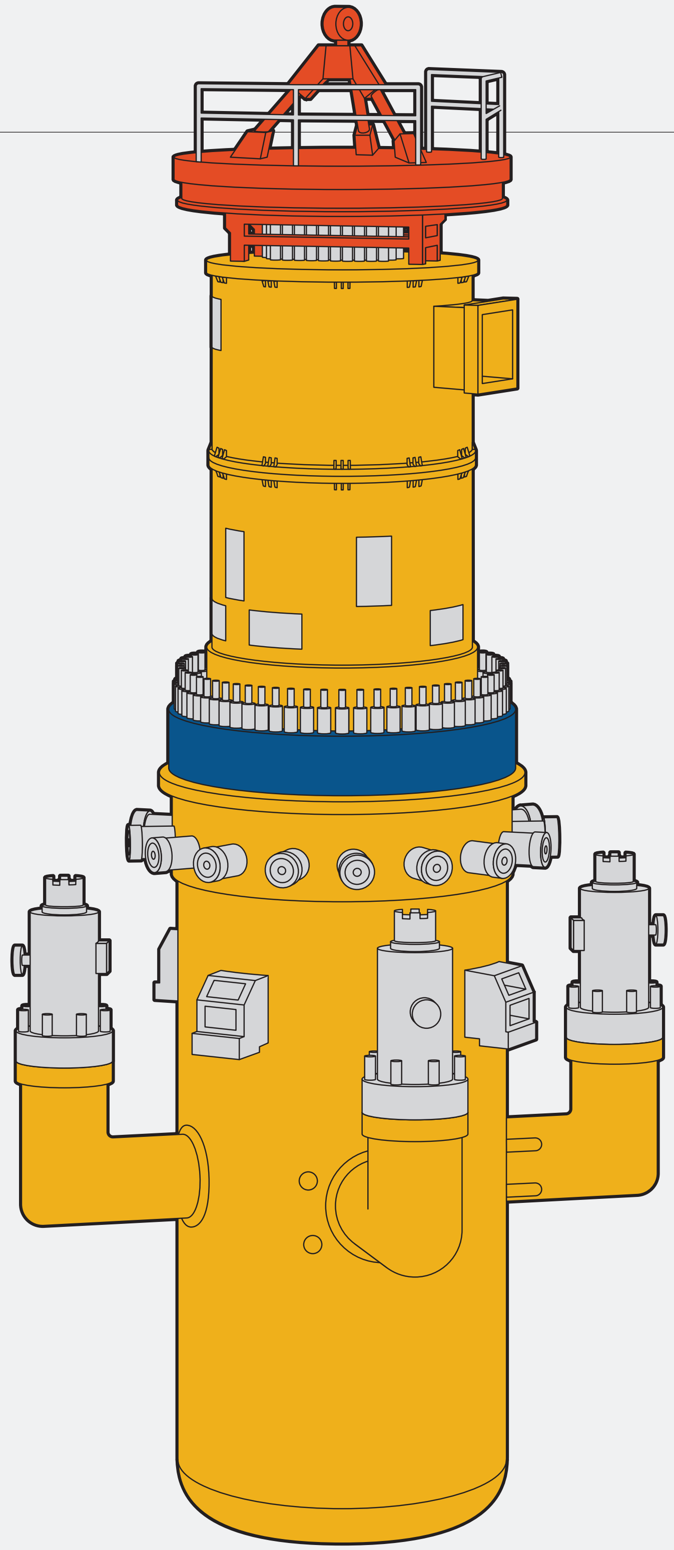

China is now developing a small modular nuclear reactor (SMR) called Linglong One, whose diminutive size (just 125 megawatts) masks an outsize ambition to spark a nuclear-fission renaissance.

Illustration by James Provost

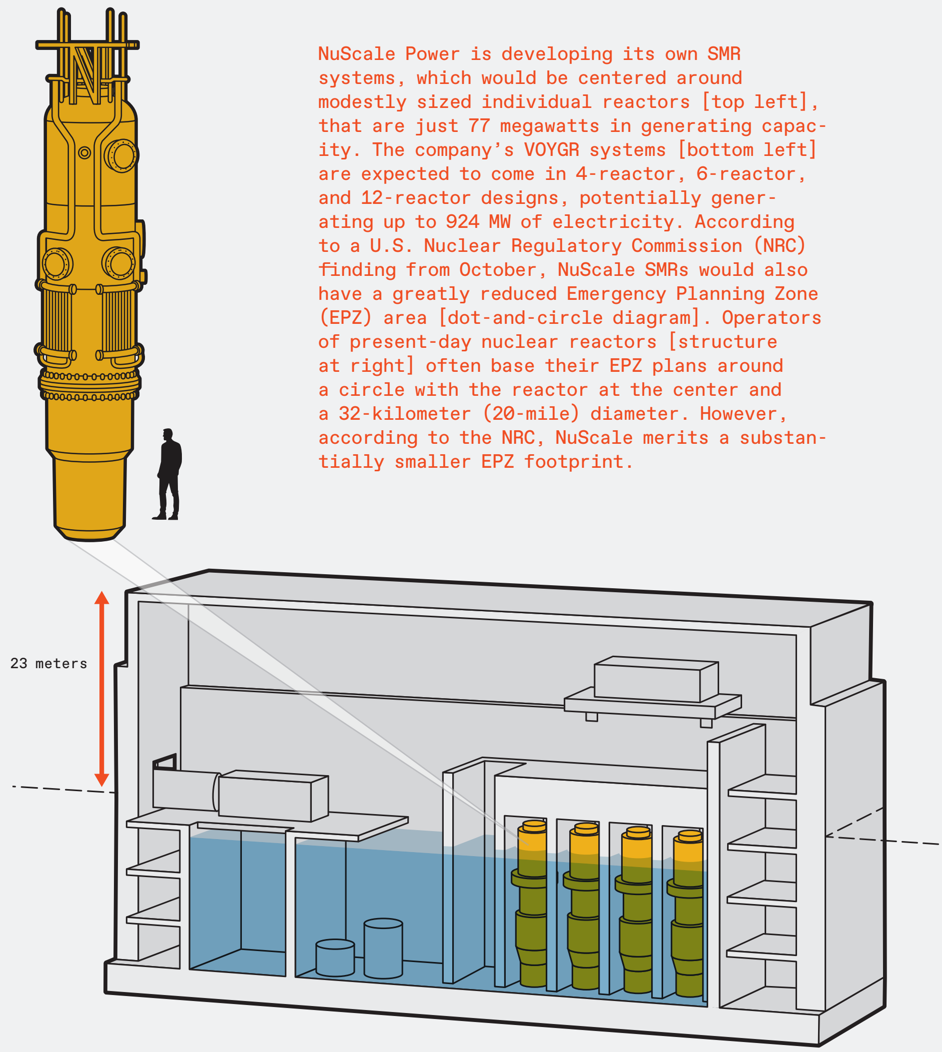

NuScale Power is developing its own SMR systems, which would be centered around modestly sized individual reactors [top left], that are just 77 megawatts in generating capacity. The company's VOYGR systems [bottom left] are expected to come in 4-reactor, 6-reactor, and 12-reactor designs, potentially generating up to 924 MW of electricity. According to a U.S. Nuclear Regulatory Commission (NRC) finding from October, NuScale SMRs would also have a greatly reduced Emergency Planning Zone (EPZ) area [dot-and-circle diagram]. Operators of present-day nuclear reactors [structure at right] often base their EPZ plans around a circle with the reactor at the center and a 32-kilometer (20-mile) diameter. However, according to the NRC, NuScale merits a substantially smaller EPZ footprint.

it's because large reactors can be expensive and prone to delays.

The twin 1,117-MW reactors at Georgia's Vogtle Electric Generating Plant, the only ones under construction in the United States, will come on line at least six years behind schedule. The 1,650-MW reactor under construction at Flamanville, in France, has experienced more than a decade of delays. Even more discouraging, nuclear's per-unit cost increased 26 percent between 2009 and 2019, while solar and wind power prices plummeted instead.

Still, a steadfast consensus remains that nuclear power isn't just desirable for a clean-energy transition—it's necessary. But some nuclear advocates feel that placing too many nuclear eggs in a single megaproject's basket is a bad idea. Instead, they think, a clean-energy transition might be better served with a fleet of smaller, more modular reactors, like Linglong One—hence the name small modular reactor (SMR).

SMRs may be smaller than today's average reactor, but they're also cheaper, less risky, and more flexible. Instead of building an airport, one analogy goes, crafting an SMR is like building an aircraft. And if nuclear reactors were aircraft, consider the SMR today a 1910s-era biplane. And it's still on the drawing board.

Once the manufacturing process scales up—if it ever scales up—makers of SMRs hope to be able to fabricate their components in a single factory, ship them out, and have them assembled on-site like flat-pack fission furniture.

Linglong One is the only one of its kind under construction today. If it's a success, China reportedly plans to use its design to power an untold number of construction projects and desalination plants. Let a thousand flowers bloom. Aside from two modified naval reactors on a boat in the Russian Arctic, every other SMR, everywhere else in the world, remains hypothetical.

But SMR plans are not in short supply. At least seven different developers plan to deploy SMRs in the United States before 2030. Most of them are demonstration reactors, not linked to the greater grid but a crucial stepping stone toward it.

Perhaps the largest SMR brand today, at least outside of China, is the U.S. startup NuScale, headquartered in Tigard, Ore. This company has developed a 77-MW SMR and envisions clumping 4, 6, and even 12 reactors together into larger power plants. NuScale has plans to build a U.S. plant in Idaho by 2030; the company is involved in the United Kingdom, Poland, and Romania as well. (It's important to recall, on the other hand, that a 2010 *IEEE Spectrum* story on the future of nuclear energy contained this projection: "NuScale is in talks with several undisclosed utilities and expects a first plant to be operational in 2018.")

In December, Rolls-Royce shortlisted three sites in England for a factory that it hopes will eventually manufacture the components for a 470-MW reactor. Rolls-Royce hopes to get the first of its reactors on the grid by 2029.

Several European countries have expressed interest in SMRs, especially in partnership with U.S. firms. The Czech state-owned energy company set aside land in the country's South Bohemian Region for an SMR project. Even France, a traditional nuclear powerhouse, plans a billion-euro investment in developing an SMR industry by 2030.

"I think, in the next 15, 20 years, there is going to be a real chance for SMRs to be commercially available and widely deployed," says Giorgio Locatelli, a nuclear project expert at Polytechnic University of Milan.

And, yet a countervailing force—regulation—has long been notorious for applying the cautionary brakes.

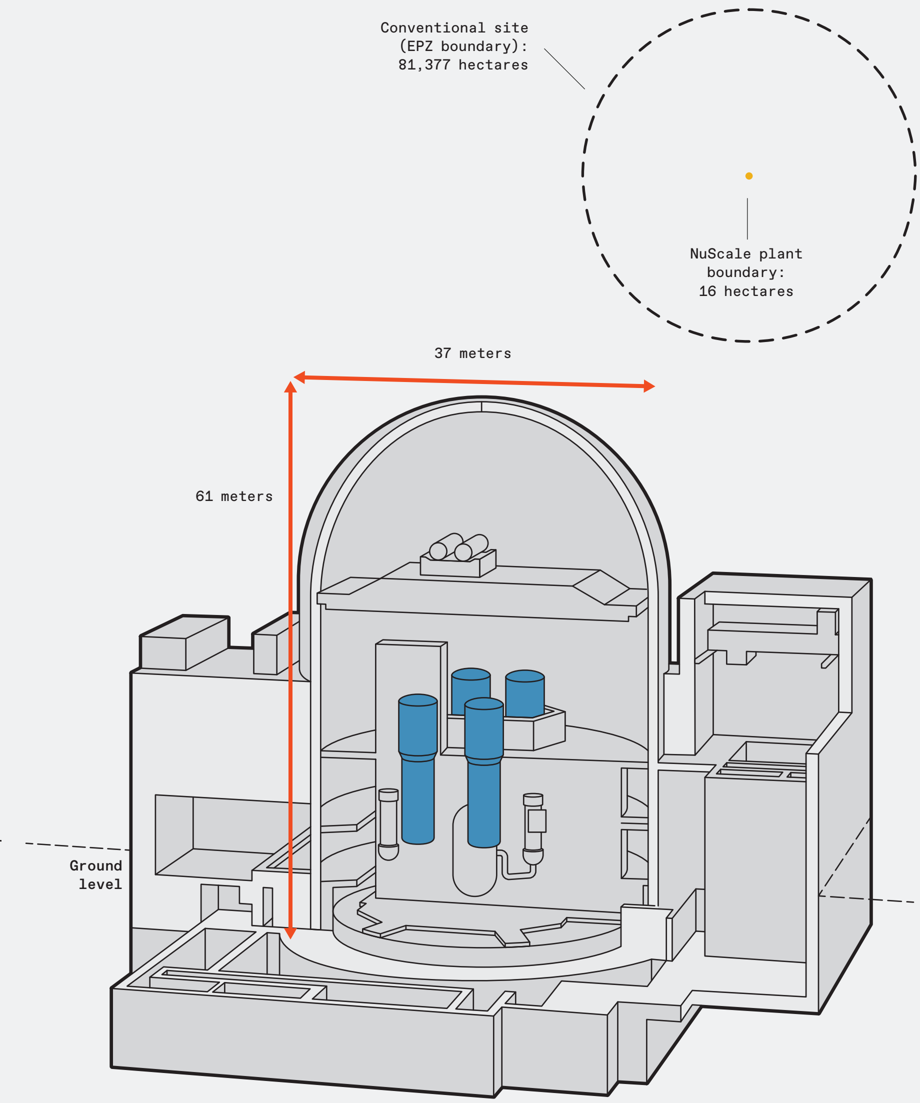

Some of those regulators' concerns come in the form of unanswered questions. Just like their larger counterparts, SMRs will produce nuclear waste. According to one recent Stanford and University of British Columbia study, SMRs yield more nuclear waste than even conventional nuclear plants. What will nuclear authorities do with that waste? No one knows, partly because every reactor design is different, and no one is sure what the SMR fleet will look like in a decade or two. Moreover, some analysts worry that bad actors could co-opt certain SMR designs to create weapons-grade plutonium.

Don't extinguish the candle on SMRs yet, though. A few nuclear regulators have at least begun to bend for SMRs.

According to Patrick White, a nuclear-regulation expert at the Nuclear Innovation Alliance, the U.S. Nuclear Regulatory Commission (USNRC) has been among the regulators who are engaging SMR developers. In 2019, President Trump signed the Nuclear Energy Innovation and Modernization Act, mandating that the USNRC create a new process specifically for future reactor designs. The result, called Part 53, is slated to become an option for SMR developers by 2027, though White says it may open up as early as 2025. What Part 53 will look like isn't yet certain.

The U.K. government, which has thrown its weight and funding behind SMRs, opened a modified regulatory approval process to SMR developers in 2021; Rolls Royce was the first to follow, and six other firms have applied. In June, French, Finnish, and Czech regulators announced they were working together to review Nuward, an SMR design backed in part by the French government; this project, they say, is a dry run for future SMR licensing.

Where, then, can SMR operators turn for fuel?

Most of today's large nuclear reactors use fuel with 3 to 5 percent uranium-235—the naturally available uranium isotope that can sustain a nuclear chain reaction. While SMR designs are diverse, many will need fuel that's more like 5 to 20 percent uranium-235. This latter type of fuel is known as high-assay low-enriched uranium (HALEU). Today, only one company commercially sells HALEU: Techsnabexport (TENEX), a subsidiary of Rosatom—Russia's state-owned nuclear-energy company.

So far, Rosatom had, as of January, avoided Western sanctions over Russian aggression in Ukraine. But TENEX has still become untenable for many of its would-be customers. For instance, TerraPower, which hoped to switch on a demonstration SMR in a deprecated coal plant in Wyoming in 2028, delayed its launch by two years due to fuel issues.

"With the exception of allies of Russia, it's not just an obstacle, it's a flat-out barrier right now," says Adam Stein, a nuclear-energy analyst at the Breakthrough Institute.

The United States has started to pierce that barrier. The 2022 Inflation Reduction Act invested US $700 million to research and develop ways of producing and transporting HALEU within the country. Ibarra, of the Nuclear Innovation Alliance, welcomes the investment, but according to him, it's a "short-to-medium-term solution." It may not be enough. One estimate suggests that the replenished HALEU won't be ready until 2028.

For many SMR-interested parties, HALEU may be a key goal in the years ahead: establishing a global HALEU supply chain that's less dependent on Russia, less susceptible to global geopolitics. It remains to be seen how Europe or the U.K. will respond—if they'll let the United States take the lead, or if they'll take action themselves. ■

ARTIFICIAL INTELLIGENCE

Memristors Tackle AI's Power Problem › All the smarts, consuming 1/800 the energy

BY CHARLES Q. CHOI

Memristive devices, which mimic the synapses between neurons, could serve as hardware for neural networks that copy the way the brain learns. Now two new studies may help solve key problems these components face with production and reliability as well as in achieving wider application.

Memristors, and related memristive devices, can remember which electrical state they were switched to after their power has been turned off. They can be used to both compute and store data, greatly reducing the energy and time that now is lost shuttling data back and forth between processors and memory. Such brain-inspired, or neuromorphic, hardware may also prove ideal for implementing neural networks—AI systems used in applications such as analyzing medical scans and guiding autonomous vehicles.

Today's memristive devices suffer from low production yields and unreliable performance. To help overcome these challenges, researchers in Israel and China have worked to achieve yields of 100 percent. They pulled it off by fabricating the memristive devices on a standard CMOS production line. The resulting silicon synapses have an energy efficiency that is 350 times as much as a high-performance Nvidia Tesla V100 graphics processing unit in handling multiply-accumulate operations, the most basic function in neural networks.

The scientists developed the new devices using the floating-gate transistor technology used in commercial flash memory. Whereas conventional floating-gate transistors have three terminals, the new components have only two. This greatly simplified fabrication and operation and reduced their size. In addition, the devices have only binary inputs and outputs, eliminating the need for the large, energy-hungry analog-to-digital and digital-to-analog converters often used in neuromorphic hardware, says study senior author Shahar Kvatinsky, an associate professor of electrical and computer engineering at the Technion–Israel Institute of Technology, in Haifa.

The new devices were able to endure more than 100,000 cycles of programming and erasing, which is done with voltage pulses. In addition, they showed only moderate device-to-device variation and are projected to retain data for more than 10 years.

The researchers use an array of some 150 of these components to implement a kind of neural network that uses only binary signals. In experiments, it could recognize handwritten digits with roughly 97 percent accuracy. Kvatinsky says that this work "is just a start—a proof of concept and not a whole integrated chip or a large neural network. Integration and scaling up is a major challenge."

In another study, a team of French researchers used memristors for Bayesian reasoning, a technique in which prior knowledge helps compute the chances that an uncertain choice might be correct. Many AI computations are nearly inscrutable, but the French system is fully explainable, which means people can use it to come up with improvements. And because it can incorporate expert knowledge, it can often perform well even when there is little available data. Even so, "it is just not obvious how to compute Bayesian reasoning with memristors," says study coauthor Damien Querlioz, a research scientist at CNRS, Université Paris-Saclay.

Implementing Bayesian reasoning using conventional electronics requires complex memory patterns, "which increase exponentially with the number of observations," says neuromorphic scientist Melika Payvand at the Institute of Neuroinformatics, in Zurich, who did not participate in either study. However, Querlioz and his colleagues "found a way of simplifying this," she notes.

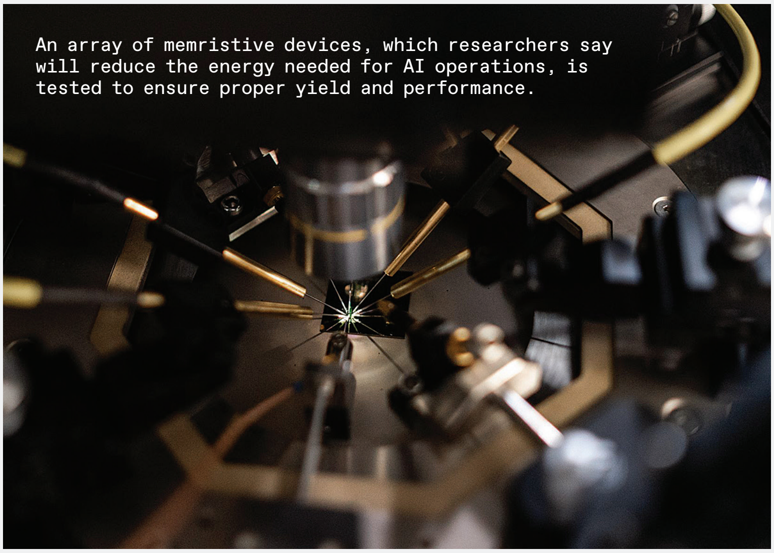

An array of memristive devices, which researchers say will reduce the energy needed for AI operations, is tested to ensure proper yield and performance.

TECHNION-ISRAEL INSTITUTE OF TECHNOLOGY

The scientists rewrote Bayesian equations so that a memristor array could perform statistical analyses that harness randomness—also known as stochastic computing. Using this approach, the array generated streams of semirandom bits at each tick of the clock. These bits were often zeros but were sometimes ones. The proportion of zeros to ones encoded the probabilities needed for the statistical calculations the array performed. This digital strategy uses relatively simple circuitry compared to that of nonstochastic methods, all of which reduce the system's size and energy demands.

The researchers fabricated a prototype circuit incorporating 2,048 hafnium oxide memristors on top of 30,080 CMOS transistors on the same chip. In experiments, the researchers had the new circuit recognize a person's handwritten signature from signals beamed from a device worn on the wrist.

Bayesian reasoning is often thought of as computationally expensive with conventional electronics. The new circuit performed handwriting recognition using 1/800 to 1/5,000 the energy of a conventional computer processor, suggesting that "memristors are a highly promising lead to provide low-energy-consumption artificial intelligence," Querlioz says.

The new device also can instantly flick on and off, which suggests it can be called upon only when needed, to conserve power. Moreover, it is also resilient to errors from random events, making it useful in extreme environments, the researchers say. All in all, the new circuit "excels in safety-critical situations, where high uncertainty is present, little data is available, and explainable decisions are required," Querlioz says. "Examples are medical sensors or circuits for monitoring the safety of industrial facilities."

A key obstacle that Bayesian systems face is how to scale them up to handle larger problems or networks. Querlioz notes that the team are now fabricating a considerably scaled-up version of their device. He notes that their circuit is currently specialized for certain types of Bayesian computation, and they want to create more adaptable designs in the future.

Both studies appeared 19 December in the journal *Nature Electronics*. ■

JOURNAL WATCH

Electric Aircraft Won't Take Off Without More Thrust

The race for all-electric planes is underway, and some early designs are making headlines. But today's early all-electric airplane models have been designed to carry no more than 30 passengers over short distances. To truly lower greenhouse gas emissions requires much larger electric planes.

One research team proposes to create one by making an all-electric version of NASA's N3-X aircraft, which in its turbojet-powered version is slated to carry 297 to 330 passengers, beginning in 2040. The researchers describe three possible electric-power-system designs in a study published in the December print issue of *IEEE Transactions on Transportation Electrification*. The lead author is Mona Ghassemi, the director of the ZEROES (Zero Emission, Realization of Optimized Energy Systems) Lab at the University of Texas at Dallas.

The researchers suggest replacing the N3-X's two turbo-electric engines with four electrochemical energy units (EEUs), which include batteries, fuel cells, and supercapacitors.

The main challenge in fully electrifying the N3-X—and large planes in general—is the immense energy demand for thrust during liftoff, which requires about 25 megawatts of power. In their paper, the researchers propose three electric-power-system designs, whereby all electric power is diverted to thrust during takeoff. Analysis shows that two of these designs could be feasible in real life, once the battery technology is advanced enough.

"With future projected advancements in...batteries, the required specific energies for the envisaged wide-body all-electric aircraft may be achieved within the next 25 years," Ghassemi notes, citing compact-fusion reactors as another option.

Consequently, more infrastructure will be needed to support the energy requirements of these all-electric aircraft. Next, Ghassemi says, her lab will design cables capable of supporting the high-voltage needs of their all-electric N3-X design.
—MICHELLE HAMPSON

NASA's proposed all-electric N3-X aircraft would carry 10 times as many passengers as today's electric planes.

NASA

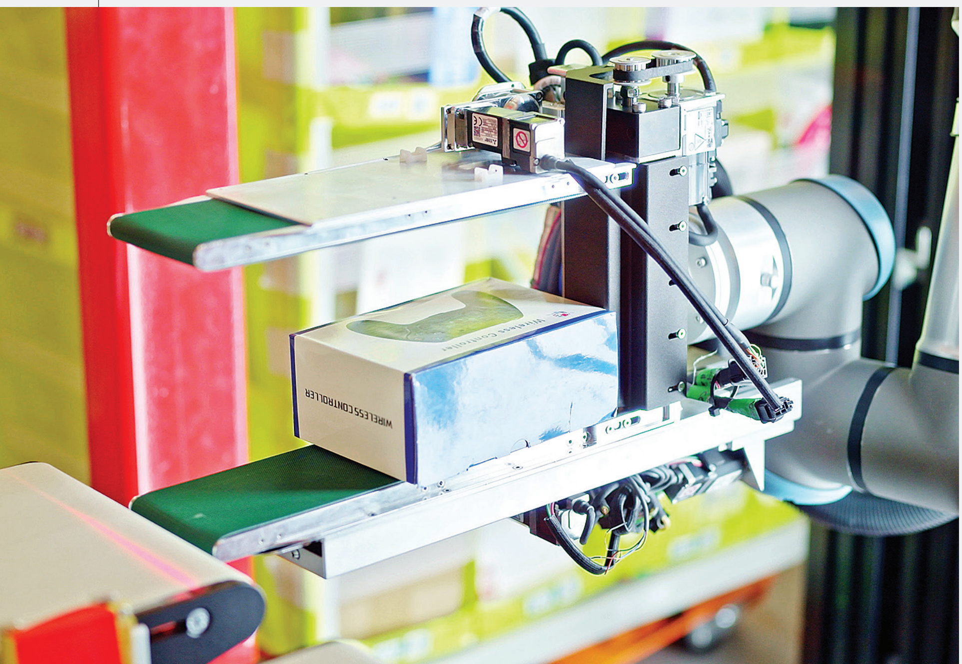

The end-of-arm tool (EOAT) is a key component of Amazon's new stowing robot.

ROBOTICS

Amazon Tells Warehouse Robots Where to Stow It

> Cramming items into warehouse pods is a surprisingly difficult problem

BY EVAN ACKERMAN

Manipulation robots in warehouses mostly take things out of bins. Such picking, as it is called, can be tricky, especially when the different types of items number in the millions. But robots are getting good at it. They're also starting to get pretty good at putting items into bins, a task called stowing.

Just a month ago, Amazon introduced Sparrow, which it describes as "the first robotic system in our warehouses that can detect, select, and handle individual products in our inventory." However, the surrounding system is doing much of the heavy lifting by presenting Sparrow with very robot-friendly bins. At highly automated warehouses everywhere, most bins either contain only identical items or just a few different items.

But robot-friendly bins are simply not the reality for the vast majority of items in an Amazon warehouse, largely because of the usual human penchant for stowing things in a way that's intuitive for us but confounds machines. Siddartha Srinivasa, professor at the Paul G. Allen School of Computer Science & Engineering at the University of Washington and, until late last year, director of Amazon Robotics AI, described the problem of stowing items as "a nightmare.... Stow fundamentally breaks all existing industrial robotic thinking." But over the past few years, Amazon Robotics researchers have put some serious work into solving it. This work focuses on the way most Amazon warehouses are actually run—with people doing most of the complex manipulation.

