3D Print a Mechanical TV › A new spin on an old television display
P. 16

The Road to Electrification › Energy transitions don't happen overnight
P. 20

Power Beaming Gets Practical › Nikola Tesla's dream may finally come true
P. 36

FOR THE TECHNOLOGY INSIDER
JUNE 2022

IEEE Spectrum

Quantum Sensors are Revolutionizing Brain Scans

IEEE

VOLUME 59 / ISSUE 6 JUNE 2022

Contents

ON THE COVER:
Photo by Daniel Stier

SSPL/GETTY IMAGES

Photo by Graphcore

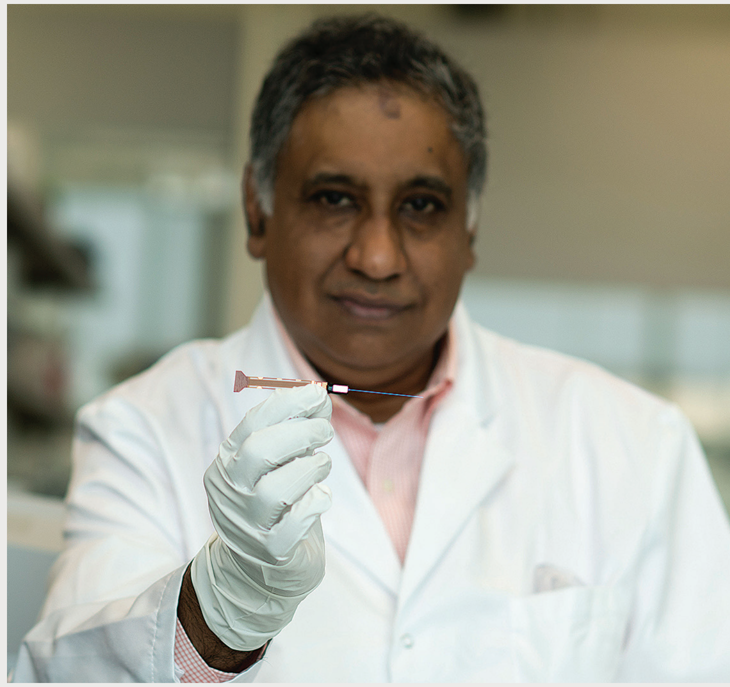

FRED LOOSEN/IMEC

BACK STORY

A Fab and a Vision

An unusual union of big tech and big science created Neuropixels, a fundamentally new technology for observing brains in action [“Eavesdropping on the Brain,” p. 30]. The alliance also created a shared global facility for collaborative neuroscience akin to the CERN particle accelerator for high-energy physics.

The partnership began with conversations between Barun Dutta [above] and Timothy D. Harris in 2010. Dutta, chief scientist at Imec, a leading nonprofit nanoelectronics R&D institute in Belgium, had use of state-of-the-art semiconductor manufacturing facilities. Harris, who directs the applied physics and instrumentation group at the Howard Hughes Medical Institute’s Janelia Research Campus, in Virginia, had connections with world-class neuroscientists who shared his vision of building a new kind of probe for neuron-level observation of living brains.

Dutta and Harris recruited allies from their institutions and beyond. They raised US $10 million from HHMI, the Allen Institute for Brain Science, the Gatsby Foundation, and the Wellcome Trust to fund the intensive R&D needed to produce a working prototype. Dutta led the effort at Imec to tap into semiconductor technology that had been inaccessible to the neuroscience community. Imec successfully delivered the first generation of Neuropixels probes to some 650 labs globally, and the second generation is due to be released in 2022.

In addition to a vibrant open-source community that sprang up to develop software to analyze the large data sets generated by these brain probes, the Allen Institute has created OpenScope, a shared brain observatory to which researchers around the world can propose experiments to test hypotheses about brain function. “Think of us as being the Intel of neuroscience,” Dutta says. “We’re providing the chips, and then labs and companies and open-source software groups around the world are building code and doing experiments with them.” ■

PAUL JAFFE

Jaffe works in the spacecraft engineering department of the U.S. Naval Research Laboratory in Washington, D.C. For many years he's been doing research on power beaming, a topic he surveys in this issue [see "Spooky Power at a Distance," p. 36]. He's particularly interested in the prospects for beaming solar energy harvested in space down to places where energy is needed on Earth. "Clean, constant, global, unlimited energy access from space could revolutionize civilization," says Jaffe.

LOUIS ROSENBERG

Rosenberg is CEO and chief scientist of Unanimous AI, a company developing AI algorithms modeled on biological swarms. In the early 1990s, he created the first interactive augmented-reality system, which he describes on page 42. Rosenberg went on to found several companies including the early virtual-reality firm Immersion. He now lives with his family on a 40-acre animal sanctuary along with hundreds of chickens, ducks, geese, turkeys, cows, pigs, goats, and sheep (real sheep, not virtual).

PETER FAIRLEY

Fairley, a contributing editor, has been tracking energy technologies and their environmental implications for *IEEE Spectrum* since 2002. In this issue [p. 8] he updates some of his reporting from 2015, when he wrote about the process of separating post-Soviet states' electricity grids from Russia's and linking them to the rest of Europe. "What surprised me was how quickly European grid operators moved to rethink the technical and market requirements they'd set for Ukraine's synchronization," he says.

PRACHI PATEL

Spectrum contributing editor Patel is a freelance science and technology writer based in Pittsburgh. In this issue [p. 12], she reports on harnessing AI to help speed the growth of algae, with the potential to bring down biofuel costs while aiding carbon capture. "I've written about AI and biofuels before, so it was exciting to see how researchers are using the former to boost the latter," she says.

IEEE Spectrum

ACTING EDITOR IN CHIEF Harry Goldstein, h.goldstein@ieee.org

FEATURES EDITOR Jean Kumagai, j.kumagai@ieee.org

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

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

PRODUCT MANAGER, DIGITAL Erico Guizzo, e.guizzo@ieee.org

EDITORIAL DIRECTOR, CONTENT DEVELOPMENT
Glenn Zorpette, g.zorpette@ieee.org

SENIOR EDITORS
Evan Ackerman (Digital), ackerman.e@ieee.org
Stephen Cass (Special Projects), cass.s@ieee.org
Samuel K. Moore, s.k.moore@ieee.org
Tekla S. Perry, t.perry@ieee.org
Philip E. Ross, p.ross@ieee.org
David Schneider, d.a.schneider@ieee.org
Eliza Strickland, e.strickland@ieee.org

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

NEWS MANAGER Mark Anderson, m.k.anderson@ieee.org

ASSOCIATE EDITORS
Willie D. Jones (Digital), w.jones@ieee.org
Michael Koziol, m.koziol@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

ADMINISTRATIVE ASSISTANT Ramona L. Foster, r.foster@ieee.org

CONTRIBUTING EDITORS Robert N. Charette, Steven Cherry, Charles Q. Choi, Peter Fairley, Edd Gent, W. Wayt Gibbs, Mark Harris, Allison Marsh, Prachi Patel, Julianne Pepitone, Lawrence Ulrich, Emily Waltz

THE INSTITUTE
EDITOR IN CHIEF Kathy Pretz, k.pretz@ieee.org
ASSISTANT EDITOR Joanna Goodrich, j.goodrich@ieee.org

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
MULTIMEDIA PRODUCTION SPECIALIST Michael Spector
ASSOCIATE ART DIRECTOR, PUBLICATIONS Gail A. Schnitzer

ADVERTISING PRODUCTION +1 732 562 6334
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org
SENIOR ADVERTISING PRODUCTION COORDINATOR
Nicole Evans Gyimah, n.gyimah@ieee.org

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

EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM*
Susan Hassler, *Chair*; Ella M. Atkins, Robert N. Charette, Francis J. Doyle III, Matthew Eisler, Shahin Farshchi, Alissa Fitzgerald, Jonathan Garibaldi, Benjamin Gross, Lawrence O. Hall, Jason K. Hui, Leah H. Jamieson, Mary Lou Jepsen, Michel M. Maharbiz, Somdeb Majumdar, Lisa May, Carmen S. Menoni, Ramune Nagisetty, Paul Nielsen, Sofia Olhede, Christopher Stiller, Wen Tong, Boon-Lock Yeo

EDITORIAL ADVISORY BOARD, *THE INSTITUTE*
Kathy Pretz, *Chair*; Qusi Alqarqaz, Stamatis Dragoumanos, Jonathan Garibaldi, Madeleine Glick, Lawrence O. Hall, Susan Hassler, Francesca Iacopi, Cecilia Metra, Mirela Sechi Annoni Notare, Shashi Raj Pandey, John Purvis, Chenyang Xu

MANAGING DIRECTOR, PUBLICATIONS Steven Heffner

DIRECTOR, BUSINESS DEVELOPMENT, MEDIA & ADVERTISING Mark David, m.david@ieee.org

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

ADVERTISING INQUIRIES Naylor Association Solutions,
Erik Albin +1 352 333 3371, ealbin@naylor.com

REPRINT SALES +1 212 221 9595, ext. 319

REPRINT PERMISSION / LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Managing Editor, *IEEE Spectrum.*

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2022 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 59, Issue No. 6. 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 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit https://www.ieee.org/about/corporate/governance/p9-26.html.

News

SENSORS

A Guide to the Quantum-Sensor Boom > Atomic scale bolsters sensing revolutions in medicine, tech, and engineering

BY CHARLES Q. CHOI

Imagine sensors that can detect the magnetic fields of thoughts, track movements at GPS accuracy without need for GPS, or detect minute quantities of a virus or other pathogen in seconds with no need for elaborate PCR assays. Just as quantum computers can theoretically find the answers to problems no classical computer could ever solve, so too can an emerging generation of quantum sensors lead to new levels of sensitivity, new kinds of applications, and new opportunities to advance a range of fields, technologies, and scientific pursuits.

Quantum technology relies on quantum effects that can arise because the universe can become a fuzzy place at its very smallest levels. For example, the quantum effect known as superposition allows atoms and other building blocks of the cosmos to essentially exist in two or more places at the same time, while another quantum effect known as entanglement can link particles so they can influence each other instantly regardless of how far apart they are.

These quantum effects are infamously fragile to outside interference. Although

Cerca Magnetics wearable magnetoencephalography (MEG) helmets use tiny lasers and even smaller clouds of rubidium atoms to sense small perturbations in magnetic signals coming from a patient's brain.

Photo by Daniel Stier

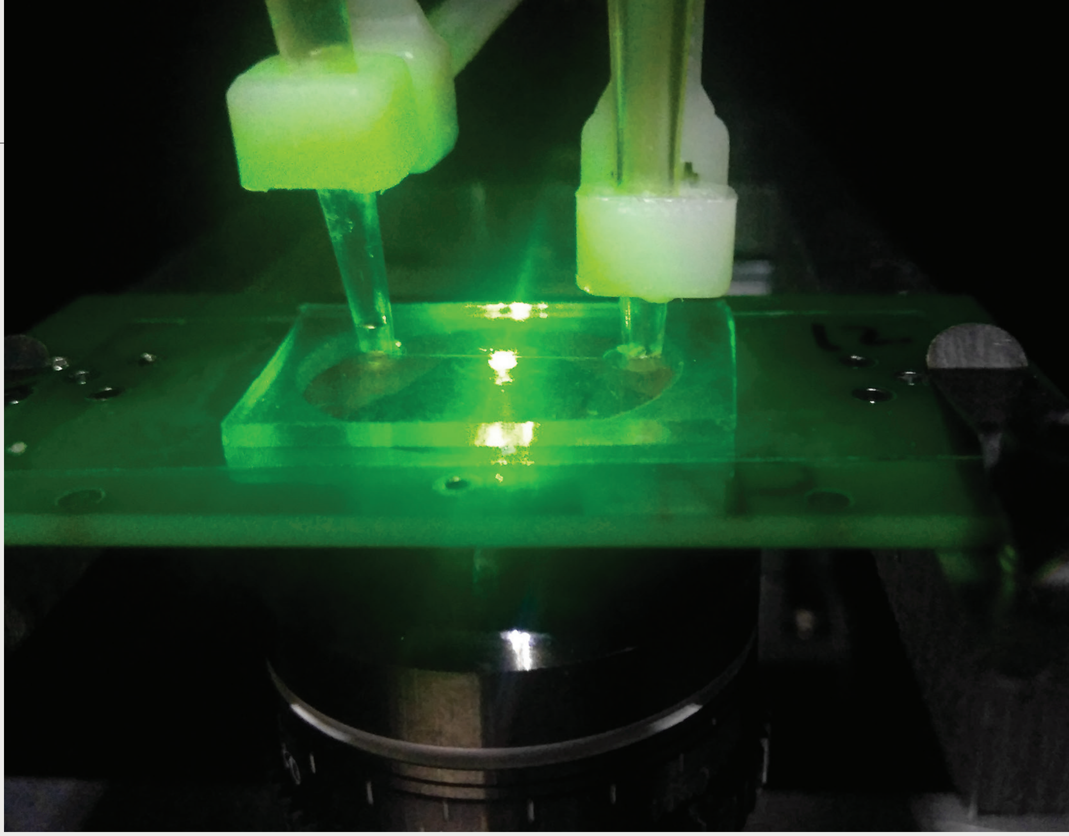

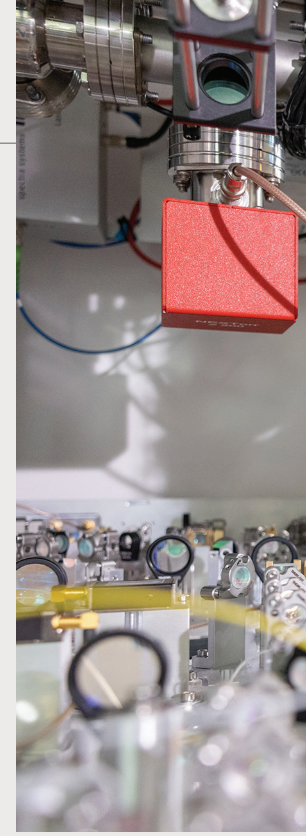

The quantum sensor for the presence of the SARS-CoV-2 virus [above] uses only low-cost materials. The devices could be scaled up, according to the researchers, to analyze a full batch of samples at once. A quantum sensor from Imperial College London and the Glasgow-based company M Squared [right] can help ships navigate even when GPS is unavailable.

quantum computers strive to overcome this weakness, quantum sensors capitalize on this vulnerability to achieve extraordinary sensitivity to the slightest disturbances in the environment. Below are just a few of the many varieties of quantum sensors being developed and deployed today.

Brain scans: Electric currents within the brain generate magnetic fields that sensors can analyze to noninvasively scan brain activity. Now quantum sensors are enabling a wearable helmet to perform such magnetoencephalography (MEG) scans with unprecedented performance and cost.

Currently MEG scans are performed with sensors known as superconducting quantum interference devices (SQUIDs). These require cooling with expensive liquid helium to -269 °C, making the scanners extremely large. In contrast, a new device from startup Cerca Magnetics, in Nottingham, England, is about the size of a Lego brick.

Each device, called an optically pumped magnetometer (OPM), contains a laser that shines a beam through a cloud of rubidium atoms at a light detector. The beam can make the magnetic fields of the rubidium atoms all line up, rendering the cloud essentially transparent. Tiny magnetic fields, such as those from brain activity, can disturb these atoms, making them capable of absorbing light, which the light detector can sense, and the laser resets the cloud so it can continue responding to magnetic disturbances.

The fact that these quantum sensors work at room temperature makes them much less bulky than SQUIDs. This means they can be placed much closer to a person's head, resulting in a signal at least two times as high and theoretically up to five times as high, for magnetic images with millimeter accuracy and millisecond resolution of surface areas of the brain, says Matthew Brookes, chairman of Cerca and professor of physics at the University of Nottingham.

As the sensors are small and lightweight, they can be mounted in a wearable helmet to let people move freely during scanning, instead of requiring them to remain still for long periods as is currently the case. In addition, the sensors can adapt to different head shapes and sizes, making it possible to scan not just adults but also children and babies. Moreover, the OPM tech costs as little as half the price of a SQUID system, Brookes says.

The Cerca scanner can help probe neurological disorders such as epilepsy, concussions, dementia, and schizophrenia, "helping shed light on many severe and debilitating conditions," he says.

Future research can aim to push these sensors closer to their theoretical limits of sensitivity, permit more freedom of movement, and add virtual reality and machine learning to boost what researchers can do with the scanners on the experimental and analytical fronts, Brookes says.

Detecting COVID: Another promising quantum sensor could lead to faster, cheaper, and more accurate tests for the SARS-CoV-2 virus behind the global pandemic. It relies on microscopic artificial diamonds with defects: A carbon atom is replaced with a nitrogen atom and the adjacent carbon atom is missing. These defects in the crystals behave like a tiny magnet whose alignment is very sensitive to magnetic fields, helping such "nitrogen-vacancy centers" serve as sensors.

The new technique, developed by researchers at MIT and the University

LEFT: CHANGHAO LI/MIT; RIGHT: IMPERIAL COLLEGE LONDON/M SQUARED

perature fluctuations and other factors lead the position estimates of conventional inertial navigation systems to drift within hours if they don't have an outside reference signal. But M Squared's device experiences negligible drift even after days, says Joseph Cotter, a research fellow at Imperial College London's Center for Cold Matter.

"The early adopters of this emerging quantum technology are likely to be those interested in long-range navigation for underwater and surface vehicles," Cotter says. "However, as the technology develops and becomes increasingly compact and lower cost, it will have wider benefits across the transportation industry through deployment on ships, trains, and aircraft."

The researchers have field tests planned for their latest device this summer. Currently the quantum accelerometer "is about the size of two washing machines," Cotter notes. "We're working to get it even more compact."

of Waterloo in Canada, involves coating nitrogen-vacancy-center diamonds roughly 25 nanometers wide with magnetic compounds that detach from the gems after they bond with the specific RNA sequence of the SARS-CoV-2 virus. When these diamonds are lit with green light, they will emit a red glow. The magnetic coating dims this glow; exposing the sensors to the virus can increase this glow.

The current gold-standard test for the SARS-CoV-2 virus takes several hours to create enough copies of the virus's genetic material to detect. Moreover, it cannot quantify the amount of virus present with high accuracy and might have false-negative rates of more than 25 percent. In contrast, computer simulations suggest that the new test can theoretically work in just a second, is sensitive enough to detect just a few hundred strands of the viral RNA, and could have false-negative rates below 1 percent.

The nanodiamonds and the other materials used in the test are cheap. In addition, this new method could be adapted to virtually any virus, including new ones that may emerge, by adjusting the magnetic coating to match the target virus. The MIT/Waterloo team is currently synthesizing and testing the sensors to see how well they actually perform. "We hope to get promising results very soon," says researcher Changhao Li, a quantum engineer at MIT.

Quantum accelerometer: The world now relies heavily on global-navigation satellite systems such as GPS, but the satellite links that enable such positioning, navigation, and timing do not work underground or underwater and are vulnerable to jamming, spoofing, and weather. Now a quantum sensor from Imperial College London and Glasgow-based M Squared can help ships navigate even when GPS is unavailable.

The quantum sensor is a kind of device called an atom interferometer. A little like the brain-scan sensor, it uses laser pulses to drive clouds of supercooled atoms into delicate states of quantum superposition. In this case, each atom's trajectory quantum mechanically interferes with itself, with its peaks and troughs augmenting or suppressing one another.

Analyzing how the phase of its atomic-wave packets shifts can reveal any acceleration or rotation the atoms experienced. With the results, the device can calculate the change in its position with time.

The resulting quantum accelerometer can help serve as the foundation of an inertial navigation system that does not rely on any outside signals. Tem-

Untold limits: Recently, scientists in Austria developed the first programmable quantum sensor, a device capable of an unprecedented level of sensitivity operating near the fundamental limits imposed by the laws of quantum mechanics.

In this work, they programmed a quantum computer to find the best settings for itself with which to measure the states of its components. They found this programmable quantum sensor could optimize itself enough to approach the fundamental sensing limit up to a factor of about 1.45. (The closer a sensor approaches the ultimate sensing limit of 1, the better its performance.) They suggest that programmable quantum sensors could find use in devices such as atomic clocks and global positioning systems, as well as magnetic and inertial sensors.

All in all, says physicist David Awschalom at Argonne National Laboratory, in Illinois, and director of the Q-Next consortium, "quantum sensors are emerging with exquisite precision to cover everything from single proteins all the way to questions in astronomy and cosmology." ■

An extended version of this article appears online as "A Quantum of Sensing—Atomic Scale Bolsters New Sensor Boom."

Smoke rises from a shelled power plant outside the town of Shchastya, near the Ukrainian city of Luhansk, on 22 February 2022.

ARIS MESSINIS/AFP/GETTY IMAGES

ENERGY

Ukraine Secures Its Grid by Plugging Into Europe's ›

The quick move kept Moscow from plunging Kyiv into darkness

BY PETER FAIRLEY

Just a few hours before massed Russian troops and missiles surged over the border into Ukraine with deadly force in February, Ukraine's grid operator opened a series of high-voltage breakers, disconnecting the nation's grid from those of Belarus, as well as Russia and the rest of the giant IPS/UPS synchronous AC power zone controlled from Moscow. It was supposed to be a 72-hour test; going ahead under those tense circumstances was a bold and risky gambit, admits Ukrenergo CEO Volodymyr Kudrytskyi.

"We heard a lot of opinions in the expert community, as well as among politicians, that it is very dangerous to disconnect from Russia and Belarus, that the Ukrainian energy system will not be able to function independently for a long time," recalled Kudrytskyi in an interview published in March by Kyiv-based news outlet *Ukrayinska Pravda*.

Ukraine's grid held, even as Russia damaged substations, transmission lines, and generators—a remarkable and heroic (though largely unheralded) achievement. As recently as November, generation shortfalls had prompted Kyiv's mayor to voice fears of rolling blackouts. Coal and gas reserves remained tight throughout the winter.

Operating as an electrical island left Ukraine more vulnerable to Russia's army, including its cyberwarriors. Like a boxer freed from an opponent's embrace, Russian forces could swing at Ukraine's newly isolated grid without fear of battle-induced blackouts cascading across into Russia.

In early March, Russia's "targeted destruction" of Ukraine's energy infrastructure prompted Ukrenergo to tag its daily social-media updates with the hashtag "EnergyFront."

But through it all, Ukrenergo and its European colleagues were quietly working on an emergency support scheme: uniting Ukraine's grid with those of the European Network of Transmission System Operators (ENTSO-E), thus making the embattled nation part of a continuous zone of synchronized AC power connecting most of continental Europe.

In mid-March, they made it happen, closing breakers to the south and synchronizing Ukraine's power system with the 50-hertz alternating current traversing ENTSO-E's wires. The move simultaneously looped in Moldova.

The result is an electrical backstop to help sustain Ukraine's defenders and its civilian population. Ukraine can now import up to 2 gigawatts of European power over links from Hungary, Poland, Romania, and Slovakia, enabling it to conserve domestic reserves of coal and hydropower.

What's more, if damage to power equipment creates an electrical shock in Ukraine, the rest of the European network—which generates roughly 35 times as much energy as Ukraine—should hold power steady.

The scale of the achievement is hard to overstate, according to Ukrenergo's CEO. "We have guaranteed ourselves an uninterrupted power supply for weeks and months," said Kudrytskyi.

"We are no longer alone," is how President Volodymyr Zelenskyy put it on the day Ukraine became part of the "energy Eurozone."

Synchronization with the European grid marked a hasty culmination of a process that began in earnest in 2017 and was not expected to be completed until 2023. Essentially, ENTSO-E members took extraordinary defensive measures—and took on risks to their own systems—to make up for as-yet-unfinished upgrades to Ukraine's infrastructure, operating procedures, and markets.

"It was almost heroic to come and save Ukraine in this way," says Tomas Jermalavičius, an Estonian security expert who has studied ongoing efforts by the former Soviet republics in the Baltics to break away from the Russian grid.

Jermalavičius says that in addition to helping Ukraine, the situation has helped the Baltic states, also on Russia's far western flank—namely, Estonia, Latvia, and Lithuania. He is one of several experts who say Russia's war on Ukraine and Ukraine's accelerated synchronization with Europe improves prospects for the Baltics' eventual separation from Russian power.

The Baltic states are NATO and European Union members. But they are not scheduled to disconnect from Russia's grid and connect to Europe's until 2025.

Jermalavičius, head of studies for the International Centre for Defence and Security, based in Tallinn, Estonia, has been expressing concerns about the political disunity that has slowed the Baltics' progress. All three countries have sizable Russian-speaking populations, and their leaders face both anti-European and pro-Russian opposition.

In a December 2021 newspaper commentary, Jermalavičius noted that a crucial test of the Baltics' ability to operate as an electrical island was scrubbed in 2019 amid intergovernmental squabbling.

Meanwhile, Russia has moved quickly to secure its own grid, adding transmission lines east of the Baltics and power plants in Kaliningrad, a Russian exclave sandwiched between Poland and Lithuania. That means it could unilaterally unplug the Baltics at any time.

"The sheer possibility of such a decision, and the technical conditions enabling it, should be sending shivers down the spines of the Baltic governments," wrote Jermalavičius, warning of "severe" impacts on security if the entire synchronization project "fell apart in the final stretches."

Fast-forward three months, however, and unity began to flower in the Baltics. On 3 March, Estonia and Latvia cut their power imports from Russia, which Lithuania had done in 2020. "With Russia's attack on Ukraine, internecine political issues are receding into the background. You see the unity of the Baltics in supporting Ukraine, and in understanding Russia's threat in all its dimensions," Jermalavičius told *IEEE Spectrum* in March.

The last big Baltic upgrade being made to prepare the way for synchronization is construction of a second connection between Poland and Lithuania (and thus between Europe and the Baltics). Last year, Polish authorities hinted that the so-called Harmony Link cable might not be ready until 2026 or 2027. Now it looks more likely that they will hit their 2025 deadline.

"Lithuania and the other Baltic states are following the plan, and it is 2025. There are technical requirements that should be completed," says Romas Švedas, a former Lithuanian diplomat and vice minister of energy.

But Jermalavičius thinks they might move faster, citing an indication from a recent meeting of regional energy ministers that acceleration of the Baltics' grid switch was on the table.

He noted that Lithuania's grid operator, Litgrid, has run promising tests without that second Polish connection. Litgrid, he said, demonstrated in December that Lithuania can operate as an isolated grid. That test also showed that the sole operating link between Lithuania and Poland—a high-voltage direct current, or HVDC, line—can deftly switch to AC mode to open a more powerful synchronous connection.

Lithuania's minister of energy, Dainius Kreivys, called it a "geopolitical turning point." ■

Journal Watch

A Smart-Home System That Respects Privacy

By sensing human activity and adjusting the environmental settings accordingly, smart-home systems could help create more energy-efficient and sustainable buildings. But there is an as-yet-unresolved conflict between these systems' need for loads of data to respond to a given environment and the fact that monitoring people's activity stirs up privacy concerns.

A new smart-home system, dubbed Chameleon, is designed to resolve the conflict. The setup, recently tested in two different environments (an office with a handful of employees and a classroom that had 15 to 20 occupants on any given day), indicated the nature and extent of human activity with 87 to 99 percent accuracy after just one week of training. "This is very valuable for scaling these devices to more buildings and use cases," says Andres Rico, a graduate research assistant at the MIT Media Lab's City Science group. Rico was part of the group that conducted the study. He and his colleagues built their system around carbon dioxide and passive infrared (PIR) sensors. The results are described in a study published 6 April in the *IEEE Internet of Things Journal*.

Carbon dioxide and infrared sensing could be useful for characterizing a variety of scenarios, such as the number of people in a room and how heavily they are breathing, which can indicate resting or active states. Each scenario has unique carbon dioxide and infrared signatures that can be analyzed using machine learning.

Those capabilities notwithstanding, says Rico, "Carbon dioxide and PIR data inherently are not intrusive," unlike systems based on cameras or other optical systems.

Rico says the team hopes to incorporate Chameleon into more real-world spaces, helping to create smarter buildings. For example, it could automatically change the HVAC settings for a room that is used for a wide variety of activities, delivering both improved energy efficiency and enhanced creature comfort. —Michelle Hampson

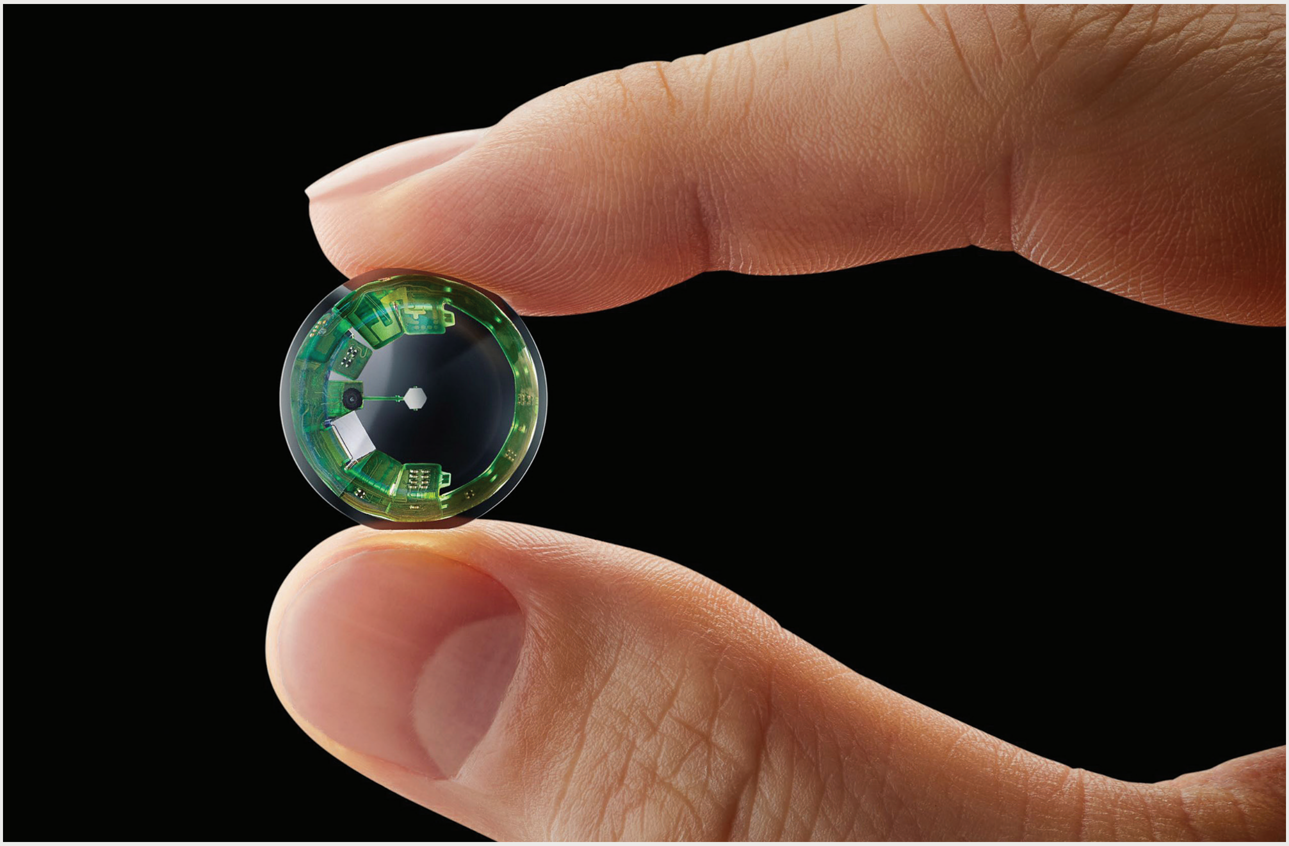

MOJO VISION

BIOMEDICAL

My Peek Through Mojo Vision's AR Contacts >

Mojo packs batteries, motion sensors, and a microLED into its latest lens

BY TEKLA S. PERRY

In March, Mojo Vision, the Saratoga, Calif.–based maker of augmented reality (AR) contact lenses, announced that it has drawn closer to realizing its goal of putting self-contained head-up displays right up against its customers' eyeballs. Though its latest AR contact lens is still just a prototype, with clinical testing and further development ahead before it can apply for the U.S. Food and Drug Administration approval needed to sell to consumers, Mojo's engineers are happy to report that they are steadily ticking off engineering milestones.

In late March, I got to peek through Mojo's newest lens. Here's what I saw, and what Steve Sinclair, Mojo senior vice president of product and marketing, had to say about the company's progress so far; Sinclair was also forthcoming about the challenges that remain.

First, the demo: I did not put the lens in my eye. This prototype is still in safety testing, and fitting a contact lens requires an eye exam. Instead, I held a lens close to one eye and peeked through. I was able to move around freely. But because holding instead of wearing the device means that it cannot track eye movements, Mojo has temporarily incorporated a tiny crosshair into the user interface to help with alignment. The nature of the demo also meant that the images I saw were flat. With a lens in each eye, the images will appear in 3D.

I tried out several applications while peering through the lens. To select them, I looked around the periphery of the real-world view in front of me, which caused icons to appear. By focusing on one (in this case, aligning the crosshairs with it), I selected it. The app I found to be the most fun to use wasn't the most complex, but it did show off the lens's onboard sensors by tagging compass headings as I turned to face different directions. I also played with several other applications, including a travel-tools app that included a high-resolution monochrome image of an incoming Uber driver, a biking app that called up heart rate and other information worth tracking during a training ride, a teleprompter app that naturally scrolled up and down to move through the text, and a monochrome video stream. All these demos took place in the prototype lens, with me fully in control of them. Mojo's team then set up a VR simulation to demonstrate eye-tracking features that normally won't work unless the lens is sitting in your eye.

This is not the first time I've seen the Mojo lens in person. But since that last demo, in 2020, the engineering team has moved from wireless power to batteries on board, increased the resolution of the display from 8,000 pixels per inch to 14,000, thinned its commercial motion sensors, developed its own radio and power management, and created several custom apps. An image sensor—a feature of earlier demos that showed off vision enhancements such as edge detection in low light—has yet to be built into the current prototype, though Sinclair says that it is in the works. What's more, he says, the company is continuing to test apps for people with low vision and still expects the visually impaired to be among the earliest customers for the device.

Here's what else Sinclair had to say about the technology as developed so far—and the path ahead:

Let's talk about the batteries.
Steve Sinclair: The battery is in the outer ring, embedded in the lens. We are partnering with a medical-grade battery company that makes implantable microbatteries for things like pacemakers, to design something safe to wear. Previously, we were using magnetic inductive coupling, and that seemed fine when the user was holding it up to their eye. But the moment you put it in your eye and start moving your head and darting your eyes around, the lens would lose that connection to the wireless field. Because it wasn't as reliable as we needed it to be, we decided about a year and a half ago to switch over to battery power.

MOJO VISION

What's New in the Mojo Lens

- Medical-grade microbatteries on board along with a custom power management IC
- High-res 14,000 pixels per inch monochrome microLED display (up from about 8,000 pixels per inch)
- Custom 5-gigahertz radio using a proprietary low-latency communications protocol
- Eye tracking using accelerometer, gyroscope, and magnetometer

Do I correctly recall that you always wanted to get batteries on board?
S.S.: We thought so. But we sped that path up, because we concluded that the other path was just not a good way to go.

Why doesn't this version have the image sensor?
S.S.: We've decided to leave out the imager, the camera, for right now; it's not critical to the use cases that we're looking at first.

Does the circuitry block vision at all?
S.S.: See the cutout on the side? Imagine I'm wearing it. It's going to be oriented like this [faces the cutout away from the nose] because I need the peripheral vision on the outside, not toward the nasal side. Ultimately, there's even more we can do to push components further toward the edges of the lens to maximize light entering the pupil.

What happens next?
S.S.: Next, we start testing [the complete prototype] on-eye and see just how well it works in different situations. We can't say exactly when that will happen. We hope it's soon, but we've got to make sure it's safe and everything's working the way we expect it to work. That testing will start with Drew Perkins, our CEO, wearing it, then probably Mike Wiemer, our CTO. And then it will go to folks like myself and others on the executive team, along with some of the key engineers who need to start evaluating things like the software and battery life.

What's the timeline for commercialization?
S.S.: Everything is predicated on eventually getting FDA certification. We don't like to presume how long that's going to take.

Can you talk about pricing?
S.S.: Like anything else, when we first bring it out, it'll be a little expensive. Our goal when we're running at volume is that it should come out somewhere close to a high-end smartphone. But factor in the fact that people are already spending [US] $500, $600, or $700 for eyewear today, subtract that out of the total price, and the premium on top of that is not huge. ■

ARTIFICIAL INTELLIGENCE

AI Speeds Algae Biofuel Growth >

New algorithms could help pond scum replace fossil fuels

BY PRACHI PATEL

Algae are terrific at counterbalancing our looming atmospheric carbon issues. Like plants, they absorb large quantities of carbon dioxide from the atmosphere and convert it into sugars via photosynthesis. Plus, the slimy green stuff grows faster than trees do, and can grow on land and in water unsuitable for cultivating food crops. All these are reasons why algae are considered an important tool to address climate change.

That is, until you try to cultivate algae on a large scale. Producing enough algae to be cost effective has proven immensely challenging; limited yield has been a deal breaker, keeping algae biofuels from reaching the market. But now, researchers have tapped into artificial intelligence to grow algae at a record pace. The result could be a dramatic drop in algae fuel costs.

By using machine learning to precisely tune the growth of blue-green algae, also known as cyanobacteria, a team of researchers at Texas A&M University has produced more than 43 grams of algae per square meter per day in an outdoor experimental setup. That's a near doubling of the 25 g/m^2/day algae-biofuel-viability threshold recently identified by the U.S. Department of Energy. It's "absolutely without any question the highest reported outdoor algae growth speed and productivity," says Joshua Yuan, professor and chair for synthetic biology and renewable products at Texas A&M.

His group's AI-aided process has the potential to bring the cost of algal biofuel below US $5 per gallon from today's exorbitant price of $33 per gallon.

That's slightly higher than corn ethanol, which is $2 to $4 per gallon (53 cents to $1.06 per liter) and is "already getting competitive with fossil fuels in California, as we can link to carbon capture to bring carbon credits," says Yuan. "Furthermore, this fuel price is [already] competitive in overseas markets, like China, where fuel prices are much higher."

In early February, the U.S. Department of Energy announced plans to appropriate $19 million to fund projects that can boost the capacity of algal systems to capture carbon dioxide. The hope is that advances in this area can serve two purposes: reduce carbon emissions and yield a bumper crop of algae for biofuels and other products. "Theoretical algae yield

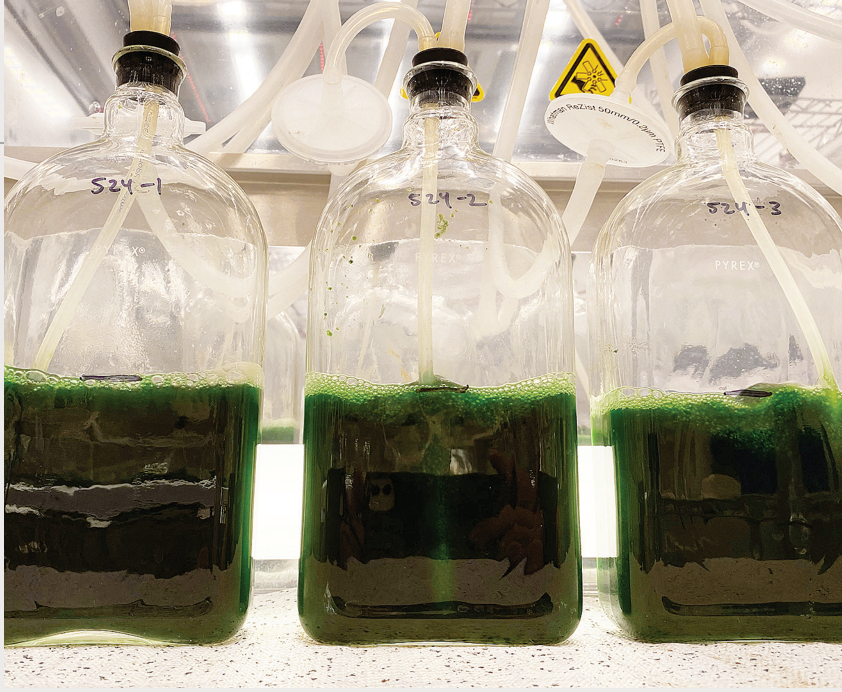

These bottles containing algae are representative of the setup for enhanced growth of the slimy green stuff aided by machine-learning algorithms.

TEXAS A&M UNIVERSITY

is huge, much better than [that of] land-based plants, but there are a few major limitations," Yuan says—namely, growing and harvesting algae.

In both closed reactors and open-pond systems, as algae grows and its concentration increases on the surface, the top layer blocks light from reaching the cells underneath, slowing their growth. Plus, heretofore, harvesting the blooms has been inefficient and costly. Traditional methods of separating algae from water, such as centrifugation, filtration, and chemical flocculation, all ratchet up the required energy input and the total cost of algae production by one-half and one-third, respectively.

The AI-powered, semiautonomous algae-cultivation system that Yuan and his colleagues reported in *Nature Communications* on 27 January overcomes these issues. It comprises two machine-learning models. One predicts, based on the intensity of the light falling on algae in water and the density of algal cells, how light will be scattered and will likely propagate through the layers of cyanobacteria. Unlike previous models, this one predicts light distribution in three dimensions. The second model predicts how this light propagation will affect the algae's growth rate. By combining the two models, Yuan says, it's possible to calculate the highest algae concentration that does not produce too much shade, allowing algae to grow at maximum speed in changing light conditions.

"Once the AI tells us [the target concentration], we control the algae growth at this concentration," he says. As soon as that level of growth is reached, the researchers manually prune part of the algae bloom and add more growth medium. In the future, they plan to make the system completely autonomous by making the pruning step robotic instead of manual. ■

Figure From Fiction

By Willie D. Jones

Many cultures through the ages have mythologized the flying, fire-breathing prowess of the dragon. Now, in the age of advanced technology, an engineer has created his own mechatronic version of the mythical beast. François Delarozière, founder and artistic director of French street-performance company La Machine, is shown riding his and his team's brainchild, called Long Ma. The 72-tonne steel-and-wood automaton can carry 50 people on a covered terrace built into its back and still walk at speeds of up to 4 kilometers per hour. It can flap its leather-and-canvas-covered wings; it will also shoot fire, smoke, or steam from its mouth, nose, eyelids, and more than two dozen other vents located along its 25-meter-long body. Long Ma spends most of its time in China, but the mechanical beast has been transported to France so it can participate in fairs there this summer. It has already been a featured attraction at the Toulouse International Fair, where it thrilled onlookers from 9 to 18 April.

PHOTOGRAPH BY ALAIN PITTON/ NURPHOTO/AP

Hands On

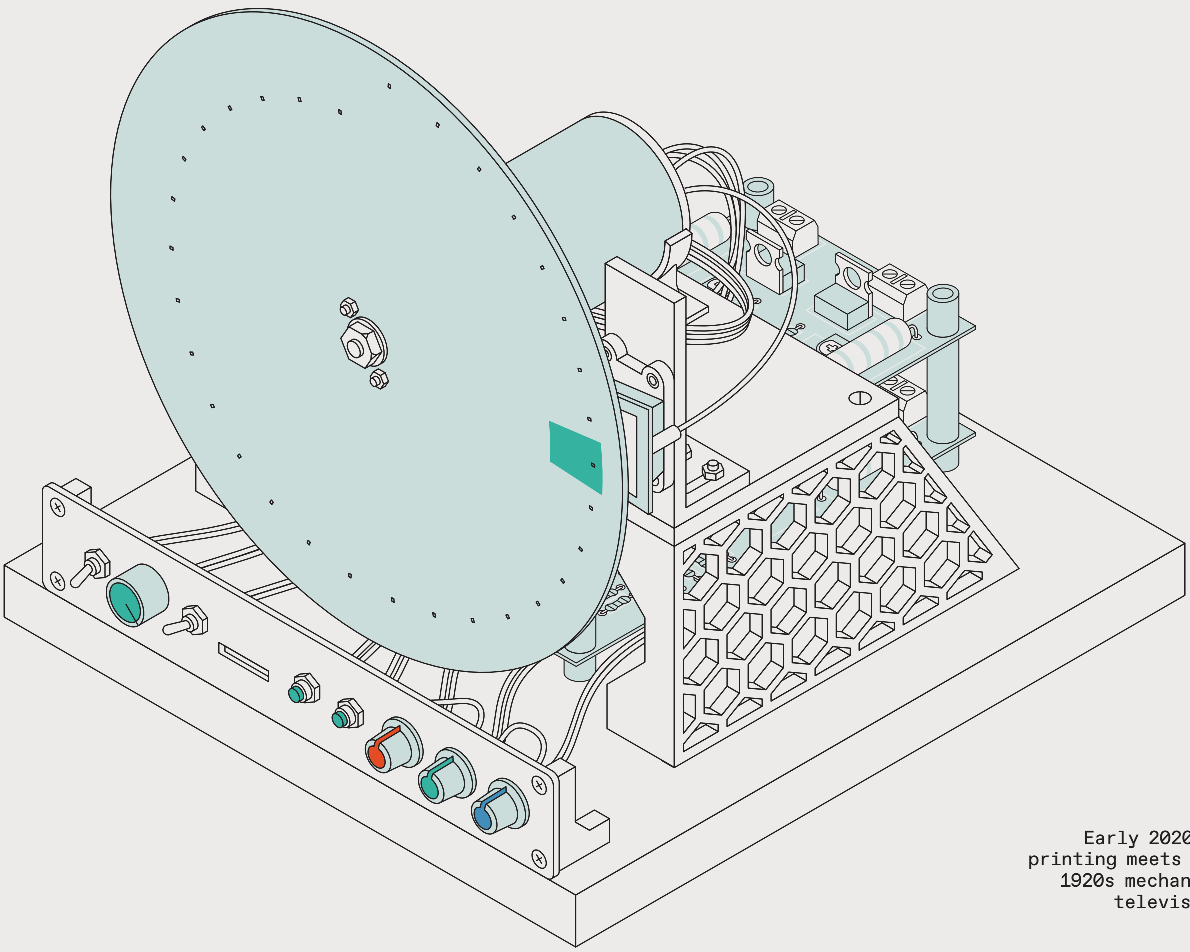

Early 2020s 3D printing meets late 1920s mechanical television.

EZ Mechanical TV ›

A new twist on the oldest type of television display

BY MARKUS MIERSE

Before flat screens, before even cathode-ray tubes, people watched television programs at home thanks to the Nipkow disk. Ninety years ago in places like England and Germany, broadcasters transmitted to commercially produced black-and-white electromechanical television sets, such as the Baird Televisor, which used these disks to produce moving images. This early programming established many of the formats we take for granted today, such as variety shows and outside broadcasts.

The size and weight of a Nipkow disk makes a display with more than a few dozen scan lines impracticable (in stark

Illustrations by James Provost

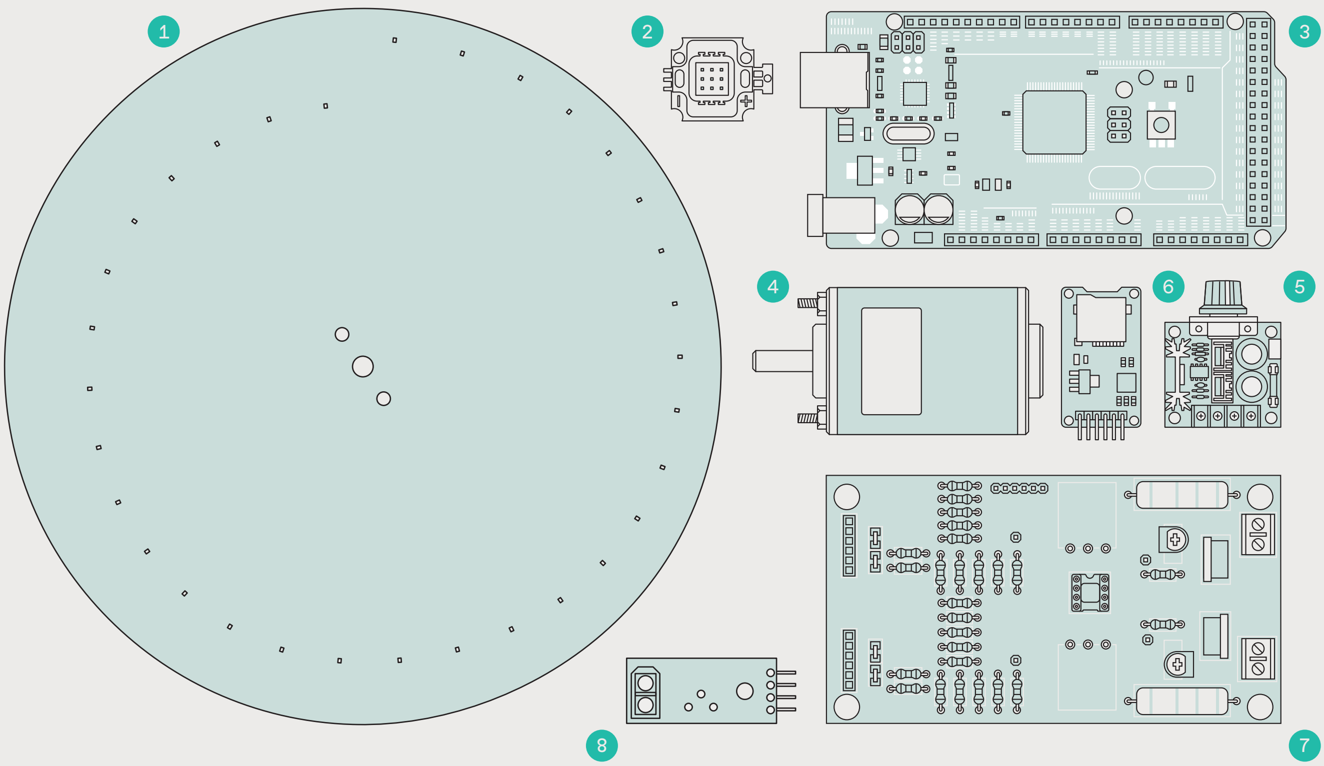

A printed 20-centimeter-diameter disk (1), complete with 32 holes, is illuminated by an RGB LED module (2). An Arduino Mega (3) controls the LEDs, while the motor's (4) speed is adjusted by a potentiometer (5). Images and movies are stored and read from an SD card (6). Digital LED data from the Mega is converted to analog voltage using a custom printed circuit board (7), and the rotation of the disk is monitored with an infrared sensor (8).

contrast to modern screens with thousands of lines). But when a mechanical TV is fed a moving image, the result is surprisingly watchable. And Nipkow displays are fascinating in their simplicity—no high voltages or complex matrices. So I wondered: What was the easiest way to build such a display that would produce a good quality image?

I'd been interested in Nipkow disks since I was a student, trying a few experiments with cardboard disks that didn't really produce anything. In more recent years, I saw that a number of people had built modern Nipkow displays—even incorporating color—but these relied on having access to pricey machine tools and materials. I set about designing an inexpensive version that could be made using a consumer-grade 3D printer.

The secret of a Nipkow disk is in its spiral of holes. A light source behind the disk illuminates a small region. A motor spins the disk, and each hole passes through the lighted region in turn, creating a series of (slightly curved) scan lines. If the illumination is varied in sync with the time it takes each hole to cross the viewing region, you can build up images in the display frame.

The first thing I had to do was figure out the disk. I chose to make the disk 20 centimeters in diameter, as that's a size most home 3D printers can manage. This dictated the resolution of my display, since there's a limit to how small you can produce precisely shaped and positioned holes. I wrote software that allowed me to generate test disks with my Prusa i3 MK3S+ printer, settling on 32 holes for 32 scan lines. The display is a trapezoid, 21.5 millimeters wide and 13.5 mm tall at one end and 18 mm tall at the other. One unexpected benefit of printing a disk with holes, rather than drilling them, was that I could make the holes square, resulting in a much sharper image than is achievable with circular holes.

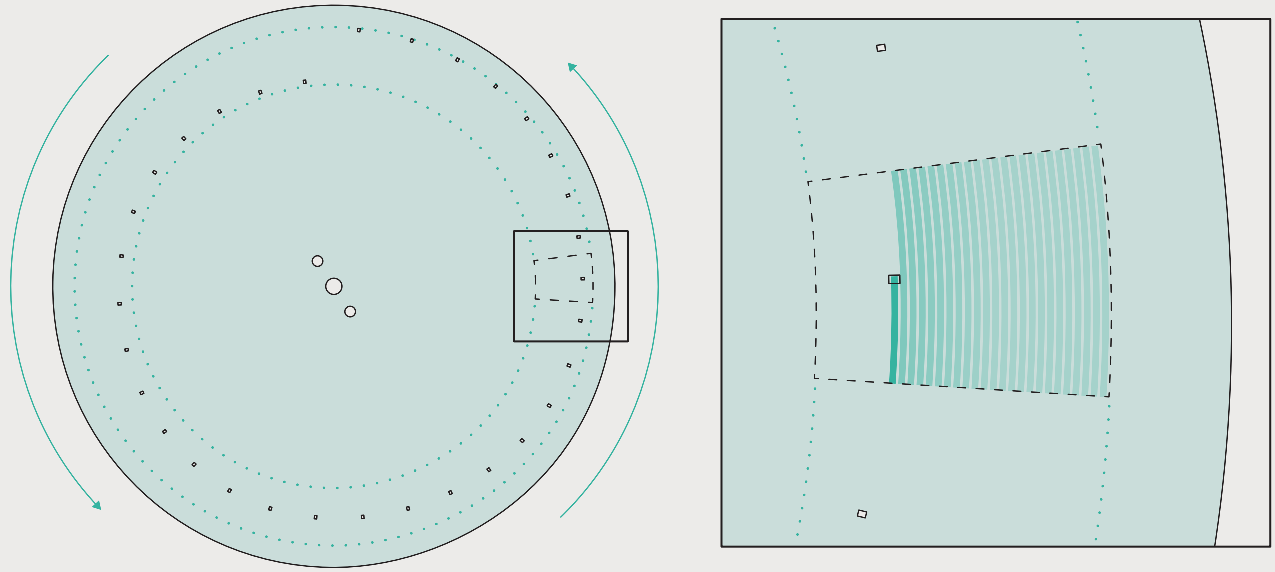

As each hole moves in front of the LED light source, the brightness of the LED is modulated to create a scan line of the image.

For a light source, I used an LED module with red, green, and blue elements placed behind a diffuser. A good picture requires a wide dynamic range of brightness and color, which means driving each element with more power and precision than a microcontroller can typically provide directly. I designed a 6-bit digital-to-analog converter circuit and had custom printed circuit boards made, each with two copies of the circuit. I stacked two PCBs on top of each other so that one copy of my DAC drives one LED color (with a spare circuit left over in case I made any mistakes populating the PCBs with components!). This gives a combined resolution of 18 bits per pixel. Three potentiometers let me adjust the brightness of each channel.

An Arduino Mega microcontroller provides the brains. The Mega has enough RAM to hold screen frames and enough input/output pins to dedicate an entire port to each color. (A port allows you to address up to eight pins simultaneously, using the bits of a single byte to turn each pin on or off.) While this did mean effectively wasting two pins per channel, the Mega has pins to spare, and addressing a port provides a considerable speed advantage over bit-banging to address each pin separately.

The Mega synchronizes its output with the disk's rotation using an infrared sensor triggered by a reflective strip attached to the back of the disk. Thanks to this sensor, I didn't have to worry about precisely controlling the speed of the disk. I used a low-noise 12-volt DC XD-3420 motor, which is easily obtained.

I connected some additional controls to the Mega—a mode button that switches between photos and videos, a play/pause button, and a skip-track button to advance to the next file. Frames printed in 3D hold everything together, mounted on a wooden base.

Because my TV uses 6 bits per channel per pixel instead of the 8 bits used by most modern image formats, I created a conversion tool that is free to download, along with all of the other supporting files for this project, from Hackster.io. Videos are treated as a collection of still images sent to the display at a rate of about 25 to 30 frames per second, depending on the exact speed of the disk. You can convert video into suitable collections using the open-source software VirtualDub, and then pass the results through my converter.

Movies and images are stored on an SD card and read into the Mega's memory using an SD module via its SPI connection. The TV simply scans the top-level directory and begins displaying images—one at a time in the case of photos, and automatically advancing to the next image in video mode. Initially, when I tried to play movies there was noticeable stuttering thanks to dropped frames. I discovered this was due to the standard Arduino SD library—it can handle data transfers at a rate of only 25 kilobytes per second or so, while at 25 frames per second the display is looking for data at a rate of 75 kB/s. The problem was solved by switching to the optimized SdFat library, which provides much faster read access.

The result is a display that's small enough to fit nicely on a desk or shelf, but produces a bright, colorful, and steady image at a frame rate fast enough for most video. All-electronic television may have ultimately triumphed over mechanical sets, but my spinning Nipkow disk is a visceral reminder that powerful forces can spring from simple origins. ■

One unexpected benefit of printing a disk was that I could make the holes square, resulting in a much sharper image.