In a typical warehouse, Amazon's drive units—mobile robots with shelves on top, called pods—drive themselves past people, like waiters bringing various delicacies to customers. The people pick items off the shelves to build up orders. Thus, the picking task is handled by robots and human beings working together. But people get the items into Amazon's warehouse workflow by stowing them on those mobile shelves in the first place. And, it turns out, people do so in a way that uses space most efficiently. This makes a lot of sense.

When an Amazon warehouse gets a new shipment of stock, the obvious thing to do might be to call up a pod with enough empty shelves to stow all of the items at once. That way, when someone places an order for an item, the newly stocked pod shows up and a human can pick the item off one of the shelves. However, if the pod gets stuck or breaks or is otherwise inaccessible, then nobody can get the items they requested. This slows down the entire system. Amazon's strategy is to instead distribute the stock across multiple pods, so that some are always available.

The process for this distributed stow is random in that a human stower might get a couple of items to put into whatever pod shows up next. It's up to the stower to decide where the items fit best, including into bins with other random items already in them, and Amazon doesn't really care as long as the inventory system keeps track of where the items end up.

Two issues immediately stood out to observers assessing the system: First, the way that Amazon products are stowed is entirely incompatible with conventional bin-picking robots. Second, it's easy to see why such stowing is a nightmare for robots. Not only must the robot carefully manipulate a jumble of objects to make room in a bin, it must also deal with elastic bands

The robot includes a tool with twin conveyor belts [in green] designed to push items into storage bins.

that keep items from spilling out of the bins but also hide things from view and impede manipulation.

"For me, it's hard, but it's not too hard—it's on the cutting edge of what's feasible for robots," says Aaron Parness, senior manager of applied science at Amazon Robotics AI. "It's crazy fun to work on."

Parness came to Amazon from Stanford and the Jet Propulsion Laboratory, where he worked on robots like StickyBot and LEMUR and was responsible for a microspine gripper, designed to grasp asteroids in microgravity. "Having robots that can interact in high-clutter and high-contact environments is super exciting because I think it unlocks a wave of applications," he adds.

What makes stowing at Amazon both cutting edge and nightmarish for robots is that it's a task that has been highly optimized for human stowers. This means that any robotic solution that would have a significant impact on the human-centered workflow is probably not going to get very far. So, Parness, along with senior applied scientist Parker Owan, had to develop hardware and software that could solve the problem as is.

Here's what they came up with: On the hardware side, there's a hook system that lifts the elastic bands out of the way to provide access to each bin. But that's the easy part; the hard part is embodied in the end-of-arm tool (EOAT), which consists of two long paddles that can gently squeeze an item to pick it up, with conveyor belts on their inner surfaces to shoot the item into the bin. An extendable thin metal spatula of sorts can go into the bin before the paddles and shift items around to make room when necessary.

To use all of this hardware requires some very complex software, because the system needs to be able to perceive the items in the bin (some of which, together with the elastic bands, may be blocking the view of other items), estimate the characteristics of each item, develop a plan for safely shoving those items around to maximize available bin space based on the object to be stowed, and then execute the motions to make it all happen. By breaking the process down this way in the lab, the Amazon researchers have been able to achieve stowing success rates of better than 90 percent.

Now, prototypes are stowing actual inventory items at an Amazon fulfillment center in Washington State. The goal is to be able to stow 85 percent of the products that Amazon stocks (millions of items). And 85 percent is good enough, because the system will be installed within the same workflow that the company's human workers use. If the system can't handle a task, it just passes it along to a person.

Regarding the other half of Amazon's warehouse problem—picking—Amazon has begun to engage with the academic community more. "My team sponsors research at MIT and at the University of Washington," says Parness. "And the team at University of Washington is actually looking at picking. Stow and pick are both really hard and really appealing problems, and in time, I hope I get to solve both!" ■

"For me, it's hard, but it's not too hard—it's on the cutting edge of what's feasible for robots."

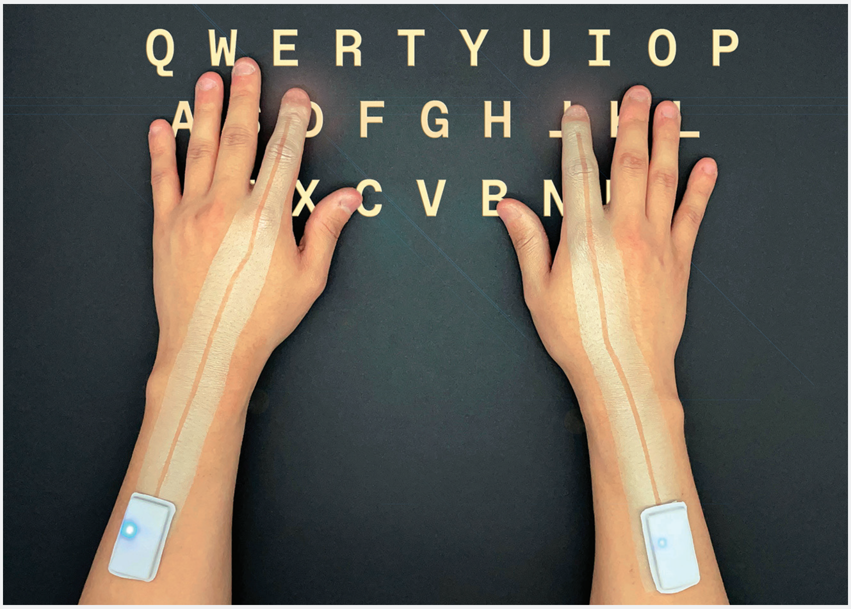

Any surface can double as a keyboard when a new spray-on sensor is tracking and translating the movements of the wearer's fingers.

KYUN KYU "RICHARD" KIM/STANFORD UNIVERSITY

SENSORS

Spray-on Smart Skin Translates Your Hand Motions > No camera, keyboard, glove, or VR motion tracker required

BY PRACHI PATEL

A **new AI-learning scheme** combined with a spray-on smart skin can recognize and decipher the movements of human hands while users type, use sign language, or even indicate the shape of simple, familiar objects. The system's developers say it can be trained quickly, with limited data, and should work for everyone.

This trick would come in handy, as it were, as a silent means of communicating with other people and with machines, including computer games operating in virtual reality. Surgeons might use gestures to remotely control medical devices, and patients might someday use them to control prosthetic arms and legs.

The point is to overcome weaknesses in existing methods for tracking gestures, says Sungho Jo, a professor in the school of computing at the Korea Advanced Institute of Science and Technology. He and his collaborators, at Seoul National University and Stanford University, reported their work in the journal *Nature Electronics*.

Today's gesture-recognition technology relies on bulky wristbands that measure electrical signals produced by muscles, or on gloves equipped with strain sensors at each joint. Another approach involves using multiple cameras that track human motion from several angles, missing important details whenever something blocks the view of a hand or some other key object. Existing software is also cumbersome, typically involving machine-learning models based on supervised learning algorithms that are computationally intense and require a lot of hand-labeled data for each new user and task.

"We tried to create a gesture-recognition system that is both lean enough in form and adaptable enough to work for essentially any user and tasks with limited data," Jo says.

The new system has two key parts. One is a mesh made of millions of silver nanowires coated with gold and embedded in a polyurethane plastic coating. The mesh, Jo says, is both durable and stretchy: "It conforms intimately to the wrinkles and folds of each human finger that wears it." A portable machine can print the mesh directly on the skin so thinly as to be almost imperceptible, he says. It is comfortable, and it can last through a few days of daily use, unless it is rubbed off with soap and water.

The team prints the mesh directly onto the back of a user's index finger. The nanowire network senses tiny changes to electrical resistance as the skin stretches underneath. As the hand moves, the nanomesh creates unique signal patterns that it wirelessly sends via a lightweight Bluetooth unit to a computer for processing.

This is where the second key part of the system—the AI—kicks in. A machine-learning algorithm maps the changing patterns in electrical conductivity to specific physical tasks and gestures. The researchers used random hand and finger motions from three different users to help the AI learn the general correlation between motions.

Based on this prior knowledge, the researchers then trained the system to distinguish between the signal patterns generated from specific tasks, such as typing on a phone, two-hand typing on a keyboard, and holding and manipulating six different shapes. Each user performed gestures related to the tasks five times to generate a small data set, which the researchers used to train the model. The algorithm learned to recognize when the user was typing a specific letter on the keyboard, for instance, or tracing the sloped surface of a pyramid. In tests, the system was then able to recognize objects being held, and sentences being typed on a virtual keyboard, by a new user.

"Our learning scheme is not only far more computationally efficient but also versatile, as it can rapidly adapt to different users and tasks with few demonstrations," Jo says.

Jo and his colleagues plan to place nanomesh devices on multiple fingers to capture a larger range of hand motions. More sensors will yield more data to be analyzed, he says, so the researchers will need to carefully consider the balance between accuracy and a reasonable computational workload. ■

INNOVATE

TRANSFORM OUR FUTURE

APPLICATIONS OPEN 20 MARCH – 5 MAY
AWARDS UP TO $500K

R&D FUNDING PROGRAM

The National Reconnaissance Office Director's Innovation Initiative (DII) funds cutting-edge scientific research in a high-risk, high-payoff environment to discover innovative concepts and creative ideas that transform overhead intelligence capabilities and systems for future national security intelligence needs. The program seeks the brightest minds and breakthrough technologies from industry, academia, national laboratories, and U.S. government agencies.

Visit the website for Broad Agency Announcement and Government Sources Sought Announcement requirements.

www.nro.gov/InnovateNRO/DII
DII@NRO.MIL

AREAS OF INTEREST

- Apertures
- Communications
- Remote Sensing
- Satellite System Design
- Sense-Making
- Other Disruptive Concepts & Technology

NATIONAL RECONNAISSANCE OFFICE
UNITED STATES OF AMERICA

Intentional Grounding

By Willie D. Jones

A team of researchers from the University of Geneva has worked out how to use lasers to improve lightning rods' ability to protect nearby structures from nature's violent electric power surges. It's long been understood that a conventional lightning rod will protect an area with a radius roughly equal to its height. Because it is impractical to make rods that extend beyond a certain length, today's sacrificial electric conductors leave some areas vulnerable. The Swiss team shot high-power laser pulses into the sky during a thunderstorm, creating an electrically conductive channel in the air. This channel drew lightning bolts from a wide area down to the tip of the metallic rod, which conveyed the electricity harmlessly into the ground.

REPORTING BY CHARLES Q. CHOI

PHOTOGRAPH BY MARTIN STOLLBERG/TRUMPF

Hands On

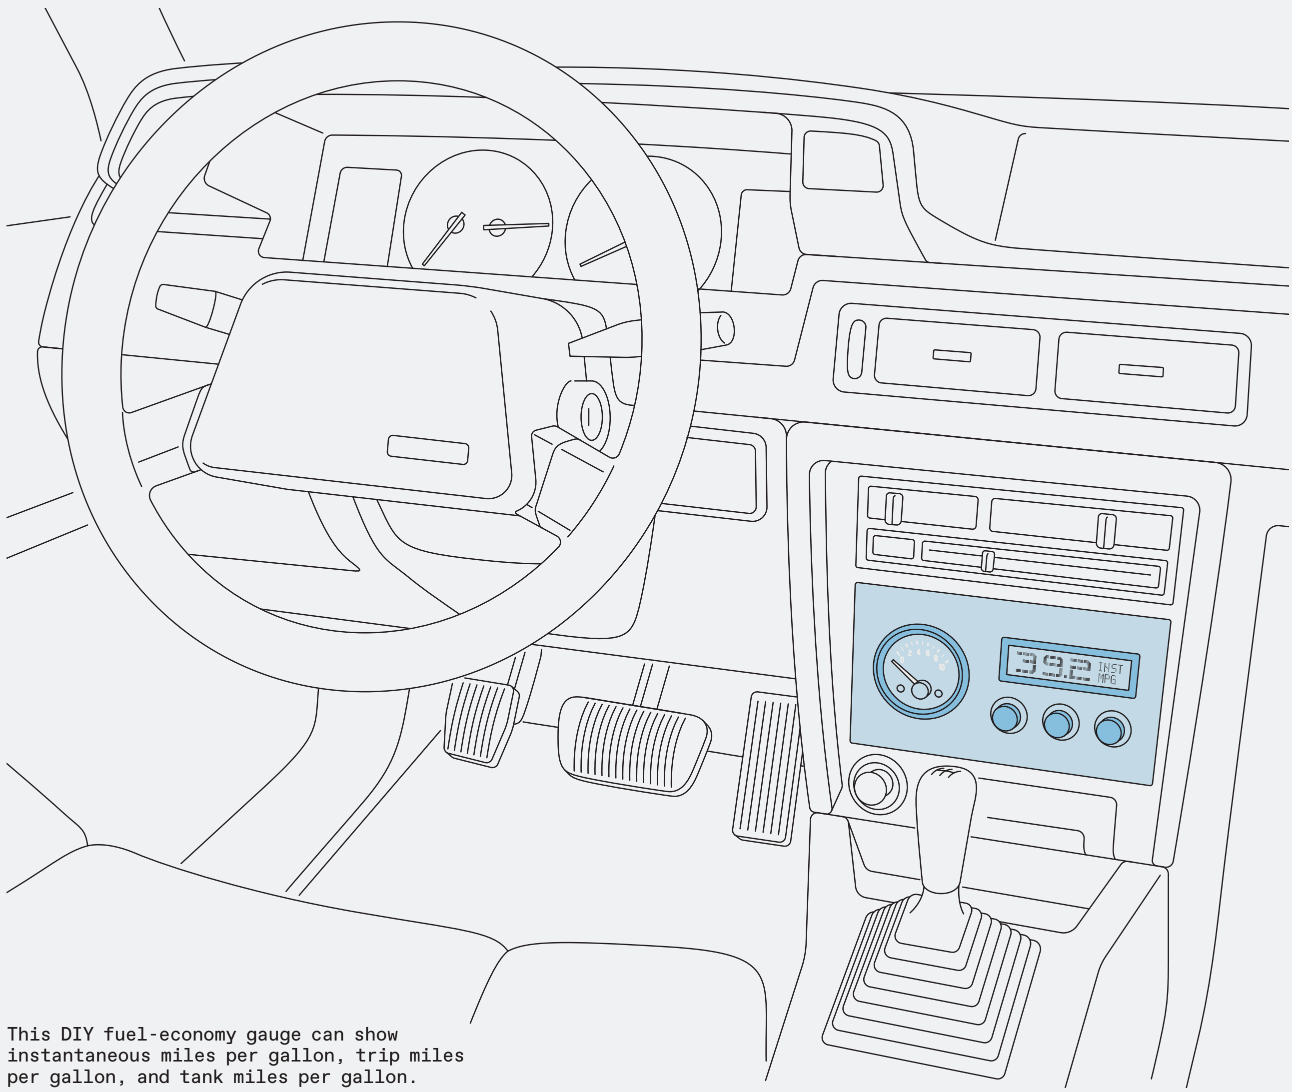

This DIY fuel-economy gauge can show instantaneous miles per gallon, trip miles per gallon, and tank miles per gallon.

A Retro Fuel-Economy Computer › This Arduino-based project outputs MPGs on a vintage dashboard gauge

BY DAVID SCHNEIDER

Russia's invasion of Ukraine and the skyrocketing fuel prices that ensued last year got me thinking about how to cut down on my consumption of gasoline. I thought briefly about purchasing a car that got better gas mileage, but it made sense first to see how much fuel economy I could eek out from the aging econobox I am now driving: a 1991 Toyota Corolla.

Possible strategies for improving its fuel economy include putting on some

Illustrations by James Provost

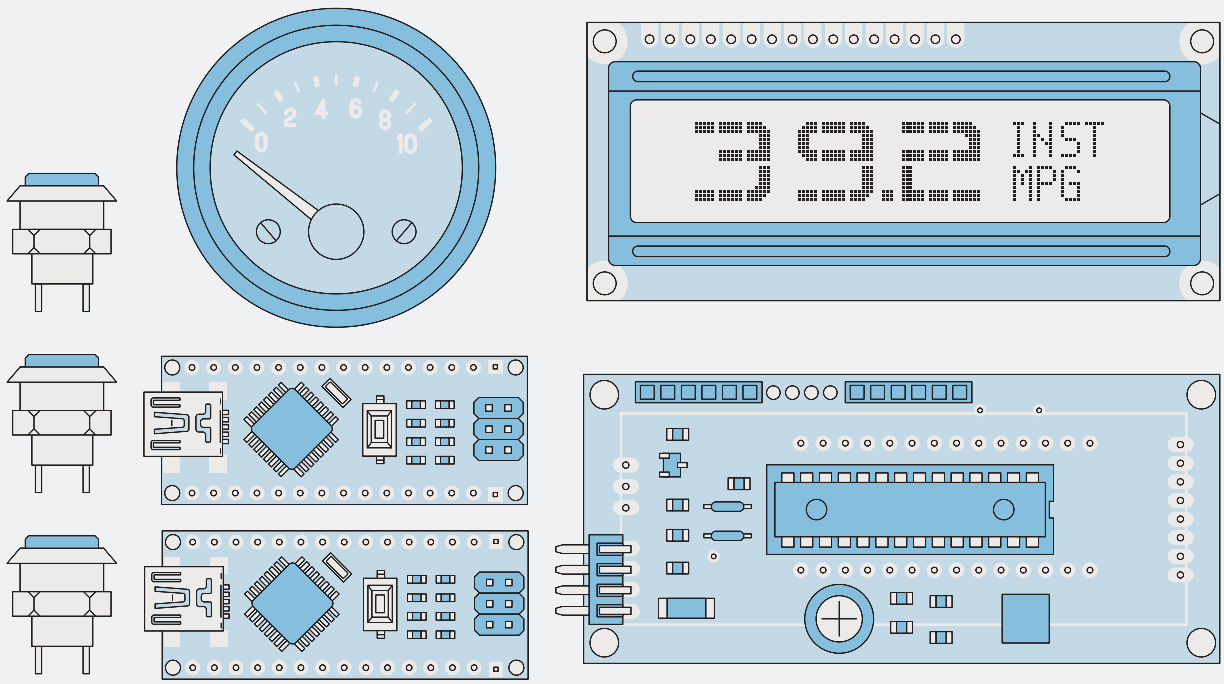

The components used in this project include two Arduino Nano boards, three push buttons, and a 16-by-2 LCD display board. The unit also includes a vintage VDO oil-pressure gauge, used to show instantaneous fuel economy in units of miles per gallon divided by 10.

low-rolling-resistance tires, adding a shroud below the engine compartment (to reduce aerodynamic drag), even removing the alternator (as some racers do) or having it operate only when braking. And just driving it differently could also help a great deal.

I knew, though, that it was going to be difficult to test out different driving techniques without more immediate feedback of their effects on fuel economy. Many newer cars show instantaneous fuel consumption on the dash. Mine doesn't. Nor would this be something easy to add.

Owners of most vehicles can purchase, for example, a ScanGauge and plug it into the OBD-II (On Board Diagnostics II) data port of their cars, which would allow them to view an estimate of instantaneous MPG while driving. But my car was manufactured years before OBD-II became standard.

So I decided to revisit a project I first explored in 2009, when I installed an Arduino-based fuel-economy computer dubbed "MPGuino" in the car that I was driving at the time. An MPGuino taps into two signals available at a car's electronic control unit (ECU): One is from the vehicle-speed sensor; the other runs the fuel injectors. With those two signals, an MPGuino can calculate how fast the car is going and how much fuel it is using, so it can show the driver instantaneous fuel economy in miles per gallon (MPG). It also records the amount of fuel used and the distance traveled over longer periods, so it can display average MPG for a single trip, or for the last tank of gas.

I recall that this gizmo was very helpful. But having since sold the car in which I had installed it, I would need now to find another MPGuino for my Corolla.

Although the software side of MPGuino is still being actively developed, hardware for it is now harder to come by. The companies listed on the MPGuino Wiki page as selling MPGuino boards or kits at one time are now out of that business. But it wasn't too difficult to cobble one together from schematics using an Arduino (I used an Arduino Nano), a 16-by-2 LCD display board, and a few discrete components.

I recalled, though, how much I didn't like having to look at that tiny LCD display while driving. So this time around I was determined to do better.

Some MPGuino enthusiasts, I discovered, were sending its output to a color flat-panel display—for the modern, glass-cockpit look. I decided to go in entirely the other direction and have mine show instantaneous gas mileage on an analog gauge.

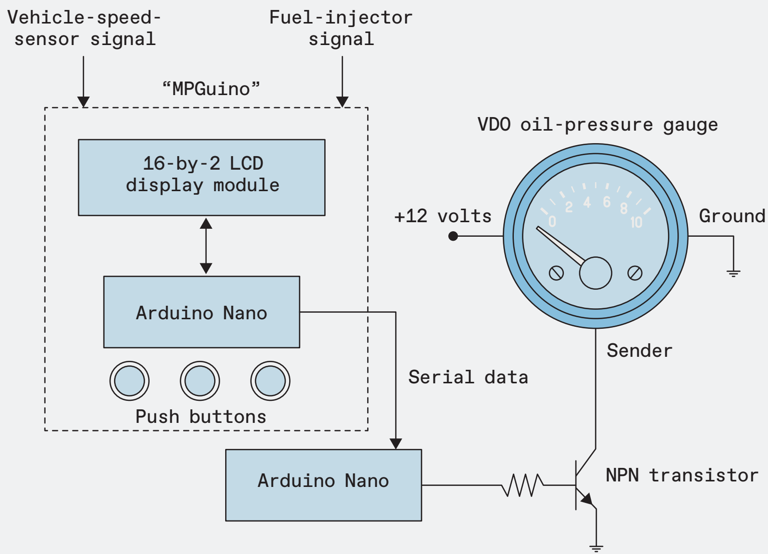

A standard MPGuino—which includes an Arduino, an LCD display and three push buttons, along with some discrete components [not shown]—is supplemented with a second Arduino and a vintage oil-pressure gauge. The variable-resistance sender unit normally used to drive that gauge is replaced here by an NPN transistor, which is switched on and off rapidly by a pulse-width-modulated signal from the second Arduino.

For that, I obtained (through eBay, naturally) a vintage VDO oil-pressure gauge, one that was salvaged from a Porsche that was not too much older than my car. In my youth, I used to drive a Porsche, but just because I am now tooling around in a tattered Corolla doesn't mean I can't still enjoy that cool look of VDO instrumentation!

Using that gauge with standard MPGuino hardware proved tricky though. While I'm pretty sure it could be done, the MPGuino code is quite complex, involving interrupts setting and clearing flags. My mind reeled at the prospect of trying to hack on it. So I decided to add a second Arduino Nano to this project, which I figured I could use initially to drive the VDO gauge and later might employ to calculate other parameters for display. It has a serial data connection with the Arduino in the MPGuino, which I found worked best with a 2016-vintage version of the MPGuino software, one tailored to output data serially in different formats.

The VDO gauge I obtained is normally connected to a sender unit that varies in resistance with changes in oil pressure. To mimic that variable resistance, I used a general-purpose NPN transistor. The base of this transistor is driven (through a 220-ohm resistor) by one of the digital output pins on the second Arduino, which outputs a pulse-width-modulated (PWM) wave of the appropriate duty cycle to drive the meter.

It took some experimentation to determine how to calibrate this analog display. Initially, I wrote code for the second Arduino that would slowly increase the duty cycle of the PWM signal, while showing the current duty cycle through the Arduino's serial monitor. I noted the PWM values that corresponded to each digit on the gauge. I could then use linear interpolation to set the PWM value for numbers between those points, coding things so that the gauge would show MPG divided by 10. (The range on this particular oil-pressure gauge goes from 0 to 10 bars.) The needle on the gauge was a little twitchy initially, so I amended my code so that the gauge would display a running average of the MPG values calculated by the MPGuino.

Tapping into the two needed signals at my car's ECU proved easy enough—I found the relevant wires without having to disassemble the dashboard, and I was able to figure out which they were from my car's factory-service manual.

But I would also need to tap into the car's 12-volt power. Normally, an MPGuino should be connected to a 12-V source that remains always on. Parasitic drain is not an issue, because the MPGuino turns off the display and goes into a low-power sleep mode soon after the car is shut down.

Connecting my device in the same way would have been problematic, because the VDO gauge and second Arduino I added would have continuously drawn power. So I added a second power input, tapping into the switched 12-V power that runs the car's radio. This powers the added Arduino (through a second 5-V regulator) and the VDO gauge, which both then turn completely off when I shut down the engine.

It's not yet clear how many more miles per gallon I'll be able to achieve just from changes in how I drive the car. But even if I'm not able to improve my gas mileage all that much, I'm finding that having the information presented on a gauge from a Porsche in my old Corolla really makes me smile. ■

I used to drive a Porsche, but just because I am now tooling around in a tattered Corolla doesn't mean I can't still enjoy that cool look of VDO instrumentation!

BY DANIEL P. DERN

Careers: Eugene H. Spafford

He's been fighting cybersecurity threats for decades

Eugene H. Spafford began working in computer cybersecurity before it was even a recognized field.

BRIAN POWELL

During **Eugene H. Spafford's** more than three decades as professor of computer sciences at Purdue University, in West Lafayette, Ind., he has made groundbreaking contributions to computer and network security. A member of the Cyber Security Hall of Fame, he is considered one of the most influential leaders in information security.

But he didn't start out aiming for a career in cybersecurity. Indeed, the field didn't really exist when he graduated from the State University of New York at Brockport with a bachelor's degree in math and computer science in 1979. Spafford then went to Georgia Tech to pursue a master's degree in information and computer science.

In the early '80s, the IEEE Fellow recalls, computer security consisted primarily of formal verification—using mathematical models and methods—and cryptography, focused on mainframes.

"We didn't have commercial networking," Spafford says. "Viruses, malware, and other cyberthreats had barely emerged. There were no tools, experts, or jobs—yet." However, computer security became a hobby of his.

"I did a lot of reading and studying on where computers might be used and where they could go wrong, as well as reading science-fiction books that explored those possibilities," he says.

Meanwhile, his graduate and postdoc work revolved around more traditional areas of computing. "The faculty [at Georgia Tech] had me design and teach a class in hardware support for operating systems," he recalls. "I loved the teaching and the investigation aspects. I ended up staying on to get a Ph.D. in 1986, researching reliable distributed computing."

His postdoc work was in software engineering: investigating how to write software that does what the developer wants it to do.

The demand for cybersecurity professionals has never been higher, given people's expanding reliance on computation and storage.

In 1987, Spafford joined Purdue's computer science faculty. A year later, he was pulled into the investigation of the Morris worm, the first high-profile cybersecurity attack.

The code had been created by a college student who allegedly intended it to be a research experiment. Also known as the Internet worm, it made headlines when it caused a major denial-of-service incident that slowed down or crashed a significant number of the computers connected to the Internet.

Spafford was part of the team charged with isolating, analyzing, and cleaning up after the worm. There was a considerable sense of urgency, he recalls, since no one knew what the worm was doing, who had written it, and what its ultimate effects might be. He put in 18-hour days dissecting the code, documenting what it did, and responding to press inquiries.

"Until the worm event, security at government agencies was primarily about mainframes and information secrecy," he says. "Now, it also was clear that the availability, even integrity, of systems could be at risk—and that we didn't have good tools for protection and analysis. Suddenly, everyone from hobbyists to Pentagon staff was concerned about securing their computers."

Spafford's early involvement in combating cybersecurity threats led him to a rewarding career as a teacher, researcher, speaker, author, consultant, and organization builder.

He wrote a conference paper, *The Internet Worm Incident*, in 1989 to capture what had happened and the lessons learned. His other security projects included developing the open-source security tools COPS and Tripwire, as well as early firewalls and intrusion-detection systems. He was one of the founders of the field of cyber forensics, which involves collecting and analyzing digital data for investigations and providing legally admissible evidence. Spafford wrote the first papers on the topic.

In 1998, Spafford founded Purdue's Center for Education and Research in Information Assurance and Security, and he became its executive director emeritus in 2016.

Just as computing and cybersecurity have evolved, so has the teaching of computing and cybersecurity, Spafford notes. "When I was starting in the field, I could describe and teach courses on how a computing system worked, from hardware to networking, and all the points along the way where security had to be put in place," he says. "Fast forward to today, and looking at any major system in use, no person alive can do the same thing. The systems have gotten so big and there are so many variables that no one person can comprehend the whole stack anymore. To do well at security, you need to understand what a stack overflow is and the timing of instructions."

Many computer science programs no longer teach assembly language or machine organization, he notes.

Spafford's work has been recognized with many awards, but the honor he's most proud of is the Purdue University Morrill Award, which he received in 2012. The award recognizes faculty who have made extraordinary contributions to the university's mission of teaching, research, and community service.

40
NUMBER OF CYBERSECURITY CAREER SPECIALIZATIONS

"It was given not only for scholarship, but also for excellence as an educator, and for my service to the community," Spafford says. "It thus represented recognition by a community of my peers for accomplishments along multiple dimensions. I value all the other recognitions I have received, but this was the one that covered the broadest scope of my work."

How well are companies doing on the security front today? Spafford says some are doing a pretty good job by partitioning their systems, hiring the right people, and doing the right kind of monitoring. But, he says, others don't understand what it means to have good security or aren't willing to spend money on securing their systems.

"We are in a marketplace where fundamental good practices are often ignored in favor of new add-ons and new features," he says. "Instead of using sound engineering principles to build strong, resilient systems, the majority of the money spent and attention paid has gone to adding yet another layer of patches and building extensions on top of fundamentally broken technologies."

Given cybersecurity's broad and still-evolving range—there are now close to 40 cybersecurity specializations—Spafford advises those contemplating a career in it to get a sense of what aspects of security they find exciting and intriguing. Once you've done that, he says, what you need to learn depends on what you will be doing.

Those interested in cybersecurity forensics, for example, will need to understand operating systems, networks, architecture, compiler design, and software engineering. "This helps you understand how systems function, how things fit together, how flaws arise, and how they are exploited," he says.

For other areas of cybersecurity, you may need to study psychology and management theory to better understand the people involved, he says. Those who want to learn about policy should get some legal background, because law enforcement calls for yet a different set of skills.

The demand for cybersecurity professionals has never been higher, given people's expanding reliance on computation and storage, and their growing digital connectivity. "All these have changed the nature of what we do with computing and have increased the attack surfaces that can be used by those who would violate security," Spafford says. "Thirty years ago, the Internet connected research centers—our homes and automobiles weren't attack surfaces. Now it's the Internet of Almost Everything." ■

Employer: Purdue University

Title: Professor of computer sciences

Education: SUNY Brockport, Georgia Tech

Publications: Spafford has authored or coauthored over 150 books, papers, and other scholarly works. He coauthored *Cybersecurity Myths and Misconceptions: Avoiding the Hazards and Pitfalls That Derail Us*, Addison-Wesley Professional, 2023, with Leigh Metcalf and Josiah Dykstra.

Government activities: Spafford has testified before the U.S. Congress nine times and has contributed to 10 major amicus curiae briefs filed with U.S. courts, including the Supreme Court.

IEEE Fellow

BY MICHAEL KOZIOL

5 Questions for Chris Taylor

Aalyria's CEO on using lasers to close the digital divide

Aalyria, a recent spinout from Google, is trying to build on other companies' near successes. The company is revamping the software platform from the ambitious-yet-failed startup Loon to optimize and reconfigure networks in real time, using whatever connections are available. Aalyria is combining that tech with free-space lasers originally developed at Lawrence Livermore National Laboratory to connect the hard-to-connect. CEO Chris Taylor answered five rapid-fire questions on the company's approach and the importance of solid engineering know-how.

Connecting remote areas has been a problem for decades—why is it so difficult to solve?
Chris Taylor: The hardest problem is how to deliver the service that solves the challenges of the digital divide, at a price point that all the people who are suffering from the digital divide can pay. This has been the age-old challenge.

Broadly, what are the two technologies Aalyria has been developing?
Taylor: Spacetime is a temporospatial, software-defined network that mashes together elements of traditional SDNs, traditional software-defined wide area networks, and then adds a digital twin of all wireless transceivers on the planet. Tightbeam is a coherent-light, free-space optics product that allows us to transfer data at high speeds, right up to the absorption lines of various atmospheric anomalies (fog, rain, snow, moisture).

Chris Taylor is the founder and CEO of Aalyria. He also is a member of the boards of Alcavio, Govini, and Ox Intel. He served in the U.S. Marine Corps for 14 years.

How are you currently improving or refining these two key technologies?
Taylor: For Spacetime, it's always about resiliency and how to ensure those data payloads are getting where they need to be as quickly as they can be there, and in the same form that they left. For Tightbeam, it's always about distance and speed and atmosphere. That's how we think about these things. Can I increase the distance? Can I increase the speed and capacity? And can I deal with the atmosphere better than anyone else to ensure that we can deliver as we said we will?

How is what Aalyria is doing different from other kinds of wireless or cell networks?
Taylor: You can autonomously find the most efficient path based on the requirements that a user has entered into the system. So we say, we need this stuff to go from A to B, to F, to Q, to R, S, T, and then to Z. And that's the path it's going to take. What we've done is said, A through Z are wired paths, if you will. But we're now going to add all the wireless transceivers on the planet. And we'll add all of the optical links that we can create with Tightbeam. Basically, Spacetime is an operating system that can do everything that you want it to do for all of your network, and Tightbeam is a killer peripheral. They don't have to be used together, but they can always be used together when necessary.

How crucial is the expertise of the former Google engineers on staff for a project like Aalyria?
Taylor: When I was growing up as a kid, we didn't yet have computers until I think I was a senior in high school, something like that. I just happened to be born at a time when the wild technological shift happened in America, and certainly in the world. Access to those technologies changed how I learned and how I saw many things in the world. It's been wildly impactful to me. If I can work with a bunch of Googlers and Metas and Amazonians, who grew up in this age, who are experts in this technology, and can help us deliver the same experience that I've had growing up—without and then with—I think that shifting from without to with is wildly impactful for any human being. It would be great if the unconnected could experience that too. ■

Photo-illustration by Stuart Bradford

By Dario Farina, Etienne Burdet, Carsten Mehring & Jaime Ibáñez

Illustrations by Chris Philpot

Roboticists Want to Give You A Third Arm

UNUSED BANDWIDTH IN NEURONS CAN BE TAPPED TO CONTROL EXTRA LIMBS

Extra limbs, controlled by wearable electrode patches that read and interpret neural signals from the user, could have innumerable uses, such as assisting on spacewalk missions to repair satellites.

What could you do with an extra limb?

Consider a surgeon performing a delicate operation, one that needs her expertise and steady hands—all three of them. As her two biological hands manipulate surgical instruments, a third robotic limb that's attached to her torso plays a supporting role. Or picture a construction worker who is thankful for his extra robotic hand as it braces the heavy beam he's fastening into place with his other two hands. Imagine wearing an exoskeleton that would let you handle multiple objects simultaneously, like Spiderman's Dr. Octopus. Or contemplate the out-there music a composer could write for a pianist who has 12 fingers to spread across the keyboard.

Such scenarios may seem like science fiction, but recent progress in robotics and neuroscience makes extra robotic limbs conceivable with today's technology. Our research groups at Imperial College London and the University of Freiburg, in Germany, together with partners in the European project NIMA, are now working to figure out whether such augmentation can be realized in practice to extend human abilities. The main questions we're tackling involve both neuroscience and neurotechnology: Is the human brain capable of controlling additional body parts as effectively as it controls biological parts? And if so, what neural signals can be used for this control?

We think that extra robotic limbs could be a new form of human augmentation, improving people's abilities on tasks they can already perform as well as expanding their ability to do things they simply cannot do with their natural human bodies. If

NEURAL CONTROL DEMONSTRATED

A volunteer exploits unused neural bandwidth to direct the motion of a cursor on the screen in front of her. Neural signals travel from her brain and spinal cord, through spinal neurons, to the tibialis anterior muscle, where they are read by an electromyography electrode array on her leg and deciphered in real time. The signals arriving at the muscle include low-frequency components [blue] that control muscle contractions, higher frequencies such as the beta band [yellow] with no known biological purpose, and noise [gray].

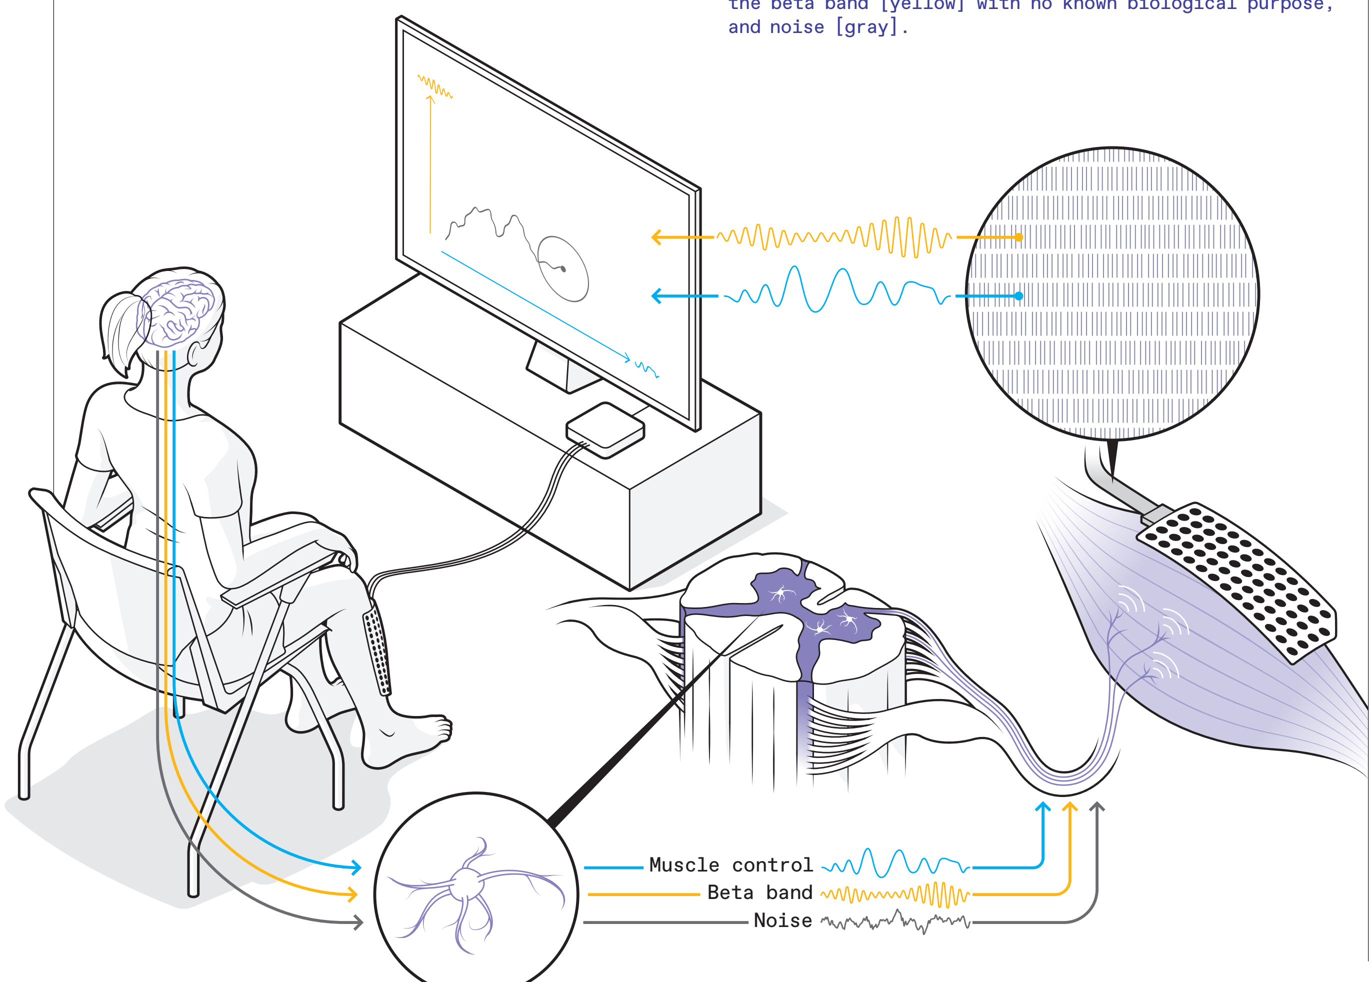

SOURCE: M. BRÄCKLEIN ET AL., JOURNAL OF NEURAL ENGINEERING

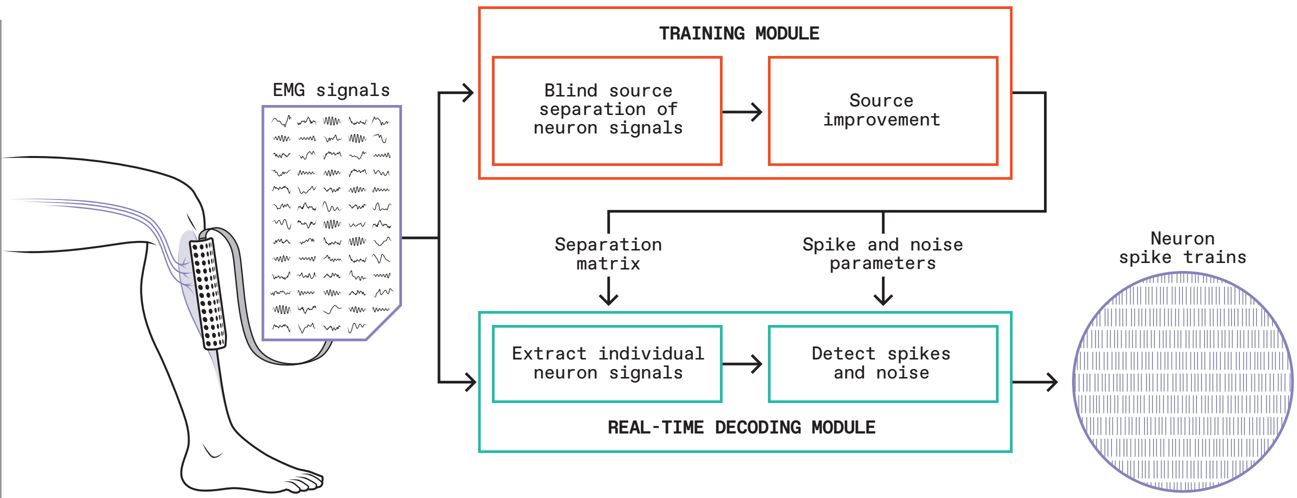

HOW THE NEURAL SIGNALS ARE DECODED

A training module [orange] takes an initial batch of electromyography (EMG) signals read by the electrode array [left], determines how to extract signals of individual neurons, and summarizes the process mathematically as a separation matrix and other parameters. With these tools, the real-time decoding module [green] can efficiently extract individual neurons' sequences of spikes, or "spike trains" [right], from an ongoing stream of EMG signals.

humans could easily add and control a third arm, or a third leg, or a few more fingers, they would likely use them in tasks and performances that went beyond the scenarios mentioned here, discovering new behaviors that we can't yet even imagine.

Robotic limbs have come a long way in recent decades, and some are already used by people to enhance their abilities. Most are operated via a joystick or other hand controls. That's how workers on manufacturing lines wield mechanical limbs that hold and manipulate components of a product. Similarly, surgeons who perform robotic surgery sit at a console across the room from the patient. While the surgical robot may have four arms tipped with different tools, the surgeon's hands can control only two of them at a time. Could we give these surgeons the ability to control four tools simultaneously?

Robotic limbs are also used by people who have amputations or paralysis. That includes people in powered wheelchairs controlling a robotic arm with the chair's joystick and those who are missing limbs controlling a prosthetic by the actions of their remaining muscles. But a truly mind-controlled prosthesis is a rarity.

The pioneers in brain-controlled prosthetics are people with tetraplegia, who are often paralyzed from the neck down. Some of these people have boldly volunteered for clinical trials of brain implants that enable them to control a robotic limb by thought alone, issuing mental commands that cause a robot arm to lift a drink to their lips or help with other simple tasks of daily life. These systems fall under the category of brain-machine interfaces (BMI). Other volunteers have used BMI technologies to control computer cursors, to type out messages, browse the Internet, and more. But most of these systems require brain surgery to insert the neural implant and include hardware that protrudes from the skull, making them suitable only for use in the lab.

Augmentation of the human body can be thought of as having three levels. The first level increases an existing characteristic, in the way that, say, a powered exoskeleton can give the wearer super strength. The second level gives a person a new degree of freedom, such as the ability to move a third arm or a sixth finger, but at a cost—if the extra appendage is controlled by a foot pedal, for example, the user sacrifices normal mobility of the foot to operate the control system. The third level of augmentation, and the least mature technologically, gives a user an extra degree of freedom without taking mobility away from any other body part. Such a system would allow people to use their bodies normally by harnessing some unused neural signals to control the robotic limb. That's the level that we're exploring in our research.

Third-level human augmentation can perhaps be achieved with invasive BMI implants, but for everyday use, we need a noninvasive way to pick up brain commands from outside the skull. For many research groups, that means relying on tried-and-true electroencephalography (EEG) technology, which uses scalp electrodes to pick up brain signals. Our groups are working on that approach, but we are also exploring another method: using electromyography (EMG) signals produced by muscles. We've spent more than a decade investigating how EMG electrodes on the skin's surface can detect electrical signals from the

BIG CELLS

We usually think of biological cells as being microscopic, but individual neurons can be a meter long, with an axon reaching from your spine to your toes. Blue whales can have axons 24 meters long.

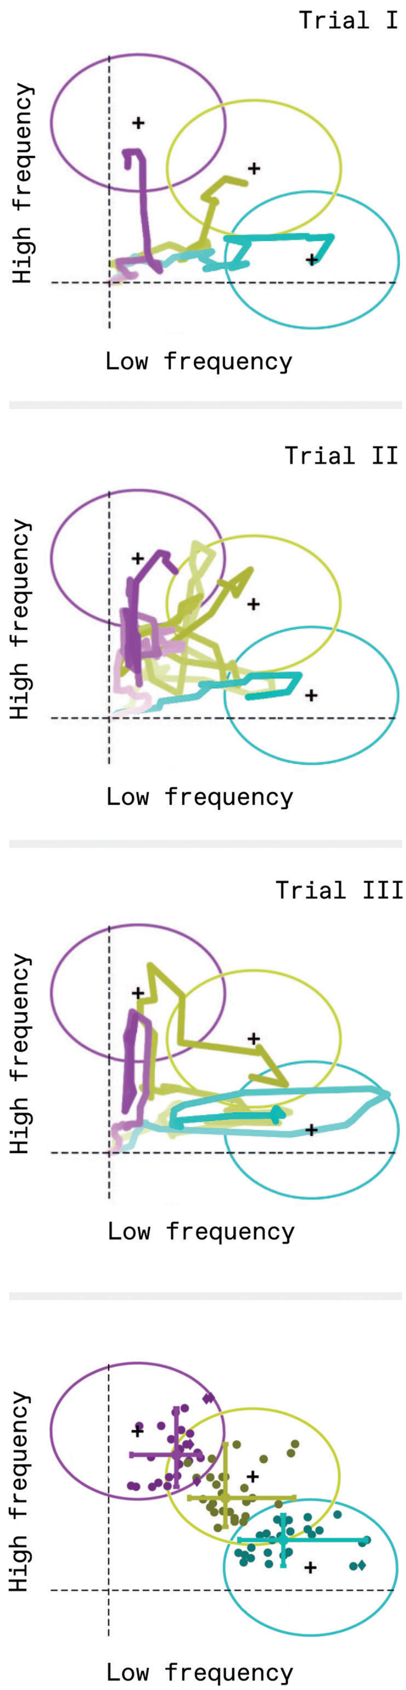

These are some results from the experiment in which low- and high-frequency neural signals, respectively, controlled horizontal and vertical motion of a computer cursor. Colored ellipses (with plus signs at centers) show the target areas. The top three diagrams show the trajectories (each one starting at the lower left) achieved for each target across three trials by one user. At bottom, dots indicate average positions achieved in successful trials, and colored crosses mark the mean positions and the range of results for each target.

SOURCE: M. BRÄCKLEIN ET AL., JOURNAL OF NEURAL ENGINEERING

muscles that we can then decode to reveal the commands sent by spinal neurons.

Electrical signals are the language of the nervous system. Throughout the brain and the peripheral nerves, a neuron "fires" when a certain voltage—some tens of millivolts—builds up within the cell and causes an action potential to travel down its axon, releasing neurotransmitters at junctions, or synapses, with other neurons, and potentially triggering those neurons to fire in turn. When such electrical pulses are generated by a motor neuron in the spinal cord, they travel along an axon that reaches all the way to the target muscle, where they cross special synapses to individual muscle fibers and cause them to contract. We can record these electrical signals, which encode the user's intentions, and use them for a variety of control purposes.

Deciphering the individual neural signals based on what can be read by surface EMG, however, is not a simple task. A typical muscle receives signals from hundreds or thousands of spinal neurons. Moreover, each axon branches at the muscle and may connect with a hundred or more individual muscle fibers distributed throughout the muscle. A surface EMG electrode picks up a sampling of this cacophony of pulses.

A breakthrough in noninvasive neural interfaces came with the discovery two decades ago that the signals picked up by high-density EMG, in which tens to hundreds of electrodes are fastened to the skin, can be disentangled, providing information about the commands sent by individual motor neurons in the spine. Such information had previously been obtained only with invasive electrodes in muscles or nerves. Our high-density surface electrodes provide good sampling over multiple locations, enabling us to identify and decode the activity of a relatively large proportion of the spinal motor neurons involved in a task. And we can now do it in real time, which suggests that we can develop noninvasive BMI systems based on signals from the spinal cord.

The current version of our system consists of two parts: a training module and a real-time decoding module. To begin, with the EMG electrode grid attached to their skin, the user performs gentle muscle contractions, and we feed the recorded EMG signals into the training module. This module performs the difficult task of identifying the individual motor neuron pulses (also called spikes) that make up the EMG signals. The module analyzes how the EMG signals and the inferred neural spikes are related, which it summarizes in a set of parameters that can then be used with a much simpler mathematical prescription to translate the EMG signals into sequences of spikes from individual neurons.

With these parameters in hand, the decoding module can take new EMG signals and extract the individual motor neuron activity in real time. The training module requires a lot of computation and would be too slow to perform real-time control itself, but it usually has to be run only once each time the

EMG electrode grid is fixed in place on a user. By contrast, the decoding algorithm is very efficient, with latencies as low as a few milliseconds, which bodes well for possible self-contained wearable BMI systems. We validated the accuracy of our system by comparing its results with signals obtained concurrently by invasive EMG electrodes inserted into the user's muscle.

Developing this real-time method to extract signals from spinal motor neurons was the key to our present work on controlling extra robotic limbs. While studying these neural signals, we noticed that they have, essentially, extra bandwidth. The low-frequency part of the signal (below about 7 hertz) is converted into muscular force, but the signal also has components at higher frequencies, such as those in the beta band at 13 to 30 Hz, which are too high to control a muscle and seem to go unused. We don't know why the spinal neurons send these higher-frequency signals; perhaps the redundancy is a buffer in case of new conditions that require adaptation. Whatever the reason, humans evolved a nervous system in which the signal that comes out of the spinal cord has much richer information than is needed to command a muscle.

That discovery set us thinking about what could be done with the spare frequencies. In particular, we wondered if we could take that extraneous neural information and use it to control a robotic limb. But we didn't know if people would be able to voluntarily control this part of the signal separately from the part they used to control their muscles. So we designed an experiment to find out.

In our first proof-of-concept experiment, volunteers tried to use their spare neural capacity to control computer cursors. The setup was simple, though the neural mechanism and the algorithms involved were sophisticated. Each volunteer sat in front of a screen, and we placed an EMG system on their leg, with 64 electrodes in a 4-by-10-centimeter patch stuck to their shin over the tibialis anterior muscle, which flexes the foot upward when it contracts. The tibialis has been a workhorse for our experiments: It occupies a large area close to the skin, and its muscle fibers are oriented along the leg, which together make it ideal for decoding the activity of spinal motor neurons that innervate it.

We asked our volunteers to contract the tibialis, essentially holding it tense and with the foot braced to prevent movement. Throughout the experiment, we looked at the variations within the extracted neural signals. We separated these signals into the low frequencies that controlled the muscle contraction and spare frequencies at about 20 Hz in the beta band, and we linked these two components respectively to the horizontal and vertical control of a cursor on a computer screen. We asked the volunteers to try to move the cursor around the screen, reaching all parts of the space, but we didn't, and indeed couldn't, explain to them how to do that. They had to rely on the visual feedback of the cursor's position and let their brains figure out how to make it move.

If humans could easily add and control a third arm, they would likely use it in new behaviors that we can't yet even imagine.

Remarkably, without knowing exactly what they were doing, these volunteers were able to perform the task within minutes, zipping the cursor around the screen, albeit shakily. Beginning with one neural command signal—contract the tibialis anterior muscle—they were learning to develop a second signal to control the computer cursor's vertical motion, independently from the muscle control (which directed the cursor's horizontal motion). We were surprised and excited by how easily they achieved this big first step toward finding a neural control channel separate from natural motor tasks. But we also saw that the control was too limited for practical use. Our next step will be to see if more accurate signals can be obtained and if people can use them to control a robotic limb while also performing independent natural movements.

We are also interested in understanding more about how the brain performs feats like the cursor control. In a variation of the cursor task, we concurrently used EEG to see what was happening in the user's brain, particularly in the area associated with the voluntary control of movements. We were excited to discover that the changes happening to the beta-band neural signals arriving at the muscles were tightly related to similar changes at the brain level. As mentioned, the beta neural signals remain something of a mystery since they play no known role in controlling muscles, and it isn't even clear where they originate. Our result suggests that our volunteers were learning to modulate brain activity that was sent down to the muscles as beta signals. This important finding is helping us unravel the mechanisms behind these beta signals.

Meanwhile, we have set up a system at Imperial College London for testing these CONTINUED ON P. 46

These are the hands of someone born with extra fingers. He has brain regions dedicated to their control, giving him exceptional dexterity. This shows that the human nervous system can control more degrees of freedom than usual.

ETIENNE BURDET

HOW DUOLINGO'S AI LEARNS WHAT YOU NEED TO LEARN

The language-learning app tries to emulate a great human tutor

BY KLINTON BICKNELL,
CLAIRE BRUST
& BURR SETTLES

Illustration by Eddie Guy

IT'S LUNCHTIME WHEN

your phone pings you with a green owl who cheerily reminds you to "Keep Duo Happy!" It's a nudge from Duolingo, the popular language-learning app, whose algorithms know you're most likely to do your 5 minutes of Spanish practice at this time of day. The app chooses its notification words based on what has worked for you in the past and the specifics of your recent achievements, adding a dash of attention-catching novelty. When you open the app, the lesson that's queued up is calibrated for your skill level, and it includes a review of some words and concepts you flubbed during your last session.