Careers

Consumer Reports' senior director of product testing, Maria Rerecich, checks the test equipment at the company's labs in Yonkers, N.Y.

ARTUR PIETRUCH

Maria Rerecich › She tests products for *Consumer Reports* to ensure they work as claimed

BY DANIEL P. DERN

For over 80 years, people have turned to *Consumer Reports* for honest and authoritative assessments of products. Today, the responsibility for managing testing and ratings for tech-containing products falls to Maria Rerecich, senior director of product testing at the independent nonprofit member organization.

Rerecich, an IEEE member, joined *Consumer Reports* in 2013 after almost three decades at Standard Microsystems Corp. (acquired in 2012 by Microchip Technology). At SMSC she was responsible for integrated circuit design, validation, and product engineering of silicon chips used in PCs.

There are obvious differences between the two jobs, Rerecich acknowledges. For starters, *Consumer Reports* rates products; it doesn't build them. "We are testing products we don't make and aren't trying to sell," Rerecich says. "My teams are testing to compare performance and check if products work as claimed, but we don't have to try to fix them if they don't."

Her 30-person department specifies tests and develops the scoring and ratings. "We work with a variety of interesting products in all types of categories and sometimes need to develop unexpected tests," she says.

One example was a situation known as Bendgate. "When the iPhone 6 came out in September 2014, there were reports that the larger model was bending when people put it in their rear pockets and sat on it," Rerecich recalls. "Lots of Bendgate videos were posted of people bending phones using their thumbs. While a little flex may be acceptable, you don't want it to deform to the point where it doesn't come back, or to break."

This concern felt like a perfect match for the organization's methodologies and test lab resources, Rerecich says.

"People's thumbs aren't calibrated. We wanted to set up a more scientific test, quantify the forces needed...and compare the iPhone 6 to other phones. We had an Instron universal testing machine, which does tensile and compression testing. We rigged it up to do bending tests on half a dozen different phones and models. We broke a lot of phones while doing testing."

Former low- or no-tech consumer products such as doorbells and dog collars are now part of the ever-increasing Internet of Things devices *Consumer Reports* is now testing. As a result, Rerecich's mission has expanded to tackle privacy and data-security concerns.

"We've seen issues in many consumer IoT products related to incomplete encryption, substandard authentication, or open vulnerabilities," she says. "For example, we've seen video doorbells and wireless home-security cameras sending email addresses, IP addresses, lists of commands, network names, and even Wi-Fi passwords as unencrypted data. When we find these [instances], we alert the manufacturers."

To date, Rerecich estimates, more than 500 products have been tested and rated for privacy and security.

Rerecich's advice to students and recent graduates includes "Learn to touch-type well and quickly, since you may spend more time at a keyboard than in a lab." Also, she urges students to get exposed to different areas of study and different types of courses. "You may not know what you like and are good at until you've tried it. For me, that was [a] chip design course at MIT—but I had also taken a course related to biomedical engineering because I thought I wanted to go into that. It's good to keep an open mind about different areas and learn as much as you can." ■

Crosstalk

Electricity's Slow Rollout

The spread of electricity has followed a protracted and characteristic pattern—slow, then fast, then slow again

One hundred forty years ago, Thomas Edison began generating electricity at two small coal-fired stations, one in London (Holborn Viaduct), the other in New York City (Pearl Street Station). Yet although electricity was clearly the next big thing, it took more than a lifetime to reach most people. Even now, not all parts of the world have easy access to it. Count this slow rollout as one more reminder that fundamental systemic transitions are protracted affairs.

Such transitions tend to follow an S-shaped curve: Growth rates shift from slow to fast, then back to slow again. I will demonstrate this by looking at a few key developments in electricity generation and residential consumption in the United States, which has reliable statistics for all but the earliest two decades of the electric period.

In 1902, the United States generated just 6 terawatt-hours of electricity, and the century-plus-old trajectory shows a clear S-curve. By 1912, the output was 25 TWh, by 1930 it was 114 TWh, by 1940 it was 180 TWh, and then three successive decadal doublings lifted it to almost 1,600 TWh by 1970. During the go-go years, the 1930s was the only decade in which gross electricity generation did not double, but after 1970 it took two decades to double, and from 1990 to 2020, the generation rose by only one-third.

As the process began to mature, the rising consumption of electricity was at first driven by declining prices, and then by the increasing variety of uses for electricity. The impressive drop in inflation-adjusted prices of electricity ended by 1970, and electricity generation reached a plateau, at about 4,000 TWh per year, in 2007.

Household electricity use remained restrained until after World War II.

The early expansion of generation was destined for industry—above all for the conversion from steam engines to electric motors—and for commerce. Household electricity use remained restrained until after World War II.

In 1900, fewer than 5 percent of all households had access to electricity; the biggest electrification jump took place during the 1920s, when the share of dwellings with connections rose from about 35 percent to 68 percent. By 1956, the diffusion was virtually complete, at 98.8 percent.

But access did not correlate strongly with use: Residential consumption remained modest, accounting for less than 10 percent of the total generation in 1930, and about 13 percent on the eve of World War II. In the 1880s, Edison light bulbs (inefficient and with low luminosity) were the first widely used indoor electricity converters. Lighting remained the dominant use for electricity in the household for the next three decades.

It took a long time for new appliances to make a difference, because there were significant gaps between the patenting and introduction of new appliances—including the electric iron (1903), the vacuum cleaner (1907), the toaster (1909), the electric stove (1912), the refrigerator (1913)—and their widespread ownership. Radio was adopted the fastest of all: Seventy-five percent of households had it by 1937.

The same dominant share was reached by refrigerators and stoves only in the 1940s—dishwashers by 1975, color TVs by 1977, and microwave ovens by 1988. Again, as expected, these diffusions followed more or less orderly S-curves.

Rising ownership of these and a range of other heavy electricity users drove the share of residential consumption to 25 percent by the late 1960s, and to about 40 percent in 2020. This share is well above Germany's 26 percent and far above China's roughly 15 percent. A new market for electricity is opening up, but slowly: So far, Americans have been reluctant buyers of electric vehicles, and, notoriously, they have long spurned building a network of high-speed electric trains, which every other affluent country has done. ■

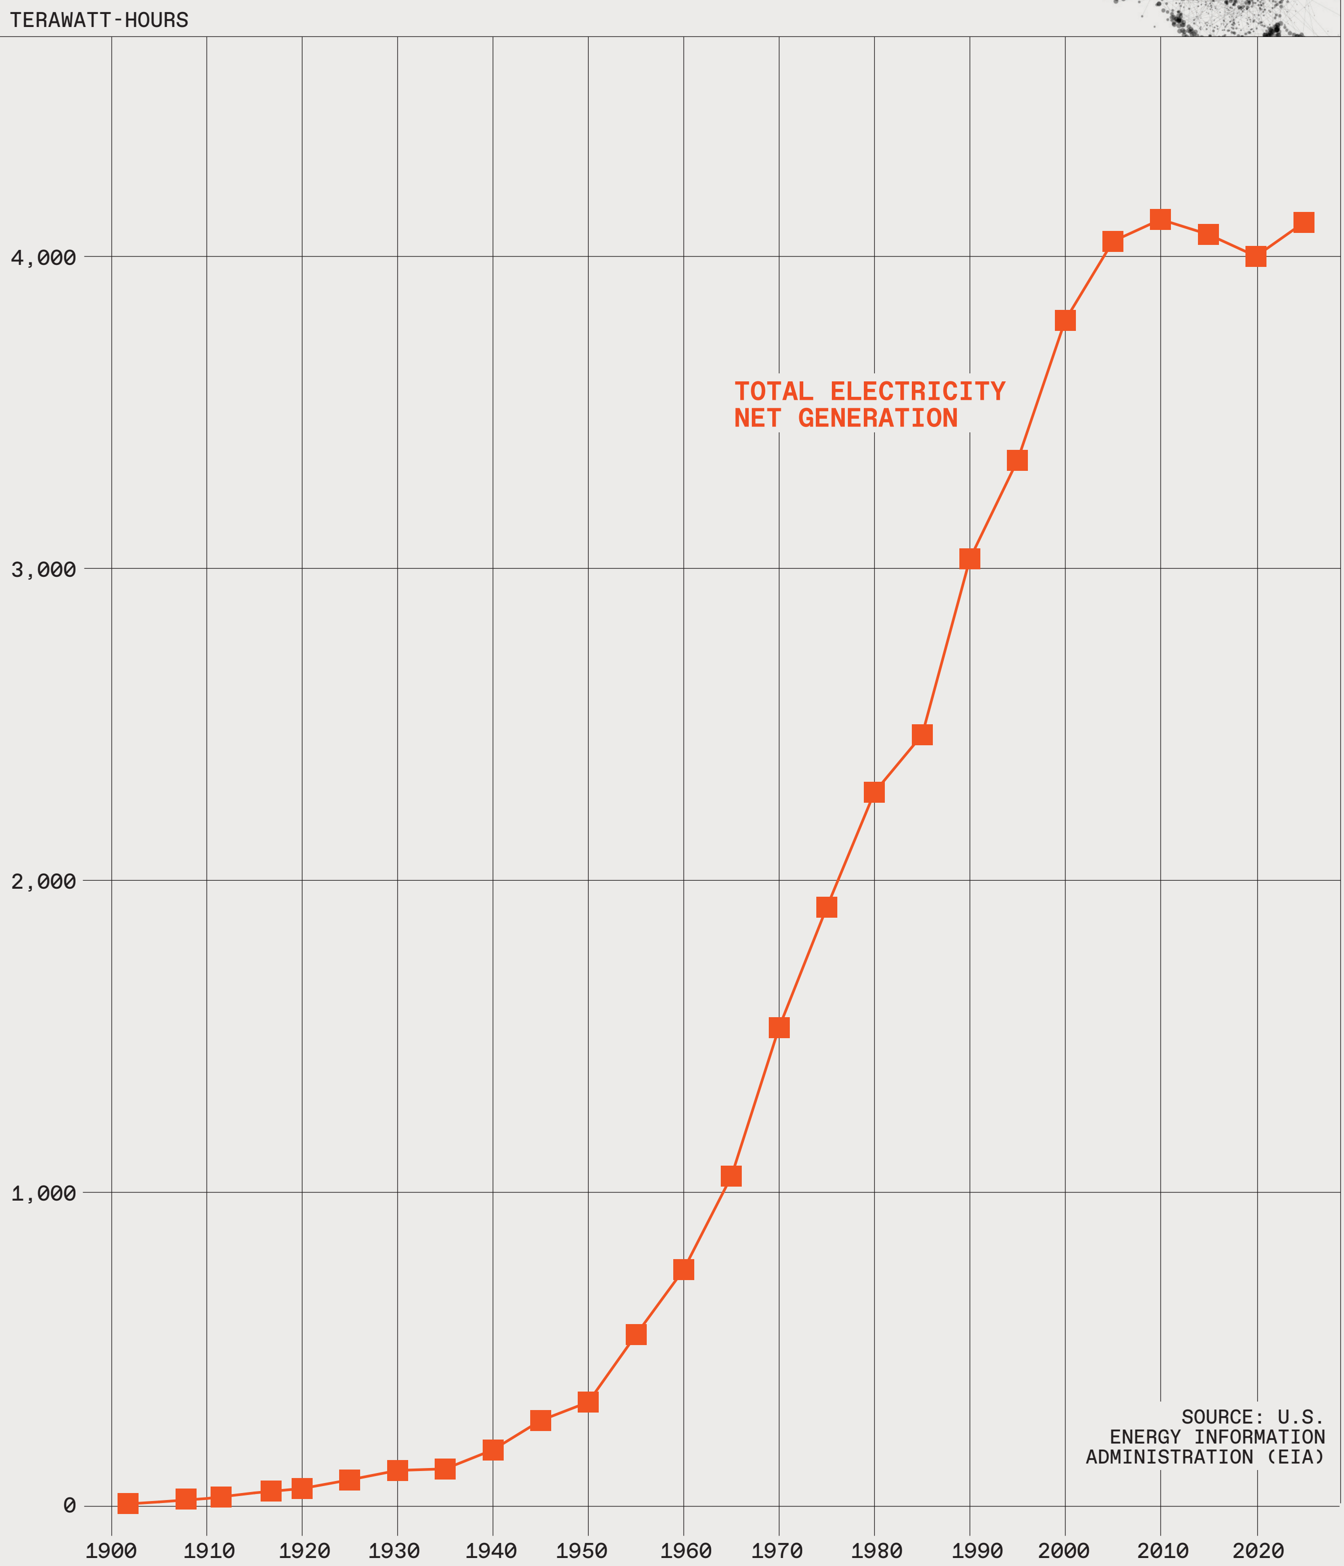
TERAWATT-HOURS
4,000
3,000
2,000
1,000
0
TOTAL ELECTRICITY NET GENERATION
1900
1910
1920
1930
1940
1950
1960
1970
1980
1990
2000
2010
2020
SOURCE: U.S. ENERGY INFORMATION ADMINISTRATION (EIA)

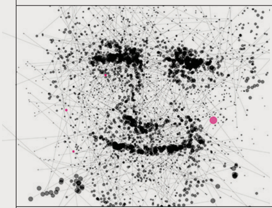

22 November 1935: The Day Globalism Was Born

Juan Trippe's Pan Am Clippers launched a new era

Today, the phrase "big tech" resonates negatively. It conjures up disturbing aspects of social media and the rise of megacorporations that seem beyond the reach of the law. And yet decades ago, big tech was associated with the glamour of motion: of speed, of power, and the thrill of exploring new frontiers.

Two leaders, Wernher von Braun and Juan Trippe, became household names as they made bold bets that paid off and enabled people to go where few thought it possible not long before. Von Braun had a troubling history: As a 30-year-old, he had convinced Adolf Hitler to fund his V-2 missiles, of which thousands were built, with slave labor. They rained down on Paris, London, and other cities, killing 9,000 people, mostly civilians.

But when von Braun's Apollo program came to fruition, in the late 1960s, huge crowds gathered every few months on the Florida coast to watch the thundering Saturn V rockets take off. It was a partylike atmosphere and a joyous time. We humans were going to the moon, making a connection that had seemed both improbable and impossible just a few years before.

The rapturous crowds in Florida gathered during a turbulent time, with popular culture dominated by sex, drugs, rock and roll, and pervasive antiestablishment sentiment. And yet, in that unsettled era, techno-optimism somehow took root. The most religious experience I have ever had was during the astronauts' live-to-Earth reading of Genesis from orbit around the moon on Christmas Eve, 1968.

Few probably realize that the big Apollo gatherings had a clear precedent in another mass outpouring of hope about large-scale human adventure. It occurred around San Francisco Bay at the height of the Great Depression, 30 years before the Apollo landings. Another indomitable spirit, Trippe, the president of Pan Am, was betting against all odds that he could open transpacific passenger service, and make it real on a timetable that no sober advisor believed possible.

On the Clipper's first return flight from Honolulu to San Francisco, disaster was very narrowly averted when the big seaplane landed in the bay with just 1 minute of fuel left in its tanks.

On 22 November 1935, the first transpacific commercial flight took to the skies. At 3:46 p.m., a Martin M-130 flying boat, the largest passenger plane built up to that point, lumbered into the air from Pan Am's base in Alameda, Calif., on the east side of San Francisco Bay. On that first trip, the plane carried only mail under contract to the U.S. post office.

Over 100,000 people had gathered around the bay to watch. Captain Edwin Musick powered the "China Clipper" northward in the bay and up over the waves. He planned to fly over the Bay Bridge, the double suspension bridge that today spans the bay and links San Francisco with Oakland. But he couldn't gain altitude quickly enough with the 4,000 gallons of fuel he was carrying. Fortunately, the roadway had not yet been hung from the catenaries, and he managed to fly under them. When he got to the Golden Gate Bridge, also under construction at the time, he just barely managed to get above it.

When Musick and his crew arrived in the Philippines six days later, there were 200,000 people cheering wildly in Manila Harbor. Musick hand-delivered a letter from President Franklin D. Roosevelt to President Manuel Quezon of the Philippines. Quezon told Musick, "You have swept away forever the distance which from the beginning of time has separated the great continent of America from the beautiful islands of the Pacific." Pan Am's landing in Manila marked the start of globalism, of our modern connectivity.

Today, as NASA's Artemis mission heralds a new era of human space exploration, it's important to remember how much difficulty and serendipity there was on the way to Manila and the moon. On the Clipper's first return flight from Honolulu to

FPG/ARCHIVE PHOTOS/GETTY IMAGES

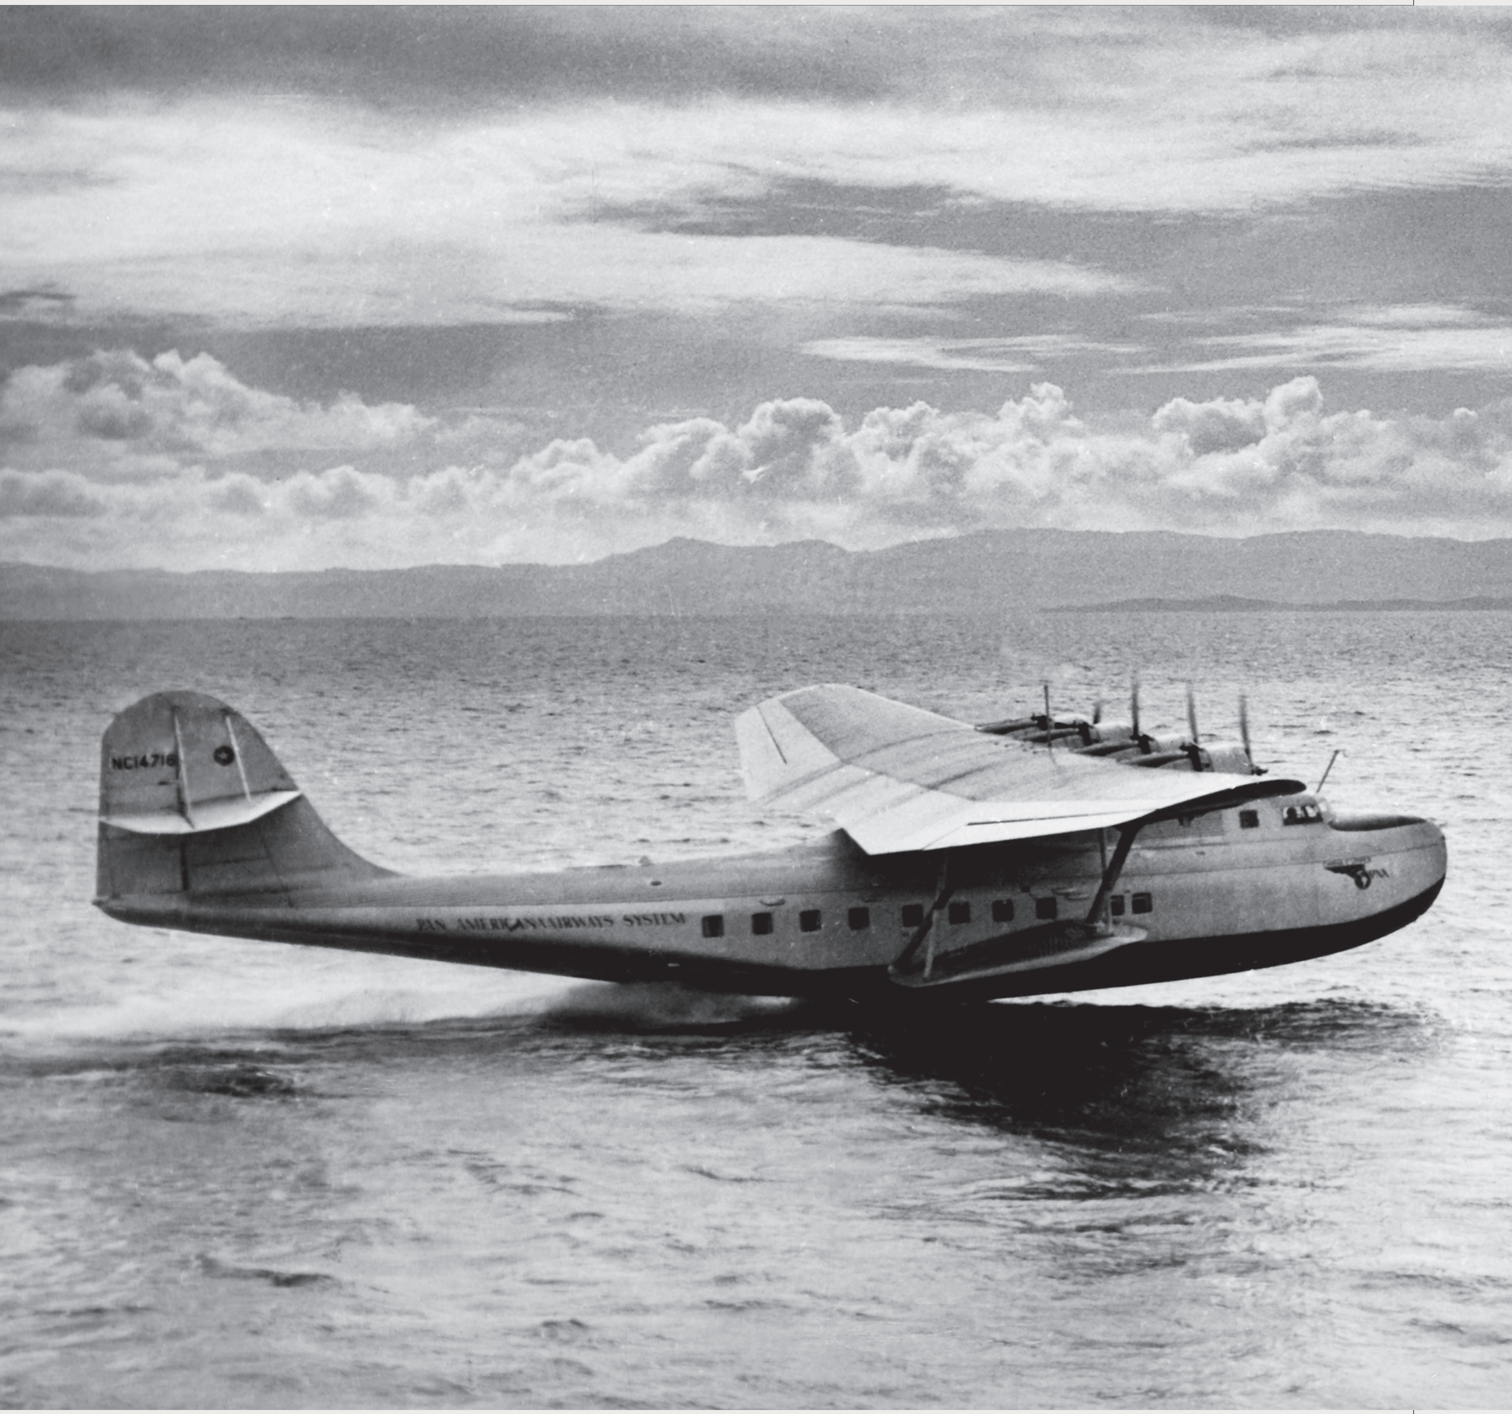

This Martin M-130 flying boat, called the China Clipper, took off on a test flight from San Francisco Bay in October 1935—one month before it made the first commercial transpacific flight.

San Francisco, disaster was very narrowly averted when the big seaplane landed in the bay with just 1 minute of fuel left in its tanks. Von Braun could have easily ended up dead or in Russia rather than in the United States.

Similarly, there will be many twists and turns on the way to the moon and Mars. The obvious successor to Trippe and von Braun would now seem to be Elon Musk—but maybe it'll be someone else. Regardless, and despite our divisions and perhaps even future pandemics, people will undoubtedly come out in droves once again to witness the takeoffs and landings. At last, big tech will again be something to celebrate. ■

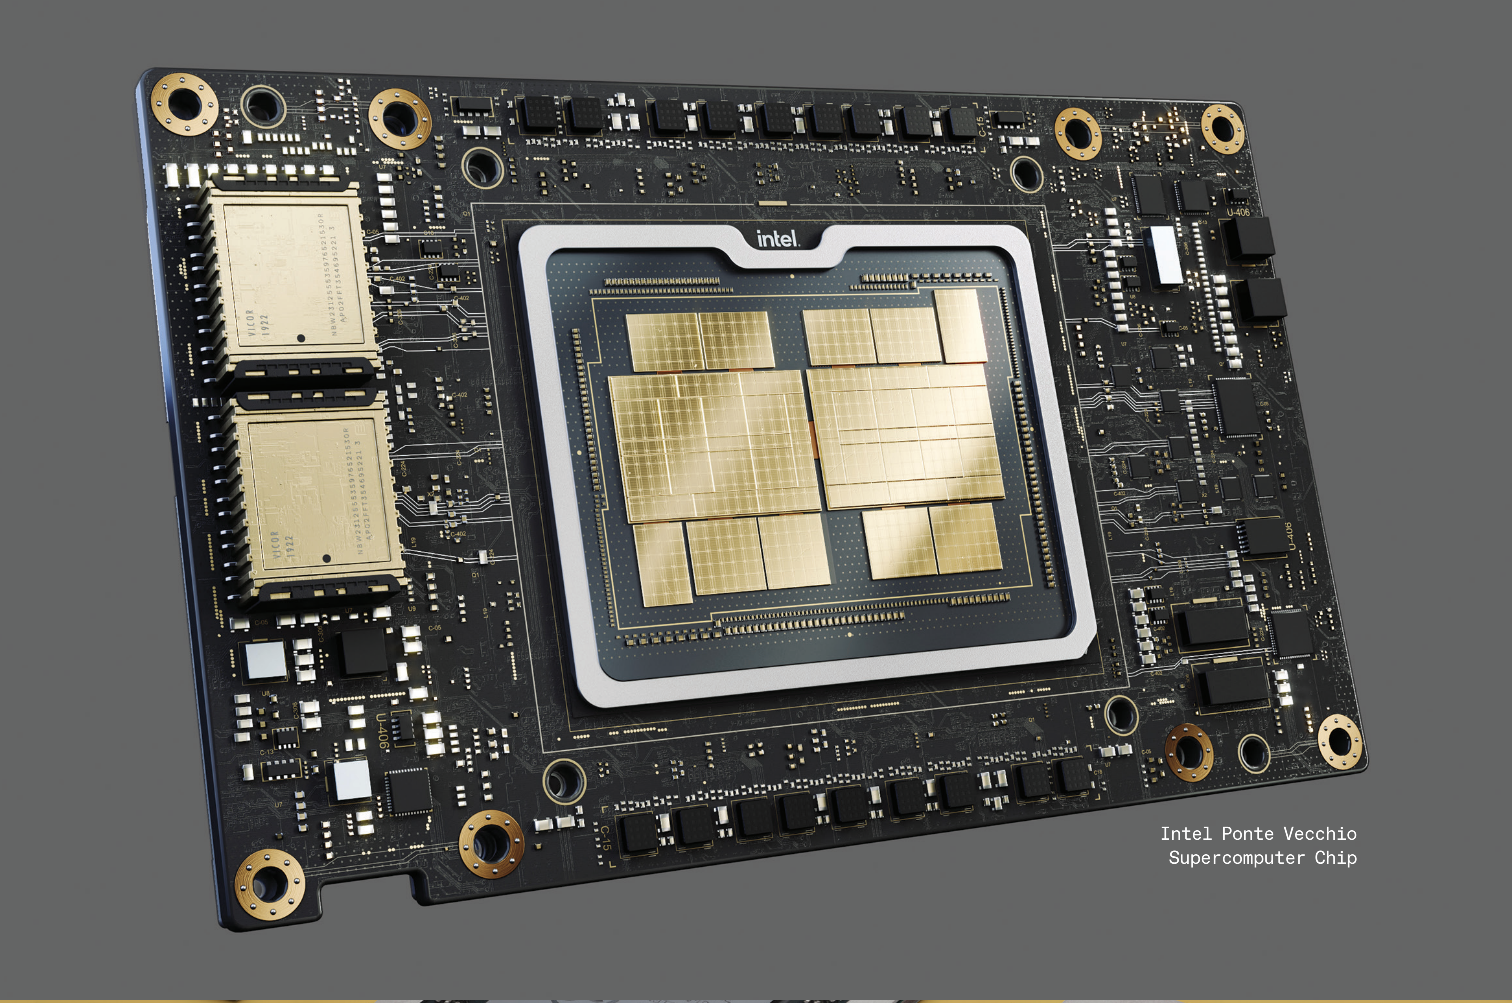

Intel Ponte Vecchio Supercomputer Chip

Graphcore Bow AI Processor

STACKING SILICON IS THE SOLUTION FOR HIGH-END COMPUTING By Samuel K. Moore

3 Paths to 3D Processors

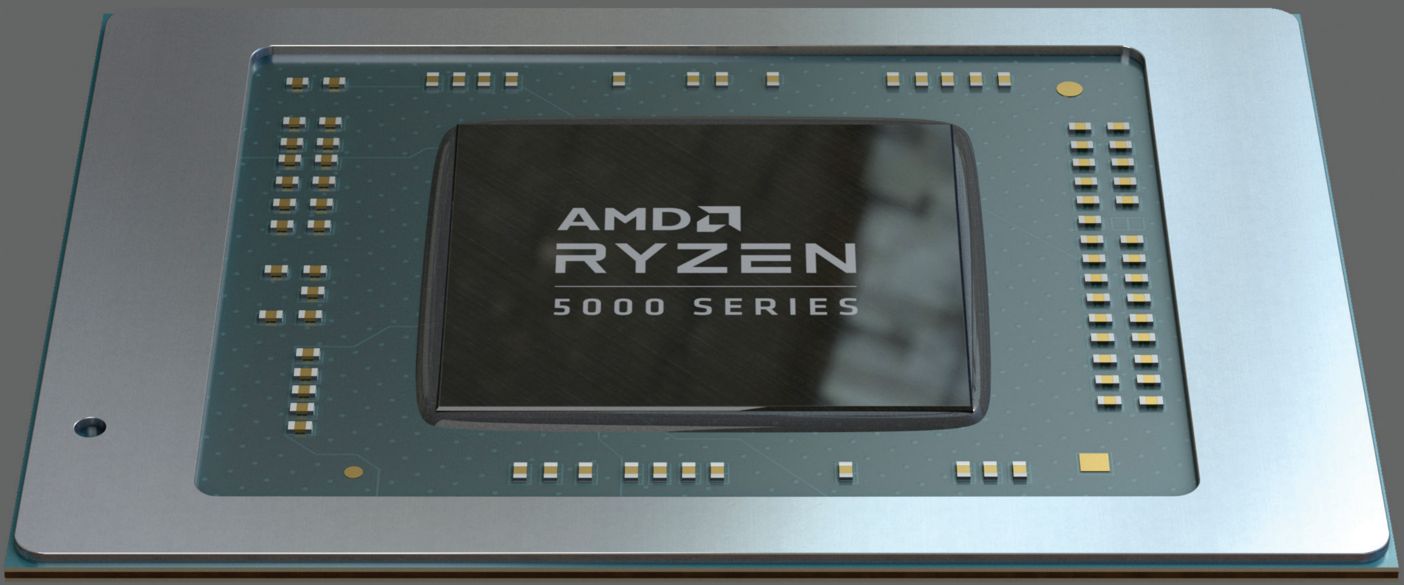

AMD Zen 3 with 3D V-Cache

A CROP OF HIGH-PERFORMANCE PROCESSORS is showing that the new direction for continuing Moore's Law is all about up. Each processor generation needs to perform better than the last, and at its most basic, that means integrating more logic onto the silicon. But there are two problems. The first is that our ability to shrink transistors and the logic and memory blocks they make up is slowing down. The second is that individual chips have reached their size limits. Photolithography tools can pattern an area no larger than about 850 square millimeters, which is about the size of a modern server GPU. ■ One answer has been to place two or more pieces of silicon side by side in the same package and stitch them together using millimeters-long, dense interconnects, so they can effectively act as a single unit. This so-called 2.5D scheme, enabled by advanced packaging technology, is already behind several top processors, which are now composed of several functional "chiplets" rather than a single IC. ■ But to sling truly huge volumes of data around as if it were all on the same chip, you need even shorter and denser connections, and that can be done only by stacking one chip atop another. Connecting two chips face-to-face in a 3D scheme can mean making hundreds or thousands of micrometers-long connections per square millimeter. These short, dense connections let data zip from one piece of silicon to another almost as rapidly and with as little energy as if the two were one chip. ■ It's taken a lot of innovation to get it to work. Engineers had to figure out how to keep heat from one chip in the stack from killing the other, decide what functions should go where and how they should be manufactured, keep the occasional bad chip from leading to a lot of costly dud systems, and deal with the resulting added complexities of doing all that at once. ■ Here are three examples that show not just how 3D chip stacking is done but also what it's good for.