Duolingo, with its gamelike approach and cast of bright cartoon characters, presents a simple user interface to guide learners through a curriculum that leads to language proficiency, or even fluency. But behind the scenes, sophisticated artificial-intelligence (AI) systems are at work. One system in particular, called Birdbrain, is continuously improving the learner's experience with algorithms based on decades of research in educational psychology, combined with recent advances in machine learning. But from the learner's perspective, it simply feels as though the green owl is getting better and better at personalizing lessons.

The three of us have been intimately involved in creating and improving Birdbrain, of which Duolingo recently launched its second version. We see our work at Duolingo as furthering the company's overall mission to "develop the best education in the world and make it universally available." The AI systems we continue to refine are necessary to scale the learning experience beyond the more than 50 million active learners who currently complete about 1 billion exercises per day on the platform.

Although Duolingo is known as a language-learning app, the company's ambitions go further. We recently launched apps covering childhood literacy and third-grade mathematics, and these expansions are just the beginning. We hope that anyone who wants help with academic learning will one day be able to turn to the friendly green owl in their pocket who hoots at them, "Ready for your daily lesson?"

B

BACK IN 1984, educational psychologist Benjamin Bloom identified what has come to be called Bloom's 2-sigma problem. Bloom found that average students who were individually tutored performed two standard deviations better than they would have in a classroom. That's enough to raise a person's test scores from the 50th percentile to the 98th.

When Duolingo was launched in 2012 by Luis von Ahn and Severin Hacker out of a Carnegie Mellon University research project, the goal was to make an easy-to-use online language tutor that could approximate that supercharging effect. The founders weren't trying to replace great teachers. But as immigrants themselves (from Guatemala and Switzerland, respectively), they recognized that not everyone has access to great teachers. Over the ensuing years, the growing Duolingo team continued to think about how to automate three key attributes of good tutors: They know the material well, they keep students engaged, and they track what each student currently knows, so they can present material that's neither too easy nor too hard.

Duolingo uses machine learning and other cutting-edge technologies to mimic these three qualities of a good tutor. First, to ensure expertise, we employ natural-language-processing tools to assist our content developers in auditing and improving our 100-odd courses in more than 40 different languages. These tools analyze the vocabulary and grammar content of lessons and help create a range of possible translations (so the app will accept learners' responses when there are multiple correct ways to say something). Second, to keep learners engaged, we've gamified the experience with points and levels, used text-to-speech tech to create custom voices for each of the characters that populate the Duolingo world, and fine-tuned our notification systems. As for getting inside learners' heads and giving them just the right lesson—that's where Birdbrain comes in.

Birdbrain is crucial because learner engagement and lesson difficulty are related. When students are given material that's too difficult, they often get frustrated and quit. Material that feels easy might keep them engaged, but it doesn't challenge them as much. Duolingo uses AI to keep its learners squarely in the zone where they remain engaged but are still learning at the edge of their abilities.

One of us (Settles) joined the company just six months after it was founded, helped establish various research functions, and then led Duolingo's AI and machine-learning efforts until last year. Early on, there weren't many organizations doing large-scale online interactive learning. The closest analogue to what Duolingo was trying to do were programs that took a "mastery learning" approach, notably for math tutoring. Those programs offered up problems around a similar concept (often called a "knowledge component") until the learner demonstrated sufficient mastery before moving on to the next unit, section, or concept.

But that approach wasn't necessarily the best fit for language, where a single exercise can involve many different concepts that interact in complex ways (such as vocabulary, tenses, and grammatical gender), and where there are different ways in which a learner can respond (such as translating a sentence, transcribing an audio snippet, and filling in missing words).

The early machine-learning work at Duolingo tackled fairly simple problems, like how often to return to a particular vocabulary word or concept (which drew on educational research on spaced repetition). We also analyzed learners' errors to identify pain points in the curriculum and then reorganized the order in which we presented the material.

Duolingo then doubled down on building personalized systems. Around 2017, the company started to make a more focused investment in machine learning, and that's when coauthors Brust and Bicknell joined the team. In 2020, we launched the first version of Birdbrain.

BEFORE BIRDBRAIN, Duolingo had made some non-AI attempts to keep learners engaged at the right level, including estimating the difficulty of exercises based on heuristics such as the number of words or characters in a sentence. But the company often found that it was dealing with trade-offs between how much people were actually learning and how engaged they were. The goal with Birdbrain was to strike the right balance.

The question we started with was this: For any learner and any given exercise, can we predict how likely the learner is to get that exercise correct? Making that prediction requires Birdbrain to estimate both the difficulty of the exercise and the current proficiency of the learner. Every time a learner completes an exercise, the system updates both estimates. And Duolingo uses the resulting predictions in its session-generator algorithm to dynamically select new exercises for the next lesson.

When we were building the first version of Birdbrain, we knew it needed to be simple and scalable, because we'd be applying it to hundreds of millions of exercises. It needed to be fast and require little computation. We decided to use a flavor of logistic regression inspired by item-response theory from the psychometrics literature. This approach models the probability of a person giving a correct response as a function of two variables, which can be interpreted as the difficulty of the exercise and the ability of the learner. We estimate the difficulty of each exercise by summing up the difficulty of its component features like the type of exercise, its vocabulary words, and so on.

The second ingredient in the original version of Birdbrain was the ability to perform computationally simple updates on these difficulty and ability parameters. We implement this by performing one step of stochastic gradient descent on the relevant parameters every time a learner completes an exercise. This turns out to be a generalization of the Elo rating system, which is used to rank players in chess and other games. In chess, when a player wins a game, their ability estimate goes up and their opponent's goes down. In Duolingo, when a learner gets an exercise wrong, this system lowers the estimate of their ability and raises the estimate of the exercise's difficulty. Just like in chess, the size of these changes depends on the pairing: If a novice chess player wins against an expert player, the expert's Elo score will be substantially lowered, and their opponent's score will be substantially raised. Similarly, here, if a beginner learner gets a hard exercise correct, the ability and difficulty parameters can shift dramatically, but if the model already expects the learner to be correct, neither parameter changes much.

To test Birdbrain's performance, we first ran it in "shadow mode," meaning that it made predictions that were merely logged for analysis and not yet used by the Session Generator to personalize lessons. Over time, as learners completed exercises and got answers right or wrong, we saw whether Birdbrain's predictions of their success matched reality—and if they didn't, we made improvements.

Once we were satisfied with Birdbrain's performance, we started running controlled tests: We enabled Birdbrain-based personalization for a fraction of learners (the experimental group) and compared their

Illustration by Eddie Guy

DEALING WITH AROUND A BILLION EXERCISES EVERY DAY REQUIRED A LOT OF INVENTIVE ENGINEERING.

learning outcomes with those who still used the older heuristic system (the control group). We wanted to see how Birdbrain would affect learner engagement—measured by time spent on tasks in the app—as well as learning, measured by how quickly learners advanced to more difficult material. We wondered whether we'd see trade-offs, as we had so often before when we tried to make improvements using more conventional product-development or software-engineering techniques. To our delight, Birdbrain consistently caused both engagement and learning measures to increase.

FROM THE BEGINNING, we were challenged by the sheer scale of the data we needed to process. Dealing with around a billion exercises every day required a lot of inventive engineering.

One early problem with the first version of Birdbrain was fitting the model into memory. During nightly training, we needed access to several variables per learner, including their current ability estimate. Because new learners were signing up every day, and because we didn't want to throw out estimates for inactive learners in case they came back, the amount of memory grew every night. After a few months, this situation became unsustainable: We couldn't fit all the variables into memory. We needed to update parameters every night without fitting everything into memory at once.

Our solution was to change the way we stored both each day's lesson data and the model. Originally, we stored all the parameters for a given course's model in a single file, loaded that file into memory, and sequentially processed the day's data to update the course parameters. Our new strategy was to break up the model: One piece represented all exercise-difficulty parameters (which didn't grow very large), while several chunks represented the learner-ability estimates. We also chunked the day's learning data into separate files according to which learners were involved and—critically—used the same chunking function across learners for both the course model and learner data. This allowed us to load only the course parameters relevant to a given chunk of learners while we processed the corresponding data about those learners.

One weakness of this first version of Birdbrain was that the app waited until a learner finished a lesson before it reported to our servers which exercises the user got right and what mistakes they made. The problem with that approach is that roughly 20 percent of lessons started on Duolingo aren't completed, perhaps because the person put down their phone or switched to another app. Each time that happened, Birdbrain lost the relevant data, which was potentially very interesting information! We were pretty sure that people weren't quitting at random—in many cases, they likely quit once they hit material that was especially challenging or daunting for them. So when we upgraded to Birdbrain version 2, we also began streaming data throughout the lesson in chunks. This gave us critical information about which concepts or exercise types were problematic.

Another issue with the first Birdbrain was that it updated its models only once every 24 hours (during a low point in global app usage, which was nighttime at Duolingo's headquarters, in Pitts-

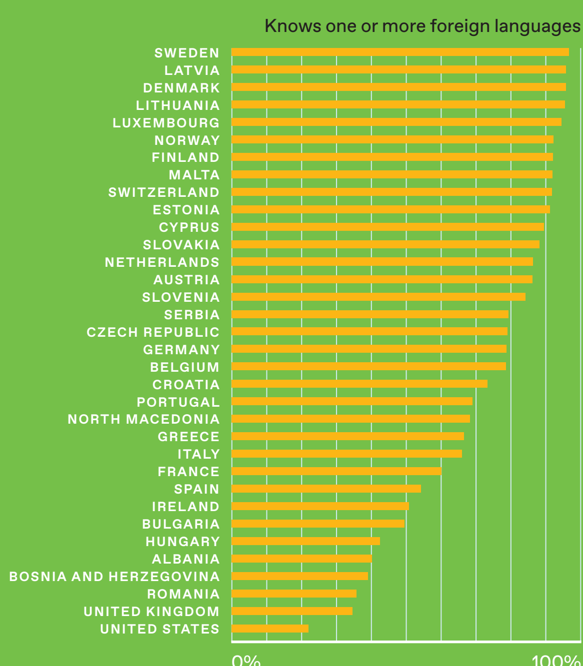

FOREIGN-LANGUAGE PROFICIENCY

People in some countries are better than others when it comes to knowing foreign languages. The United States, in particular, ranks far behind European nations in this regard.

SOURCE: EUROSTAT; U.S. CENSUS BUREAU

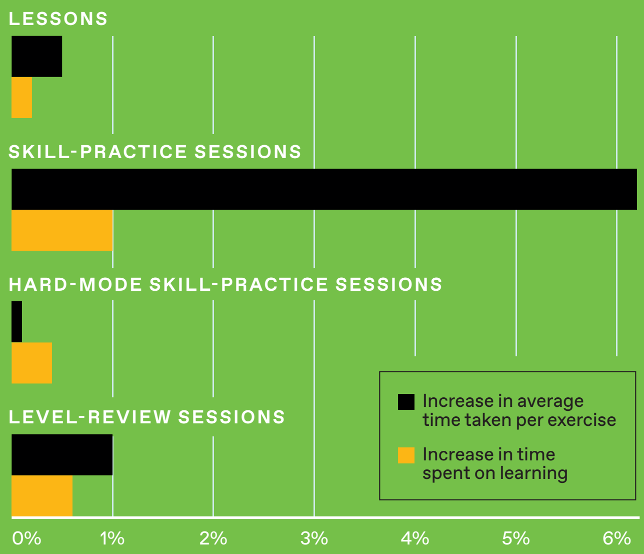

BETTER AI, BETTER LEARNING

Learner performance improved with the transition from version 1 of Birdbrain to version 2. In particular, there were increases in the time spent on learning and in the average time taken per exercise, a proxy for the difficulty of the lesson.

burgh). With Birdbrain V2, we wanted to process all the exercises in real time. The change was desirable because learning operates at both short- and long-term scales; if you study a certain concept now, you'll likely remember it 5 minutes from now, and with any luck, you'll also retain some of it next week. To personalize the experience, we needed to update our model for each learner very quickly. Thus, within minutes of a learner completing an exercise, Birdbrain V2 will update its "mental model" of their knowledge state.

In addition to occurring in near real time, these updates also worked differently because Birdbrain V2 has a different architecture and represents a learner's knowledge state differently. Previously, that property was simply represented as a scalar number, as we needed to keep the first version of Birdbrain as simple as possible. With Birdbrain V2, we had company buy-in to use more computing resources, which meant we could build a much richer model of what each learner knows. In particular, Birdbrain V2 is backed by a recurrent neural-network model (specifically, a long short-term memory, or LSTM, model), which learns to compress a learner's history of interactions with Duolingo exercises into a set of 40 numbers—or in the lingo of mathematicians, a 40-dimensional vector. Every time a learner completes another exercise, Birdbrain will update this vector based on its prior state, the exercise that the learner has completed, and whether they got it right. It is this vector, rather than a single value, that now represents a learner's ability, which the model uses to make predictions about how they will perform on future exercises.

The richness of this representation allows the system to capture, for example, that a given learner is great with past-tense exercises but is struggling with the future tense. V2 can begin to discern each person's learning trajectory, which may vary considerably from the typical trajectory, allowing for much more personalization in the lessons that Duolingo prepares for that individual.

Once we felt assured that Birdbrain V2 was accurate and stable, we conducted controlled tests comparing its personalized learning experience with that of the original Birdbrain. We wanted to be sure we had not only a better machine-learning model but also that our software provided a better user experience. Happily, these tests showed that Birdbrain V2 consistently caused both engagement and learning measures to increase even further. In May 2022, we turned off the first version of Birdbrain and switched over entirely to the new and improved system.

MUCH OF WHAT we're doing with Birdbrain and related technologies applies outside of language learning. In principle, the core of the model is very general and can also be applied to our company's new math and literacy apps—or to whatever Duolingo comes up with next.

Birdbrain has given us a great start in optimizing learning and making the curriculum more adaptive and efficient. How far we can go with personalization is an open question. We'd like to create adaptive systems that respond to learners based not only on what they know but also on the teaching approaches that work best for them. What types of exercises does a learner really pay attention to? What exercises seem to make concepts click for them?

Those are the kinds of questions that great teachers might wrestle with as they consider various struggling students in their classes. We don't believe that you can replace a great teacher with an app, but we do hope to get better at emulating some of their qualities—and reaching more potential learners around the world through technology. ■

MAINSPRING'S LINEAR GENERATOR CAN RUN ON ALMOST ANY FUEL

The Omnivorous

Linear generators are currently being manufactured at one of Mainspring Energy's facilities in Menlo Park, Calif. Clockwise from top left, technicians work on the frame of a linear generator core; a technician routes compressed air from an air spring to a generator's bearings; a gantry moves a completed core across the manufacturing floor for final assembly; another technician prepares copper wires to be wound into coils; and a stator is positioned in preparation for assembly. TOP ROW: CREATIVE SHOT; BOTTOM ROW: MAINSPRING

Generator

By MATT SVRCEK

It's January 2030 and your electric heat pump is warming the house while your electric car charges in the garage, all powered by solar panels on your roof and by wind and solar generators at your local utility. It doesn't matter that it's been raining for two weeks because your utility is tapping into ammonia produced with last summer's sunshine. It's consuming that ammonia in a linear generator.

The linear generator can quickly switch between different types of green (and not-so-green, if need be) fuel, including biogas, ammonia, and hydrogen. It has the potential to make the decarbonized power system available, reliable, and resilient against the vagaries of weather and of fuel supplies. And it's not a fantasy; it's been developed, tested, and deployed commercially.

The cofounders of Mainspring Energy, of which I am one, spent 14 years developing this technology, and in 2020 we began rolling it out commercially. It is currently installed at tens of sites, producing 230 to 460 kilowatts at each. We expect linear generators at many more locations to come on line within the next year.

THE STORY OF the linear generator began nearly two decades ago at Stanford University's Advanced Energy Systems Laboratory, when mechanical engineering professor Christopher Edwards asked some of us Ph.D. students a simple question: "What is the most efficient and practical way possible to convert chemical-bond energy into useful work?"

We started by considering fuel cells, since they can be very efficient. But fuel cells use catalysts to trigger the chemical reactions that release energy, and catalysts typically cost a lot, degrade over time, and respond poorly to rapid changes in load. So we began looking for an alternative.

We knew that we could trigger the release of energy simply by compressing a mixture of air and fuel. Here's how that would work.

First, fuel and air enter a closed chamber with movable end walls. Next, those end walls move toward each other, compressing the mixture of fuel and air. As this happens, the molecules within the mixture collide faster and faster, until they at last break apart and re-form into different molecules, releasing the energy stored in their chemical bonds. That energy causes the new molecules to collide even faster and more often, not just with themselves but also with the walls of the chamber, raising the pressure in the chamber. It all happens without a spark or any other ignition source.

The pressure pushes the walls outward with more force than that needed to push them inward at the beginning of the cycle. Once these walls reach their initial position, and the pressure within the chamber reverts to its initial state, a new batch of fuel and air flows in, pushing the molecules created by the previous cycle out of the chamber and starting the process all over. That's the theory.

To test it out, in 2008 we constructed an apparatus capable of compressing through a volume 100 times that of the starting value, then expanding back again. We used a metal tube two meters long and 50 millimeters in diameter, with a closed wall on one end and a metal slug as the moving wall. This arrangement works like a piston that compresses a gas inside a cylinder in an engine, although that's where the similarities end—the "piston" in our device was not attached to a crankshaft, or to anything at all. I'll discuss in a moment the limitations of this type of engine architecture for this kind of reaction, and how we solved them with a new type of machine. But it was a good place to start.

Our first device was very simple—it could run only one "shot" at a time, and it did not produce electricity; that is, we did not harvest the energy produced. But we could use it to measure the efficiency of the reaction, meaning the extra push that must be applied to the moving wall during expansion relative to how much fuel was used. And the results were excellent: The device was as efficient as a fuel cell, just as we had hoped.

Now we had to build a version that could generate electricity and run for years at a reasonable cost. In 2010, Shannon Miller, Adam Simpson, and I incorporated Mainspring Energy to build a real-world system. Khosla Ventures provided our initial seed money; to date

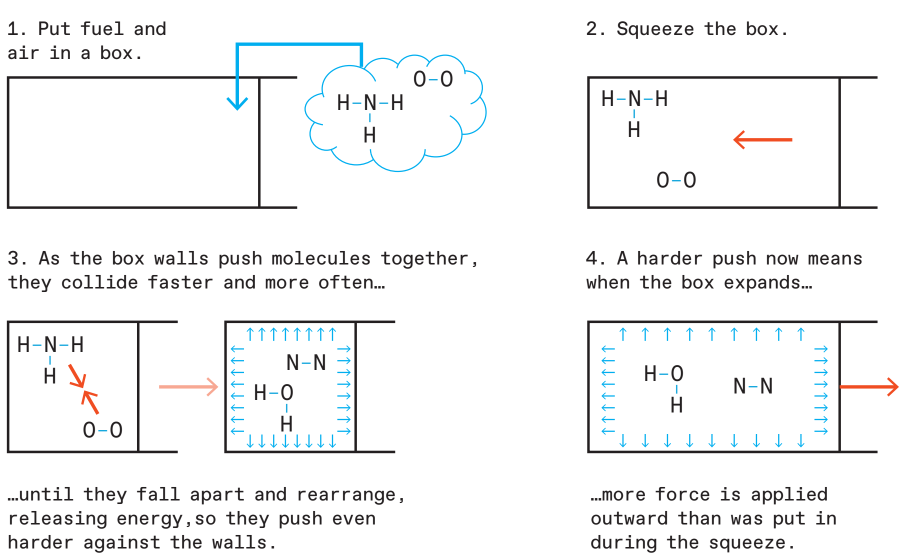

The efficient, clean, flameless reaction at the heart of the Mainspring generator works with nearly any fuel, including carbon-free ammonia as shown here. The ammonia reacts with oxygen in air to produce nitrogen gas and water, and the resulting force pushes against the walls of the box.

In Mainspring's linear generator, two translators move within a center reaction zone located between two outer air springs. A set of stationary copper coils surrounds each translator, forming a linear electromagnetic machine (LEM). A cycle begins with the introduction of air and fuel into the center reaction zone. Energy stored in the air springs from a previous cycle compresses the mixture until a flameless reaction occurs. The reaction drives the translators, to which magnets are attached, back through the copper coils, producing electricity. The force of this motion also compresses the air springs, readying the system for the next cycle.

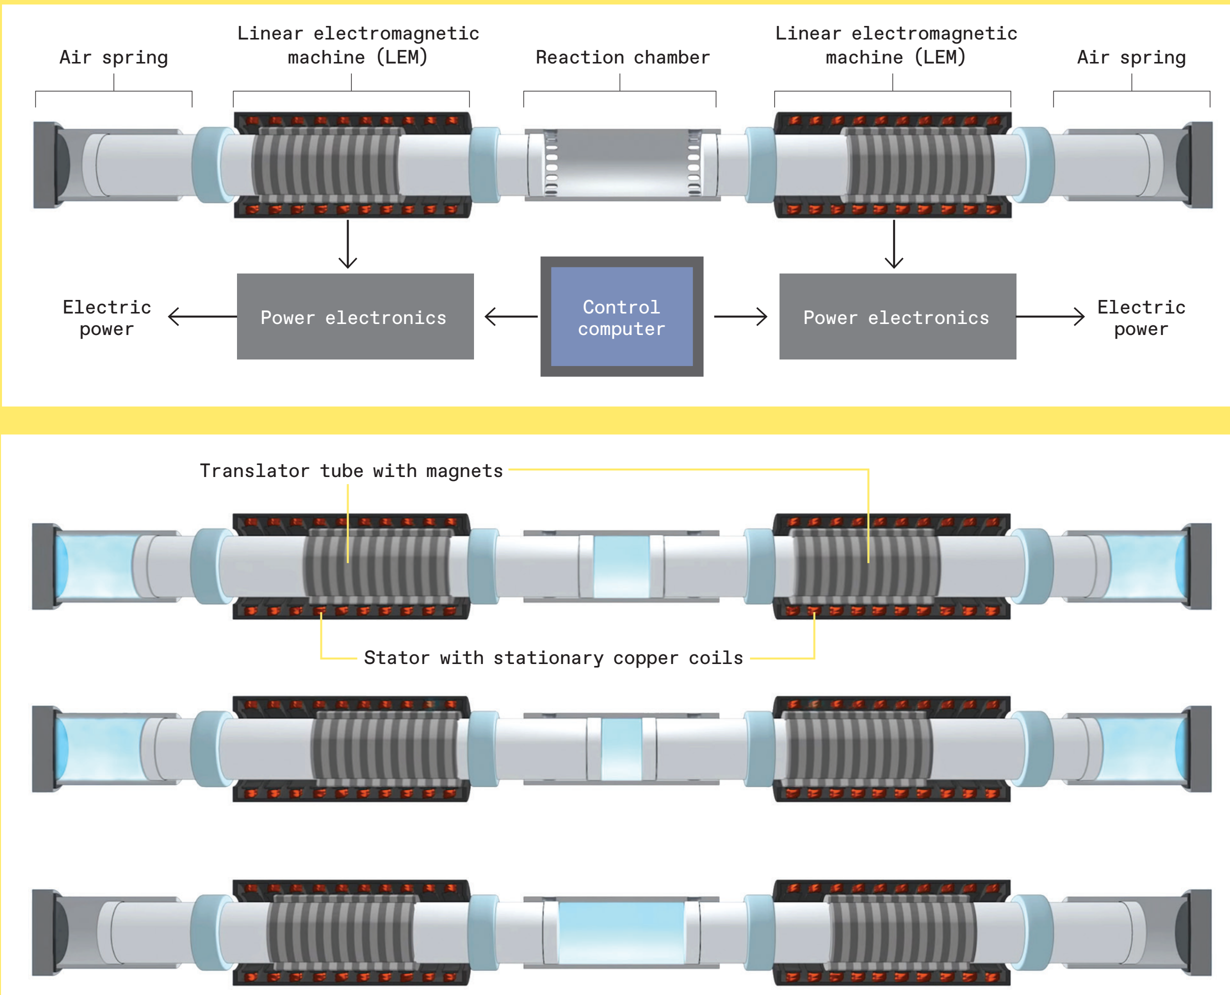

PREVIOUS PAGE AND ABOVE: MAINSPRING

we've raised more than US $500 million from a range of investors, including Khosla, American Electric Power, Bill Gates, and NextEra Energy.

Generators that use the flameless compression reaction had been built before in research labs, based on a conventional combustion-engine architecture, but they were limited by the difficulty of controlling the reaction in this type of apparatus. To be efficient, the mixture needs to be compressed just enough to initiate the reaction. If compression continues after the reaction happens, it fights against the pressure generated by the reaction, wasting energy. If compression stops too soon, the reaction never happens.

This optimal compression varies with conditions, beginning with the choice of fuel: Hydrogen, for example, reacts with less compression than ammonia. Running at a partial power output instead of full power or running on a hot day versus a cold one also changes the optimal compression.

A conventional engine harvests energy when the extra pressure generated from the reaction pushes on a piston, which pushes on a connecting rod to rotate a crankshaft. The crankshaft geometry constrains the piston to always follow the same motion, and therefore the same amount of compression, no matter what. Such an engine can't adapt to changes in the required compression, and that makes it hard to control the reaction.

So rather than mimicking an engine, we designed a new machine that ties the compression and expansion motion directly to the generation of electricity, and in doing so provides the necessary reaction control. This machine ended up looking completely different from—and having almost no parts in common with—a conventional engine. So we felt a new name was needed, and we called it the linear generator.

PICTURE A SERIES of five cylindrical assemblies arranged in a line, held within a boxlike frame. The central tube is the reaction chamber; it's where the fuel and air go. On either side of it sits a linear electromagnetic machine (LEM) that converts the push from pressure directly into electric power. At each end of the generator is an air-filled cylindrical chamber that acts as a spring to bounce the moving part of the LEM back to the center. The whole arrangement—two air springs, two LEMs, and a reaction chamber—forms a linear generator core. It's long and skinny: A machine rated at 115 kW is about 5.5 meters long and about 1 meter high and wide.

The LEM, in principle, is an electric motor that has been unrolled to form a line instead of a circle. It consists of a moving part—the translator—and a stationary part—the stator. The translator is a long, straight tube with an array of neodymium permanent magnets attached to the perimeter, near the center. An end plate caps each translator tube and seals to the inner surface of the reaction chamber. The capped end of the translator does the actual compression, as the piston would in an engine, but it is wildly different in design. The stator is a series of copper coils. As the translator moves back and forth in a straight line inside the coils, the magnets generate current that feeds an 800-volt DC bus.

It works rather like regenerative braking. An electric car's motor acts in reverse, as a generator, to convert the car's motion into electricity, to feed the batteries. Here, the LEM converts the translator's kinetic energy into electricity.

$38 BILLION

Projected global market for integrated large-scale and distributed renewables over the next five years

Our control computer immediately adjusts the flow of current through the coils via an array of power-switching transistors to make the LEM apply more or less force. The LEM can hit a desired turnaround position within about 1/10th of a millimeter, then target and hit a different turnaround position on the next cycle. The system determines a turnaround position at which the level of compression triggers the reaction just before the end of the stroke, the most efficient point.

This ability to automatically and rapidly adjust compression is remarkable in two ways.

First, the generator maintains the optimum reaction process throughout the entire load range, from idle all the way to full power, in order to follow demand. For example, if power demand drops, the fuel will flow more slowly and the fuel molecules will thus be a little more dilute; they'll need a little more compression, and our system will provide just the right amount.

One real-world example of the system working this way pairs our generators with a 3.3-megawatt rooftop solar array. When the sun is shining, our generators turn off, and when the sun goes down or goes behind a cloud, our generators automatically turn on within seconds, immediately providing precisely as much power as the building requires.