PREVIOUS PAGES, CLOCKWISE FROM TOP: INTEL; AMD; GRAPHCORE

3D Technologies

HYBRID BONDING binds copper pads at the top of a chip's interconnect stack directly to the copper pads on a different chip. In hybrid bonding, the pads are in small recesses, surrounded by oxide insulator. The insulator is chemically activated and instantly bonds when pressed to its opposite at room temperature. Then, in an annealing step, the copper pads expand and bridge the gap to form a low-impedance link.

Hybrid bonding offers a high density of connections—in the range of 10,000 bonds per square millimeter, many more than in microbump technology, which offers about 400–1,600/mm² [chart]. The pitch—the distance from the edge of one interconnect to the far edge of the next—achievable today is about 9 micrometers, but tighter geometries are in the works. The technology will likely hit its limits around a pitch of 3 μm or so, says Lihong Cao, director of engineering and technical marketing at packaging-technology com-

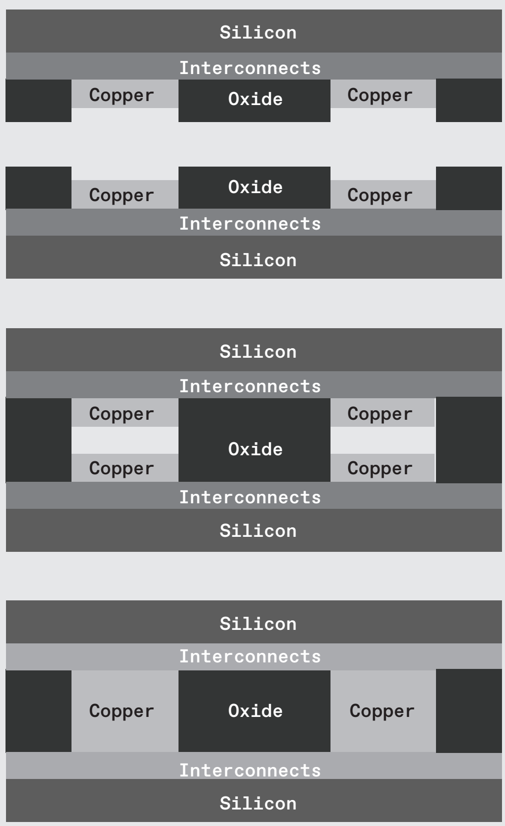

Hybrid bonding starts by forming recessed copper ports at the top "face" of the chip [top]. The surrounding oxide dielectric is chemically activated, so when the two chips are pressed together at room temperature they instantly bond [middle]. The bound chips are annealed, expanding the copper to form a conductive connection [bottom].

Density per millimeter

3D microbumps	3D hybrid bonding
400-1,600	>10,000

Through-silicon vias [pillars, bottom half] and hybrid bonding [rectangles, center] connect an AMD compute chiplet [bottom] to an SRAM chiplet [top].

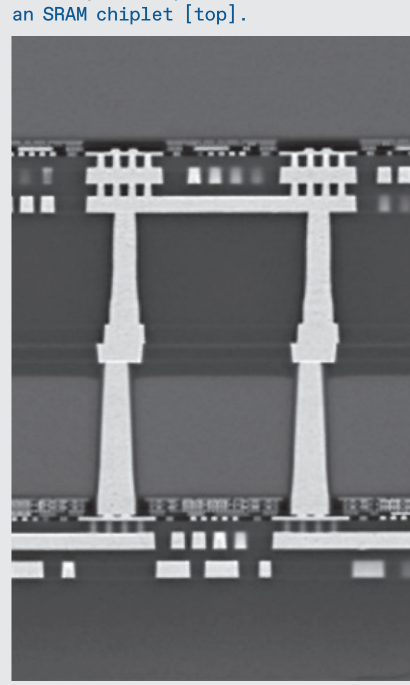

AMD's Zen 3 with 3D V-Cache

Personal computers have long come with the option to add more memory, to speed up extra-large applications and data-heavy work. Thanks to 3D chip stacking, AMD's next-generation CPU chiplet comes with that option, too. It's not an aftermarket add-on, of course, but if you're looking to build a computer with some extra oomph, ordering a processor with an extra-large cache memory could be the way to go.

The new Zen 3 compute chiplet and its predecessor are both made using the same TSMC manufacturing process—and therefore have the same size transistors, interconnects, and everything else. But AMD made so many architectural alterations to Zen 3 that even without the extra cache memory, its clock runs 6 percent faster and it demonstrates a 19 percent performance improvement on average.

3D Technologies: TSMC's chip-on-wafer hybrid bonding, through-silicon vias

On top of those architectural gems is the inclusion of a set of through-silicon vias (TSVs), vertical interconnects that burrow straight down through most of the silicon. The TSVs are built within the Zen 3's last-level cache, blocks of static RAM called L3. The cache sits in the middle of the compute chiplet and is shared across all eight of its cores.

In processors destined for data-heavy workloads, the backside of the Zen 3 wafer is thinned down until the TSVs are exposed. Then a 64-megabyte SRAM chiplet is attached to those exposed TSVs using what's called hybrid bonding [see sidebar, "3D Technologies"]. The latter technique can form connections between CPU core and cache memory every 9 micrometers. Finally, for structural stability and heat conduction, blank silicon chiplets are attached to cover the remainder of the Zen 3 CPU die.

Adding the extra cache by setting it beside the CPU die in a 2.5D arrangement was not an option,

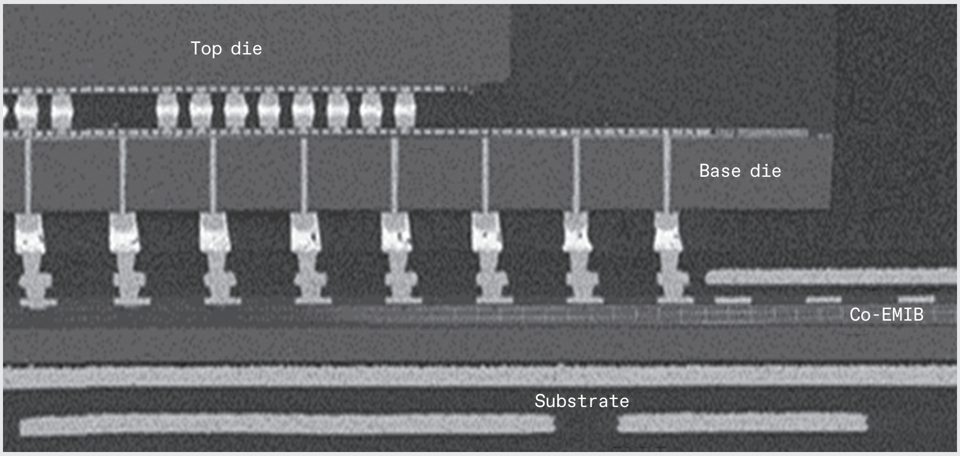

Intel's Ponte Vecchio uses microbumps [short pillars, top left] to connect a compute chiplet to a base die. Through-silicon vias [thin vertical lines, center] send signals from both the base die and the compute tile to the package. The base die connects horizontally to a second base die using a silicon bridge [horizontal bar, right].

pany ASE Group. The most critical step to improve hybrid bonding is keeping wafers from warping and reducing the surface roughness of each side to nanometer-scale perfection, she says.

MICROBUMPS are essentially an extremely scaled-down version of a standard packaging technology called flip-chip. In flip-chip, bumps of solder are added to the end points of the interconnects at the top (face) of a chip. The chip is then flipped onto a package substrate with a matching set of interconnects, and the solder is melted to form a bond. To stack two chips with this technology, one chip must have short copper pillars protruding from the surface. These are then capped with a "microbump" of solder, and the two chips are joined face-to-face by melting the solder.

The minimum distance from the start of one connection to the far edge of the next, the pitch, can be less than 50 μm when using microbumps. Intel used a 36-μm-pitch version of its Foveros 3D integration technology in the Ponte Vecchio supercomputer chip. Samsung says its microbump technology, called 3D X-Cube, is available with a 30-μm pitch. The technology cannot match hybrid bonding's density [above]. However, its requirements for alignment and planarization are not as strict as hybrid bonding's, making it easier to stack multiple chips that are made using different manufacturing technologies onto a single base chip.

THROUGH-SILICON VIAS (TSVs) are interconnects that descend vertically down through a chip's silicon. They don't extend through a wafer's entire depth, so the backside of the silicon has to be ground away until the TSV is exposed. They are often necessary in 3D stacked chips because the chips are bonded together so that their interconnects are face-to-face. In that case, the TSVs provide the stack with access to power and data. They have been in wide use for years in high-bandwidth dynamic RAM, which stacks multiple memory chips vertically. But with 3D chip stacking, the technology has come to logic chips, too. ■

PHOTOS: LEFT: AMD; RIGHT: INTEL

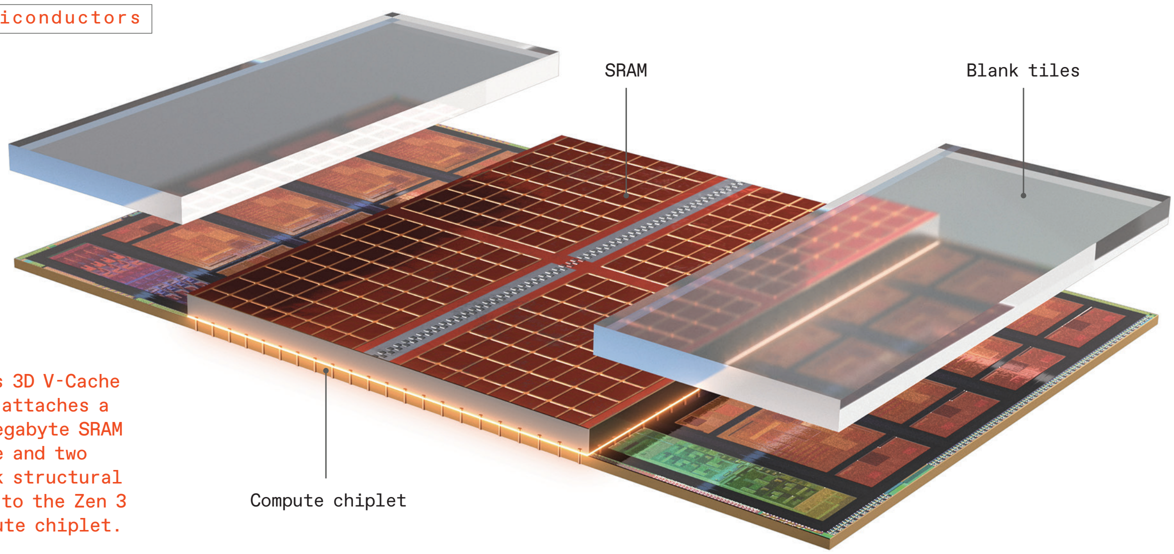

AMD's 3D V-Cache tech attaches a 64-megabyte SRAM cache and two blank structural dies to the Zen 3 compute chiplet.

because data would take too long to get to the processor cores. "Despite tripling the L3 [cache] size, 3D V-Cache only added four [clock] cycles of latency—something that could only be achieved through 3D stacking," John Wuu, AMD senior fellow design engineer, told virtual attendees of the IEEE International Solid-State Circuits Conference (ISSCC) in February.

The bigger cache made its mark in high-end games. Using the desktop Ryzen CPU with 3D V-Cache sped games rendered at 1080p by an average of 15 percent. It was good for more serious work as well in the Epyc server CPU, shortening the run time for difficult semiconductor-design verification work by 66 percent.

The industry's ability to shrink SRAM is slowing compared to how well it can shrink logic, Wuu pointed out. So future SRAM expansion packs will likely continue to be made using more established manufacturing processes while the compute chiplets are driven down to Moore's Law's bleeding edge.

3D Technologies: TSMC's wafer-on-wafer hybrid bonding, through-silicon vias

Graphcore's Bow AI Processor

3D integration can speed computing even if one chip in the stack doesn't have a single transistor on it. England-based AI computer company Graphcore managed a huge increase to its systems' performance just by attaching a power-delivery chip to its AI processor. The addition of the power-delivery silicon means the combined chip, called Bow, can run faster—1.85 gigahertz versus 1.325 GHz—and at lower voltage than its predecessor. That translates to computers that train neural nets up to 40 percent faster with as much as 16 percent less energy compared to its previous generation. Importantly, users get this improvement with no change to their software at all.

The power-management die is packed with a combination of capacitors and TSVs. The latter are just to deliver power and data to the processor chip. It's the capacitors that really make the difference. Like the bit-storing components in dynamic RAM, these capacitors are formed in deep, narrow trenches in

Graphcore's Bow AI Processor is a 3D stack made up of a compute chip and a power delivery die with integrated capacitors.

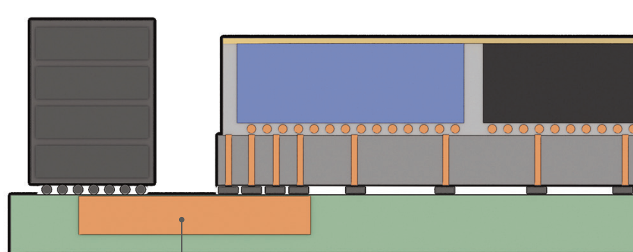

TOP: AMD; BOTTOM: GRAPHCORE

the silicon. Because these reservoirs of charge are so close to the processor's transistors, power delivery is smoothed out, allowing the processor cores to run faster at lower voltage. Without the power-delivery chip, the processor would have to increase its operating voltage above its nominal level to work at 1.85 GHz, consuming a lot more power. With the power chip, it can reach that clock rate and consume less power, too.

The manufacturing process used to make Bow is unique but not likely to stay that way. Most 3D stacking is done by bonding one chip to the other while one of them is still on the wafer, called chip-on-wafer. Instead Bow used TSMC's wafer-on-wafer technology, where an entire wafer of one type is bonded to an entire wafer of the other, then diced up into chips. Bow is the first chip on the market to use the technology, which led to a higher density of connections between the two dies than could be achieved using a chip-on-wafer process, says Simon Knowles, Graphcore's chief technical officer and cofounder.

Although Graphcore's power-delivery chip has no transistors, those might be coming. Using the technology only for power delivery "is just the first step for us," says Knowles. "It will go much further than that in the near future."

Intel's Ponte Vecchio Supercomputer Chip

The Aurora supercomputer is designed to become one of the first U.S.-based high-performance computers to pierce the exaflop barrier—a billion billion high-precision floating-point calculations per second. To get Aurora to those heights, Intel's Ponte Vecchio packs more than 100 billion transistors across 47 pieces of silicon to make a single processor. Using both 2.5D and 3D technologies, Intel squeezed 3,100 square millimeters of silicon—nearly equal to four Nvidia A100 GPUs—into a 2,330-mm^2 footprint.

Intel Fellow Wilfred Gomes told engineers at ISSCC that the processor pushed Intel's 2D and 3D integration technologies to the limits.

3D Technologies: Intel's Foveros 3D integration (microbumps), through-silicon vias

Each Ponte Vecchio is really two mirror-image sets of chiplets tied together using Intel's 2.5D integration technology, Co-EMIB. Short for co-embedded multi-die interconnect bridge, Co-EMIB forms a silicon bridge of high-density interconnects between two 3D stacks of chiplets. Co-EMIB dies also connect high-bandwidth memory and an I/O chiplet to the "base tile," the chiplet on which the rest are stacked.

The base tile uses Intel's 3D integration technology, called Foveros, to stack compute and cache-memory chiplets atop it. Foveros uses microbumps—short copper pillars each topped with a micrometers-wide ball of solder—to make vertical connections a few dozen micrometers apart. Signals and power get into this stack by means of TSVs.

Eight compute tiles, four cache tiles, and eight blank "thermal" tiles meant to remove heat from the processor are all attached to the base tile. The base itself provides cache memory and a network that allows any compute tile to access any memory element.

Needless to say, none of this was easy. It took innovations in yield management, clock circuits, thermal regulation, and power delivery, Gomes told conference attendees. For example, Intel engineers chose to supply the processor with a higher-than-normal voltage (1.8 volts) so that current would be low enough to simplify the package. Circuits in the base tile reduce the voltage to something closer to 0.7 V for use on the compute tiles, and each compute tile had to have its own power domain in the base tile. Key to this ability were new high-efficiency components called coaxial magnetic integrated inductors. Because these are built into the package substrate, the circuit actually snakes up and down between the base tile and the package before supplying the voltage to the compute tile.

It's taken 14 years to go from the first petaflop supercomputer in 2008 to exaflops this year, Gomes said. Advanced packaging, such as 3D stacking, is among the technologies that could help shorten the next thousandfold computing improvement to just six years, Gomes predicted. ■

INTEL

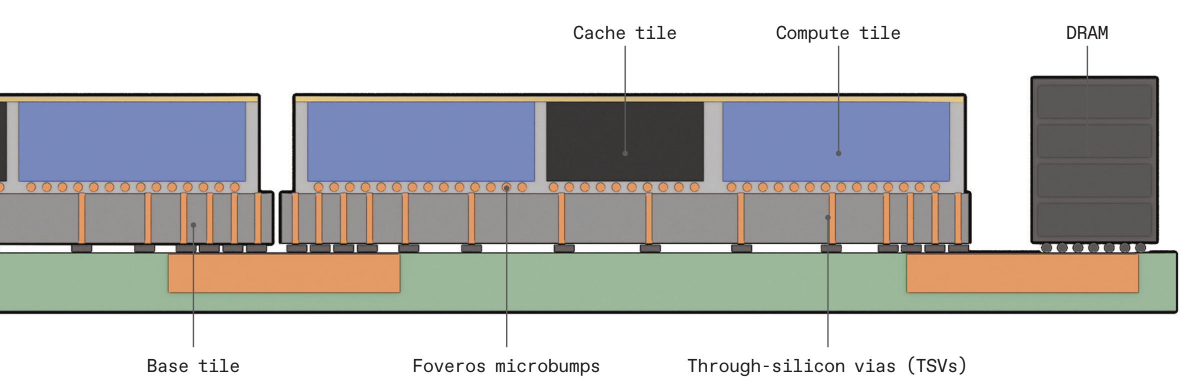

The different parts of Ponte Vecchio are made using different manufacturing processes. Intel's Foveros technology creates the 3D interconnects, and its Co-EMIB makes the horizontal connections.