Providing the compression that's needed, just when it's needed, also unlocks the capability to operate efficiently using fuels that have widely different properties. For example, hydrogen reacts with little compression, but ammonia requires a lot. The linear generator is fuel agnostic—it can run on a wide range of fuels, including natural gas, biogas, hydrogen, ammonia, syngas, and even alcohols without compromising performance.

That's the LEM. The remaining pieces of the architecture came about as we worked to maintain the inherent efficiency of the reaction in a real machine that has minimal losses like friction and heat transfer while running reliably for billions of cycles.

One of the biggest choices we had to make was the overall layout of the machine. We knew that the pressurized gas had to push on a moving wall directly connected to an electromagnetic force, but there were several ways to make that happen. In the first year or so we founders, together with seven other engineers, spent many hours at a whiteboard considering our options. Ultimately, we chose a symmetrical layout with two translators meeting in a single, central cylinder. Our fuel-air mixture, slightly pressurized, enters through the holes on one end. When the translators move away from that end, these holes are uncovered, and because the fresh mix-

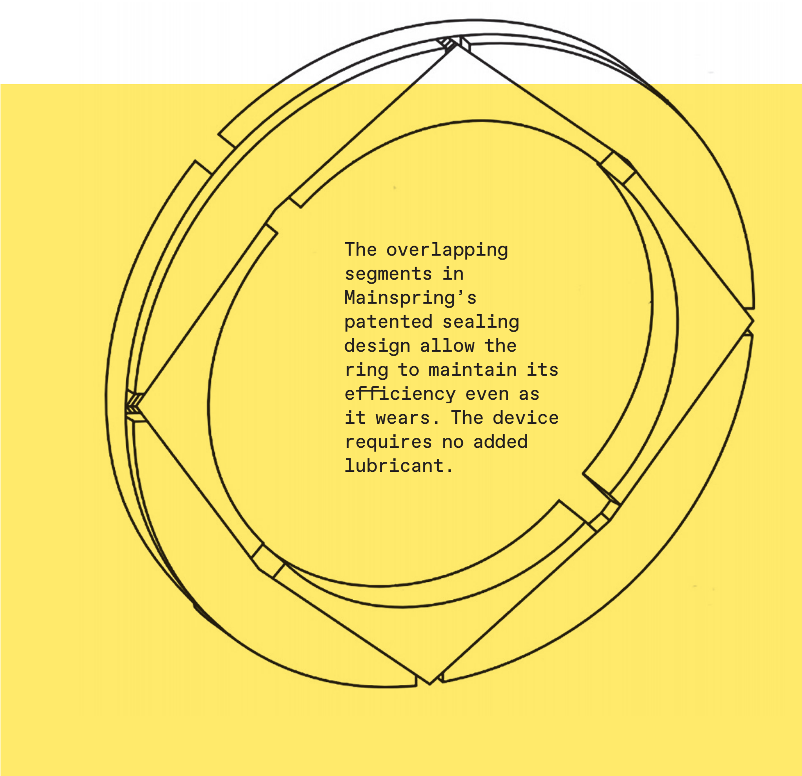

The overlapping segments in Mainspring's patented sealing design allow the ring to maintain its efficiency even as it wears. The device requires no added lubricant.

ture is at slightly higher pressure, it flows into the cylinder, pushing the used materials out the holes on the other end.

This choice replaces the conventional engine valve train—valves, seats, guides, seals, springs, rockers, camshaft, bearings, timing chain, and oil lubrication—with a simple set of holes in the cylinder wall. Another advantage of combining two translators in a single cylinder is the reduction of heat-transfer losses by nearly half.

Our last major design choice was to add an air chamber to either end of the generator. As the translators move outward during the expansion portion of the cycle, the outer ends of the translators compress plain air in the outer chambers, thus storing a fraction of the reaction energy. This stored energy is recovered afterward, when the compressed air pushes the translators back toward the center to start the next compression cycle. It's the same idea as storing energy by compressing and releasing a mechanical spring. This way, the LEMs can apply their braking forces and generate power in both directions, allowing us to cut their size by half.

We also let a small amount of this pressurized air out of our system to feed air bearings. Compared with oil-lubricated bearings, air bearings have lower friction and simpler seals. They work just like an air hockey game, where an array of small holes creates a pressurized film of air on which the puck floats.

IN 2012, ABOUT A YEAR and a half after our initial round of $10 million, we completed the first prototype that generated power. It put out only 1 kW.

A couple of days after we had gotten it to work for the first time, one of our investors let us know that he was planning to drop by our Menlo Park, Calif., headquarters to see it run. The engineer who had done most of the electrical design realized that, for a demo, we needed a way to see it making power, so he ran out to a nearby hardware store, bought a couple of halogen work lights, and plugged them directly into the electrical bus. Though barely more impressive than a school science project in which a potato powers a light bulb, it proved that our design worked.

But the output was a long way from our commercial target, 200 kW, a number we had picked because it would provide enough power for a typical retail store.

Mainspring Energy's first commercial product contains two linear generator cores. This unit, installed outside a store in Northern California, can produce up to 230 kilowatts of power.

OUR NEXT MILESTONE came in late 2013, when we built a 50-kW machine. And…it didn't work at all.

It had a teething problem not uncommon with large power equipment. An array of coils switching high voltage at a relatively high frequency generates a lot of electrical noise. In our device, it fed back to our position sensor and caused the LEM to vibrate, creating a sound we called "the crunchies." Our electrical and controls engineers were able to work through the problem and eliminate it.

But then we hit a wall—literally: The side of the translator would scrape along the cylinder wall whenever we tried to produce more than a few kilowatts.

To explain what happened, I need to describe one more component of our linear generator: the seal between the translator and the cylinder wall. This seal exists to keep the pressurized gas from escaping while still allowing the translator to slide.

Typically, you would use a layer of liquid oil between the two parts to avoid friction. But remember, we're adding fresh air and fuel into the cylinder through holes in the cylinder wall, and if we used a liquid lubricant in this arrangement, it would be nearly impossible to keep it from getting into the fuel mix and burning during the reaction process, creating noxious emissions.

So we decided to develop an oil-free sealing system. It worked well in our 1-kW device, and so we scaled the same design up for the 50-kW model. But though the machine got larger, the clearance requirements stayed the same in an absolute sense and thus were tighter in a relative sense. That allowed tiny distortions in the components to create points of friction, causing further distortions, ending with a runaway scraping problem.

After months of trying various tweaks, we still couldn't run beyond around 20 percent of full power without scraping. So we threw out the old sealing design and started over. We ended up inventing a unique carbon sealing-ring assembly that floats independently from the translator, one that can expand as it wears down, thus maintaining its seal.

This fixed the problem, and within a few more months we were running at full power for hundreds of hours. The next big scaling step—from 50 kW to 100 kW—was less difficult and culminated in our first official prototype, which we installed in the parking lot behind our building.

CONTINUED ON P. 46

CodeController(ParagraphEditor)>>processHeader
CodeController(ParagraphEditor)>>processSubHeader
CodeController(ParagraphEditor)>>processByline

The Machine That Transformed Computing

The Xerox Alto combined windows, WYSIWYG, menus, and networking to make the modern PC

By David C. Brock

User In

(Form readFrom: 'TheMachineThatTransformed(

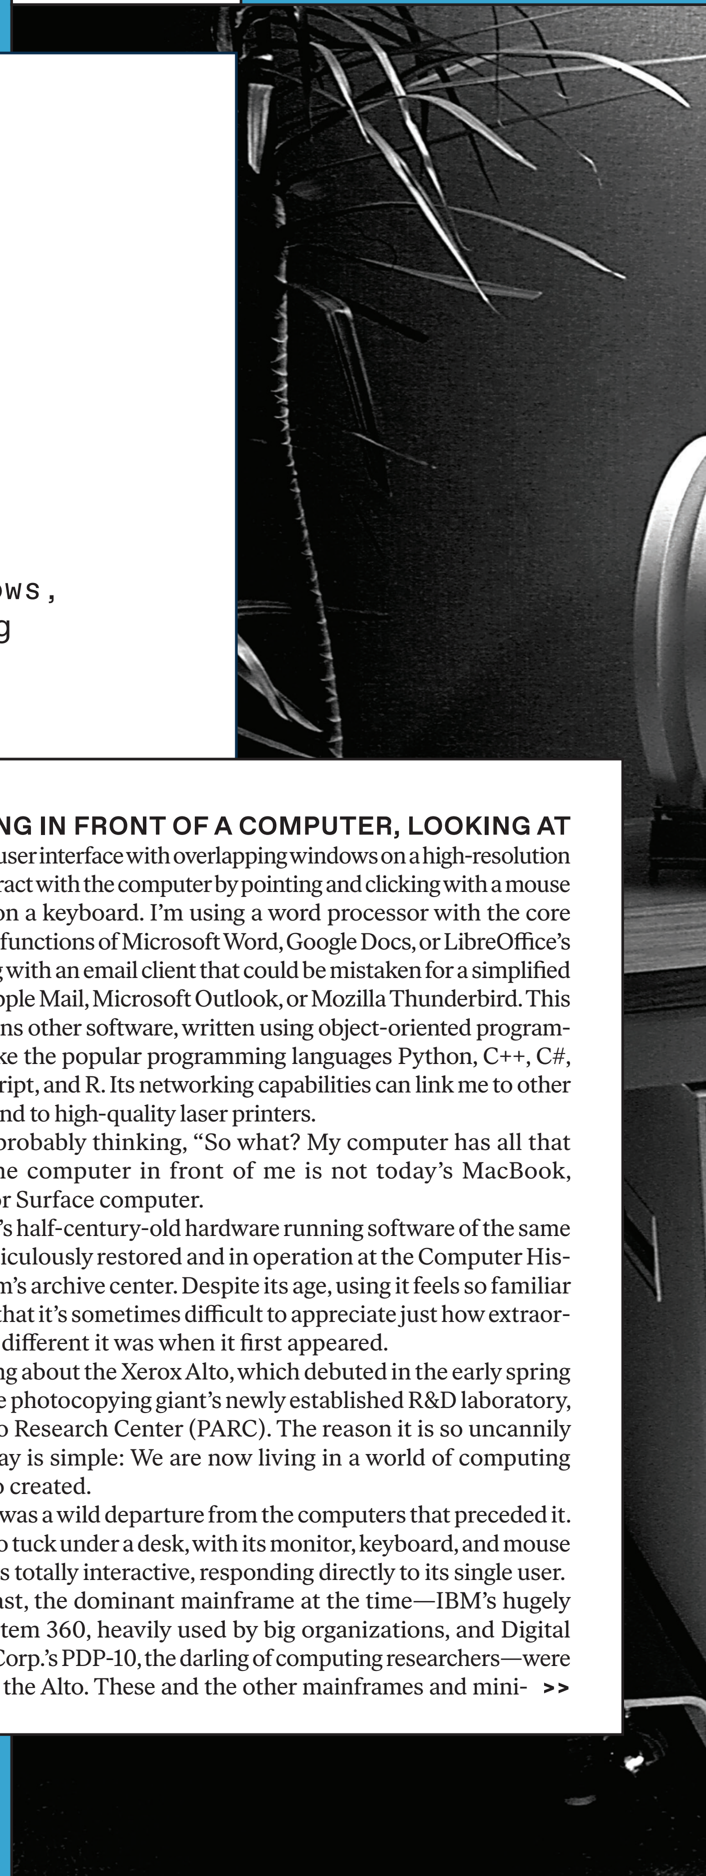

PARC, A XEROX COMPANY

Form Editor

I'M SITTING IN FRONT OF A COMPUTER, LOOKING AT its graphical user interface with overlapping windows on a high-resolution screen. I interact with the computer by pointing and clicking with a mouse and typing on a keyboard. I'm using a word processor with the core features and functions of Microsoft Word, Google Docs, or LibreOffice's Writer, along with an email client that could be mistaken for a simplified version of Apple Mail, Microsoft Outlook, or Mozilla Thunderbird. This computer runs other software, written using object-oriented programming, just like the popular programming languages Python, C++, C#, Java, JavaScript, and R. Its networking capabilities can link me to other computers and to high-quality laser printers.

You are probably thinking, "So what? My computer has all that too." But the computer in front of me is not today's MacBook, ThinkPad, or Surface computer.

Rather, it's half-century-old hardware running software of the same vintage, meticulously restored and in operation at the Computer History Museum's archive center. Despite its age, using it feels so familiar and natural that it's sometimes difficult to appreciate just how extraordinary, how different it was when it first appeared.

I'm talking about the Xerox Alto, which debuted in the early spring of 1973 at the photocopying giant's newly established R&D laboratory, the Palo Alto Research Center (PARC). The reason it is so uncannily familiar today is simple: We are now living in a world of computing that the Alto created.

The Alto was a wild departure from the computers that preceded it. It was built to tuck under a desk, with its monitor, keyboard, and mouse on top. It was totally interactive, responding directly to its single user.

In contrast, the dominant mainframe at the time—IBM's hugely popular System 360, heavily used by big organizations, and Digital Equipment Corp.'s PDP-10, the darling of computing researchers—were nothing like the Alto. These and the other mainframes and mini- >>

By 1975, dozens of Xerox PARC's researchers had personal Altos in their offices and used them daily. The large cabinet contained a CPU, memory, and a removable disk pack. On the desk are additional disk packs and the Alto's vertical display, mouse, and keyboard.

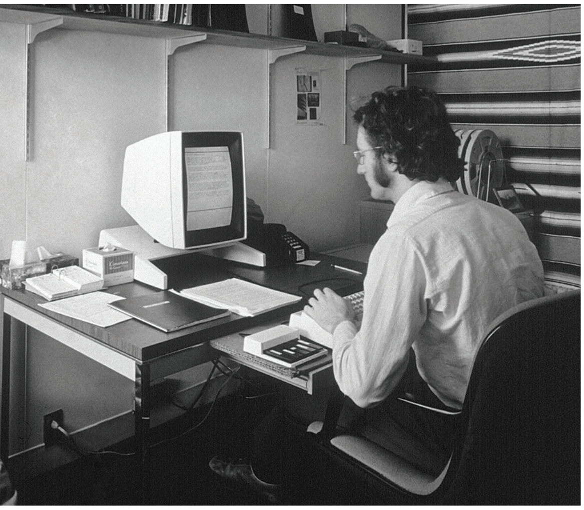

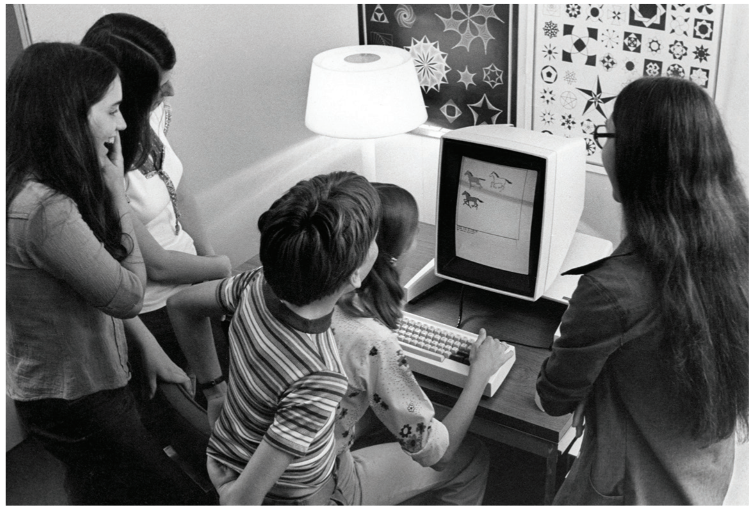

>> computers of the era were room-size affairs, almost always located somewhere away from the user and almost always under the control of someone else. The many simultaneous users of one such computer shared the system as a common resource. They typically connected to it with a teletypewriter, though the most avant-garde users may have employed simple text-only video terminals.

The people who developed the Alto came to Xerox PARC from universities, industrial labs, and commercial ventures, bringing with them diverse experiences and skills. But these engineers and programmers largely shared the same point of view. They conceived and developed the Alto in a remarkable burst of creativity, used it to develop diverse and pathbreaking software, and then moved out of Xerox, taking their achievements, design knowledge, and experiences into the wider world, where they and others built on the foundation they had established.

Broadly speaking, the PARC researchers set out to explore possible technologies for use in what Xerox had tagged "the office of the future." They aimed to develop the kind of computing hardware and software that they thought could be both technologically and economically possible, desirable, and, perhaps to a lesser extent, profitable in about 10 to 15 years.

The type of computing they envisioned was thoroughly interactive and personal, comprehensively networked, and completely graphical—with high-resolution screens and high-quality print output.

This vision wasn't entirely new or limited to Xerox PARC. Rather, the elements of it had developed over the previous decade among a community of computing and artificial-intelligence researchers at a handful of academic institutions and firms. The U.S. Department of Defense's Advanced Research Projects Agency (ARPA, now DARPA), the primary funder for academic computing and AI research in the United States from 1961 to 1970 (and for decades beyond), had curated these researchers for generous support contracts. Naturally enough, these ARPA contractors became leaders in U.S. computing.

So the broad vision for computing's future was well in place when, at the close of the 1960s, Xerox decided to create a new corporate R&D laboratory. Xerox had established a patent position that gave it a virtual monopoly in photocopying, controlling 90 percent or more of the US $1.7 billion U.S. market in 1972—about $12 billion in today's dollars. The riches that resulted made it easy for the company to bankroll this new R&D laboratory focused on forward-looking projects.

Oddly, at the time, an expensive new laboratory was also immediately financially attractive: R&D expenditures were frequently counted as assets instead of business expenses, all with Wall Street's approval. The more you spent, the better your balance sheet looked.

The new laboratory was to include sections intended to extend Xerox's lock

> Remarkably, the researchers at PARC had actually succeeded in making the computer of the future.

on patents surrounding the physical production of documents in offices, that is, materials like photoconductors—which become more electrically conductive when exposed to light—along with toners and optics, all key parts of copier technology. Other sections were to be devoted to computing, including the Computer Science Laboratory led by Jerry Elkind and Bob Taylor. The two were disciples of J.C.R. Licklider, long a prominent evangelist for personal, interactive, networked computing through his roles at MIT, Bolt Beranek & Newman (a contract research company deeply involved in building the Internet, now Raytheon BBN), ARPA, IBM, and numerous professional organizations.

This wasn't Xerox's only big bet on computing. By the late 1960s, Xerox executives had begun to see information technology making an impact in the office environment. Computers were generating reports, inventories, and analyses on paper for the use of office workers. IBM, which was the dominant producer of the quintessential office machine, the electric typewriter, had become a hugely profitable giant in computing. So in March 1969, Xerox agreed to acquire computer-manufacturer SDS—which had made some inroads into the timeshared-computing market with minicomputers—for the eyebrow-raising price of $900 million (a value of more than $7 billion today).

PARC, A XEROX COMPANY

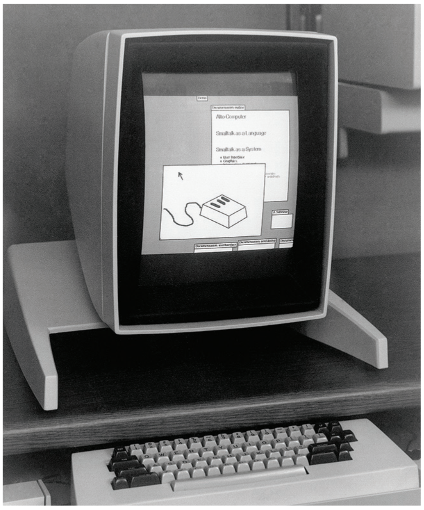

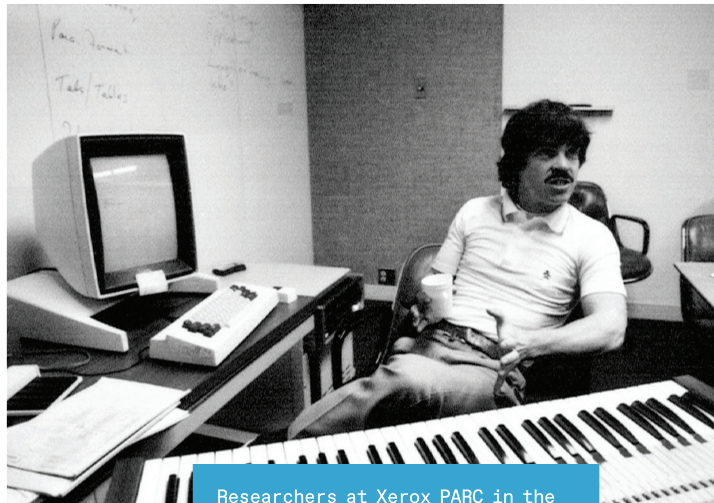

Researchers at Xerox PARC in the 1970s found a huge variety of uses for their Alto computers. For example, Larry Tesler [far left, opposite page] developed modeless word-processing software; researchers conducted learning experiments with children; and Alan Kay [above] attached the keyboard of an electric organ to his Alto. Bob Taylor [center], comanager of PARC's Computer Science Laboratory, gave the computer its name.

Just a few months later, Xerox's leadership was debating where to locate their new laboratory. In a memo to the CEO in June of 1969, Jacob Goldman, Xerox's chief scientist, who was responsible for setting up the new laboratory, outlined some possible locations and some of the problems with them:

New Haven? "Traditional Yale faculty snobbery." Princeton? "Community not a hospitable one." Boston? "Job hopping is a way of life." Southern California? "Attractive, smog-free residential areas too remote." The late Chuck Thacker, perhaps the key designer of the Alto hardware, kept a photocopy of this memo, which now resides in his papers at the Computer History Museum.

Initially, the executives deemed Palo Alto too far from other Xerox locations, but they quickly overcame that reservation, and the Xerox Palo Alto Research Center was born.

To populate PARC's computing research sections, Taylor harvested throughout the ARPA community, enticing many of its promising young researchers by offering them a new, well-heeled venue for pursuing their collective vision for interactive, networked, graphical, personal computing. Into PARC flowed many former ARPA contractors, including some from Stanford's Artificial Intelligence Laboratory, the University of Utah's computer graphics operation, Doug Engelbart's group at SRI, and BBN's artificial-intelligence efforts. Taylor also drew researchers from the Berkeley Computer Corp. (BCC), a struggling startup that had emerged from ARPA-funded Project Genie's attempt to build an interactive programming environment on a timeshared computer at the University of California, Berkeley.

A **Alan Kay, who had been** known in ARPA circles first as a graduate student at the University of Utah and then as a young researcher at Stanford, joined PARC soon after it opened its doors in 1970. There, Kay established the Learning Research Group, centered around his conception of a hand-held, screen-oriented, graphical, networked personal computer, which he called the Dynabook. If you've ever used an iPad or other tablet computer, the similarity to Kay's Dynabook idea is astonishing.

Kay envisioned a new sort of software environment and programming language for the Dynabook. In this digital world, children and adults alike would be able to create their own tools, models, and simulations; share them and build on one another's work; and exchange the resulting knowledge. The key to accomplishing all that would be a new approach to coding that came to be known as object-oriented programming.

Alan Kay's 1968 cardboard mock-up of the Dynabook foreshadowed generations of portable computers.

Not long after Kay had articulated his vision in an internal PARC memo and had attracted researchers like Dan Ingalls, and later Adele Goldberg, to work toward it with him, he was approached by Thacker and Butler Lampson. These two PARC researchers also wanted to build a small single-user computer.

Lampson and Thacker had come to PARC from the ashes of BCC to work in PARC's Computer Science Laboratory. At BCC, and previously at Project Genie, the pair had made significant contributions to timeshared computing. The essence of timesharing was just what it sounds like—sharing the capabilities and resources of a larger computer among multiple simultaneous users. Lampson and Thacker now suspected that the future might not lie in timesharing but rather in small computers, each used exclusively by an individual and networked together for communication and for sharing files of all sorts.

Lampson and Thacker proposed to Kay that he give them most of his group's annual budget to build one of these small

MARK RICHARDS/COMPUTER HISTORY MUSEUM

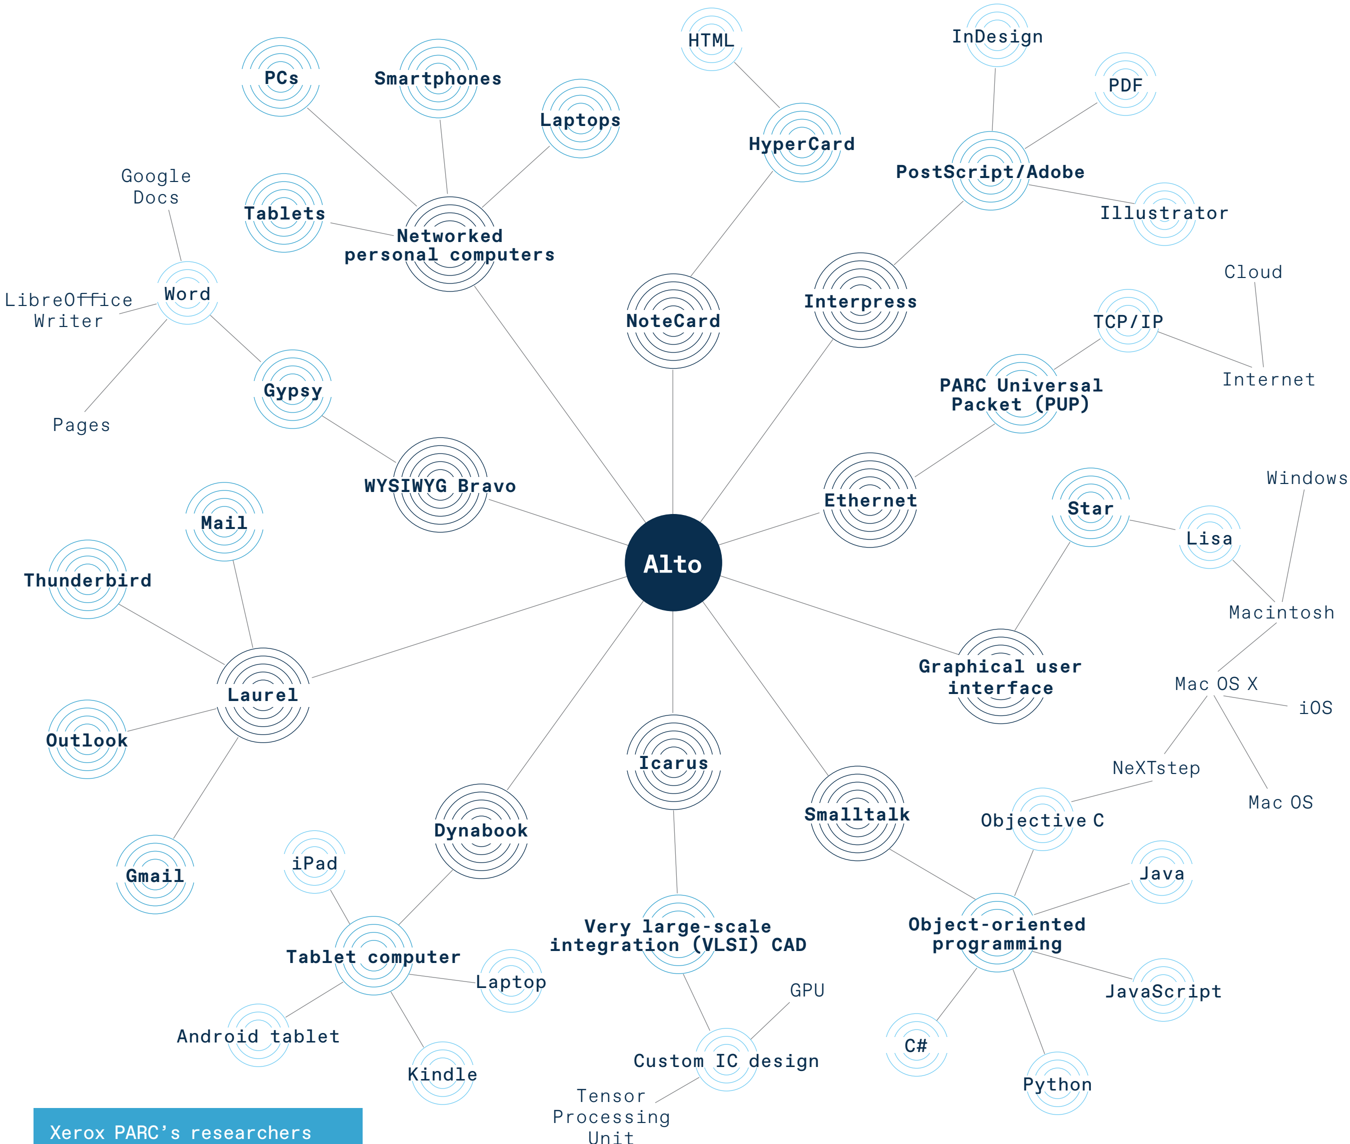

Xerox PARC's researchers took the ideas implemented in the Alto and sent them out into the world, where they are reflected in software and hardware being used today.

computers quickly and relatively inexpensively. It would be ready in just a couple of months and have many of the characteristics of the Dynabook. But instead of being portable, this small computer would fit under a desk. And instead of being able to sketch on a screen with a stylus, as Kay imagined for the Dynabook, it would use a mouse (an earlier invention from Engelbart's SRI laboratory) as a pointer to navigate a cursor on a high-resolution graphical display.

Kay gambled his budget on Lampson and Thacker's proposal, calling it the "Interim Dynabook." Taylor, the charismatic comanager of the Computer Science Laboratory, quickly gave the small computer its official name: the Alto.

PARC's collective bet on the future of integrated circuits shaped Thacker's approach to designing the Alto. The computer would require a large amount of expensive main memory to display its graphics on-screen, but PARC's researchers reasoned that the trend for integrated circuits of increasing complexity to slash the cost of computing (soon to be known as Moore's Law) was real and would continue. So although the cost of the Alto's memory using the new 1103 DRAM chips from Intel would be huge at first, in a decade or so, the researchers figured, the cost of memory would be exponentially less.

In December 1972, Lampson penned an internal memo, "Why Alto?," that argued for PARC creating large numbers of Alto computers. Of course, Kay's group would need a dozen or so for the development of their inventive software environment and programming language (which would soon become known as Smalltalk) and their planned learning experiments with children. But, Lampson argued, the Alto would also be perfect for much broader experiments in personal computing and networking.

While Elkind, the other comanager of the Computer Science Laboratory, was skeptical of the proposal, he became a convert after seeing the prototype Alto in April 1973. By 1974, as historian Leslie Berlin has noted, Elkind was promoting the Alto in the strongest possible terms to top company executives as Xerox's route to the future of computing.

He wasn't the only convert. Once researchers at PARC tried out the Alto, they wanted it. Over the years, hundreds of these machines were produced, proliferating throughout PARC, into Xerox more widely, and eventually outside of the firm in select locations, including university laboratories and even the White House.

It was Bob Metcalfe, then a young networking expert straight out of graduate school at Harvard, who came up with the local-networking approach for the Altos that would eventually become critical in computer networking writ large.

On 22 May 1973, Metcalfe wrote a memo describing his "ETHER Network." His design built on networking technology from the famed Arpanet, then being constructed, along with an experimental digital radio network developed at the University of Hawaii called ALOHAnet. By November 1973, Metcalfe and another PARC researcher, David Boggs, had developed a network that began to come to life inside the research center.

Metcalfe left PARC in 1979 to found 3Com, which, along with other startups, commercialized Ethernet. Ethernet soon became the dominant local-networking standard and remains a vitally important wired-networking technology used today.

Remarkably, the researchers at PARC had actually succeeded in making the computer of the future. The networked Alto machines—and the astonishing array of software developments they enabled—wove together the key elements of personal computing that surround us to this day.

But in no way was the Alto the last—or the complete—word on personal computing. Xerox's attempt to turn the Alto into a true commercial product, 1981's Xerox Star, introduced a common graphical interface on the metaphor of a desktop, with graphical icons for files, folders, printers, and the like. Nevertheless, the true success of the Alto's computing idiom required that the technology escape the confines of the monopolist firm that had given rise to it. The approach pioneered at PARC only truly thrived in the more open, horizontal, competitive market of the early personal-computer industry. Success required a larger community.

Members of Xerox PARC's Learning Research Group, including Alan Kay [with mustache, farthest from camera] and Adele Goldberg [left of Kay, auburn hair, leaning forward] meet in the fabled "beanbag room."

ALAN KAY

Steve Jobs and a whole team from Apple toured PARC in 1979. The visit was arranged as a quid pro quo for allowing Xerox to invest in Jobs's exciting new personal-computer company. Viewing Alto's graphical user interface, Jobs had what he later described as an epiphany, one that reoriented his efforts at Apple forever after. He quickly hired Larry Tesler from PARC. Tesler had made key contributions to the Alto's software, including programs for document editing, printing, and Smalltalk.

Many other PARC researchers would join Apple and help it bring the graphical user interface into prominence. Meanwhile, PARC researcher Charles Simonyi, who had developed the Alto's extremely popular Bravo "What you see is what you get" (WYSIWYG) word processor, left for Microsoft, where he would work to turn Bravo into Word and launch Microsoft into the world of application software.

Smalltalk was eventually commercialized by a PARC spin-off at the end of the 1980s, accelerating the profound effects that Smalltalk and object-oriented programming were already having on the development of software and coding.

Also in the early 1980s, an entire group of researchers left PARC to start Adobe, aiming to commercialize the approaches to computer printing and digital typography they had explored at PARC. Adobe's technology was essential to what became a multibillion-dollar market in desktop publishing and, later, to the now ubiquitous PDF.

While these and many other companies made computing the Alto way the industry standard, Xerox, too, directly benefited from the research at PARC. Laser printing, invented by Gary Starkweather at PARC not long after the opening of the laboratory, paid Xerox handsome dividends as it slowly but surely displaced other technologies for document duplication and printing.

As should now be apparent, how the Alto came to shape our lives with computers a half century later isn't the story of any one individual. In our culture, however, the history of technology is habitually presented as a sequence of remarkable individual achievements. But this is wrong. Innovation is the work of groups, of communities. These provide the context and the medium for the actions of the individual. Leadership is a meaningless concept outside of a group.

The remarkable story of the Alto is the story of such communities. It is a story of how a broad research community developed a shared vision for interactive, networked, graphical, personal computing. It is a story of how a smaller group of talented individuals came together in a new laboratory to realize that vision and to experiment with it. And it is a story of this group moving on, finding new colleagues and organizations in the rapidly expanding personal-computer industry, and working for decades to bring the Alto way of computing to the world. ■

The Computer History Museum holds extensive collections about the Alto and the community that created it, including a remarkable digital archive that will be released soon. For news, visit https://computerhistory.org/art-of-code/

Roboticists Want to Give You a Third Arm CONTINUED FROM P. 27

new technologies with extra robotic limbs, which we call the MUlti-limb Virtual Environment, or MUVE. Among other capabilities, MUVE will enable users to work with as many as four lightweight wearable robotic arms in scenarios simulated by virtual reality. We plan to make the system open for use by other researchers worldwide.

C**onnecting our control technology** to a robotic arm or other external device is a natural next step, and we're actively pursuing that goal. The real challenge, however, will not be attaching the hardware, but rather identifying multiple sources of control that are accurate enough to perform complex and precise actions with the robotic body parts.

We are also investigating how the technology will affect the neural processes of the people who use it. For example, what will happen after someone has six months of experience using an extra robotic arm? Would the natural plasticity of the brain enable them to adapt and gain a more intuitive kind of control? A person born with six-fingered hands can have fully developed brain regions dedicated to controlling the extra digits, leading to exceptional abilities of manipulation. Could a user of our system develop comparable dexterity over time? We're also wondering how much cognitive load will be involved in controlling an extra limb. If people can direct such a limb only when they're focusing intently on it in a lab setting, this technology may not be useful. However, if a user can casually employ an extra hand while doing an everyday task like making a sandwich, then that would mean the technology is suited for routine use.

Other research groups are pursuing similar neuroscience questions with different types of control mechanisms. Domenico Prattichizzo and colleagues at the University of Siena, in Italy, have demonstrated a wrist-mounted soft robotic sixth finger. It enables a user with a hand weakened by a stroke to grip objects securely. Users wear a cap with EMG electrodes and send commands to the finger by raising their eyebrows. Harry Asada's group at MIT has experimented with many types of extra robotic limbs, including a wearable suit that used EMG to detect muscle activity in the torso to control extra limbs.

Other groups are experimenting with control mechanisms involving scalp-based EEG or neural implants. It is early days for movement augmentation, and researchers around the world have just begun to address the most fundamental questions of this emerging field.

Two practical questions stand out for our system: Can we achieve neural control of extra robotic limbs concurrently with natural movement, and can the system work without the user's exclusive concentration? If the answer to either question is no, we won't have a practical technology, but we'll still have an interesting new tool for studying the neuroscience of motor control. If the answer to both questions is yes, we may be ready to enter a new era of human augmentation. For now, our (biological) fingers are crossed. ■

The Omnivorous Generator

CONTINUED FROM P. 39

WE STILL NEEDED to make the linear generator affordable. The technology had the advantage of using fewer parts than engines or turbines and lacking the expensive catalyst of fuel cells. But we had to figure out package design, engineering for high-volume manufacturing, and the supply chain for a product we decided would consist of two side-by-side linear generators for a total power of 230 kW. We made a few mistakes along the way.

One big one involved our efforts to reduce the cost of physically attaching the magnet array to the outside of the translator tube. In the prototypes, we secured the magnets against the tube by winding resin-impregnated Kevlar fiber around the outside of the glued-on magnets. In our first attempt at cost reduction, we switched to an impregnated cloth wrap that went on more quickly and easily. But after building a couple of units with this approach we discovered that magnets were coming loose under the wrap. So we went back to the wound-Kevlar approach, and eventually reduced its cost by developing an automated winding process.

FINALLY, IN JUNE 2020, in the thick of the COVID pandemic, a crew pulled a flatbed truck up to our Silicon Valley headquarters, loaded up the first-in-the-world production linear generator, and drove it 30 kilometers to a paying customer's site—part of a national retail chain. A couple of days later we flipped the switch, and we were in business! A few months later we delivered our second unit to a Kroger store in Southern California, and shortly after that a pair of units went to a Lineage Logistics cold-storage facility.

When we started the company, we optimized the first generator for natural gas because it was then most widely available, least expensive, and relatively clean. Even though it does produce carbon emissions, our system's efficiency makes it greener than the traditional generators that it replaces.

We see our linear generator as the cornerstone of a zero-carbon grid because of its unique flexibility: It can handle nearly any scale of power, from single units to grid-connected arrays; it's easily permitted and installed wherever power is needed; and it runs on almost any fuel. We have run one of our stock units on hydrogen and on anhydrous ammonia. We have a customer project running on renewable biogas in a landfill. We plan to start operating other biogas projects at wastewater treatment plants and dairy waste digesters this year. We are getting ready to deploy arrays of up to dozens of generators for large-scale operations, like electric-truck charging. And we are now designing larger, utility-scale versions in the megawatt output range. These will all use the same core technology without any radical design changes.

And yes, Professor Edwards, we think we have answered that question you posted some 20 years ago: "What is the most efficient and practical way possible to convert chemical-bond energy into useful work?"

It's the linear generator. ■

THE INSTITUTE

NEWS OF THE IEEE
VOLUME 47 / ISSUE 1

How Diana Gamzina's Startup Is Slashing the Cost of Satellite Amplifiers P. 62

Nvidia's CTO on the Future of Supercomputing P. 50

Internet Pioneer Vint Cerf Receives Medal of Honor P. 52

Photography by Gabriela Hasbun

Making IEEE a Force for Change

AS 2023 IEEE president, it is my goal to work with all members—particularly our students, young professionals, and affinity committee members—to make IEEE a more successful and resilient global technical organization and for there to be global recognition of IEEE as a force for change.

For more than 40 years, IEEE has been an integral part of my pursuit of excellence in my professional life. It has given me the opportunity to speak at more than 200 IEEE events, where I have had the pleasure of engaging with academics, young professionals, and midcareer engineers in industry and government, some of whom are women and underrepresented minorities. Such engagements at the grassroots level have provided me deeper insights into understanding the needs of members from diverse communities and how IEEE can best provide them with opportunities to advance their career.

Through proactive outreach, we can demonstrate the benefits of membership and IEEE's relevance to current and prospective members. These offerings include unparalleled networking opportunities at more than 2,000 international IEEE conferences—both face to face and virtual—as well as access to the finest technical literature, massive resources for upskilling, and innovative collaboration facilities with colleagues worldwide.

I appreciate both the global nature of IEEE as well as the rich diversity of the engineering and technical fields in which our members excel. We must continue to look for opportunities to provide more value and services to professionals working in industry, academia, and public service, while paying keen attention to underrepresented communities such as women and entrepreneurs, as well as our students and young professionals, who are our future technologists.

Realizing IEEE's mission

IEEE's mission is to advance technology for the benefit of humanity. Each of us can play a critical role in improving living conditions for society. IEEE members are responsible citizens and therefore seek to help address global challenges using technological solutions.

Today the world faces its largest modern-day threat: climate change. We see the effects that extreme weather—such as changes in weather patterns, crop losses, and rising sea levels—has on communities worldwide. It's natural for us to want to contribute.

As the world's largest organization of technical professionals, IEEE has both the opportunity and the responsibility to assist in organizing the response of engineers, scientists, and technical professionals across the world to address the causes, mitigate the impact, and adapt to climate change. IEEE recognizes this global crisis and is committed to helping alleviate the effects of climate change through pragmatic, accessible technical solutions for sustainability, and providing engineers and technologists with a neutral space for discussion and action.

IEEE has significant ongoing, relevant efforts and expertise that can be brought to bear on this issue. Our scholarly publications, conference proceedings, technical standards, and other materials can help foster the exchange of technical knowledge and information for the critical climate issues that our planet faces. Two significant resources are the IEEE Climate Change website (ieeeclimatechange.org) and the IEEE Climate Change Collection in the IEEE Xplore Digital Library (ieeexplore.ieee.org). The collection is a repository of more than 7,000 IEEE articles addressing the causes of climate

FOLLOW ME ON SOCIAL MEDIA FOR UPDATES AND INFORMATION FROM ACROSS IEEE

Facebook: @ieeepresident Instagram: @ieeepresident
Twitter: @IeeePresident LinkedIn: linkedin.com/showcase/ieeepresident

IEEE

change along with strategies to mitigate and adapt to it.

IEEE also has been reaching out to all major engineering and technology societies in Asia, Africa, Europe, North America, and South America to identify areas of collaboration and portfolios of solutions to address climate change.

Our organization has the credibility to bring experts to the table to present diverse viewpoints and seek common solutions. For example, IEEE had a significant presence at the U.N. Climate Change Convention (COP27) in Egypt in November, where I spoke at events organized by two engineering organizations: the Global Alliance of Universities on Climate and the International Renewable Energy Agency.

I believe IEEE's impact starts at the local level, beginning with building awareness of the multitude of resources membership provides. With more than 340 sections around the world, local engineers, technologists, and computer scientists can deliver technology solutions for their unique communities. Our focus should be on identifying IEEE members and their colleagues at the local level who can be the points of contact for those looking for solutions to these complex issues. Another focus should be bringing academics and industry practitioners together to share their ideas and experiences with one another and with the next generation of engineers.

I welcome all IEEE members and volunteers around the globe to contribute their experiences and insights so that together we can explore and discover solutions to address climate change. I encourage our community to be contributors and influencers for climate action through research, knowledge sharing, technology advancements, solutions development, and more.

Let us employ the opportunities that IEEE offers—such as international cooperation, continuing professional development, and ambitious technology development—to build a more sustainable future.

—SAIFUR RAHMAN
IEEE president and CEO

Please share your thoughts with me at president@ieee.org.

EDITOR'S NOTE

The Enduring Impact of the Internet

Vint Cerf and other tech pioneers are connecting the world

EVERY TIME YOU connect to the Internet, you're relying on technology that this year's IEEE Medal of Honor recipient, Vint Cerf, helped create [page 52].

As a developer of the TCP/IP protocols, Cerf led the effort to build the early computer network that eventually became the Internet. These days, he is vice president and chief Internet evangelist for Google. He is receiving IEEE's highest honor for helping to create "the Internet architecture and providing sustained leadership in its phenomenal growth in becoming society's critical infrastructure."

Of course, many engineers over the years have played key roles in designing today's Internet. You probably use products every day that Michael Kagan had a hand in developing. As a cofounder and former chief technology officer at Mellanox Technologies, Kagan helped develop the company's high-speed InfiniBand products, which are found in most of the world's supercomputers. Mellanox's Ethernet products are in nearly all cloud data centers.

Kagan became chief technology officer of Nvidia after that company acquired Mellanox in 2020. On page 50, find out what the IEEE senior member is working on now. One technology he is excited about is Nvidia's Omniverse, a real-time graphics collaboration used by carmakers to test autonomous vehicles. Kagan is also helping to build Earth-2, a climate modeling initiative designed to let people simulate mitigation techniques to combat global warming.

Cerf and Kagan share with other engineers a drive to develop better products and, in the process, improve people's lives. Building a better wheelchair became Rory Cooper's mission after a spinal-cord injury left him paralyzed from the waist down. He designed an electric-powered wheelchair that helped the user stand up. Since then, the life Fellow has devoted his career to developing assistive technology for others. On page 54, learn about his inventions that have made the world more accessible for wheelchair users.

Diana Gamzina is helping to drastically lower the cost of satellite communications. Her startup, Elve, builds millimeter-wave power amplifiers, the vacuum-electronics devices that enable communication with satellite networks and create long-distance ground-to-ground links.

Amplifiers on the market now cost as much as US $1 million apiece and take up to a year to produce using conventional manufacturing processes. On page 62, read how the IEEE senior member is using advanced materials and new manufacturing technologies to make an amplifier in a week and sell it for just 10 percent of the usual price. Her ultimate goal is to have her amplifiers installed on cellphone towers, enabling high-data-rate communication in remote and rural locations.

The engineers featured in this issue of *The Institute* are committed to connecting the world by removing barriers.

—KATHY PRETZ
Editor in chief, *The Institute*

For updates about IEEE and its members, visit us at spectrum.ieee.org/the-institute

PROFILE

Michael Kagan on the Future of High-Performance Computing

Nvidia's CTO builds a supercomputer to model climate change

BY KATHY PRETZ

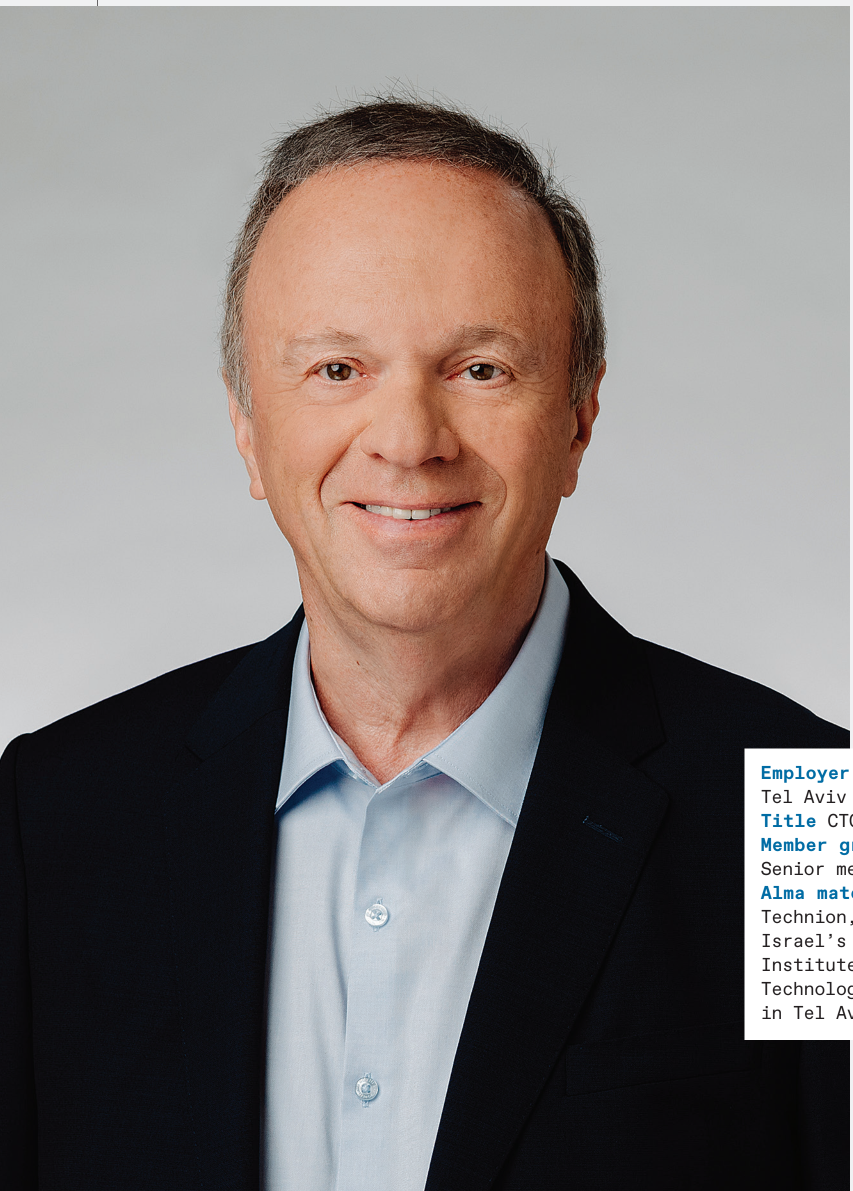

Employer Nvidia, Tel Aviv
Title CTO
Member grade Senior member
Alma mater Technion, Israel's Institute of Technology in Tel Aviv

NVIDIA (2)

IN 2019 MICHAEL Kagan was leading the development of accelerated networking technologies as chief technology officer at Mellanox Technologies, which he and eight colleagues had founded two decades earlier. Then in April 2020 Nvidia acquired the company for US $7 billion, and Kagan took over as CTO of the tech goliath—his dream job.

At Mellanox, based in Yokneam Illit, Israel, Kagan had overseen the development of high-performance networking for computing and storage in cloud data centers. The company made networking equipment such as adapters, cables, and high-performance switches, as well as a new type of processor, the DPU (data-processing unit). The company's high-speed InfiniBand and Ethernet products are in most of the world's fastest supercomputers and cloud data centers, Kagan says.

He now focuses on integrating a wealth of Nvidia technologies to build accelerated computing platforms, whose foundation are three chips: the CPU, the GPU, and the DPU. The DPU can support the ability to offload, accelerate, and isolate data center workloads, reducing CPU and GPU workloads.

"At Mellanox we worked on the data center interconnect, but at Nvidia we are connecting state-of-the-art computing to become a single unit of computing: the data center," Kagan says. Interconnects are used to link multiple servers and combine the entire data center into one, giant computing unit.

"I have access and an open door to Nvidia technologies," he says. "That's what makes my life exciting and interesting. We are building the computing of the future."

From Intel to Mellanox

Kagan was born in St. Petersburg, Russia—then known as Leningrad. After he graduated high school, his family moved to Israel. As with many budding engineers, his curiosity led him to disassemble and reassemble things to figure out how they worked. And, with several engineers in the family, he says, pursuing an engineering career was an easy decision.

He joined Intel in Haifa, Israel, as a design engineer and eventually relocated to the company's offices in Hillsboro, Ore., where he worked on the 80387 floating-point coprocessor.

A year later, after returning to Israel, Kagan served as an architect of the i8060XP vector processor and then managed the design of the Pentium MMX microprocessor.

During his 16 years there, he worked his way up to chief architect. In 1999, while preparing to move his family to California to lead a high-profile project, a former coworker at Intel, Eyal Waldman, asked Kagan to join him and five other acquaintances to form Mellanox.

The team of cofounders and vision drew him in. During his nearly 22 years there, he said, he had no regrets. "It was one of the greatest decisions I've ever made," he says.

InfiniBand, the startup's breakout product, was designed for what today is known as cloud computing, Kagan says. More than half the machines at the top 500 computer companies use the Mellanox interconnect, now the Nvidia interconnect.

"Most of the cloud providers use Nvidia's networking and compute technologies," Kagan says, adding that the partnership between Mellanox and Nvidia was "natural," as the two companies had been doing business together for nearly a decade.

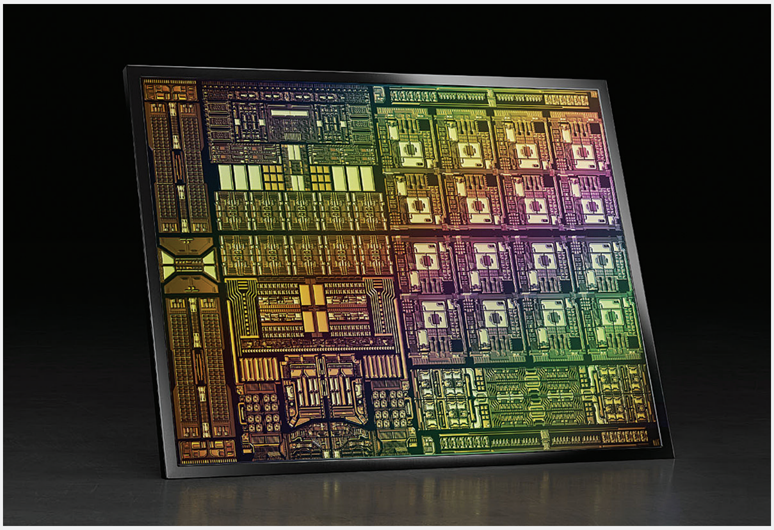

Nvidia's BlueField DPU offloads, accelerates, and isolates a variety of networking, storage, and security services.

Nvidia's supercomputers

As CTO of Nvidia, Kagan has shifted his focus from pure networking to the integration of multiple Nvidia technologies including building BlueField data-processing units and the Omniverse real-time graphics collaboration platform.

He says Nvidia's vision for the data center of the future is based on its three chips: CPU, DPU, and GPU.

"These three pillars are connected with a very efficient and high-performance network that was originally developed at Mellanox and is being further developed at Nvidia," he says.

Development of the BlueField DPU is now a key priority for the company. It is a data center infrastructure on a chip, optimized for high-performance computing. It also offloads, accelerates, and isolates a variety of networking, storage, and security services.

"In the data center, you have no control over who your clients are," Kagan says. "You're better off isolating yourself and other customers from each other by having a segregated or different computing platform run the operating system, which is basically the infrastructure management, the resource management, and the provisioning."

Kagan is particularly excited about the Omniverse, a product that uses Pixar's Universal Scene Description software for creating virtual worlds—what has become known as the metaverse. Kagan describes the 3D platform as "creating a world by collecting data and making a physically accurate simulation of the world."

Car manufacturers are using the Omniverse to test-drive autonomous vehicles. Instead of physically driving a car, data about the virtual world can be generated to train the AI models.

"You can create situations that the car has to handle in the real world but that you don't want it to meet in the real world, like a car crash," Kagan says.