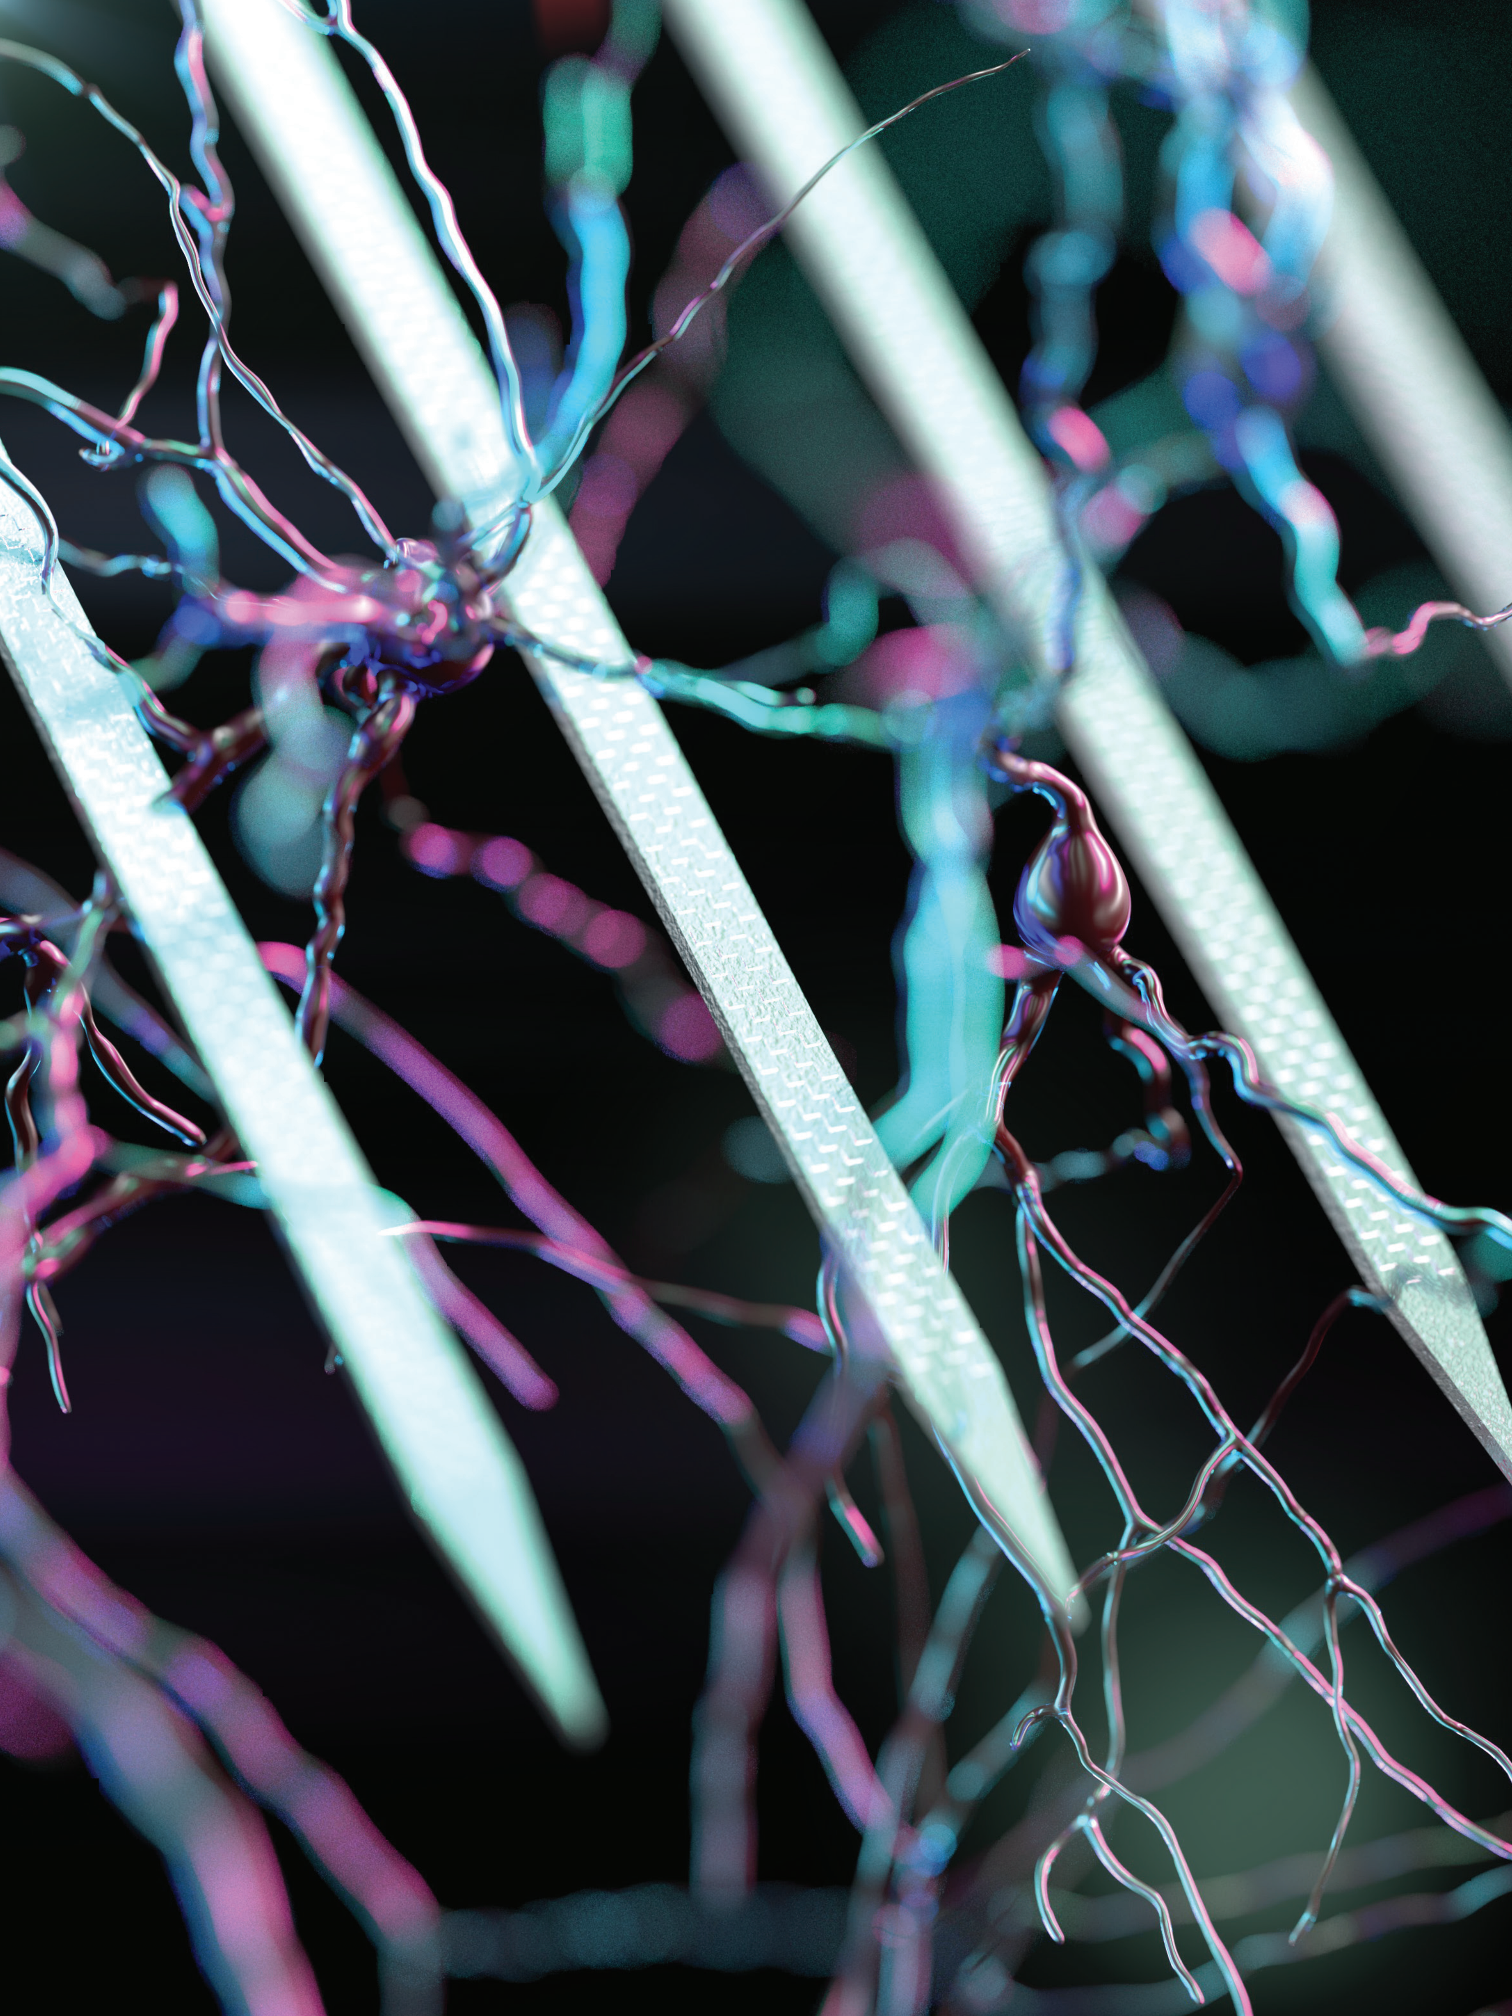

EAVESDROPPING ON THE BRAIN

With 10,000 electrodes, this neural implant senses more than ever before

By BARUN DUTTA

Illustration by ANATOMYBLUE

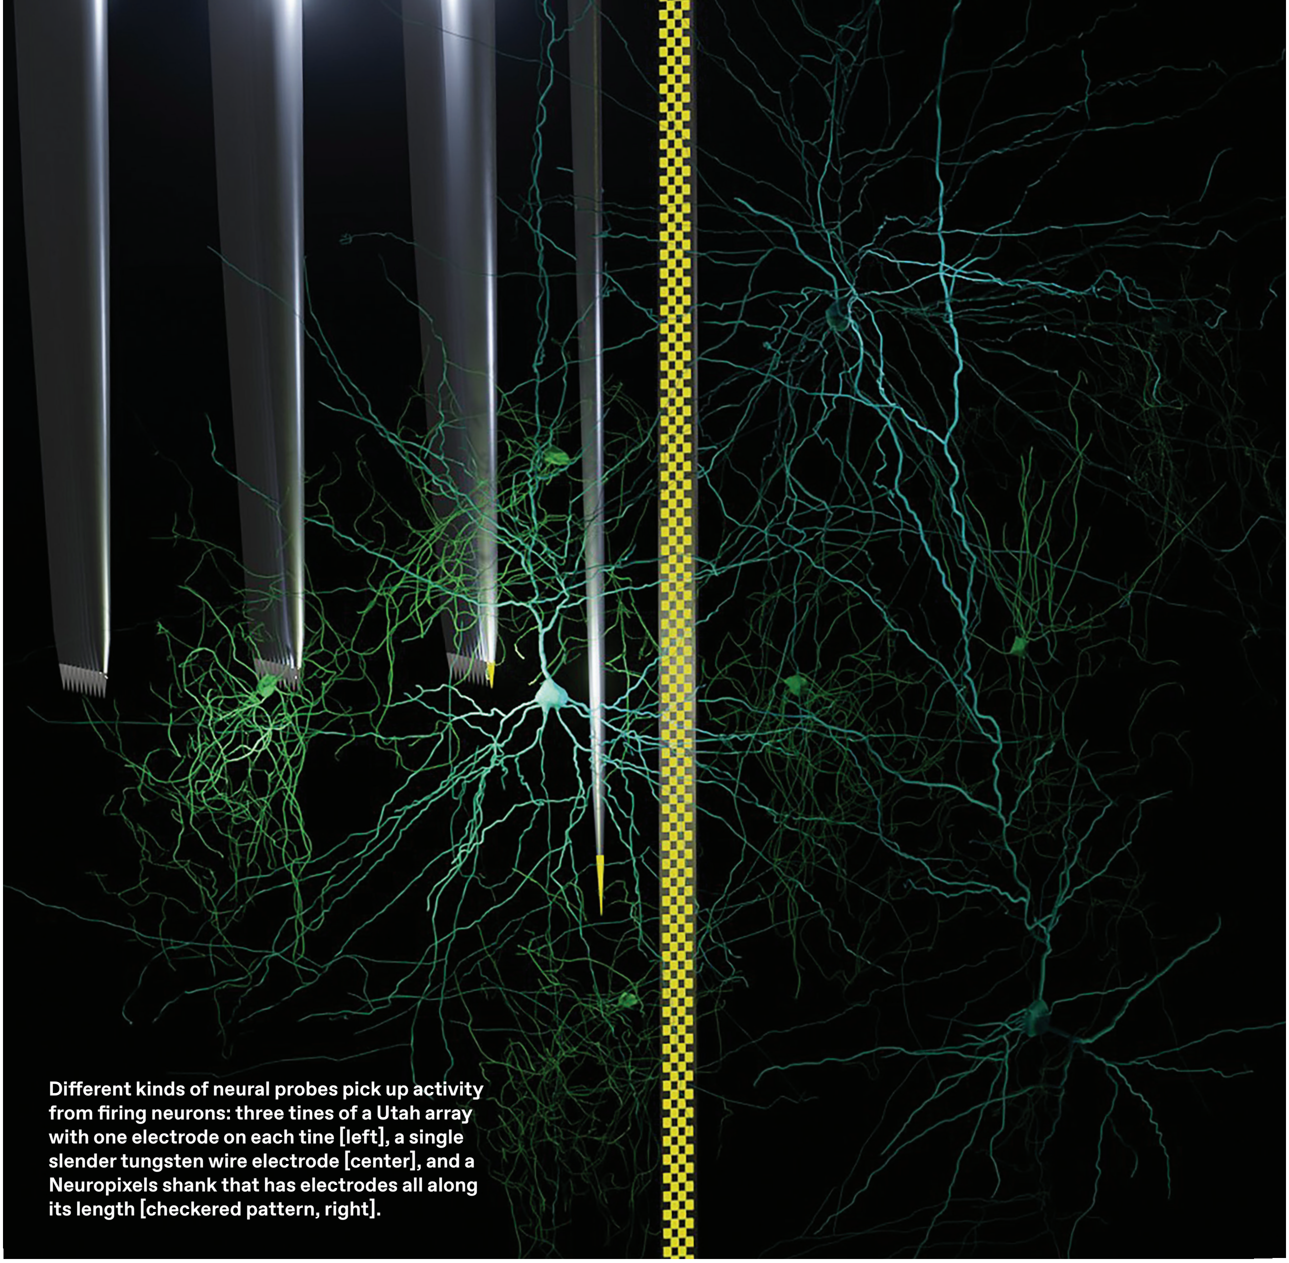

Different kinds of neural probes pick up activity from firing neurons: three tines of a Utah array with one electrode on each tine [left], a single slender tungsten wire electrode [center], and a Neuropixels shank that has electrodes all along its length [checkered pattern, right].

MASSACHUSETTS GENERAL HOSPITAL/IMEC/NATURE NEUROSCIENCE

I

IMAGINE A PORTABLE COMPUTER built from a network of 86 billion switches, capable of general intelligence sophisticated enough to build a spacefaring civilization—but weighing just 1.2 to 1.3 kilograms, consuming just 20 watts of power, and jiggling like Jell-O as it moves. There's one inside your skull right now. It is a breathtaking achievement of biological evolution. But there are no blueprints.

Now imagine trying to figure out how this wonder of bioelectronics works without a way to observe its microcircuitry in action. That's like asking a microelectronics engineer to reverse engineer the architecture, microcode, and operating system running on a state-of-the-art processor without the use of a digital logic probe, which would be a virtually impossible task.

So it's easy to understand why many of the operational details of humans' brains (and even the brains of mice and much simpler organisms) remain so mysterious, even to neuroscientists. People often think of technology as applied science, but the scientific study of brains is essentially applied sensor technology. Each invention of a new way to measure brain activity—including scalp electrodes, MRIs, and microchips pressed into the surface of the cortex—has unlocked major advances in our understanding of the most complex, and most human, of all our organs.

The brain is essentially an electrical organ, and that fact plus its gelatinous consistency pose a hard technological problem. In 2010, I met with leading neuroscientists at the Howard Hughes Medical Institute (HHMI) to explore how we might use advanced microelectronics to invent a new sensor. Our goal: to listen in on the electrical conversations taking place among thousands of neurons at once in any given thimbleful of brain tissue.

Timothy D. Harris, a senior scientist at HHMI, told me that "we need to record every spike from every neuron" in a localized neural circuit within a freely moving animal. That would mean building a digital probe long enough to reach any part of the thinking organ, but slim enough not to destroy fragile tissue on its way in. The probe would need to be durable enough to stay put and record reliably for weeks or even months as the brain guides the body through complex behaviors.

For an electrical engineer, those requirements add up to a very tall order. But more than a decade of R&D by a global, multidisciplinary team of engineers, neuroscientists, and software designers has at last met the challenge, producing a remarkable new tool that is now being put to use in hundreds of labs around the globe. As chief scientist at Imec, a leading independent nanoelectronics R&D institute, in Belgium, I saw the opportunity to extend advanced semiconductor technology to serve broad new swaths of biomedicine and brain science. Envisioning and shepherding the technological aspects of this ambitious project has been one of the highlights of my career.

We named the system Neuropixels because it functions like an imaging device, but one that records electrical rather than photonic fields. Early experiments already underway—including some in humans—have helped explore age-old questions about the brain. How do physiological needs produce motivational drives, such as thirst and hunger? What regulates behaviors essential to survival? How does our neural system map the position of an individual within a physical environment?

Successes in these preliminary studies give us confidence that Neuropixels is shifting neuroscience into a higher gear that will deliver faster insights into a wide range of normal behaviors and potentially enable better treatments for brain disorders such as epilepsy and Parkinson's disease.

Version 2.0 of the system, demonstrated last year, increases the sensor count by about an order of magnitude over that of the initial version produced just four years earlier. It paves the way for future brain-computer interfaces that may enable paralyzed people to communicate at speeds approaching those of normal conversation. With version 3.0 already in early development, we believe that Neuropixels is just at the beginning of a long road of exponential Moore's Law–like growth in capabilities.

I

IN THE 1950S, researchers used a primitive electronic sensor to identify the misfiring neurons that give rise to Parkinson's disease. During the 70 years since, the technology has come far, as the microelectronics revolution miniaturized all the components that go into a brain probe: from the electrodes that pick up the tiny voltage spikes that neurons emit when they fire, to the amplifiers and digitizers that boost signals and reduce noise, to the thin wires that transmit power into the probe and carry data out.

By the time I started working with HHMI neuroscientists in 2010, the best electrophysiology probes, made by NeuroNexus and Blackrock Neurotech, could record the activity of roughly 100 neurons at a time. But they were able to monitor only cells in the cortical areas near the brain's surface. The shallow sensors were thus unable to access deep brain regions—such as the hypothalamus, thalamus, basal ganglia, and limbic system—that govern hunger, thirst, sleep, pain, memory, emotions, and other important perceptions and behaviors. Companies such as Plexon make probes that reach deeper into the brain, but they are limited to sampling 10 to 15 neurons simultaneously. We set

The team realized that they could mount two Neuropixels 2.0 probes on one headstage, the board that sits outside the skull, providing a total of eight shanks with 10,240 recording electrodes.

IMEC

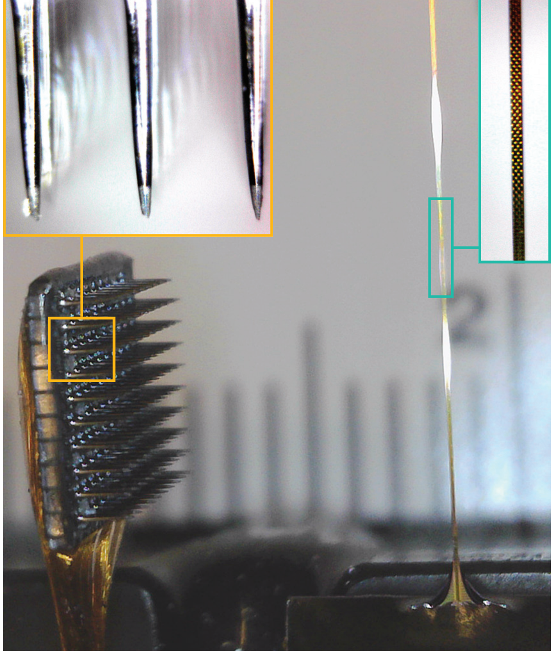

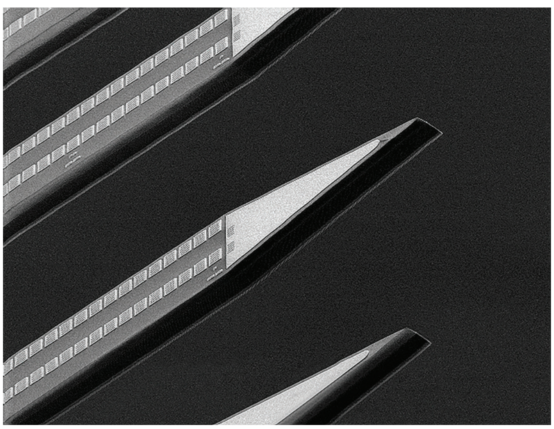

The most common neural recording device today is the Utah array [top image, at left], which has one electrode at the tip of each of its tines. In contrast, a Neuropixels probe [top image, at right] has hundreds of electrodes along each of its long shanks. An image taken by a scanning electron microscope [bottom] magnifies the tips of several Neuropixels shanks.

MASSACHUSETTS GENERAL HOSPITAL/IMEC/NATURE NEUROSCIENCE

for ourselves a bold goal of improving on that number by one or two orders of magnitude.

To understand how brain circuits work, we really need to record the individual, rapid-fire activity of hundreds of neurons as they exchange information in a living animal. External electrodes on the skull don't have enough spatial resolution, and functional MRI technology lacks the speed necessary to record fast-changing signals. Eavesdropping on these conversations requires being in the room where it happens: We needed a way to place thousands of micrometer-size electrodes directly in contact with vertical columns of neurons, anywhere in the brain. (Fortuitously, neuroscientists have discovered that when a brain region is active, correlated signals pass through the region both vertically and horizontally.)

These functional goals drove our design toward long, slender silicon shanks packed with electrical sensors. We soon realized, however, that we faced a major materials issue. We would need to use Imec's CMOS fab to mass-produce complex devices by the thousands to make them affordable to research labs. But CMOS-compatible electronics are rigid when packed at high density.

The brain, in contrast, has the same elasticity as Greek yogurt. Try inserting strands of angel-hair pasta into yogurt and then shaking them a few times, and you'll see the problem. If the pasta is too wet, it will bend as it goes in or won't go in at all. Too dry, and it breaks. How would we build shanks that could stay straight going in yet flex enough inside a jiggling brain to remain intact for months without damaging adjacent brain cells?

Experts in brain biology suggested that we use gold or platinum for the electrodes and an organometallic polymer for the shanks. But none of those are compatible with advanced CMOS fabrication. After some research and lots of engineering, my Imec colleague Silke Musa invented a form of titanium nitride—an extremely tough electroceramic—that is compatible with both CMOS fabs and animal brains. The material is also porous, which gives it a low impedance; that quality is very helpful in getting currents in and clean signals out without heating the nearby cells, creating noise, and spoiling the data.

Thanks to an enormous amount of materials-science research and some techniques borrowed from microelectromechanical systems (MEMS), we are now able to control the internal stresses created during the deposition and etching of the silicon shanks and the titanium nitride electrodes so that the shanks consistently come out almost perfectly straight, despite being only 23 micrometers thick. Each probe consists of four parallel shanks, and each shank is studded with 1,280 electrodes. At 1 centimeter in length, the probes are long enough to reach any spot in a mouse's brain. Mouse studies published in 2021 showed that Neuropixels 2.0 devices can collect data from the same neurons continuously for over six months as the rodents go about their lives.

THE THOUSANDFOLD DIFFERENCE in elasticity between CMOS-compatible shanks and brain tissue presented us with another major problem during such long-term studies: how to keep track of individual neurons as the probes inevitably shift in position relative to the moving brain. Neurons are 20 to 100 µm in size; each square pixel (as we call the electrodes) is 15 µm across, small enough that it can record the isolated activity of a single neuron. But over six months of jostling activity, the probe as a whole can move within the brain by up to 500 µm. Any particular pixel might see several neurons come and go during that time.

The 1,280 electrodes on each shank are individually addressable, and the four parallel shanks give us an effectively 2D readout, which is quite analogous to a CMOS camera image, and the inspiration for the name Neuropixels. That similarity made me realize that this problem of neurons shifting relative to pixels is directly analogous to image stabilization. Just like the subject filmed by a shaky camera, neurons in a chunk of brain are correlated in their electrical behavior. We were able to adapt knowledge and algorithms developed years ago for fixing camera shake to solve our problem of probe shake. With the stabilization software active, we are now able to apply automatic corrections when neural circuits move across any or all of the four shanks.

The first Neuropixels device [top] had one shank with 966 electrodes. Neuropixels 2.0 has four shanks with 1,280 electrodes each. Two probes can be mounted on one headstage.

IMEC

We needed a way to place thousands of micrometer-size electrodes directly in contact with vertical columns of neurons, anywhere in the brain.

Version 2.0 shrank the headstage—the board that sits outside the skull, controls the implanted probes, and outputs digital data—to the size of a thumbnail. A single headstage and base can now support two probes, each extending four shanks, for a total of 10,240 recording electrodes. Control software and apps written by a fast-growing user base of Neuropixels researchers allow real-time, 30-kilohertz sampling of the firing activity of 768 distinct neurons at once, selected at will from the thousands of neurons touched by the probes. That high sampling rate, which is 500 times as fast as the 60 frames per second typically recorded by CMOS imaging chips, produces a flood of data, but the devices cannot yet capture activity from every neuron contacted. Continued advances in computing will help us ease those bandwidth limitations in future generations of the technology.

IN JUST FOUR YEARS, we have nearly doubled the pixel density, doubled the number of pixels we can record from simultaneously, and increased the overall pixel count more than tenfold, while shrinking the size of the external electronics by half. That Moore's Law–like pace of progress has been driven in large part by the use of commercial-scale CMOS and MEMS fabrication processes, and we see it continuing.

A next-gen design, Neuropixels 3.0, is already under development and on track for release around 2025, maintaining a four-year cadence. In 3.0, we expect the pixel count to leap again, to allow eavesdropping on perhaps 50,000 to 100,000 neurons. We are also aiming to add probes and to triple or quadruple the output bandwidth, while slimming the base by another factor of two.

Just as was true of microchips in the early days of the semiconductor industry, it's hard to predict all the applications Neuropixels technology will find. Adoption has skyrocketed since 2017. Researchers at more than 650 labs around the world now use Neuropixels devices, and a thriving open-source community has appeared to create apps for them. It has been fascinating to see the projects that have sprung up: For example, the Allen Institute for Brain Science in Seattle recently used Neuropixels to create a database of activity from 100,000-odd neurons involved in visual perception, while a group at Stanford University used the devices to map how the sensation of thirst manifests across 34 different parts of the mouse brain.

We have begun fabricating longer probes of up to 5 cm and have defined a path to probes of 15 cm—big enough to reach the center of a human brain. The first trials of Neuropixels in humans were a success, and soon we expect the devices will be used to better position the implanted stimulators that quiet the tremors caused by Parkinson's disease, with 10-µm accuracy. Soon, the devices may also help identify which regions are causing seizures in the brains of people with epilepsy, so that corrective surgery eliminates the problematic bits and no more.

Future generations of the technology could play a key role as sensors that enable people who become "locked in" by neurodegenerative diseases or traumatic injury to communicate at speeds approaching those of typical conversation. Every year, some 64,000 people worldwide develop motor neuron disease, one of the more common causes of such entrapment. Though a great deal more work lies ahead to realize the potential of Neuropixels for this critical application, we believe that fast and practical brain-based communication will require precise monitoring of the activity of large numbers of neurons for long periods of time.

An electrical, analog-to-digital interface from wetware to hardware has been a long time coming. But thanks to a happy confluence of advances in neuroscience and microelectronics engineering, we finally have a tool that will let us begin to reverse engineer the wonders of the brain. ■

SPOOKY POWER AT A DISTANCE

Researchers have beamed substantial amounts of energy at distances over 1 kilometer

BY PAUL JAFFE

A power-beaming system developed by PowerLight Technologies conveyed hundreds of watts of power during a 2019 demonstration at the Port of Seattle.

Wires have a lot going for them when it comes to moving electric power around, but they have their drawbacks too. Who, after all, hasn't tired of having to plug in and unplug their phone and other rechargeable gizmos? It's a nuisance.

Wires also challenge electric utilities: These companies must take pains to boost the voltage they apply to their transmission cables to very high values to avoid dissipating most of the power along the way. And when it comes to powering public transportation, including electric trains and trams, wires need to be used in tandem with rolling or sliding contacts, which are troublesome to maintain, can spark, and in some settings will generate problematic contaminants.

Many people are hungry for solutions to these issues—witness the widespread adoption over the past decade of wireless charging, mostly for portable consumer electronics but also for vehicles. While a wireless charger saves you from having to connect and disconnect cables repeatedly, the distance over which energy can be delivered this way is quite short. Indeed, it's hard to recharge or power a device when the air gap is just a few centimeters, much less a few meters. Is there really no practical way to send power over greater distances without wires?

To some, the whole notion of wireless power transmission evokes images of Nikola Tesla with high-voltage coils spewing miniature bolts of lightning. This wouldn't be such a silly connection to make. Tesla had indeed pursued the idea of somehow using the ground and atmosphere as a conduit for long-distance power transmission, a plan that went nowhere. But his dream of sending electric power over great distances without wires has persisted.

Guglielmo Marconi, who was Tesla's contemporary, figured out how to use "Hertzian waves," or electromagnetic waves, as we call them today, to send signals over long distances. And that advance brought with it the possibility of using the same kind of waves to carry energy from one place to another. This is, after all, how all the energy stored in wood, coal, oil, and natural gas originally got here: It was transmitted 150 million kilometers through space as electromagnetic waves—sunlight—most of it millions of years ago.

Can the same basic physics be harnessed to replace wires today? My colleagues and I at the U.S. Naval Research Laboratory, in Washington, D.C., think so, and here are some of the reasons why.

THERE HAVE BEEN sporadic efforts over the past century to use electromagnetic waves as a means of wireless power transmission, but these attempts produced mixed results. Perhaps the golden year for research on wireless power transmission was 1975, when William Brown, who worked for Raytheon, and Richard Dickinson of NASA's Jet Propulsion Laboratory (now retired) used microwaves to beam power across a lab with greater than 50 percent end-to-end efficiency. In a separate demonstration, they were able to deliver more than 30 kilowatts over a distance of about a mile (1.6 kilometers).

These demonstrations were part of a larger NASA and U.S. Department of Energy campaign to explore the feasibility of solar-power satellites, which, it was proposed, would one day harvest sunlight in space and beam the energy down to Earth as microwaves. But because this line of research was motivated in large part by the energy crisis of the 1970s, interest in solar-power satellites waned in the following decades, at least in the United States.

Although researchers revisit the idea of solar-power satellites with some regularity, those performing actual demonstrations of power beaming have struggled to surpass the high-water mark for efficiency, distance, and power level reached in 1975. But that situation is starting to change, thanks to various recent advances in transmission and reception technologies.

Most early efforts to beam power were confined to microwave frequencies, the same part of the electromagnetic spectrum that today teems with Wi-Fi, Bluetooth, and various other wireless signals. That choice was, in part, driven by the simple fact that efficient microwave transmitting and receiving equipment was readily available.

But there have been improvements in efficiency and increased availability of devices that operate at much higher frequencies. Because of limitations imposed by the atmosphere on the effective transmission of energy within certain sections

PREVIOUS PAGES: POWERLIGHT TECHNOLOGIES

During a 2019 demonstration at the Naval Surface Warfare Center in Bethesda, Md., this laser beam safely conveyed 400 watts over a distance of 325 meters.

of the electromagnetic spectrum, researchers have focused on microwave, millimeter-wave, and optical frequencies. While microwave frequencies have a slight edge when it comes to efficiency, they require larger antennas. So, for many applications, millimeter-wave or optical links work better.

For systems that use microwaves and millimeter waves, the transmitters typically employ solid-state electronic amplifiers and phased-array, parabolic, or metamaterial antennas. The receiver for microwaves or millimeter waves uses an array of elements called rectennas. This word, a portmanteau of *rectifier* and *antenna*, reflects how each element converts the electromagnetic waves into direct-current electricity.

Any system designed for optical power transmission would likely use a laser—one with a tightly confined beam, such as a fiber laser. The receivers for optical power transmission are specialized photovoltaic cells designed to convert a single wavelength of light into electric power with very high efficiency. Indeed, efficiencies can exceed 70 percent, more than double that of a typical solar cell.

AT THE U.S. NAVAL Research Laboratory, we have spent the better part of the past 15 years looking into different options for power beaming and investigating potential applications. These include extending the flight times and payload capacities of drones, powering satellites in orbit when they are in darkness, powering rovers operating in permanently shadowed regions of the moon, sending energy to Earth's surface from space, and distributing energy to troops on the battlefield.

You might think that a device for sending large amounts of energy through the air in a narrow beam sounds like a death ray. This gets to the heart of a critical consideration: power density. Different power densities are technically possible, ranging from too low to be useful to high enough to be dangerous. But it's also possible to find a happy medium between these two extremes. And there are also clever ways to permit beams with high power densities to be used safely. That's exactly what a team I was part of did in 2019, and we've successfully extended this work since then.

One of our industry partners, PowerLight Technologies, formerly known as LaserMotive, has been developing laser-based power-beaming systems for more than a decade. Renowned for winning the NASA Power Beaming Challenge in 2009, this company has not only

To underscore how safe the system was, the host of the BBC science program "Bang Goes the Theory" stuck his face fully into a power beam.

U.S. NAVAL RESEARCH LABORATORY

The 400 watts we were able to transmit was, admittedly, not a huge amount, but it was sufficient to brew us some coffee.

achieved success in powering robotic tether climbers, quadcopters, and fixed-wing drones, but it has also delved deeply into the challenges of safely beaming power with lasers. That's key, because many research groups have demonstrated laser power beaming over the years—including teams at the Naval Research Laboratory, Kindai University, the Beijing Institute of Technology, the University of Colorado Boulder, JAXA, Airbus, and others—but only a few have accomplished it in a fashion that is truly safe under every plausible circumstance.