The company is also building what it calls the most powerful AI supercomputer for climate science: Earth-2, a digital twin of the planet. By design, Earth-2 will continuously run models to predict climate and weather events at both the regional and global levels.

Kagan says the climate modeling technology will enable people to try mitigation techniques for global warming and see what their impact is likely to be in 50 years.

"We are actually moving forward at a fairly nice pace," he says. "But the thing is that you always need to reinvent yourself and do the new thing faster and better, and basically win with what you have and not look for infinite resources. This is what commitment means." ■

An IEEE senior member, Kagan says membership gives him access to information about technical topics that would otherwise be challenging to obtain. He likes connecting with members who are working on similar projects because he always learns something new.

"It inspires you to do your job in a different way," he says.

THE ROYAL SOCIETY

NEWS

Medal of Honor Goes to Vint Cerf

BY JOANNA GOODRICH

IEEE LIFE FELLOW Vinton "Vint" Cerf, widely known as the "father of the Internet," is the recipient of this year's IEEE Medal of Honor. He is being recognized "for co-creating the Internet architecture and providing sustained leadership in its phenomenal growth in becoming society's critical infrastructure."

The IEEE Foundation sponsors the annual award.

While working as a program manager at the U.S. Defense Advanced Research Projects Agency (DARPA) Information Processing Techniques Office in 1974, Cerf and IEEE Life Fellow Robert Kahn designed the Transmission Control Protocol and the Internet Protocol.

TCP manages data packets sent over the Internet, making sure they don't get lost, are received in the proper order, and are reassembled at their destination correctly. IP manages the addressing and forwarding of data to and from its proper destinations. Together they make up the Internet's core architecture and enable computers to connect and exchange traffic.

"Cerf's tireless commitment to the Internet's evolution, improvement, oversight, and evangelism throughout its history has made an indelible impact on the world," said one of the endorsers of the award. "It is largely due to his efforts that we even have the Internet, which has changed the way society lives.

"The Internet has enabled a large part of the world to receive instant access to news, brought us closer to friends and loved ones, and made it easier to purchase products online," the endorser said. "It's improved access to education and scientific discourse, made smartphones useful, and opened the door for social media, cloud computing, video conferencing, and streaming. Cerf also saw early on the importance of decentralized control, with no one company or government completely in charge."

Since 2005, Cerf has been vice president and chief Internet evangelist at Google in Reston, Va., spreading the word about adopting the Internet in service to public good. He is responsible for identifying new technologies and enabling policies that support the development of advanced, Internet-based products and services.

Enhancing the World Wide Web

Cerf left DARPA in 1982 to join Microwave Communications Inc. (now part of WorldCom), headquartered in Washington, D.C., as vice president of its digital information services division. A year later, he led the development of MCI Mail, the first commercial email service on the Internet.

In 1986 he left the company to become vice president of the newly formed Corporation for National Research Initiatives, also in Reston. He worked alongside Kahn at the not-for-profit organization, developing digital libraries, gigabit speed networks, and knowledge robots (mobile software agents used in computer networks).

He returned to MCI in 1994 and served as a senior vice president for 11 years before joining Google.

Together with Kahn, Cerf founded the nonprofit Internet Society in 1992. The organization helps set technical standards, develops Internet infrastructure, and helps lawmakers set policy.

Cerf served as its president from 1992 to 1995 and was chairman of the board of the Internet Corp. for Assigned Names and Numbers from 2000 to 2007. ICANN works to ensure a stable, secure, and interoperable Internet by managing the assignment of unique IP addresses and domain names. It also maintains tables of registered parameters needed for the protocol standards developed by the Internet Engineering Task Force.

Cerf has received several recognitions for his work, including the 2004 Turing Award from the Association for Computing Machinery. The honor is known as the Nobel Prize of computing. Together with Kahn, he was awarded a 2013 Queen Elizabeth Prize for Engineering, a 2005 U.S. Presidential Medal of Freedom, and a 1997 U.S. National Medal of Technology and Innovation. ■

NEWS

Microwave Society Has a New Name

BY KATHY PRETZ

Antennas like these in the Atacama Large Millimeter Array telescope are some of the technologies that members of the IEEE Microwave Theory and Technology Society work on.

THE IEEE SOCIETY focused on advancing microwave theory and its applications—including RF, microwave, millimeter-wave, and terahertz technologies—has changed its name to the IEEE Microwave Theory and Technology Society. The one-word change from *techniques* to *technology* took effect in January.

The new name is more appealing to younger engineers, who had a difficult time relating to the term *microwave techniques*, according to the society. Many members working in industry reported that they considered the term to be too academic and disconnected from the technology.

In the early 1940s, the word *microwave* was rarely seen in print because most microwave engineers worked on classified projects for the U.S. Department of Defense. But by the early 1950s several technical societies including the American Institute of Electrical Engineers and the Institute of Radio Engineers—IEEE's predecessor societies—began publishing articles on the technology. The IRE started holding technical talks on microwaves at its yearly convention.

The IRE approved the formation of the Professional Group for Microwave Electronics in March 1952. It was later renamed the Professional Group on Microwave Theory and Techniques, and in 1974 the name was changed to the IEEE Microwave Theory and Techniques Society.

The society held its first symposium in 1952. Today MTT-S technically or financially sponsors more than 50 conferences, workshops, and other events each year.

In 1953 the society launched the *IEEE Transactions on Microwave Theory and Techniques*. The publication became a quarterly, and in 1966 it became a monthly. The society's publications include the *IEEE Journal of Microwaves*, the fifth-most-downloaded journal in the IEEE Xplore Digital Library, according to the society's website.

Its Women in Microwaves group is working to increase diversity in the microwave field. It aims to attract more female graduate students and professionals to join. The group holds events that help enhance the professional and personal growth of female members and raise their visibility.

The society, which celebrated its 70th anniversary last year, has more than 11,000 members in its 190 chapters worldwide. ■

New Board to Oversee IEEE's Humanitarian Activities

BY JOANNA GOODRICH

THE IEEE HUMANITARIAN Activities Committee (HAC) was elevated to the IEEE Humanitarian Technology Board (HTB), effective 1 January. The IEEE Board of Directors approved the change in November.

The new board will oversee all humanitarian activities across IEEE's groups, regions, and societies and is responsible for fostering new collaborations. It also will fund related projects and activities.

The change provides a systematic process for creating programs and managing existing ones including EPICS in IEEE, IEEE MOVE, the IEEE Power & Energy Society Scholarship Plus Initiative, IEEE Smart Village, and the IEEE Special Interest Group on Humanitarian Technology.

HTB will focus on raising awareness of how technically trained people can contribute to society and provide guidance on how to do that. The board also will continue HAC's work building a sustainable development community within and beyond IEEE by partnering with other organizations.

HTB will continue to host HAC's annual Global Summit. The virtual event connects IEEE members from around the world who are using or want to use their technical expertise to address sustainable development challenges.

"Bringing together IEEE groups to collaborate on humanitarian projects and working with outside organizations will help make IEEE a world leader in sustainable development," says IEEE Senior Member Sampathkumar Veeraraghavan, the 2021–2022 HAC chair. "By doing so, IEEE can have a more graduated and aggregated global impact." ■

RECOGNITION OF RECORD GROWTH AND IMPACT

HAC was honored with the 2022 American Society of Association Executives' Power of Associations Summit Award in the global development category. It was the first time IEEE received the award for its humanitarian activities.

JOHN MOORE/GETTY IMAGES

PROFILE

His Wheelchair Tech Makes the World More Accessible

Rory Cooper develops customized controls and wheelchairs for rough terrain

BY JOANNA GOODRICH

Cooper [seated] develops tools to help people with disabilities get around with more ease and comfort.

FOR MORE THAN 25 years, Rory Cooper has been developing technology to improve the lives of people with disabilities. His inventions have helped countless wheelchair users get around with more ease and comfort.

Technologies that the IEEE life Fellow has developed include the SmartWheel and the VCJ-CA, a variable-compliance joystick with compensation algorithms. The SmartWheel attaches to a manual wheelchair to measure the force of pushes, push frequency, stroke length, smoothness, and speed of both the push and the wheelchair. Wheelchair athletes use the data to optimize their performance. It is also helpful in determining adjustments to minimize stress injuries for more typical users.

The VCJ-CA lets users customize the driving controls of electric-powered wheelchairs and is used today in just about every such chair.

These days, Cooper and his team at the University of Pittsburgh's Human Engineering Research Laboratories are working to develop advancements including a wheelchair that can travel on rough terrain. Cooper founded the HERL in collaboration with the U.S. Department of Veterans Affairs.

For those and other "extensive contributions to wheelchair technology that have expanded mobility and reduced secondary injuries for millions of people with disabilities," Cooper received the 2022 IEEE Biomedical Engineering Award.

Employer Human Engineering Research Laboratories at the University of Pittsburgh
Title Director
Member grade Life Fellow
Alma mater California Polytechnic State University, in San Luis Obispo

He says that although technology, medicine, and society have evolved significantly in the way they can help people with disabilities, "there's still a lot of opportunity for technology to further improve people's lives and health." As HERL director and a professor of bioengineering, physical medicine, rehabilitation, and orthopedic surgery at the University of Pittsburgh, he plans to develop more helpful tools.

Changing the course of his career

Cooper was left paralyzed from the waist down due to a bicycle accident

UNIVERSITY OF PITTSBURGH

in 1980 that happened while he was in the U.S. Army. He left soon after and earned a bachelor's degree in electrical engineering from California Polytechnic State University, in San Luis Obispo. He went on to receive a master's degree from Cal Poly in the same subject, taking classes while working as an instrumentation and control engineer at Pacific Gas and Electric in Diablo Canyon, Calif.

During his graduate studies he took a biomedical engineering class and fell in love with the field, he says. He also had started teaching apprentices at PG&E the basics of control systems and electronics—which provided another type of inspiration.

Educating the apprentices "was a great thing for me and perhaps a mistake for PG&E because I found that I really enjoyed teaching," Cooper says, laughing.

Thinking he'd rather teach than pursue an industry career, he headed to the University of California, Santa Barbara, for a Ph.D. There he began developing a device that came to be called the SmartWheel. The mechanical instrument has a complex set of sensors integrated with a single-board computer with wireless communication and is mounted onto wheelchairs.

"I started to develop the technology because I wanted to try to win a medal in the Paralympics," Cooper says.

The SmartWheel measures the forces and torques applied by athletes to the push rim (the part on the chair individuals use to turn the wheels). An encoder measures the wheel's speed and orientation. Athletes can use the data to optimize their performance by adjusting their body position, customizing the design of their chair, and positioning and orienting their wheels with respect to their shoulders.

It worked for Cooper: He received a bronze Paralympic medal in wheelchair racing in 1988.

Cooper [second from the left] and his colleagues—David Constantine, Andrin Vuthaj, and Jorge Candiotti[from left to right]—at the University of Pittsburgh's Human Engineering Research Laboratories working on the MeBot.

ABIGAIL ALBRIGHT

But he hadn't perfected the device when, in 1989, he joined California State University in Sacramento as a faculty member. That same year he joined the Department of Veterans Affairs as a postdoctoral researcher.

After five years, he left Cal State but continued working part time at the VA. In 1994 he joined the University of Pittsburgh as a professor, establishing the HERL that year to develop and enhance technology that promotes people's mobility, function, and inclusion. One of those projects was continuing development of the SmartWheel, which became commercially available in 2000.

Making wheelchairs inclusive

Another technology Cooper is proud of is the variable-compliance joystick with compensation algorithms. Before the VCJ-CA was invented, the controls of electric-powered wheelchairs were analog, making it difficult to customize the controls, he says. If the user had even the slightest tremor or tic, the wheelchair could move unintentionally.

"There were a lot of people who were reliant on others to push their wheelchair or to operate its controls for them," Cooper says. "But these wheelchair users wanted independent mobility, so I began studying how to make this possible."

The joystick's hardware and software can be customized to fit each user's needs. For example, individuals with restricted hand or arm movement can tailor the stiffness of the joystick according to their reach, strength, and control. The algorithms allow individuals to customize their wheelchair's speed, braking, and turning capabilities. The algorithms also can adapt to a user's tremor, range of motion, ability to generate motion, and ability to control the direction of their arm, hand, or finger.

"The VCJ-CA is now used in almost every electric-powered wheelchair in the world—which is pretty cool," Cooper says.

"The VCJ-CA is now used in almost every electric-powered wheelchair in the world—which is pretty cool."

Bringing stability and safety to users

The most common cause of emergency-room visits by wheelchair users is falling from the chair or tipping over, Cooper says.

"This often happens when the individual's wheelchair hits thresholds in doorways, drives off small curbs, or transitions from a sidewalk to a ramp," he says.

Since 2013, he and his team have been working on the Mobility Enhancement Robotic Wheelchair to minimize such injuries.

Known as the MEBot, the wheelchair can climb curbs up to 20 centimeters high and can self-level as it drives over uneven terrain. It does so thanks to six wheels that move up and down plus two sets of smaller omnidirectional wheels in the front and back. The wheelchair's larger, powered wheels can reposition themselves to simulate front-, mid-, or rear-wheel drive.

User trials were completed in 2021, and Cooper predicts the MEBot will become available within the next five years. ■

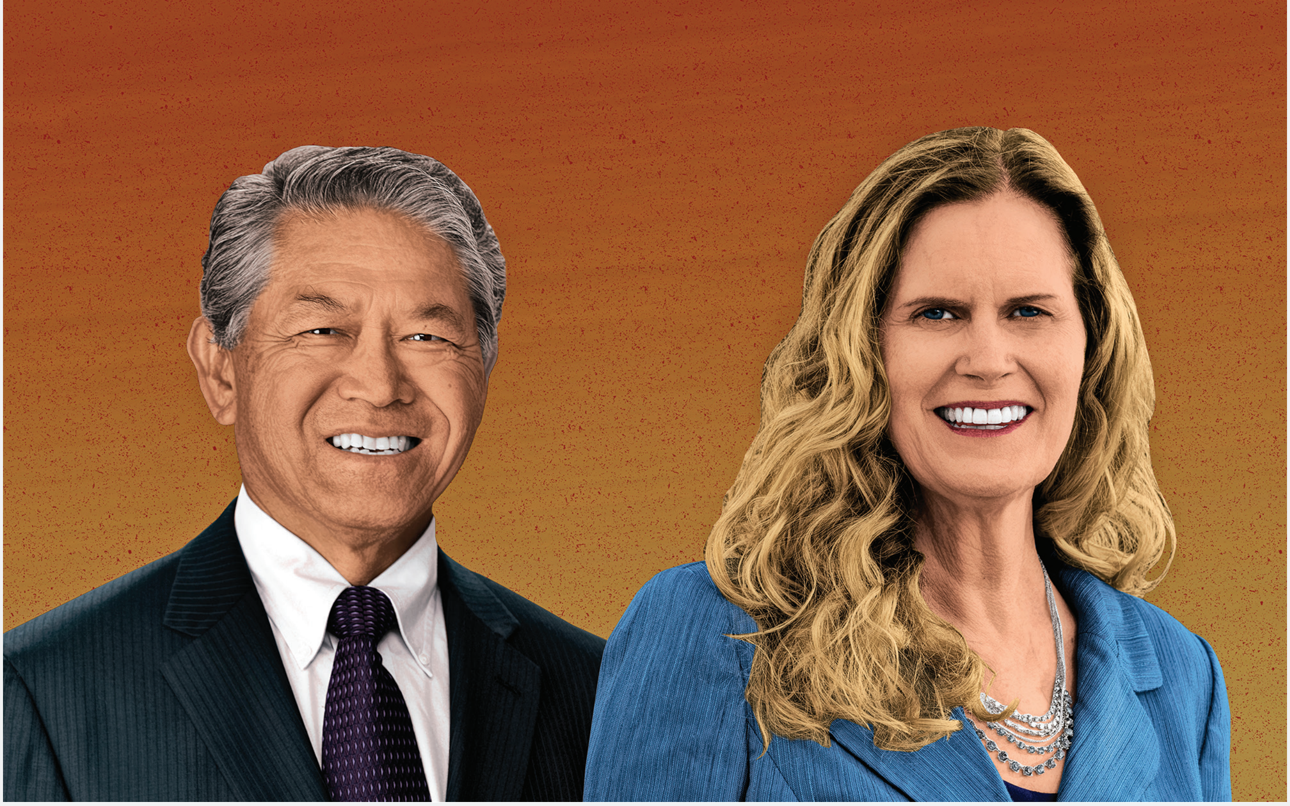

NEWS

Meet the Members Running for 2024 President-Elect

BY JOANNA GOODRICH

THE IEEE BOARD of Directors has nominated Life Fellow Roger Fujii and Senior Member Kathleen Kramer as candidates for IEEE president-elect.

The winner of this year's election will serve as IEEE president in 2025. For more information about the election, president-elect candidates, and petition process, visit the IEEE election website (ieee.org/election).

Roger Fujii

LIFE FELLOW

FUJII IS PRESIDENT of Fujii Systems of Rancho Palos Verdes, Calif., which designs critical systems. Before starting his company, Fujii was vice president at Northrop Grumman's engineering division in San Diego (US $1.086 billion).

His area of expertise is certifying critical systems. He has been a guest lecturer at California State University, the University of California, and Xiamen University.

An active IEEE volunteer, Fujii most recently chaired the IEEE financial transparency reporting committee and the IEEE ad hoc committee on IEEE in 2050. The ad hoc committee envisioned scenarios to gain a global perspective of what the world might look like in 2050 and beyond and what potential futures might mean for IEEE.

He was 2016 president of the IEEE Computer Society, 2021 vice president of the IEEE Technical Activities Board, and 2012–2014 director of Division VIII.

Fujii received the 2020 IEEE Richard E. Merwin Award, the IEEE Computer Society's highest-level volunteer service award.

Kathleen Kramer

SENIOR MEMBER

KRAMER IS A PROFESSOR of electrical engineering at the University of San Diego, where she served as chair of the EE department and director of engineering from 2004 to 2013. As director she provided academic leadership for engineering programs and developed new programs.

Her areas of interest include multisensor data fusion, intelligent systems, and cybersecurity in aerospace systems.

Kramer has worked for companies including Bell Communications Research, Hewlett-Packard, and Viasat.

She is a distinguished lecturer for the IEEE Aerospace and Electronic Systems Society and leads the society's technical panel on cybersecurity.

Kramer is chair of the IEEE ad hoc committee on innovating funding models. She was also the 2019–2021 IEEE secretary and the 2017–2018 director of IEEE Region 6.

Kramer has contributed to several advances for cybersecurity and robotics in ABET, the global accrediting organization for academic programs in applied science, computing, engineering, and technology.

Photo-illustration by Gluekit

Tech Predictions for 2023 and Beyond

TO UNDERSTAND KEY technology trends for this year and beyond, IEEE conducted a global survey of 350 CIOs, CTOs, IT directors, and other technology leaders. Here's a look at what they said.

98% of respondents said AI-powered autonomous collaborative software and mobile robots will automate processes and tasks, allowing workers to be more efficient.

Which industry sector will be the most impacted by technology in 2023?

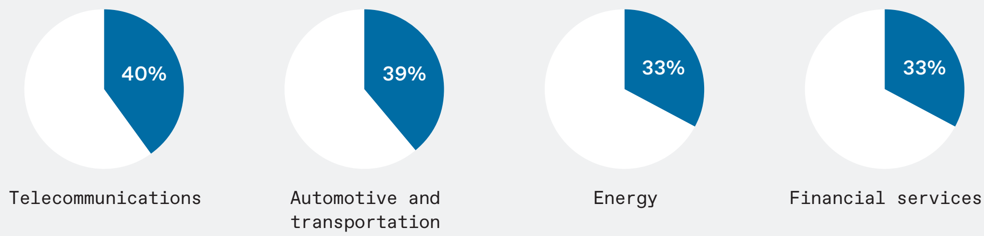

What innovations will advance the development of the metaverse?

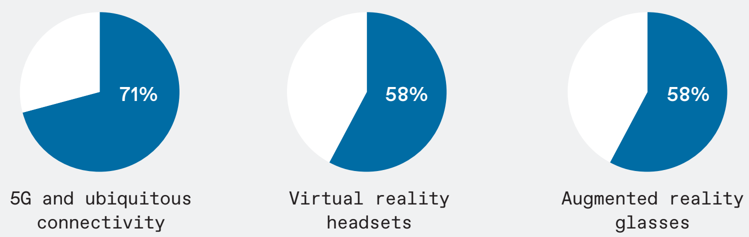

What areas will benefit the most from 5G?

Remote learning and education	56%
Telemedicine	51%
Entertainment, sports, and live-event streaming	51%
Personal and professional daily communication	49%
Transportation and traffic control	29%
Manufacturing and assembly	25%
Carbon footprint reduction and energy efficiency	23%

SOURCE: *THE IMPACT OF TECHNOLOGY IN 2023 AND BEYOND: AN IEEE GLOBAL STUDY*

What will be the most important technology in 2023?

Cloud computing	40%
5G	38%
The metaverse	37%
Electric vehicles	35%
Industrial Internet of Things	33%

95% of respondents said incorporating technologies that would help their organization become more sustainable is a top priority.

HISTORY

The Birth of Random-Access Memory

Britain's Manchester "Baby" was the first electronic digital computer to store a program

BY JOANNA GOODRICH

WHETHER YOU'RE STREAMING a movie on Netflix, playing a video game, or just looking at digital photos, your computer is regularly dipping into its memory for instructions. Without random-access memory, a computer today can't even boot up.

Over the years, memory has been made up of vacuum tubes, glass tubes filled with mercury and, most recently, semiconductors.

But the first computers didn't have any reprogrammable memory at all. Until the late 1940s, every time a machine needed to change tasks, it had to be physically reprogrammed and rewired, according to the Science and Industry Museum in Manchester, England.

The first electronic digital computer capable of storing instructions and data in a read/write memory was the Manchester Small Scale Experimental Machine, known as the Manchester "Baby." It successfully ran a program from memory in June 1948.

Computing pioneers Frederic C. Williams, Tom Kilburn, and Geoff Tootill developed and built the machine and its storage system—the *Williams-Kilburn tube*—at the University of Manchester.

"The Baby was very limited in what it could do, but it was the first-ever real-life demonstration of electronic stored-program computing, the fast and flexible approach used in nearly all computers today," James Sumner, a lecturer on the history of technology at the University of Manchester, said in an interview with the *Manchester Evening News*.

How Baby came to remember

After World War II, research groups around the world began investigating ways to build computers that could perform multiple tasks from memory. One such researcher was British engineer F.C. Williams, a radar pioneer who worked at the Telecommunications Research Establishment (TRE), in Malvern, England.

Williams had an impressive background in radar systems and electronics research. He helped develop the "identification: friend or foe" system, which used radar pulses to distinguish Allied aircraft during the war.

Because of his expertise, in 1945 the TRE tasked Williams with editing and contributing content to a series of books on radar techniques. As part of his research, he traveled to Bell Labs in Murray Hill, N.J., to learn about work being done to remove ground echoes from the radar traces on CRTs. Williams came up with the idea of using two CRTs and storing the radar trace by passing it back and forth between the two.

> The computer was composed of metal racks, hundreds of valves and vacuum tubes, and a panel of vertically mounted hand-operated switches.

Williams returned to the TRE and began to investigate the idea, realizing that the approach also could be used to store digital data with just one CRT. Kilburn, a scientific officer at the TRE, joined Williams in his research.

A CRT uses an electron gun to send a focused beam of electrons toward a phosphor-laden screen. The phosphors glow where the beam strikes; the glow eventually fades until struck again by the electron beam.

To store digital data, Williams and Kilburn used a more powerful electron beam. When it hit the screen, it knocked a few electrons aside, briefly creating a positively charged spot surrounded by a negative halo. Reading the data involved writing to each data spot on that plate and decoding the pattern of current generated in a nearby metal plate—which would depend on whether there was something written at that spot previously.

It turned out that the electron charges leaked away over time (just as phosphors on a TV screen fade) and didn't allow the tube to keep storing data, according to an entry about the Baby on the Engineering and Technology History Wiki. To maintain the charge, the electron beam had to repeatedly read the data stored on the phosphor and regenerate the associated charge pattern. Such refreshing is also used in the DRAM present in today's computers.

In 1946 the men demonstrated a device that could store 1 bit. It is now called the *Williams-Kilburn tube*; sometimes just the *Williams tube*.

Also in 1946, Williams joined the University of Manchester as chair of its electrotechnology department. The TRE temporarily assigned Kilburn to work with him there, and the two continued their research at

Weighing almost a tonne, the Manchester “Baby” computer was 5 meters long and 2 meters tall.

SCIENCE MUSEUM GROUP

the university’s Computing Machine Laboratory. A year later Williams recruited computer scientist Tootill to join the team. And in 1947 they successfully stored 2,048 bits using a Williams-Kilburn tube.

Building the prototype

To test the reliability of the Williams-Kilburn tube, in 1948 Tilburn and Tootill, with guidance from Computing Machine Laboratory founder Max Newman and computer scientist Alan Turing, built a small-scale experimental machine. It took them six months, using surplus parts from WWII-era code-breaking machines. And the Manchester Baby was born.

The Baby took up an entire room in the laboratory building. It was 5 meters long, 2 meters tall, and weighed almost a tonne. The computer consisted of metal racks, hundreds of valves and vacuum tubes, and a panel of vertically mounted hand-operated switches. Users entered programs into memory, bit by bit via the switches, and read the output directly off the face of the Williams-Kilburn tube.

On 21 June 1948, Baby ran its first program. Written by Kilburn to find the highest factor of an integer, it consisted of 17 instructions. The machine ran through 3.5 million calculations in 53 minutes before getting the correct answer.

By 1953, 17 pioneering computer design groups worldwide had adopted the Williams-Kilburn RAM technology.

The IEEE commemorated Baby as an IEEE Milestone in June. The IEEE United Kingdom and Ireland Section sponsored the nomination. Administered by the IEEE History Center and supported by donors, the Milestone program recognizes outstanding technical developments around the world. ■

> “The Baby was the first-ever real-life demonstration of electronic stored-program computing, and the fast and flexible approach is used in nearly all computers today.”

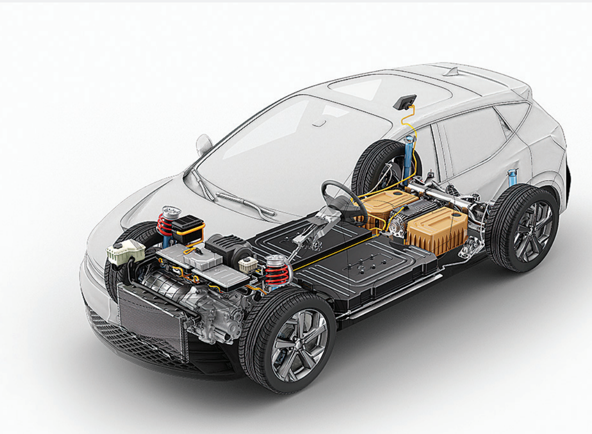

IEEE STANDARDS

Four Standards Found Under the Hood

EVERY YEAR CAR manufacturers release new and redesigned models including electric vehicles. You might not realize it, but IEEE standards can be found inside every new car.

These are just a few that help make vehicles safer and more reliable.