Perhaps the most dramatic demonstration of safe laser power beaming prior to our team's effort was by the company Lighthouse Dev in 2012. To underscore how safe the system was, the host of the BBC science program "Bang Goes the Theory" stuck his face fully into a power beam sent between buildings at the University of Maryland. This particular demonstration took advantage of the fact that some infrared wavelengths are an order of magnitude safer for your eyes than other parts of the infrared spectrum.

That strategy works for relatively low-power systems. But as you push the level higher, you soon get to power densities that raise safety concerns regardless of the wavelength used. What then? Here's where the system we've demonstrated sets itself apart. While sending more than 400 watts over a distance that exceeded 300 meters, the beam was contained within a virtual enclosure, one that could sense an object impinging on it and trigger the equipment to cut power to the main beam before any damage was done. Other testing has shown how transmission distances can exceed a kilometer.

Careful testing (for which no BBC science-program hosts were used) verified to our satisfaction the functionality of this feature, which also passed muster with the U.S. Navy's Laser Safety Review Board. During the course of our demonstration, the system further proved itself when, on several occasions, birds flew toward the beam, shutting it off—but only momentarily. You see, the system monitors the volume the beam occupies, along with its immediate surroundings, allowing the power link to automatically reestablish itself when the path is once again clear. Think of it as a more sophisticated version of a garage-door safety sensor, where the interruption of a guard beam triggers the motor driving the door to shut off.

For our demonstrations, observers in attendance were able to walk around between the transmitter and receiver without needing to wear laser-safety eyewear or take any other precautions. That's because, in addition to designing the system so that it can shut itself down automatically, we took care to consider the possible effects of reflections from the receiver or the scattering of light from particles suspended in the air along the path of the beam.

The 400 watts we were able to transmit was, admittedly, not a huge amount, but it was sufficient to brew us some

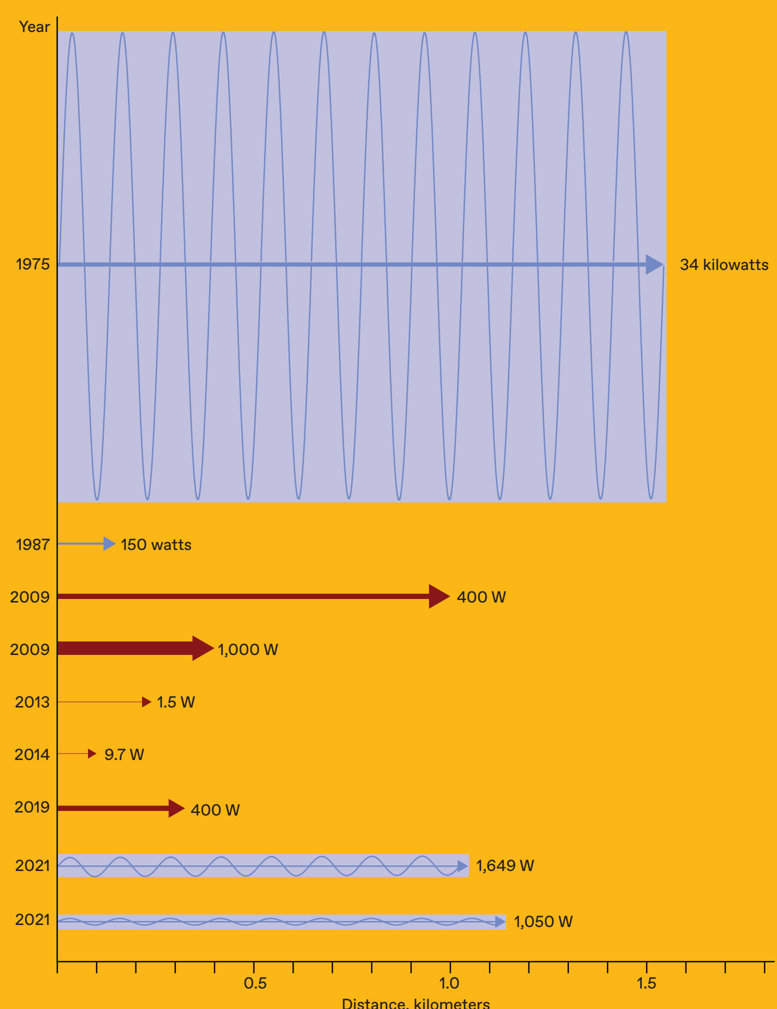

There have been many demonstrations of power beaming over the years, using either microwaves [blue] or lasers [red], with the peak-power record having been set in 1975 [top]. In 2021, the author and his colleagues took second and third place for the peak-power level achieved in such experiments, having beamed more than a kilowatt over distances that exceeded a kilometer, using much smaller antennas.

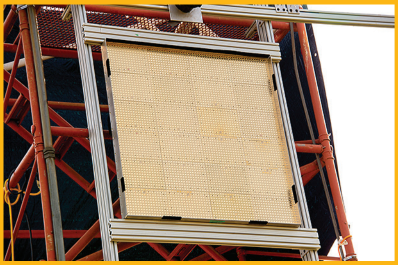

Last year, the author and his colleagues carried out a demonstration at the U.S. Army's Blossom Point test facility south of Washington, D.C. They used 9.7-gigahertz microwaves to send 1,649 watts (peak power) from a transmitter outfitted with a 5.4-meter-diameter parabolic dish [top] over a distance of 1,046 meters to a 2-by-2-meter "rectenna" [middle] mounted on a tower [right], which transformed the beam into usable electric power.

U.S. NAVAL RESEARCH LABORATORY

coffee, continuing what's become de rigueur in this line of experimentation: making a hot beverage. (The Japanese researchers who started this tradition in 2015 prepared themselves some tea.)

Our next goal is to apply power beaming, with fully integrated safety measures, to mobile platforms. For that, we expect to increase the distance covered and the amount of power delivered.

But we're not alone: Other governments, established companies, and startups around the world are working to develop their own power-beaming systems. Japan has long been a leader in microwave and laser power beaming, and China has closed the gap if not pulled ahead, as has South Korea.

At the consumer-electronics level, there are many players—Powercast, Ossia, Energous, GuRu, and Wi-Charge among them. And the multinational technology giant Huawei expects power beaming for smartphone charging within "two or three [phone] generations."

For industrial applications, companies like Reach Labs, TransferFi, MH GoPower, and MetaPower are making headway in employing power beaming to solve the thorny problem of keeping batteries for robots and sensors, in warehouses and elsewhere, topped off and ready to go. At the grid level, Emrod and others are attempting to scale power beaming to new heights.

On the R&D front, our team demonstrated within the past year safe microwave wireless power transmission of 1.6 kilowatts over a distance of a kilometer. Companies like II-VI Aerospace & Defense, Peraton Labs, Lighthouse Dev, and others have also recently made impressive strides. Today, ambitious startups like Solar Space Technologies, Solaren, Virtus Solis, and others operating in stealth mode are working hard to be the first to achieve practical power beaming from space to Earth.

As such companies establish proven track records for safety and make compelling arguments for the utility of their systems, we are likely to see whole new architectures emerge for sending power from place to place. Imagine drones that can fly for indefinite periods and electrical devices that never need to be plugged in—ever—and being able to provide people anywhere in the world with energy when hurricanes or other natural disasters ravage the local power grid. Reducing the need to transport fuel, batteries, or other forms of stored energy will have far-reaching consequences. It's not the only option when you can't string wires, but my colleagues and I expect, within the set of possible technologies for providing electricity to far-flung spots, that power beaming will, quite literally, shine. ■

HOW A PARACHUTE ACCIDENT HELPED JUMP-START AUGMENTED REALITY

IN 1992, HARDWARE FOR THE FIRST INTERACTIVE AR SYSTEM LITERALLY FELL FROM THE SKIES

BY LOUIS ROSENBERG

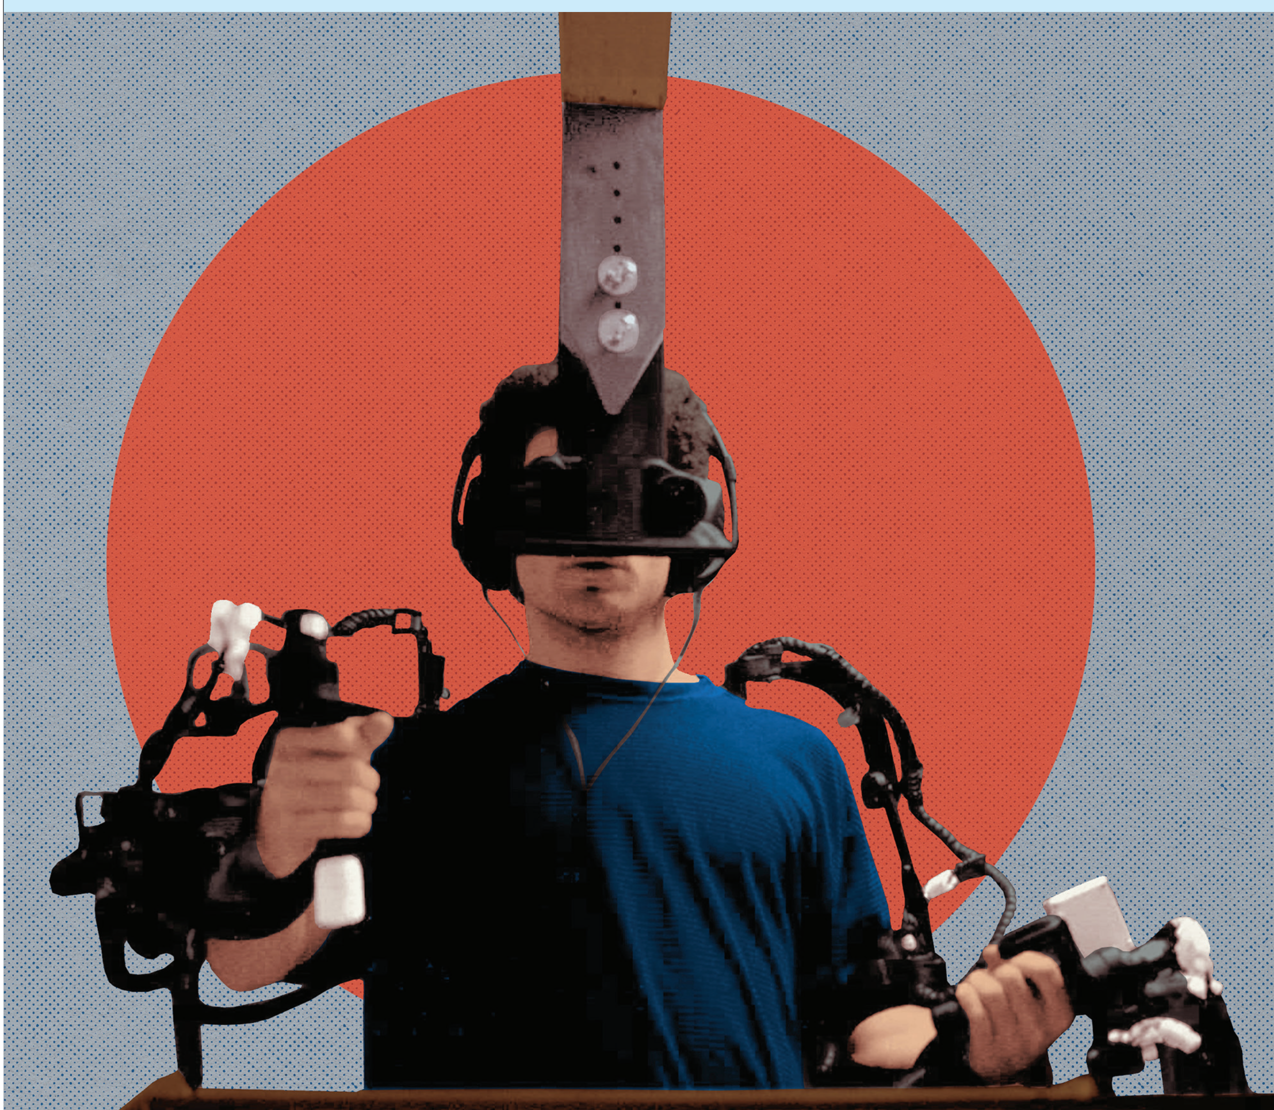

I climb into an upper-body exoskeleton that's covered in sensors, motors, gears, and bearings, and then lean forward, tilting my head up to press my face against the eyepieces of a vision system hanging from the ceiling. In front of me, I see a large wooden board, painted black and punctuated by a grid of metal holes. The board is real. So is the peg in my hand that I'm trying to move from one hole to another, as fast as I can. When I begin to move the peg, a virtual cone appears over the target hole, along with a virtual surface easing toward it. I can feel the surface as I slide the peg along it toward the cone and into the hole.

This was the Virtual Fixtures platform, which was developed in the early 1990s to test the potential of "perceptual overlays" to improve human performance in manual tasks that require dexterity. And it worked.

These days, virtual-reality experts look back on the platform as the first interactive augmented-reality system that enabled users to engage simultaneously with real and virtual objects in a single immersive reality.

The project began in 1991, when I pitched the effort as part of my doctoral research at Stanford University. By the time I finished—three years and multiple prototypes later—the system I had assembled filled half a room and used nearly a million dollars' worth of hardware. And I had collected enough data from human testing to definitively show that augmenting a real workspace with virtual objects could significantly enhance user performance in precision tasks.

Given the short time frame, it might sound like all went smoothly, but the project came close to getting derailed many times, thanks to a tight budget and substantial equipment needs. In fact, the effort might have crashed early on, had a parachute—a real one, not a virtual one—not failed to open in the clear blue skies over Dayton, Ohio, during the summer of 1992.

Before I explain how a parachute accident helped drive the development of augmented reality, I'll lay out a little of the historical context.

Thirty years ago, the field of virtual reality was in its infancy, the phrase itself having only been coined in 1987 by Jaron Lanier, who was commercializing some of the first headsets and gloves. His work built on earlier research by Ivan Sutherland, who pioneered head-mounted display technology and head-tracking, two critical elements that sparked the VR field. Augmented reality (AR)—that is, combining the real world and the virtual world into a single immersive and interactive reality—did not yet exist in a meaningful way.

Louis Rosenberg tests Virtual Fixtures [left], the first interactive augmented-reality system, which he developed at Wright-Patterson Air Force Base, in 1992.

Rosenberg [below] spent some of his time working in the Advanced Displays and Spatial Perception Laboratory of the Ames Research Center as part of his research in augmented reality.

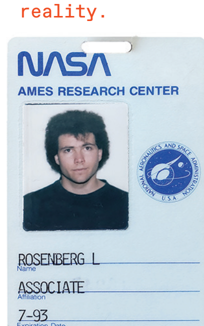

ORIGINAL PHOTOS: LOUIS ROSENBERG. PHOTO-MANIPULATION: GLUEKIT

Back then, I was a graduate student at Stanford University and a part-time researcher at NASA's Ames Research Center, interested in the creation of virtual worlds. At Stanford, I worked in the Center for Design Research, a group focused on the intersection of humans and technology that created some of the very early VR gloves, immersive vision systems, and 3D audio systems. At NASA, I worked in the Advanced Displays and Spatial Perception Laboratory of the Ames Research Center, where researchers were exploring the fundamental parameters required to enable realistic and immersive simulated worlds.

Of course, knowing how to create a quality VR experience and being able to produce it are not the same thing. The best PCs on the market back then used Intel 486 processors running at 33 megahertz. Adjusted for inflation, they cost about US $8,000 and weren't even a thousandth as fast as a cheap gaming computer today. The other option was to invest $60,000 in a Silicon Graphics workstation—still less than a hundredth as fast as a mediocre PC today. So, though researchers working in VR during the late '80s and early '90s were doing groundbreaking work, the crude graphics, bulky headsets, and lag so bad it made people dizzy or nauseous plagued the resulting virtual experiences.

I was conducting a research project at NASA to optimize depth perception in early 3D-vision systems, and I was one of those people getting dizzy from the lag. And I found that the images created back then were definitely virtual but far from reality.

Still, I wasn't discouraged by the dizziness or the low fidelity, because I was sure the hardware would steadily improve. Instead, I was concerned about how enclosed and isolated the VR experience made me feel. I wished I could expand the technology, taking the power of VR and unleashing it into the real world. I dreamed of creating a merged reality where virtual objects inhabited your physical surroundings in such an authentic manner that they seemed like genuine parts of the world around you, enabling you to reach out and interact as if they were actually there.

I was aware of one very basic sort of merged reality—the head-up display in use by military pilots, enabling flight data to appear in their lines of sight so they didn't have to look down at cockpit gauges. I hadn't experienced such a display myself, but became familiar with them thanks to a few blockbuster 1980s hit movies, including *Top Gun* and *Terminator*. In *Top Gun* a glowing cross hair appeared on a glass panel in front of the pilot during dogfights; in *Terminator*, cross hairs joined text and numerical data as part of the fictional cyborg's view of the world around it.

Neither of these merged realities was the slightest bit immersive, presenting images on a flat plane rather than connected to the real world in 3D space.

But they hinted at interesting possibilities. I thought I could move far beyond simple cross hairs and text on a flat plane to create virtual objects that could be spatially registered to real objects in an ordinary environment. And I hoped to instill those virtual objects with realistic physical properties.

I needed substantial resources—beyond what I had access to at Stanford and NASA—to pursue this vision. So I pitched the concept to the Human Sensory Feedback Group of the U.S. Air Force's Armstrong Laboratory, now part of the Air Force Research Laboratory.

To explain the practical value of merging real and virtual worlds, I used the analogy of a simple metal ruler. If you want to draw a straight line in the real world, you can do it freehand, going slow and using significant mental effort, and it still won't be particularly straight. Or you can grab a ruler and do it much quicker with far less mental effort. Now imagine that instead of a real ruler, you could grab a virtual ruler and make it instantly appear in the real world, perfectly registered to your real surroundings. And imagine that this virtual ruler feels physically authentic—so much so that you can use it to guide your real pencil. Because it's virtual, it can be any shape and size, with interesting and useful properties that you could never achieve with a metal straightedge.

Of course, the ruler was just an analogy. The applications I pitched to the Air Force ranged from augmented manufacturing to surgery. For example, consider a surgeon who needs to make a dangerous incision. She could use a bulky metal fixture to steady her hand and avoid vital organs. Or we could invent something new to augment the surgery—a virtual fixture to guide her real scalpel, not just visually but physically. Because it's virtual, such a fixture would pass right through the patient's body, sinking into tissue before a single cut had been made. That was the concept that got the military excited, and their interest wasn't just for in-person tasks like surgery but for distant tasks performed using remotely controlled robots. For example, a technician on Earth could repair a satellite by controlling a robot remotely, assisted by virtual fixtures added to video images of the real worksite. The Air Force agreed to provide enough funding to cover my expenses at Stanford along with a small budget for equipment. Perhaps more significantly, I also got access to computers and other equipment at Wright-Patterson Air Force Base near Dayton, Ohio.

And what became known as the Virtual Fixtures Project came to life, with the goal of building a prototype that could be rigorously tested with human subjects. And I became a roving researcher, developing core ideas at Stanford, fleshing out some of the underlying technologies at NASA Ames, and assembling the full system at Wright-Patterson.

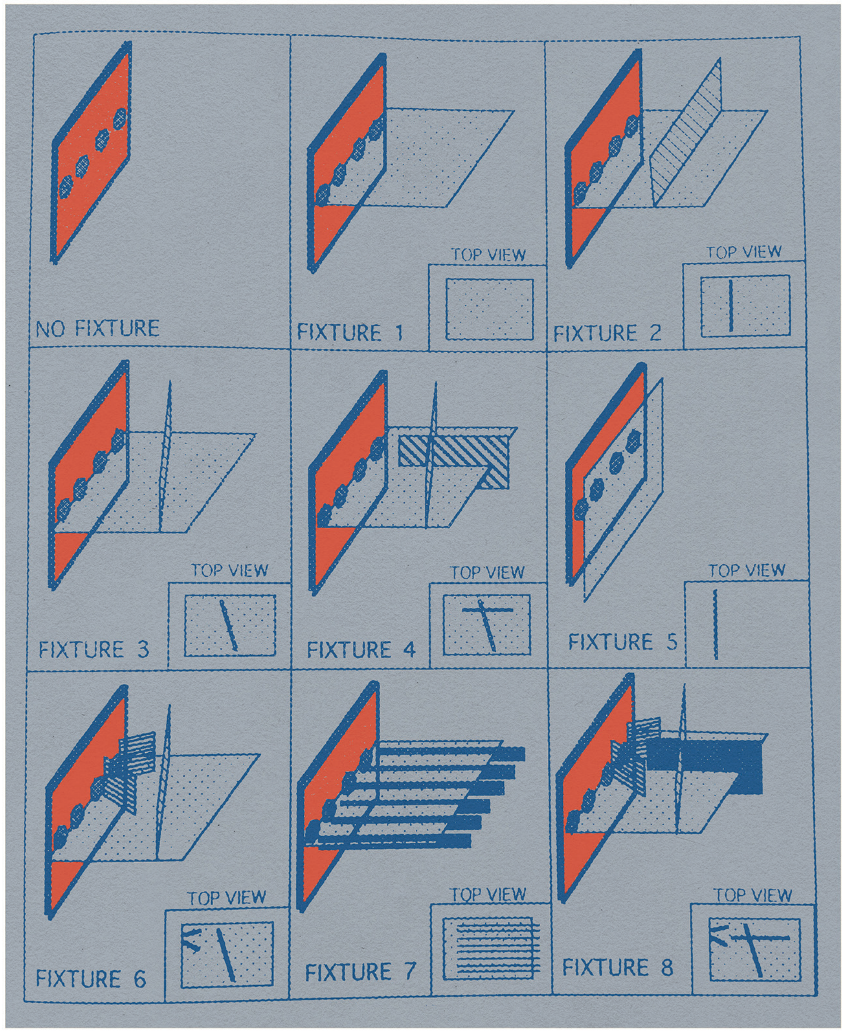

These early drawings of a real pegboard combined with virtual overlays generated by a computer—an early version of augmented reality—were created by Louis Rosenberg as part of his Virtual Fixtures project.

Now about those parachutes.

As a young researcher in my early twenties, I was eager to learn about the many projects going on around me at these various laboratories. One effort I followed closely at Wright-Patterson was a project designing new parachutes. As you might expect, when the research team came up with a new design, they didn't just strap a person in and test it. Instead,

The Fitts's Law peg-insertion task involves having test subjects quickly move metal pegs between holes. The pegboard shown here [left] was real, while the cones that helped guide the user to the correct holes were virtual. In his sketch of his augmented-reality system [right], Louis Rosenberg shows a user of the Virtual Fixtures platform wearing a partial exoskeleton and peering at a real pegboard augmented with cone-shaped virtual fixtures.

they attached the parachutes to dummy rigs fitted with sensors and instrumentation. Two engineers would go up in an airplane with the hardware, dropping rigs and jumping alongside so they could observe how the chutes unfolded. Stick with my story and you'll see how this became key to the development of that early AR system.

Back at the Virtual Fixtures effort, I aimed to prove the basic concept—that a real workspace could be augmented with virtual objects that feel so real, they could assist users as they performed dexterous manual tasks. To test the idea, I wasn't going to have users perform surgery or repair satellites. Instead, I needed a simple repeatable task to quantify manual performance. The Air Force already had a standardized task it had used for years to test human dexterity under a variety of mental and physical stresses. It's called the Fitts's Law peg-insertion task and it involves having test subjects quickly move metal pegs between holes on a large pegboard.

So I began assembling a system that would enable virtual fixtures to be merged with a real pegboard, creating a mixed-reality experience perfectly registered in 3D space. I aimed to make these virtual objects feel so real that bumping the real peg into a virtual fixture would feel as authentic as bumping into the actual board.

I wrote software to simulate a wide range of virtual fixtures, from simple surfaces that prevented your hand from overshooting a target hole, to carefully shaped cones that could help a user guide the real peg into the real hole. I created virtual overlays that simulated textures and had corresponding sounds, even overlays that simulated pushing through a thick liquid as if it were virtual honey.

Of the various applications for Virtual Fixtures that we considered at the time, the most commercially viable back then involved manually controlling robots.

For more realism, I modeled the physics of each virtual element, registering its location accurately in three dimensions so it lined up with the user's perception of the real wooden board. Then, when the user moved a hand into an area corresponding to a virtual surface, motors in the exoskeleton would physically push back, an interface technology now commonly called haptics. It indeed felt so authentic that you could slide along the edge of a virtual surface the way you might move a pencil against a real ruler.

To accurately align these virtual elements with the real pegboard, I needed high-quality video cameras. Video cameras at the time were far more expensive than they are today, and I had no money left in my budget to buy them. This was a frustrating barrier: The Air Force had given me access to a wide range of amazing hardware, but when it came to simple cameras, they couldn't help. It seemed like every research project needed them, most of far higher priority than mine.

Which brings me back to the skydiving engineers testing experimental parachutes. These engineers came into the lab one day to chat; they mentioned that their chute had failed to open, their dummy rig plummeting to the ground and destroying all the sensors and cameras aboard.

This seemed like it would be a setback for my project as well, because I knew if there were any extra cameras in the building, the engineers would get them.

But then I asked if I could take a look at the wreckage from their failed test. It was a mangled mess of bent metal, dangling circuits, and smashed cameras. Still, though the cameras looked awful with cracked cases and damaged lenses, I wondered if I could get any of them to work well enough for my needs.

By some miracle, I was able to piece together two working units from the six that had plummeted to the ground. And so, the first human testing of an interactive augmented-reality system was made possible by cameras that had literally fallen out of the sky and smashed into the earth.

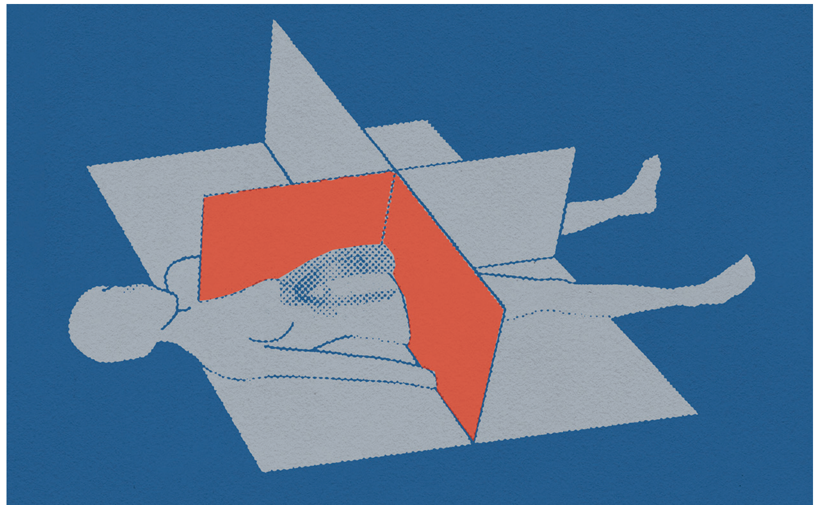

One imagined use for augmented reality at the time of its creation was in surgery. Today, augmented reality is routinely used for surgical training, and surgeons are beginning to use it in the operating room.

To appreciate how important these cameras were to the system, think of a simple AR application today, like *Pokémon Go*. If you didn't have a camera on the back of your phone to capture and display the real world in real time, it wouldn't be an augmented-reality experience; it would just be a standard video game.

The same was true for the Virtual Fixtures system. But thanks to the cameras from that failed parachute rig, I was able to create a mixed reality with accurate spatial registration, providing an immersive experience in which you could reach out and interact with the real and virtual environments simultaneously.

As for the experimental part of the project, I conducted a series of human studies in which users experienced a variety of virtual fixtures overlaid onto their perception of the real task board. The most useful fixtures turned out to be cones and surfaces that could guide the user's hand as they aimed the peg toward a hole. The most effective involved physical experiences that couldn't be easily manufactured in the real world but were readily achievable virtually. For example, I coded virtual surfaces that were "magnetically attractive" to the peg. For the users, it felt as if the peg had snapped to the surface. Then they could glide along it until they chose to yank free with another snap. Such fixtures increased speed and dexterity in the trials by more than 100 percent.

Of the various applications for Virtual Fixtures that we considered at the time, the most commercially viable back then involved manually controlling robots in remote or dangerous environments—for example, during hazardous waste cleanup. If the communications distance introduced a time delay in the telerobotic control, virtual fixtures became even more valuable for enhancing human dexterity.

Today, researchers are still exploring the use of virtual fixtures for telerobotic applications with great success, including in satellite repair and robot-assisted surgery.

I went in a different direction, pushing for more mainstream applications for augmented reality. That's because the part of the Virtual Fixtures project that had the greatest impact on me personally wasn't the improved performance in the peg-insertion task. Instead, it was the big smiles that lit up the faces of the human subjects when they climbed out of the system and effused about what a remarkable experience they'd had. Many told me, without prompting, that this type of technology would one day be everywhere.

I agreed with them. I was convinced we'd see this type of immersive technology go mainstream by the end of the 1990s. In fact, I was so inspired by the enthusiastic reactions people had when they tried those early prototypes, I founded a company in 1993—Immersion—with the goal of pursuing mainstream consumer applications. Of course, it hasn't happened nearly that fast.

At the risk of being wrong again, I sincerely believe that virtual and augmented reality, now commonly referred to as the metaverse, will become an important part of most people's lives by the end of the 2020s. Based on the recent surge of investment by major corporations into improving the technology, I predict that by the early 2030s augmented reality will replace the mobile phone as our primary interface to digital content.

And no, none of the test subjects who experienced that early glimpse of augmented reality 30 years ago knew they were using hardware that had fallen out of an airplane. But they did know that they were among the first to reach out and touch our augmented future. ■

THE INSTITUTE

NEWS OF THE IEEE
VOLUME 46 / ISSUE 2

AMD's Lisa Su Breaks Through the Silicon Ceiling P. 50

Find Volunteers Faster with IEEE's Recruitment Website P. 52

GeeRemit App Moves Money and Saves It Too P. 61

A High Ideal

DURING MY TRIP to Mexico City this year for the IEEE Region 9 annual meeting, I had the chance to visit the Memory and Tolerance Museum, whose mission is to increase the understanding of cultural diversity and eradicate hatred and genocide around the world. It was a humbling experience to see exhibits of the horrible events against humanity, which continue to happen in almost every corner of the world. It made me wonder: Why do such atrocities still happen?

Could it be that because we speak different languages, we don't understand each other; we follow different cultural and social norms, so we behave differently; we have different physical appearances, so we do not look like each other? These differences create perceived distance between us.

IEEE's mission is to bring together individuals with a common interest in advancing technology for the benefit of humanity. The mission helps IEEE build its international community of technologists—regardless of nationality—to encourage friendship, support peaceful cooperation, and further the exchange of ideas. Our activity in IEEE demonstrates the ability of technologists to work together for the greater global public good without boundaries.

Maintaining high ideals

This is a high ideal, and IEEE's apolitical, neutral stance allows us to engage with our members in every part of the world—those living in democracies, monarchies, and dictatorships as well as in countries undergoing civil upheaval, and in nations in conflict. IEEE does not suspend our commitment to support this global scientific community in challenging times. Instead, we aspire to strengthen relationships within our international technical community to nurture peaceful coexistence and collaboration and to continue to build relations among individuals across national boundaries, ideologies, and politics that may help foster a more peaceful world.

IEEE's relationship is with our members—not their countries, employers, or institutions. Across IEEE, we value our members as individuals and recognize and celebrate their technical efforts and contributions to the field. We are open, diverse, and inclusive, and we hold our fellow members personally accountable for their behavior, and not for the actions of others.

At different times throughout our organization's history, there have been calls for IEEE to expel members due to the behavior of governments, institutions, or employers, or to limit access to IEEE resources, journals, and events based on nationality or work affiliation. Nothing in IEEE's bylaws or policies permits this to occur. In fact, our core policies clearly assert that we do not discriminate based on nationality, race, religion, gender, disability, age, or sexual orientation.

IEEE's continued ability to engage researchers and members in so many countries is contingent on IEEE's continual recognition as an apolitical organization that maintains its neutrality on matters outside its technical scope.

Protecting freedoms

As an international organization operating in more than 160 countries, IEEE supports the free and open exchange of scholarly and academic work and the global advancement of science and technology. IEEE is committed to enabling an environment of international cooperation and the

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

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

IEEE

sharing of our members' wealth of knowledge to drive innovation. Science and engineering are fundamental enterprises for which openness, international collaboration, the free flow of ideas, and the involvement of all talented individuals are essential for ongoing advancement. Science, engineering, and humanity prosper when there is freedom of association and communication.

Significant work still needs to be done to ensure that scientists and engineers have the right to pursue their careers without discrimination. IEEE is committed to the realization and maintenance of such freedoms.

These freedoms are challenged in many countries. As I write this column, the ongoing invasion of Ukraine and the ensuing atrocities and humanitarian crisis are assaults on the conscience. Around the world, other countries are also in conflict stemming from political tensions, drought, disease, and more.

When leaving the museum in Mexico City, a thought kept coming back to me: To meet IEEE's mission to advance technology for the benefit of humanity, we must bring together our members who share the common fundamental value of promoting worldwide peace.

The world will be a better place only when everywhere is a better place.

A sense of belonging

IEEE has been my professional home for more than 37 years. My membership in IEEE has fostered within me a sense of belonging and made me feel like I was connected to something bigger than myself.

I hope that IEEE can continue to offer you and all technical professionals around the world the same opportunities that it has offered me: the free exchange of ideas, learning and skill-building, career advancement opportunities and leadership training and, most importantly, lifelong friendships and networking with like-minded colleagues. After all, IEEE is our professional home!

—K.J. RAY LIU
IEEE president and CEO

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

EDITOR'S NOTE

IEEE Fellows Who Broke Barriers

The challenges they overcame shaped their careers

I OFTEN MARVEL at the fantastic problem-solving skills of technologists.

Consider the remarkable people in this issue: a pioneering woman who's rebuilding one of the world's leading semiconductor companies, a computational neuroscientist who's devising a better AI model, engineers who refused to give up on their idea for a universal plug-and-play communication system, and a fintech startup founder who figured out how to let users of her money-transfer app put aside savings. Each one, confronted with a significant challenge, dove right in and overcame it.

As CEO, president, and chairwoman of Advanced Micro Devices, Lisa Su is doing what few women in tech have had the opportunity to do: lead a large semiconductor company. She brought AMD back from the brink of bankruptcy and is transforming it into a leading, high-performance company.

Su, an IEEE Fellow, is the first woman to receive one of the semiconductor industry's most prestigious honors, the IEEE Robert N. Noyce Medal. Read about her climb up the leadership ladder on page 50.

Stephen Grossberg is another member trying to transform an industry: artificial intelligence. The IEEE Fellow wants to help companies that use AI solve problems associated with deep learning, which can be unsuitable for some applications because it can experience catastrophic forgetting.

On page 54, Grossberg describes an alternative—adaptive resonance theory—for both biological and artificial intelligence based on cognitive and neural research he has conducted.

Today we can easily connect smartphones, printers, game controllers, and other external devices to our computers thanks to the Universal Serial Bus standard [page 58]. Released in 1996 by Intel, USB was recently commemorated with an IEEE Milestone. Ajay Bhatt, one of the engineers who helped develop it, said in an interview with CNN that the technology had exceeded his wildest of imaginations.

Mobile apps have become popular ways for people in the United States to send money to aid those in sub-Saharan Africa. But many of the apps don't have a separate account for savings. They also lack a way to send money to entities with no cellphone number, such as philanthropic institutions. On page 61, find out how IEEE Fellow Sandra K. Johnson addressed the issues with her GeeRemit app. She says its features allow users to help transform lives.

Johnson isn't the only one taking on challenges related to a digital product and service. IEEE is working to transform how digital-service providers design products and services for children [page 56], because protecting their online privacy is becoming a top priority around the world. Its new Standard for an Age Appropriate Digital Services Framework Based on the 5Rights Principles for Children provides practical steps that project teams, suppliers, process assessors, and others can take to ensure their offerings are age-appropriate.

—KATHY PRETZ
Editor in chief, *The Institute*

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

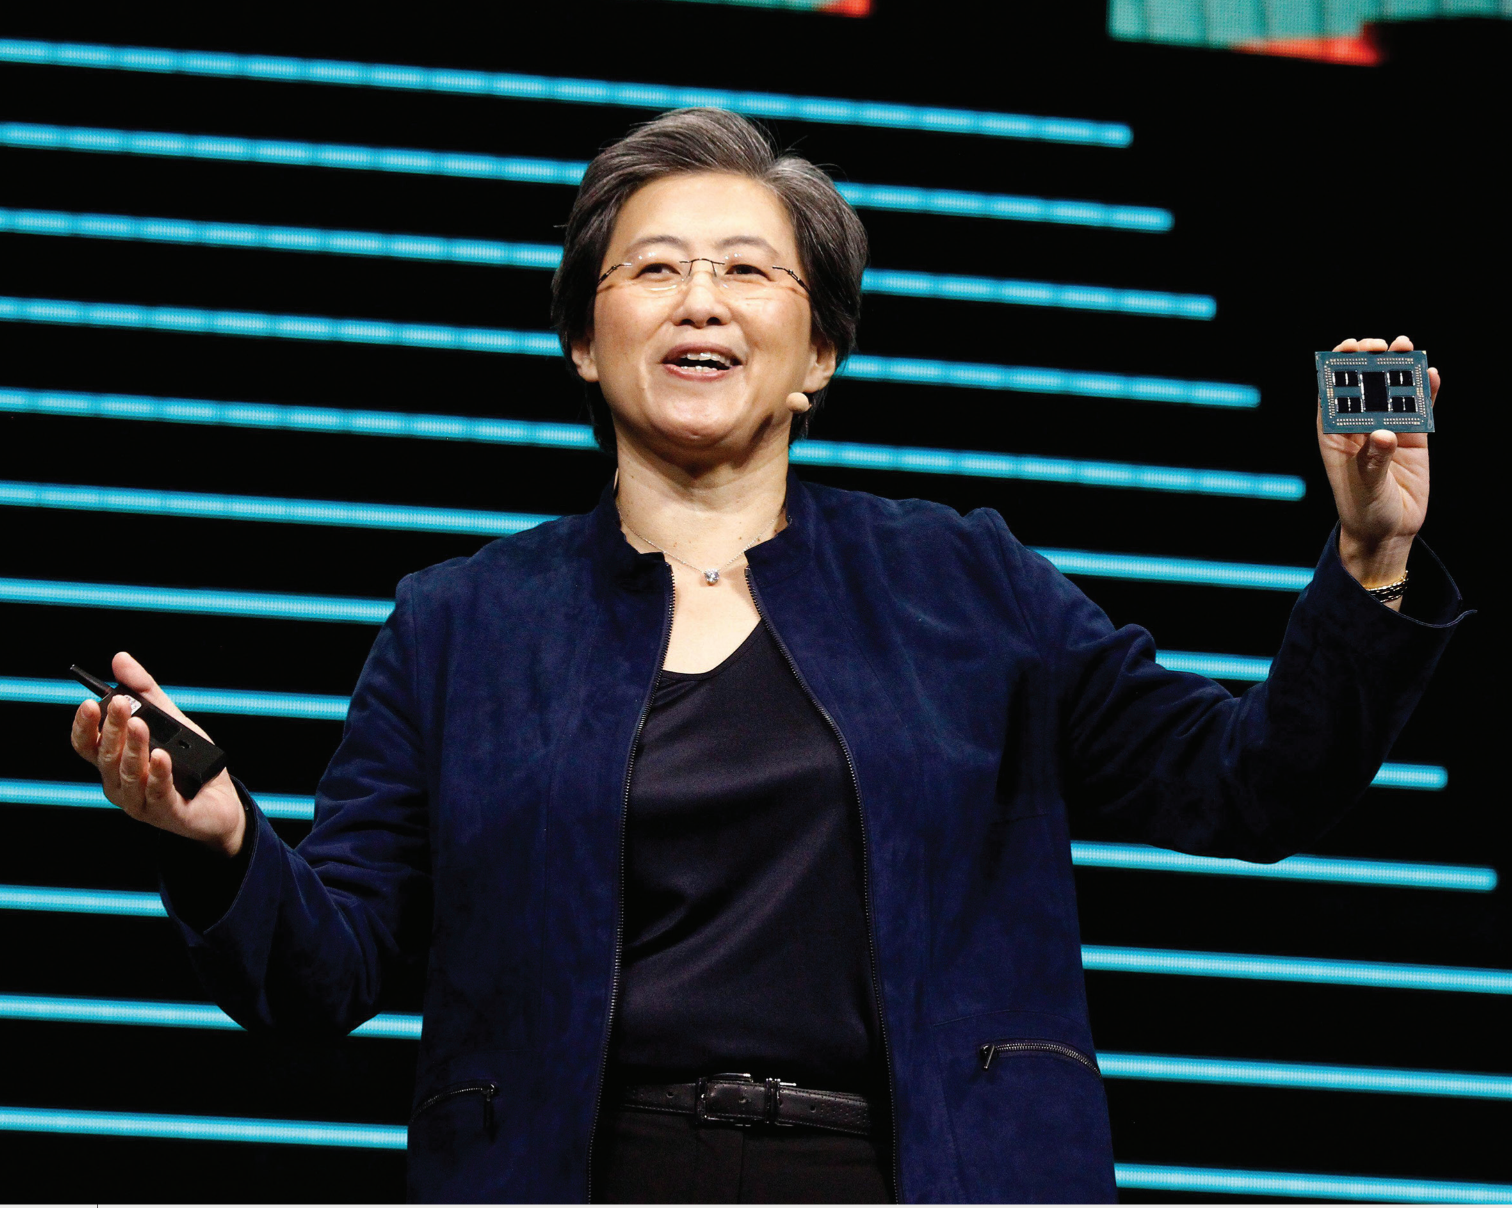

IEEE Fellow Lisa Su unveiled AMD's new 64-core Threadripper 3990WX chip at the 2020 International Consumer Electronics Show in Las Vegas.

ALAMY

PROFILE

Lisa Su Breaks Through the Silicon Ceiling

AMD's CEO is the first woman to receive IEEE's highest semiconductor award

BY KATHY PRETZ

WHEN LISA SU became CEO of Advanced Micro Devices in 2014, the company was on the brink of bankruptcy. Since then, AMD's stock has soared—from less than US $2 per share to more than $110. The company is now a leader in high-performance computing.

Su has received many accolades for spearheading AMD's turnaround. She added another honor in 2021: the IEEE Robert N. Noyce Medal. Su is the first woman to receive the award, which recognizes her "leadership in groundbreaking semiconductor products and successful business strategies that contributed to the strength of the microelectronics industry." Sponsored by Intel, the Noyce Medal is considered to be one of the semiconductor industry's most prestigious honors.

"To be honest, I would have never imagined that I would receive the Noyce award," the IEEE Fellow says. "It's an honor of a lifetime. To have that recognition from my peers in the technical community is a humbling experience. But I love what I do and being able to contribute to the semiconductor industry."

Employer Advanced Micro Devices
Title CEO
Member grade Fellow
Alma mater MIT

Climbing the leadership ladder

Su decided to study electrical engineering, she says, because she was drawn to the prospect of building hardware.

"I felt like I was actually building and making things," she says.

"It might surprise people that my parents would have preferred that I became a medical doctor," she says, laughing. "That was the most well-respected profession when I was growing up. But I never really liked the sight of blood. I ended up getting a Ph.D., which I guess was the next best thing."

Su has spent most of her career working on semiconductor projects for large companies. Along the way, she evolved from researcher to manager to top executive. Looking back, Su divides her career path into two parts. The first 20 or so years she was involved in R&D; for the past 16 years, she has worked on the business side.

> "I really do believe that you can be trained to be a good leader."

Her first job was with Texas Instruments, in Dallas, where she was a member of the technical staff at the company's semiconductor process and device center. She joined in 1994, but after a year, she left for IBM, in New York. There, she was a staff member researching device physics. In 2000 she was assigned to be the technical assistant for IBM's chief executive. She later was promoted to director of emerging projects.

She made the switch to management in 2006, when she was appointed vice president of IBM's semiconductor research and development center in New York.

Su says she doesn't agree with the notion that leadership is an innate ability.

"I really do believe that you can be trained to be a good leader," she says. "As engineers transition into business or management, you have to think about a different set of challenges that are not necessarily 'How do you make your transistor go faster?' but [instead] 'How do you motivate teams?' or 'How do you understand more about what customers want?' I've made my share of mistakes in those transitions, but I've also learned a lot."

One of the first places she got a chance to put her training into action was at Freescale Semiconductor, in Austin, Texas. In 2007 she took over as CTO and oversaw the company's R&D efforts. She was promoted to senior vice president and general manager of Freescale's networking and multimedia group. In that role, she was responsible for global strategy, marketing, and engineering for the embedded communications and applications processor business.

She left in 2012 to join AMD, also in Austin, as senior vice president, overseeing the company's global business units. Two years later she was appointed president and CEO, the first woman to run a Fortune 500 semiconductor company.

It took more than leadership skills to get to the top, she says.

"It's a little bit of you have to be good [at what you do], but you also have to be lucky and be in the right place at the right time," she says. "I was fortunate in that I had a lot of opportunities throughout my career."

AMD's business is booming, and Su is credited with expanding the market for the company's chips beyond PCs to game consoles and embedded devices. AMD released products in 2017 with its Ryzen desktop processors and Epyc server processors for data centers. They are based on its Zen microarchitecture, which enabled the chips to quickly process more instructions than the competition. The Radeon line of graphics cards for gaming consoles debuted in 2000.

AMD is currently focused on building the next generation of supercomputers—which Su says will be "important in many aspects of research going forward."

In 2020 the company announced its advanced CPUs, GPUs, and software will be powering Lawrence Livermore National Laboratory's El Capitan exascale-class supercomputer. Predicted to be the world's fastest when it goes into service in 2023, El Capitan is expected to expand the use of artificial intelligence and machine learning.

Important association

Su joined IEEE while a student.

"I think IEEE is still the foremost organization for bringing researchers together to share their findings, to network, and to develop and build relationships," she says. "I've met many people through my IEEE connections, and they continue to be close colleagues. It's just a great organization to move the industry forward." ■

IEEE PRODUCTS

Volunteering Platform Gets New Features

BY JOANNA GOODRICH

SEVERAL FEATURES HAVE been added to the IEEE Volunteering Platform. The website enables members to search for opportunities across the organization, be it short- or long-term, local or remote. Those looking for helpers can post positions they need to fill.

Although the IEEE Young Professionals group leads the development of the platform, the website helps all IEEE groups increase their volunteer pipeline.

"We believe volunteering is a member benefit," says Member Emre Ayranci, the IEEE Young Professionals committee chair. "We are happy to help IEEE members find volunteering opportunities as well as enable volunteer leaders to have access to more volunteers who would like to contribute to their projects and initiatives."

New features

- All members can now post an opportunity, whereas before, only members with a leadership position had that ability.
- Those posting an opportunity can now choose how many hours per week, month, or year it should take to complete a project.
- Those looking for volunteers now can designate a so-called co-owner to provide backup on the project as well as assist with such tasks as recruiting people and communicating with the volunteers. The names of both the owner and co-owner are displayed.
- The applicants' names now appear in a sorted list under the opportunity on the creator's and co-owner's dashboard. The list is broken down into applicants who have applied to the opportunity, those who were invited to apply, and those who declined the invitation.
- Creators and co-owners can use the platform to email all applicants at once as well as download every applicant's contact information.
- To make volunteering opportunities more visible, they can now be posted from the platform directly to Facebook, LinkedIn, and Twitter.

IEEE NEWS

Humanitarian SIGHT Program Reports Record Growth

The IEEE Humanitarian Activities Committee's Special Interest Group on Humanitarian Technology program focuses on raising awareness of how technically trained people can contribute to society. Last year it celebrated reaching several milestones. Here's a look at two of them.

MEMBERSHIP GROWTH

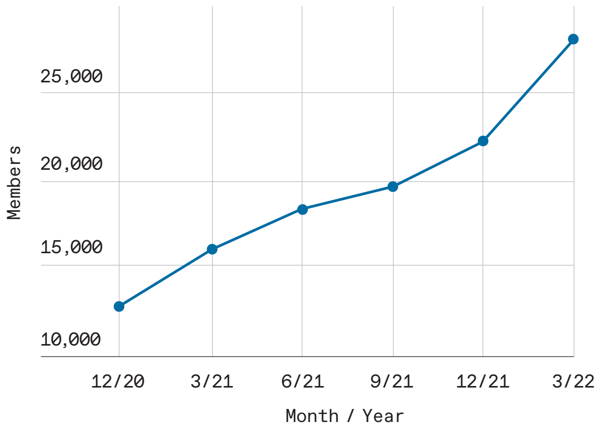

PROJECT PROPOSAL SUBMISSIONS

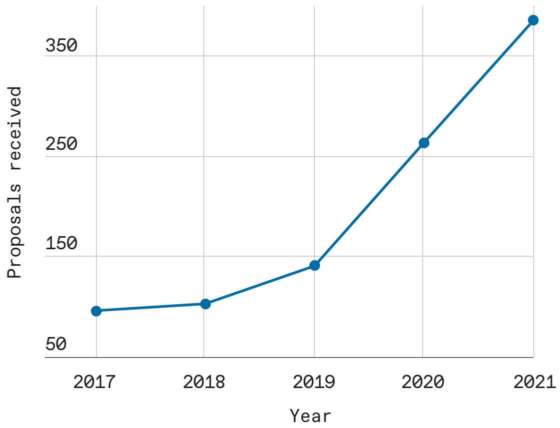

SOURCE: THE IEEE HUMANITARIAN ACTIVITIES COMMITTEE

ARTIFICIAL INTELLIGENCE

Don't Trust Deep Learning, Says Brain Modeler

Stephen Grossberg maintains his approach to intelligence is better

BY KATHY PRETZ

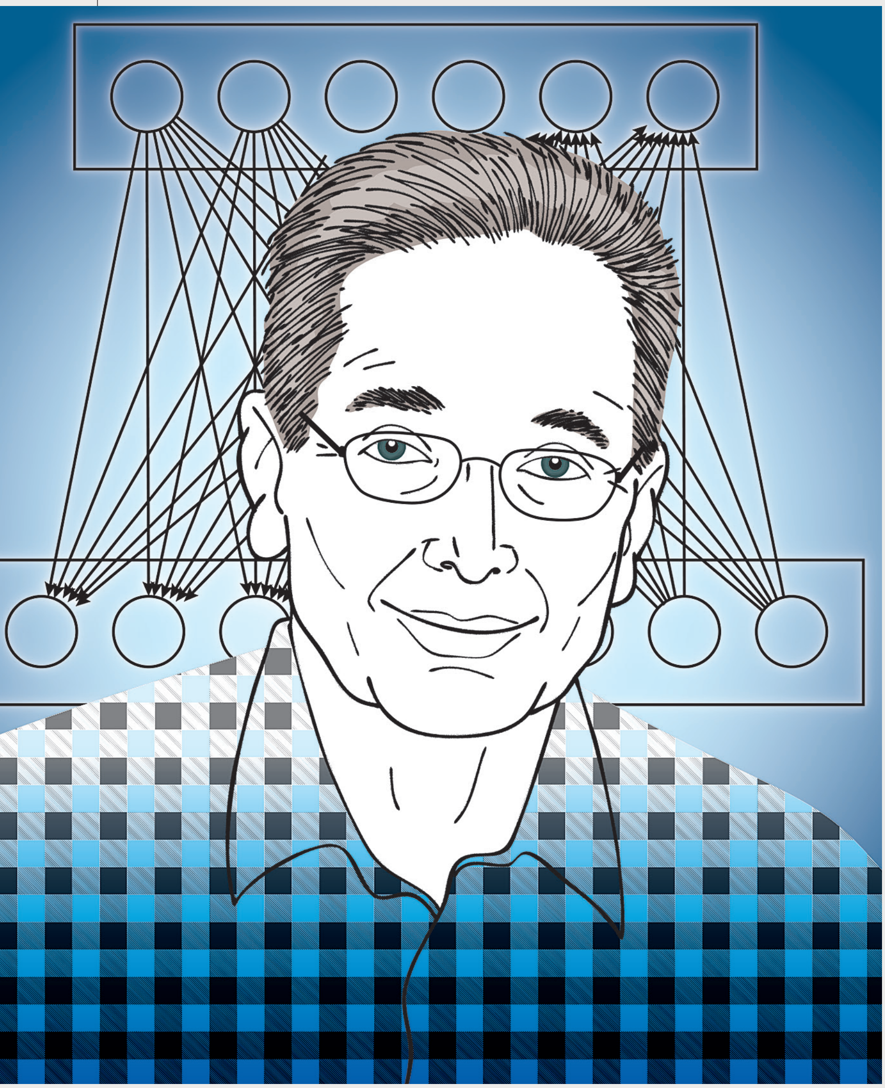

DURING THE PAST 20 years, deep learning has come to dominate artificial intelligence research and programs through a series of useful commercial applications. But underneath the dazzle are some deep-rooted problems that threaten the technology's ascension.

The inability of a typical deep learning program to perform well on more than one task, for example, severely limits application of the technology to specific tasks in rigidly controlled environments. More seriously, it has been claimed that deep learning is untrustworthy because it is not explainable—and unsuitable for some applications because it can experience catastrophic forgetting.

Said more plainly, if the algorithm does work, it may be impossible to fully understand why. And while the tool is slowly learning a new database, an arbitrary part of its learned memories can suddenly collapse. It might therefore be risky to use deep learning on any life-or-death application such as a medical one.

In his new book, IEEE Fellow Stephen Grossberg argues that an entirely different approach is needed. *Conscious Mind, Resonant Brain: How Each Brain Makes a Mind* describes an alternative model for both biological and artificial intelligence based on cognitive and neural research Grossberg has been conducting for

Photo-illustration by MCKIBILLO

decades. He calls his model Adaptive Resonance Theory, or ART.

Grossberg—an endowed professor of cognitive and neural systems, and of mathematics and statistics, psychological and brain sciences, and biomedical engineering at Boston University—based ART on his theories about how the brain processes information.

"Our brains learn to recognize and predict objects and events in a changing world that is filled with unexpected events," he says.

Based on that dynamic, ART uses supervised and unsupervised learning methods to solve such problems as pattern recognition and prediction. Algorithms using the theory have been included in large-scale applications such as classifying sonar and radar signals, detecting sleep apnea, recommending movies, and computer-vision-based driver-assistance software.

ART can be used with confidence because it is explainable and does not experience catastrophic forgetting, Grossberg says. He adds that ART solves what he has called the *stability-plasticity dilemma*: how a brain or other learning system can autonomously learn quickly (plasticity) without experiencing catastrophic forgetting (stability).

Grossberg, who formulated ART in 1976, is a pioneer in modeling how brains become intelligent. He is the founder and director of Boston University's Center for Adaptive Systems and the founding director of the Center of Excellence for Learning in Education, Science, and Technology. Both centers have sought to understand how the brain adapts and learns, and to develop technological applications based on their findings.

Grossberg attempts to explain in his nearly 800-page book how "the small lump of meat that we call a brain" gives rise to thoughts, feelings, hopes, sensations, and plans. In particular, he describes biological neural models that attempt to explain how that happens. The book also covers the underlying causes of conditions such as Alzheimer's disease, autism, amnesia, and post-traumatic stress disorder.

> **Understanding how brains give rise to minds is also important for designing smart systems in computer science, engineering, and tech, including AI and smart robots.**

"Understanding how brains give rise to minds is also important for designing smart systems in computer science, engineering and tech, including AI and smart robots," he writes. "Many companies have applied biologically inspired algorithms of the kind that this book summarizes in multiple engineering and technological applications."

The theories in the book, he says, are not only useful for understanding the brain but also can be applied to the design of intelligent systems that are capable of autonomously adapting to a changing world. Taken together, the book describes the fundamental process that enables people to be intelligent, autonomous, and versatile.

The beauty of ART

Grossberg writes that the brain evolves to adapt to new challenges. There is a common set of brain mechanisms that control how humans retain information without forgetting what they have already learned, he says.

"We retain stable memories of past experiences, and these sequences of events are stored in our working memories to help predict our future behaviors," he says. "Humans have the ability to continue to learn throughout their lives, without new learning washing away memories of important information that we learned before."

One of the problems faced by classical AI, he says, is that it often built its models on how the brain *might* work, using concepts and operations that could be derived from introspection and common sense.

"Such an approach assumes that you can introspect internal states of the brain with concepts and words people use to describe objects and actions in their daily lives," he writes. "It is an appealing approach, but its results were all too often insufficient to build a model of how the biological brain really works."

The problem with today's AI, he says, is that it tries to imitate the results of brain processing instead of probing the mechanisms that give rise to the results. People's behaviors adapt to new situations and sensations "on the fly," Grossberg says, thanks to specialized circuits in the brain.