- **Wi-Fi**
  IEEE 802.11
- **Image Quality of Vehicle Camera**
  IEEE 1858-2016
- **Rechargeable Battery**
  IEEE 1725
- **DC Quick Charger**
  IEEE 2030.1.1

SOURCE: IEEE STANDARDS ASSOCIATION

IEEE PRODUCTS

Two Universities Sign Open-Access Agreements

BY KATHY PRETZ

THE UNIVERSITY OF California and the Conference of Italian University Rectors (a consortium of state and non-state universities known as the CRUI) each recently signed what's known as a read-and-publish agreement with IEEE. The agreement lets the institutions' researchers access documents behind the paywall in the IEEE Xplore Digital Library, like traditional subscribers.

The agreement also allows them to publish their open-access articles in IEEE periodicals without paying out of pocket.

IEEE open-access articles are supported by article-processing charges instead of subscriptions. In such arrangements, the APCs are paid for by the organizations where the researchers work or study.

IEEE publishes more than 200 leading journals and magazines, and the digital library contains more than 5 million documents including scientific journals, conference proceedings, and standards.

The four-year agreement with the University of California lets researchers from its 10-campus system publish open-access articles in IEEE's fully open gold and traditional hybrid journals. *Gold open access* refers to periodicals in which all content is freely available once published. These publications are normally supported by APCs. The university's libraries will cover the APCs, according to a news release about the agreement.

The CRUI signed an unlimited read-and-publish open-access agreement with IEEE. The three-year deal gives researchers from the consortium's 54 institutions the right to publish an unrestricted number of open-access articles in IEEE's hybrid journals and fully open-access journals, according to the news release. The cost of accessing IEEE Xplore content and the APCs will be covered by license fees paid by consortium members.

There is a provision in the agreement that publishing articles in IEEE open-access journals will be done with a Creative Commons attribution license unless otherwise requested by the author, the release said. The license lets others distribute, modify, and add to the article as long as the author is credited for the original version.

Visit open.ieee.org to learn more about the IEEE open-access options for authors and institutions and to see the list of more than 300 institutions that already have an agreement. ■

NEWS

Notice to Membership

THE IEEE ETHICS and Member Conduct Committee (EMCC) received a complaint against Dr. Tao Li, a member with the grade of Fellow of the IEEE. Following an EMCC investigation, a hearing board appointed by the IEEE Board of Directors found cause that Dr. Li had violated Section II, Subsections 7 and 8 of the IEEE Code of Ethics. The IEEE Board of Directors sustained these findings and imposed the sanction of Expulsion from IEEE Membership on Dr. Li, in accordance with IEEE Bylaw I-110.5. The IEEE Board of Directors also determined that this notification to the IEEE membership should be made.

ISTOCK

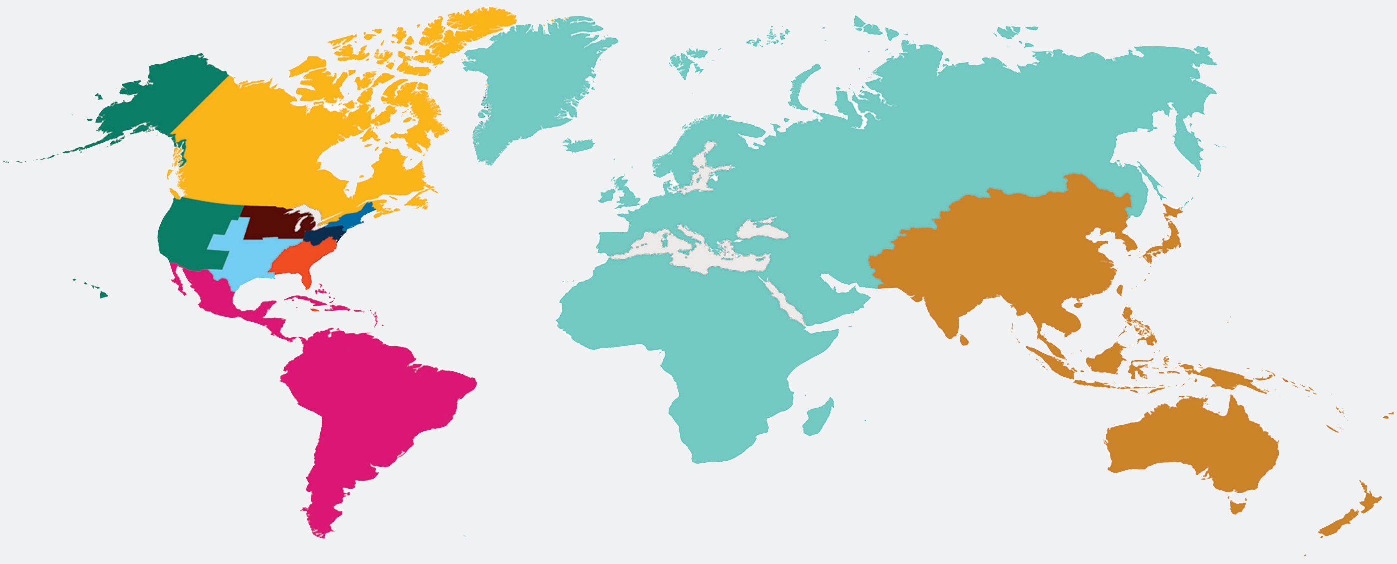

NEWS

IEEE is Realigning Its Geographic Regions

A MAJOR STEP in restructuring IEEE's regions was taken during the November 2022 IEEE Board of Directors meeting, in which the Board approved the proposed region realignment as outlined on page 55 of the September 2022 edition of *The Institute*.

IEEE has been working to restructure its 10 geographic regions to provide a more equitable representation across its global membership.

Some of the factors that were considered when evaluating the realignment were membership counts, geographic locations, time zones, and providing the best overall experience for members.

During the IEEE Board meeting, the 2022 vice president of Member and Geographic Activities, David Koehler, presented MGA's progress and the proposed region realignment.

The plan's approval is the first of many steps. It allows for the consolidation of the six U.S.-based regions into five, joining together current IEEE Region 1 (Northeastern U.S.) and Region 2 (Eastern U.S.) and the splitting of current IEEE Region 10 (Asia and Pacific) into two regions.

Those regional changes will be effective in January 2028.

In addition, the Board formally approved the concept of zones and zone representatives. A zone is a substructure within a region with a significant number of members. In Region 8 (Europe, Middle East, and Africa) and Region 10, four zones in total were approved in November. Two zones were formed in IEEE Region 8 and two in Region 10, which became effective in January 2023.

In those larger regions, zone representatives can assist in the region and provide an additional voice for members within the region as well as represent them on the MGA board.

MGA is continuing its work on the planned region realignment to ensure all the appropriate steps are taken to make the transition. ■

Among the factors considered when evaluating the realignment were geographic location and the best overall experience for members.

Founder Diana Gamzina holds Elve's millimeter-wave power amplifier system, which uses 100 watt E-band traveling-wave tubes.

STARTUP

How Elve Slashed the Cost of Power Amps for Spacecraft

Additive manufacturing cuts production cost by 90 percent

BY KATHY PRETZ

DIANA GAMZINA IS on a mission to drastically reduce the price of millimeter-wave power amplifiers. The vacuum-electronics devices are used for communication with distant space probes and for other applications that need the highest data rates available.

An amplifier can cost as much as US $1 million because it's made using costly, high-precision manufacturing and manual assembly. Gamzina's startup, Elve, is using advanced materials and new manufacturing technologies to lower the unit price.

It can take up to a year to produce one amplifier using conventional manufacturing processes, but Elve is already making about one per week, Gamzina says. Elve's process enables sales at about 10 percent of the usual price, making large-volume markets more accessible.

Launched in June 2020, the startup produces affordable systems for wireless connections that deliver optical fiber quality, or what Gamzina calls *elvespeed connectivity*. The company's name, she says, refers to atmospheric emission of light and very low frequency perturbations due to electromagnetic pulse sources. Elves can be seen as a flat ring glowing in Earth's upper atmosphere. They appear for just a few milliseconds and can grow to be as much as 320 kilometers wide.

Based in Davis, Calif., Elve employs 12 people as well as a handful of consultants and advisors.

For her work with amplifiers, Gamzina, an IEEE senior member, was recognized with the 2022 Vacuum Electronics Young Scientist Award from the IEEE Electron Devices Society. She received the award at last year's IEEE International Vacuum Electronics Conference, in Monterey, Calif.

"Dr. Gamzina's innovation and contributions to the industry are remarkable," IEEE Member Jack Tucek, general chair of the conference, said in a news release about the award.

Expanding the market

In addition to running her company, Gamzina works as a staff scientist at the SLAC National Accelerator Laboratory, in Menlo Park, Calif.—a U.S. Department of Energy national lab operated by Stanford. It was in this role that she began speaking to industry representatives about how to expand the market for high-performance millimeter-wave amplifiers.

From those discussions, she found that lowering the price was key.

"Customers are paying between $250,000 and $1 million for each individual device," she says. "When you hear these numbers, you realize why people don't think this price is anywhere close to affordable for growing the market."

Founded 2020
Headquarters Davis, Calif.
Employees 12

Elve's millimeter-wave power amplifier system weighs 4 kilograms, measures 23 centimeters by 15 cm by 8 cm, and is powered by a 28-volt

Photography by Gabriela Hasbun

DC input bus. The heart of the amplifier system is a set of traveling-wave tubes. TWTs are a subset of vacuum electronics that can amplify electromagnetic signals by more than a hundredfold over a wide bandwidth of frequencies. They're commonly used for communications and radar imaging applications, Gamzina says.

The TWTs take a millimeter-wave signal and make it "interact with a high-energy electron beam so that the signal steals the energy from the electron beam. That's how it gets amplified," Gamzina says. "In some ways, it's a simple concept, but it requires excellence in RF design, manufacturing, vacuum science, electron emission, and thermal management to make it all work."

Elve's millimeter-wave amplifiers enable communication with satellite networks and create long-distance ground-to-ground links.

"This allows enormous amounts of data to be sent, leading to higher data rates that are comparable to fiber or laser-type communication devices," Gamzina says. "Another advantage is that the amplifier can operate in most inclement weather."

The initial market for Elve's amplifiers is the communications field, she says, especially for backhaul: the connections between base stations and the core network infrastructure.

"Our ultimate goal is to be part of the terrestrial market such that the amplifiers are installed on cellphone towers," Gamzina says, "enabling high-data-rate communication in remote and rural locations all over the world."

The company plans to develop millimeter-wave amplifiers for imaging and radar applications as well.

> It can take up to a year to produce one amplifier using conventional manufacturing processes, but Elve is making about one per week.

Learning to run a startup

Gamzina's work at Elve is related to research she is conducting at the SLAC lab as well as previous work she did as a development engineer with the millimeter-wave research group at the University of California, Davis. There she studied vacuum devices that operate at terahertz frequencies for use as power sources in particle accelerators, broadcast transmitters, industrial heating, radar systems, and communication satellites.

She also led research programs in additive manufacturing techniques. She holds three U.S. patents related to the manufacture of vacuum-electronics devices.

To learn how to launch a startup, she took a class at Stanford on the subject and she attended a 10-week training course for founders offered by 4thly, a global startup accelerator. She also says she learned a few things as a youngster while helping out at her father's hydraulic equipment manufacturing company.

Raising funds to keep one's business afloat is difficult for many startups, but it wasn't the case for Elve, Gamzina says.

"We have good market traction, a good product, and a good team," she says. "There's been a lot of interest in what we're doing."

What has been a concern, she says, is that the company is growing so rapidly, it might be difficult to scale up production.

"Making hundreds or thousands of these has been a challenge to even comprehend," she says. "But we are ready to go from the state we're in to the next big jump." ■

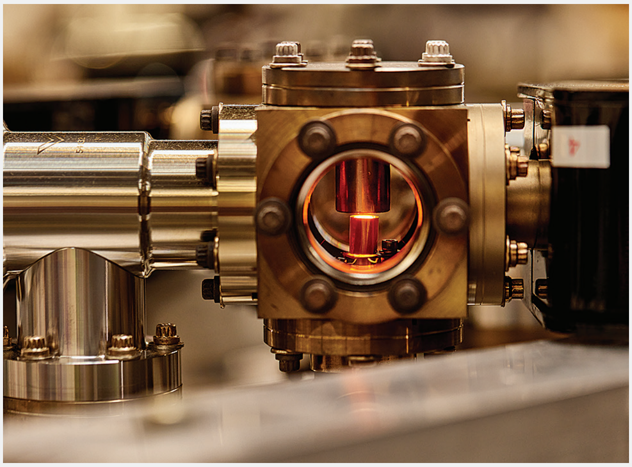

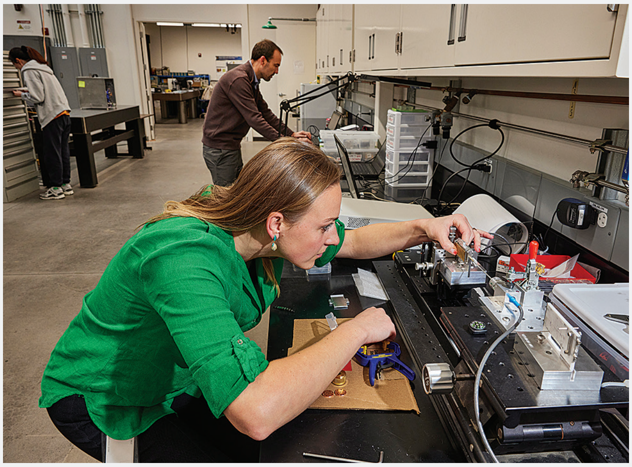

Elve's nano-composite scandate tungsten cathode [left] glowing at 1,000 °C. Production manager Arnela Sinevod and CTO Rich Kowalczyk perform magnetic measurements at the company's facilities [right].

IEEE ELECTION

Countdown to the IEEE Annual Election

ON 1 MAY the IEEE Board of Directors is scheduled to announce the candidates to be placed on this year's ballot for the annual election of officers—which begins on 15 August.

The ballot includes IEEE president-elect candidates and other officer positions up for election.

The Board of Directors has nominated IEEE Life Fellow Roger U. Fujii and IEEE Senior Member Kathleen A. Kramer as candidates for 2024 IEEE president-elect. Visit the IEEE elections page (ieee.org/election) to learn about the candidates.

The ballot also includes nominees for delegate-elect/director-elect openings submitted by division and region nominating committees, IEEE Technical Activities vice president-elect, IEEE-USA president-elect, IEEE Standards Association president-elect, IEEE Women in Engineering committee chair-elect, and IEEE Standards Association board of governors members-at-large.

IEEE members who want to run for an office but who have not been nominated need to submit their petition intention to the IEEE Board of Directors by 15 April. Petitions should be sent to the IEEE Corporate Governance staff: elections@ieee.org.

Those elected take office on 1 January 2024.

To ensure voting eligibility, members are encouraged to review and update their contact information (ieee.org/go/my_account) and communication preferences (ieee.org/election-preferences) by 30 June.

Given ever-changing global conditions, members might wish to vote electronically instead of by mail.

For more information about the offices up for election, the process of getting on the ballot, and deadlines, visit the IEEE elections page or write to elections@ieee.org.

2022 Election Results

HERE IS THE IEEE Tellers Committee tally of votes counted in the 2022 annual election and approved in November by the IEEE Board of Directors.

Proposed constitutional amendment
For: 24,248
Against: 6,542

IEEE president-elect, 2023
Thomas M. Coughlin: 10,908
Kathleen A. Kramer: 10,769
Kazuhiro Kosuge: 8,682
Maike Luiken: 4,365

IEEE division delegate-elect/director-elect, 2023

Division I
Yong Lian: 1,641
Samar K. Saha: 1,443

Division III
Stefano Bregni: 1,591
Vincent W.S. Chan: 1,248
Wahab Almuhtadi: 563

Division V
Christina M. Schober: 2,489
Dejan S. Milojicic: 2,150

Division VII
Christopher E. Root: 2,323
Lalit K. Goel: 1,342

Division IX
Aylin Yener: 1,524
Ahmed Tewfik: 1,207
Walter "Walt" D. Downing: 1,061

IEEE region delegate-elect/director-elect, 2023—2024

Region 2
Felicia Harlow: 1,533
Rhonda L. Farrell: 1,009

Region 4
Constance "Connie" A. Kelly: 943
Tarek Lahdhiri: 547
Hamid Vakilzadian: 334

Region 6
Joseph C. Wei: 2,926
Gora Datta: 1,415

Region 8
Michael G. Hinchey: 4,829
Adeel Sultan: 2,085

Region 10
Takako Hashimoto: 5,804
Ling Chuen "Michael" Ong: 2,437

IEEE Standards Association board of governors member-at-large, 2023—2024
David T. Chen: 744
Kishik Park: 618

IEEE Standards Association board of governors member-at-large, 2023—2024
Dorothy V. Stanley: 1,108
Sha Wei: 288

IEEE Technical Activities vice president-elect, 2023
Manfred "Fred" J. Schindler: 11,893
Rakesh Kumar: 9,684

IEEE-USA president-elect, 2023
Keith A. Moore: 9,886
James R. Look: 6,456

The University of Bonn calls for applications for a

Hertz Chair at the interface Mathematics, Modelling and Simulation of Complex Systems

(W3, full professorship)

starting as soon as possible.

Establishing distinguished Hertz Chairs at the interface of outstanding research areas at the University of Bonn is a central measure following the success of the University of Bonn within the German Excellence Strategy, a funding program of the Federal and State Governments to support cutting-edge research at universities. Under this scheme, the University of Bonn establishes an exceptionally endowed Hertz Chair within the Transdisciplinary Research Area (TRA) “Mathematics, Modelling and Simulation of Complex Systems”.

The Hertz Chair will serve to further develop the profile of the TRA by contributing to the research fields and methods in the TRA’s areas of scientific excellence and potential. Ideal candidates for this Hertz Chair will have a strong track record in mathematics, computer science or computational sciences, and a proven interest in application fields such as life science, economics or geoscience. The holder is expected to make significant contributions regarding future Clusters of Excellence.

Beyond outstanding scientific achievements within the TRA’s profile, we are looking for a candidate who is experienced in acquiring third party funding, demonstrates leadership potential and organizational skills and is interested in participating in interdisciplinary teaching programs.

The Hertz Chair will be integrated into a vibrant and active research environment. It will be positioned at the interface between highly visible and established research areas in Bonn. These include outstanding expertise in the fields of mathematics, machine learning, robotics, life sciences and quantitative economics with three Clusters of Excellence (HCM – Hausdorff Center for Mathematics, ImmunoSensation2 – The Immune Sensory System, PhenoRob – Robotics and Phenotyping for Sustainable Crop Production), as well as affiliated extra-university research centers (Max Planck Institute for Mathematics, Fraunhofer Institute for Algorithms and Scientific Computing SCAI, Fraunhofer Institute for Intelligent Analysis and Information Systems IAIS). The holder of this Hertz Chair will have access to excellent resources and research structures.

Formal requirements are regulated by § 36 of the Higher Education Act of North Rhine-Westphalia (Hochschulgesetz des Landes Nordrhein-Westfalen).

The University of Bonn is committed to diversity and equal opportunities. It is certified as a family- friendly university and offers a dual career service. Increasing the proportion of women among scientific staff and promoting women’s careers in particular remain priorities of the University. We therefore strongly encourage women with relevant qualifications to apply. The University also encourages suitable candidates with severe disabilities to apply. Applications are processed in accordance with the Equal Opportunities Act of North Rhine-Westphalia (Landesgleichstellungsgesetz).

Please find further information on the Transdisciplinary Research Area “Mathematics, Modelling and Simulation of Complex Systems” here: www.tra1.uni-bonn.de

For further information on this Hertz Chair, please contact the spokespersons of the TRA, Prof. Dr. Heiko Röglin (roeglin@cs.uni-bonn.de) or Prof. Dr. Anton Bovier (bovier@uni-bonn.de).

Applications received by **March 15, 2023** are guaranteed full consideration. The appointment committee reserves its right to also consider applications submitted after this date.

Applicants must submit the usual documents (cover letter, CV including teaching experience, research interests and research plan, copies of university degrees and certificates, list of publications) in English via our online portal: https://berufungsportal.uni-bonn.de

The Global Institute of Future Technology (GIFT hereafter) at Shanghai Jiao Tong University sincerely invites applications for tenure-track or tenured professoriate and research scientist positions, mainly at Associate and Full levels, related to Sustainable Energy.

GIFT manages endowments from a variety of sources, including a recent donation of more than 200 million US dollars. GIFT’s new building will be completed within two years, adding 50,000 square meters to the current space. With strong government and university support, we will rapidly expand our faculty, including high-level hiring (e.g. NAE/NAS members and Fellows of multiple professional societies). At present, five research centers have been established: Large-scale Energy Storage and Reliability Center, Solid-state Battery Center, Advanced Sustainable Power Systems Center, Green Energy and Modern Agriculture Center, and Perovskite Photovoltaics Center.

Applicants must hold a doctorate degree in a relevant field. Candidates are expected to establish a vigorous research program and contribute to undergraduate and graduate education. For full consideration, please send a CV, teaching plan and research plan, copies of three representative publications and contact information of five referees as a single PDF file to the GIFT Search Committee (email: gift-facultysearch@sjtu.edu.cn). For more information, please visit **http://gift.sjtu.edu.cn**.

Intensive Wireless Communications Course Series

Register for all four courses in one transaction and use promo code IWCS2023 to save 25% on the registration fee.

Fundamentals of Wireless Communication
22-23 March 2023

Network and Service Architecture
29-30 March 2023

Cellular Networks
12-13 April 2023

Non-Cellular Wireless Systems
19 April 2023

INTENSIVE WIRELESS COMMUNICATIONS SERIES 2023 IEEE ComSoc

Successfully complete all courses in this **Intensive Wireless Communications** course series and earn a digital badge.

Interested in purchasing this training for your employees? Contact iwc-mktg@ieee.org for more details.

BY DAVID C. BROCK

Past Forward

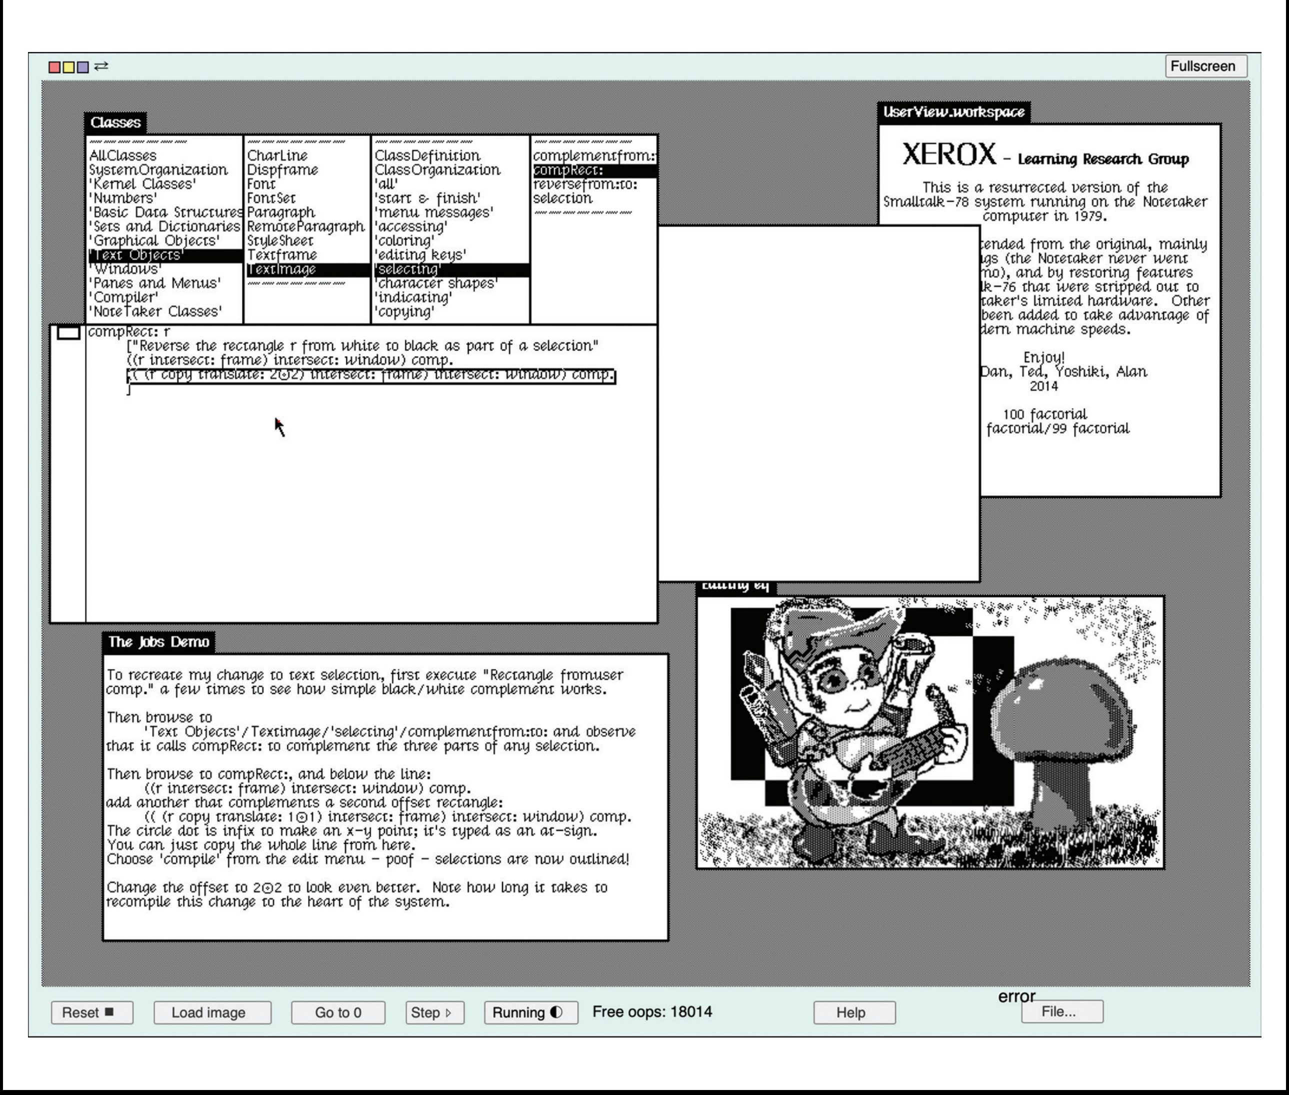

Steve Jobs's Object Lesson

Late in 1979, Steve Jobs and engineers from Apple visited the Xerox Palo Alto Research Center (PARC). There they were introduced to the experimental Alto computer and the Smalltalk language and computing environment, developed by Alan Kay's Learning Research Group. Smalltalk was a breakthrough in object-oriented programming, an approach that's used in many of the most popular programming languages today. During his visit, Jobs was taken with Smalltalk's graphical user interface, and it reshaped his subsequent approach at Apple. This screen capture from the version of Smalltalk demoed for Jobs shows a critical moment that impressed him. Dan Ingalls, one of the developers of Smalltalk, was able to change the user interface live, switching the appearance of selected text from white text on a black background to a bounding rectangular box (as in the upper left window).

Smalltalk-78 had a graphical user interface that greatly impressed Steve Jobs when he saw it demoed in 1979.

More recently, Ingalls created the Smalltalk Zoo, a collection of Smalltalk emulators that you can run in your browser to experience for yourself Jobs's aha moment. ■

FOR MORE ON THE HISTORY OF SMALLTALK AND THE ALTO, SEE "THE MACHINE THAT TRANSFORMED PERSONAL COMPUTING," IN THIS ISSUE.

HANSEN HSU/COMPUTER HISTORY MUSEUM


LIVE
LIFE ON
YOUR
TERMS.
INSURANCE
THAT FITS
YOU.
Group Term Life Insurance
engineered for you.

Semantic segmentation for wildlife conservation.