People can learn from new situations, he adds, and unexpected events are integrated into their collected knowledge and expectations about the world.

ART's networks are derived from thought experiments on how people and animals interact with their environment, he adds: "ART circuits emerge as computational solutions of multiple environmental constraints to which humans and other terrestrial animals have successfully adapted." This fact suggests that ART designs might in some form be embodied in all future autonomous adaptive intelligent devices, whether biological or artificial.

"The future of technology and AI will depend increasingly on such self-regulating systems," Grossberg concludes. "It is already happening with efforts such as designing autonomous cars and airplanes. It's exciting to think about how much more may be achieved when deeper insights about brain designs are incorporated into highly funded industrial research and applications." ■

IEEE STANDARDS

Standard Seeks to Protect Kids in Cyberspace

BY KATHY PRETZ

TO HELP DIGITAL-SERVICE providers do a better job of designing products and services with children in mind, the IEEE Standards Association (IEEE SA) recently published IEEE 2089-2021.

The Age Appropriate Digital Services Framework Working Group developed the standard under the auspices of the emerging technology standards committee of the IEEE Consumer Technology Society.

The IEEE Standard for an Age Appropriate Digital Services Framework Based on the 5Rights Principles for Children provides practical steps that project teams, suppliers, process assessors, and others can take to ensure their online products and services are safer for minors.

The 5Rights Foundation established the principles as part of its mission to make "systematic changes to the digital world to ensure it caters [to] children and young people, by design and default."

The five principles are: presenting information in an age-appropriate way, upholding children's rights, offering fair terms for children, recognizing childhood, and putting children ahead of commercial interests and ahead of platform status.

The standard also embodies the U.N. Convention on the Rights of the Child, an agreement that establishes the civil, political, economic, social, and cultural rights of all children.

The term *age appropriateness* covers a variety of values that support children, including sustainability, privacy, usability, convenience, controllability, accountability, and inclusivity, according to the 55-page standard.

IEEE SA said it believes the standard will encourage organizations to design their services with children in mind, demonstrate commitment to social responsibility, and encourage adherence to local regulatory requirements.

"While there are localized efforts to address children's rights and safety in digital products and services, there has never been consensus-driven guidance applicable on a global scale," Konstantinos Karachalios, managing director of IEEE SA, said in a news release about the standard. "This standard provides organizations a framework to practically orient design processes for age-appropriate digital services toward responsible technological innovation inclusive of children."

Processes and principles

The IEEE 2089-2021 working group looked at a variety of ways information about children is collected and stored. They include search engines that expose young people to advertising, messaging systems used in online gaming such as chat apps, in-game purchasing, video-streaming services, and social media platforms. The group also reviewed augmented-reality applications, role-playing games, multimedia content, and the use of microphones, cameras, and other items that are part of Internet-connected devices.

Unlike other standards, IEEE 2089-2021 is not a protocol, a technology, or a specification, says IEEE Senior Member Katina Michael, the standard's working group chair.

"IEEE 2089 is a design practice," she says. "It's about acknowledging and identifying risks and addressing them before you deploy a new service."

Michael is a professor with Arizona State University's School for the Future of Innovation in Society, with a joint appointment in the school of computing and augmented intelligence. She is the director of the Society Policy Engineering Collective.

The standard outlines 11 processes to follow to mitigate and manage risks throughout the life cycle of development, delivery, and distribution of digital products and services. The framework includes recognizing child users and meeting their needs, upholding children's rights, taking a child-centered approach to data use, and writing published terms in age-appropriate formats. For each process, the standard defines the purpose; outcomes; activities and tasks; and inputs and outputs. It also identifies key roles for project teams such as an age-appropriate lead and a children's-rights advocate.

IEEE SA has made the standard available at no cost because of its expected impact. ■

ISTOCKPHOTO

IEEE PRODUCTS

Learn In-Demand IoT Concepts From New IEEE Academy

BY JOHANNA PEREZ

SURVEYS OF IEEE members regularly show "continuing professional education" as one of the primary reasons they join. To deliver on that need, the IEEE ad hoc committee on lifelong learning and continuing education initiated efforts to provide increased value for IEEE members, technical professionals, and engineers.

The committee developed the IEEE Academies, which are designed to teach in-demand, technical concepts to members.

The program's learning-path format helps people understand technical concepts even if they do not have a deep background in a specific technology. The learner is guided through a logical, continuous path that ties the concepts and materials together.

The first one to launch—the IEEE Academy on Internet of Things—uses existing content from across IEEE and combines it with new learning tools developed by industry experts.

The demand for the Internet of Things is steadily rising. There are more than 10 billion IoT devices already, with the new connections expected to increase by 20 times in the next three years, according to Datamation. With the IoT job market expected to grow to nearly US $1.6 trillion by 2026, according to Datamation, it is crucial for working technical professionals to keep their skills up to date.

Each learning path offers 0.9 continuing-education units or nine professional development hours. Upon completion, the learner earns a certificate.

The two learning paths in the IoT Academy are:

Communications Standards

Communication technology is an essential part of the IoT, as it allows devices to be interconnected. This learning path covers the basic principles of the technology and practical usage of standardized communication.

Computing Platforms

IoT computing platforms are important to the development and deployment of applications. This learning path provides an overview of current and future trends.

Visit the IEEE Learning Network for member and nonmember pricing.

Johanna Perez is a former digital marketing specialist for IEEE Educational Activities.

Machine Learning Mastery in 5 Courses

BY BRITNEY DO

AS BUSINESSES GRAPPLE with increasing amounts of data and search for ways to use it effectively, they're turning more and more to machine learning and deep learning. Both models use statistics to make predictions, but there are differences.

Machine learning employs algorithms to identify patterns and make predictions. When the algorithmic model makes a wrong prediction, a programmer must troubleshoot. Deep learning functions similarly, but its artificial neural network enables it to problem-solve more like a human. It can correct itself in the case of a bad prediction.

Advanced knowledge of mathematics, statistics, data analysis, and programming is fundamental for a machine learning engineer. To help technical professionals better understand the technology, IEEE Educational Activities created a five-course program, Machine Learning: Predictive Analysis for Business Decisions.

The five courses are:

Machine Learning in the Age of Enterprise Big Data

Examines the fundamental types of machine learning that drive business insights and reviews advanced computational intelligence for business processes.

Machine Learning in a Data-Driven Business Environment

This course can help you comprehend diverse sources that allow businesses to collect, store, organize, and interpret data.

Sound Business Practices for Data Mining and Predictive Analysis

Explore tools to measure business performance. This course explains how predictive and prospective analytics can deliver insights.

Machine Learning Algorithms, Models, and Systems Integration

This course will provide a solid understanding of available software and best practices for machine learning model integration.

Machine Learning Platforms, Technology, and Tools

Study the computational infrastructure that is necessary for enabling machine learning with big data. This course explains the concepts and techniques necessary for deploying scalable machine learning.

IEEE Educational Activities also offers Enhancing Business Operations With Machine Learning, an on-demand virtual event.

Visit the IEEE Learning Network for member and nonmember pricing.

Britney Do is a former digital marketing intern for IEEE Educational Activities.

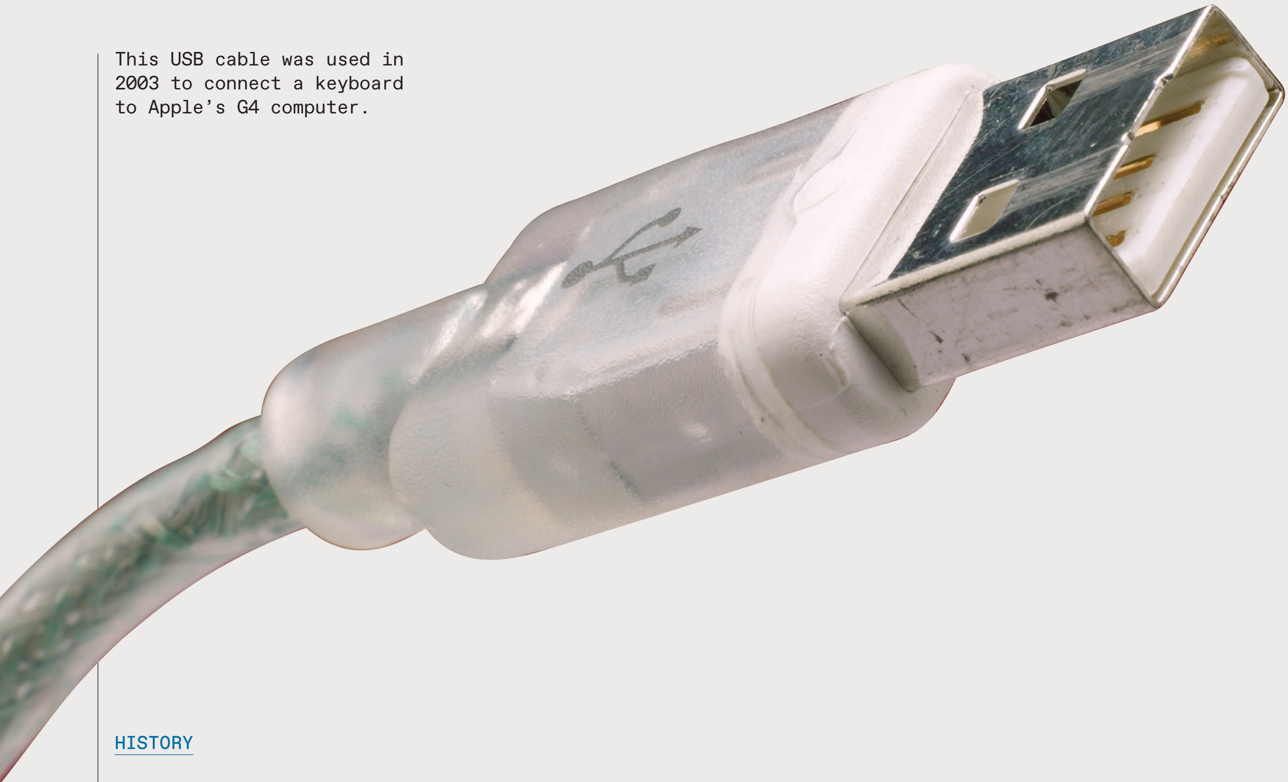

This USB cable was used in 2003 to connect a keyboard to Apple's G4 computer.

SSPL/GETTY IMAGES

HISTORY

How USB Came to Be

The plug-and-play tech made it easy to connect everything

BY JOANNA GOODRICH

BEFORE THE DEVELOPMENT of USB (Universal Serial Bus), attaching a camera, printer, or other accessory to a computer was not so simple. Users sometimes needed to open up their computer and add hardware to give them the communications port they needed.

The USB, which was released in 1996 by Intel, simplified things. USB ports now are standard on personal computers and are built into many other electronic devices such as smartphones, eBook readers, and game consoles.

The ubiquitous USB standard has been commemorated with an IEEE Milestone.

Collaboration is key

Many of the problems consumers encountered when they tried to attach peripherals to their computer in the 1990s arose because of the lack of standard practices among the industry's many suppliers, as noted in the Milestone's entry on the Engineering and Technology History Wiki. Another problem was that most PCs had a limited number of input ports, and adding more could be difficult.

Ajay Bhatt, one of the engineers at Intel who helped develop USB, says that even as a technologist, he struggled with upgrading his PC.

"I looked at the architecture, and I thought, You know what? There are better ways of working with computers, and this is just too difficult," he said in a 2019 interview with *Fast Company*.

In the early 1990s, Bhatt told his boss about his idea of developing a universal plug-and-play communication system—something the user didn't need to adjust. His manager wasn't interested.

Bhatt was passionate about his idea, though, so he decided to join a different research team at Intel. And there he was given the green light.

127
devices could be connected to a PC at once using USB

In 1992 Bhatt visited the Jones Farm Conference Center, in Hillsboro, Ore., where he met with engineers from different tech companies who also were looking into developing a plug-and-play scheme. It was there that engineers from Compaq, Digital Equipment Corp. (DEC), IBM, Intel, Microsoft, NEC, and Nortel formed an alliance.

"The industry as a whole recognized that it had a big problem that needed to be solved," Jim Pappas said in an Intel article on the USB interface. At the time of the meeting, Pappas was an engineering manager at DEC, but he eventually joined Intel as a program manager for its USB-development team.

From dream to reality

Before the group began developing USB, it explored what was already available. It looked at Ethernet-like technologies, audio interfaces, Apple's GeoPort, and IEEE 1394—also known as the Firewire standard. But none had all the traits the team sought. In particular, the engineers wanted something that was inexpensive, user-friendly, able to power peripherals, and offered a lot of bandwidth.

To keep manufacturing costs down, the engineers designed USB to work with a slender, four-conductor cable that could be as long as 5 meters. One end of the cable had an A connector, which plugged into the computer; the B connector on the other end plugged into the external device.

At the time, computers didn't typically provide power for such external devices. Most peripherals had to be plugged into an outlet while connected to a PC. But USB allowed a computer to supply sufficient power for some peripherals.

Another advantage of USB was that, in principle, it allowed as many as 127 peripherals to be connected to one PC at a time. A single computer was unlikely to have 127 USB ports, but the number of ports available could be increased by adding USB hubs.

The team announced its first design in 1995. At 12 megabits per second, USB 1.0 was "faster than anything else that normally would come at the back of the PC," Pappas told *Fast Company*.

Ajay Bhatt was one of the engineers who helped develop USB.

PRADEEP GAUR/MINT/GETTY IMAGES

The team encountered a problem, though: 12 Mb/s was too fast for computer mouses, joysticks, keyboards, and other accessories with unshielded cables. The engineers solved the problem by arranging for USB 1.0 to support communications at 1.5 Mb/s as well.

That approach allowed USB to work at low speed for low-cost peripherals with unshielded cables and at high speed for devices with shielded cables, such as printers and floppy-disk drives.

USB 1.1, released in 1996, did not become popular until 1998, after it was showcased at the COMDEX trade show in Las Vegas.

At a news conference there, an Intel team attached 127 peripherals to one PC. The engineers hired Bill Nye to plug in the last of the devices. In Pappas's 2019 *Fast Company* interview, he said that once Nye did so, the team sent a document to various destinations to print. "We had a whole stage full of different printers!" Pappas said.

The delay between the release of USB 1.1 in 1996 and when it caught on is understandable because Microsoft Windows 98, which was released in June 1998, was the first operating system to support USB. Two months later, Apple released its iMac, which lacked a floppy-disk drive but did have a pair of USB ports. Although Apple was not among the companies that worked on the USB project, it helped make the technology mainstream.

Since then, three more USB generations have emerged. The most recent, USB4, was released in 2019.

"Who would have thought that a connector that we had defined in the early '90s would still be usable today?" Bala Cadambi, who worked on the USB development team, said in the *Fast Company* interview. "That's very rare. We had cost constraints, performance constraints. It was designed for a desktop, not a smartphone. Looking back at it, it was wonderful that we accomplished what we did, that it withstood the test of time."

The IEEE Oregon Section sponsored the Milestone nomination. Administered by the IEEE History Center and supported by donors, the Milestone program recognizes outstanding technical developments around the world. ■

> "I couldn't imagine where USB has gone or where it will continue to go. This has exceeded the wildest of my imaginations."

IEEE NEWS

Tech Predictions for 2022 and Beyond

TO UNDERSTAND KEY technology trends, priorities, and predictions for 2022 and beyond, IEEE conducted a global study of 350 technology leaders including CTOs, CIOs, and IT directors. The leaders also shared how the coronavirus pandemic affected the adoption of technologies at their companies. Here is a look at what they said.

50%

Nearly half of technology leaders predicted that the significant and rapid increase of devices connected to their company's business this year would become unmanageable.

WHAT WILL BE THE MOST IMPORTANT TECHNOLOGY IN 2022?

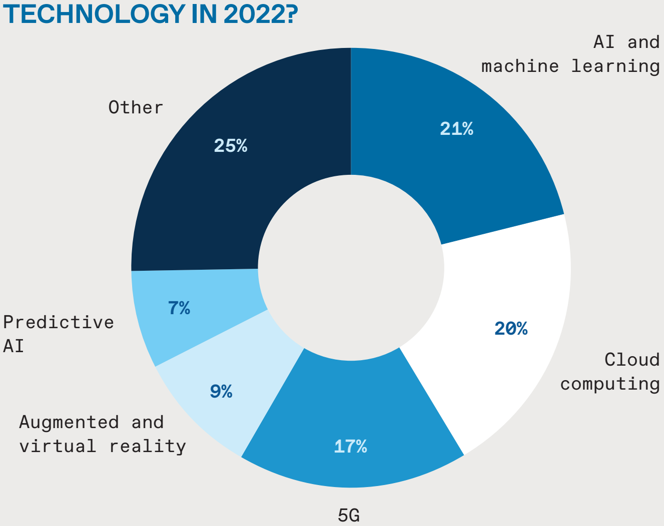

WHICH INDUSTRY SECTOR WILL BE MOST IMPACTED BY TECHNOLOGY IN 2022?

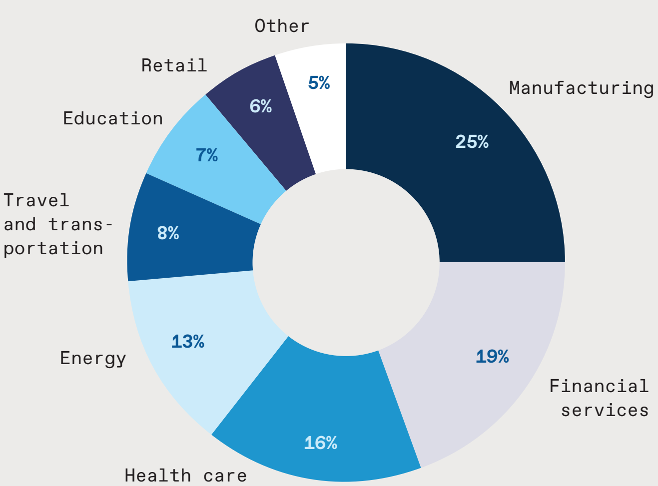

DUE TO ROUNDING, NUMBERS MAY APPEAR TO NOT ADD UP TO 100 PERCENT

TOP 10 TECHNOLOGIES COMPANIES ACCELERATED THE ADOPTION OF IN 2021 DUE TO COVID-19

- Cloud computing
- AI and machine learning
- 5G
- Augmented, virtual, and mixed reality
- Blockchain
- Teleconferencing
- Internet of Things
- Robotics
- Predictive AI
- Edge computing

150%

There was a 150 percent increase in the number of devices—such as tablets, robots, and drones—used at tech companies due to remote work.

SOURCE: THE IMPACT OF TECHNOLOGY IN 2022 AND BEYOND: AN IEEE GLOBAL STUDY

STARTUP

This Member's App Moves Money and Saves It Too

Money-transfer tool aims to help people fund worthy causes

BY KATHY PRETZ

MOBILE MONEY-TRANSFER APPS are now bigger in developing countries than just about anywhere else. In Kenya, for example, mobile transactions last year were up by about 20 percent over 2020, according to Capacity.

For a fee, people send money to friends and family, pay bills, and buy groceries using a mobile wallet app. The transactions are made through cellular service providers, although the funds held in a mobile money account are protected by local financial regulations.

That's important, because most of the apps' users don't have a bank account.

Although mobile money is fast and easy to use, it has a few drawbacks: The transfer service doesn't provide a separate account for savings, and there's no way to send money to an entity that lacks a cellphone number, such as a charitable organization.

Sandra K. Johnson [left] is working to change those shortcomings with her fintech startup, Global Mobile Finance, which is based in Morrisville, N.C.

The IEEE Fellow developed the geeRemit app to let users in the United States send money to people in sub-Saharan Africa. A small portion of the fee can be donated to a charitable organization or deposited into a financial account. GeeRemit is currently in the pilot stage.

Photo-illustration by Max-o-matic

Johnson says the app's features allow users to lend a helping hand to their friends and family living abroad.

"GeeRemit provides a social benefit by transforming lives," she says. "The more money you send, the more you can save over time. The recipient can use that money to pay for a life-changing event such as a wedding, schooling, or starting a business. This social component is attractive for many people who send money, because they already have the mindset to help others."

For her development of geeRemit, *Inc.* magazine included her on its 2020 Top 100 Female Founders list.

Johnson held a number of technical leadership positions in IBM offices around the world during her more than 25 years with the company. She left IBM in 2014. In 2018 she launched the fintech company and is currently its only employee.

She also has a consulting company, SKJ Visioneering, in Research Triangle Park, N.C. In addition, she is a visiting scholar at North Carolina Agricultural and Technical State University, in Greensboro.

Eureka moment

After a 2007 visit to the Elmina Slave Castle, a former slave-trading depot in Ghana, Johnson had what she calls a "eureka moment." She vowed to use her technical skills to improve the living conditions of sub-Saharan Africans.

She got an opportunity to do so in 2012, when she took the job of CTO in IBM's Central, East, and West Africa regional office, in Nairobi, Kenya. Part of Johnson's job required her to travel as the company's technical leader in the region. She noticed that people were using mobile money to pay for just about everything.

"Even though this is a cash-based society, no one uses cash," Johnson says. "They all pay through mobile money. Kenya was where mobile money became a success."

During her trips back to the United States she sometimes transmitted money to Nairobi to pay bills, but the process wasn't easy.

"I experienced the challenges of slow transfer times and very high fees," she says. "I coupled the two experiences together and decided to create this money-transfer app."

> "GeeRemit provides a social benefit by transforming lives."

Data is gold

Johnson is looking for processes to analyze the data she's collecting so she can discover new ways to allow customers to send more money or provide some other benefit.

"The issue with a money-transfer protocol is that it needs to be fast and low-cost. But it also needs to require as little data as possible," she says.

One benefit might be to send the app's users reminders of upcoming birthdays to prompt them to send a cash gift to friends and relatives. She also would like to use analytics to determine whether a user needs to save money for college, say, or for starting a business. The app would then suggest the person set aside money.

"There's a certain segment of the market that's attracted to the additional functions and features," she says. "We've done some market research, and potential customers are willing to pay up to 2 percent more for this service."

> "I experienced the challenges of slow transfer times and very high fees. I coupled the two experiences together and decided to create this money transfer app."

Overcoming hurdles

Johnson says her main challenge is finding funding. She has spent the past few years raising money, mostly from friends and family but also through crowdfunding websites such as IFundWomen, which helps female founders. Johnson also received money through angel investors and a government-regulated fundraising program. She applied for several grants including one from an American Express program for businesses owned by Black women.

Johnson says she's stuck in a Catch-22 situation. To sell investors on her app, she says, "we need market traction, and to get market traction, we need funds to get to the market."

Another hurdle is that the business of transmitting money is heavily regulated in the United States. Each state requires a money-transmitter license, and fees for the license can range from a few thousand to a few hundred thousand dollars. Because Georgia's fee is the least expensive, she's running the pilot program for the app in the Atlanta area.

The early adopters who are using the app are so far transmitting money only to Ghana and Kenya. She doesn't need to obtain licenses in those countries, she says, because she has partnered with a global payments network that already has licenses there. GeeRemit has been integrated on the network's front end, and the delivery of funds is handled on the back end.

"Right now, we have a license and a focus," Johnson says. "As we generate revenue from those using the service in Georgia, we have a plan in place to obtain licenses in other states." ■

BY ALLISON MARSH

Past Forward

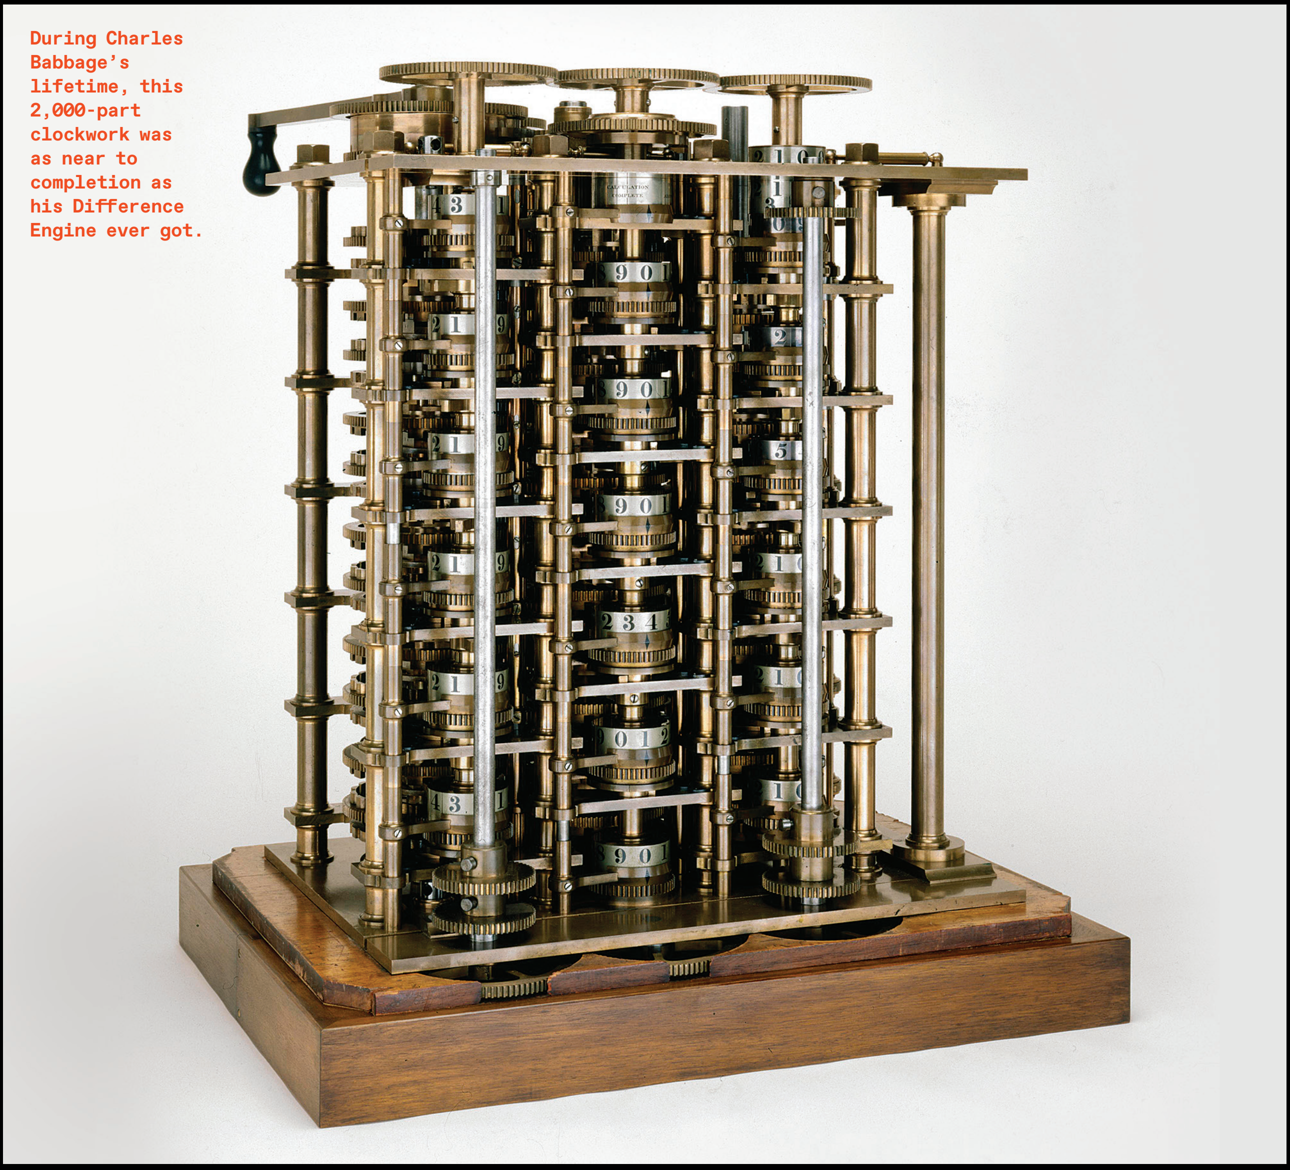

During Charles Babbage's lifetime, this 2,000-part clockwork was as near to completion as his Difference Engine ever got.

The Clockwork Computer

Two hundred years ago this month, Charles Babbage presented to the Royal Astronomical Society a prototype of a remarkable machine that used a clockwork mechanism to solve polynomial equations. Today, Babbage's Difference Engine is often considered the first automatic computer. Babbage then did the almost unimaginable—he convinced the British government to fund the building of a full-scale version. The government saw the value in a machine that could calculate the many numerical tables used for navigation, construction, finance, and engineering, thereby reducing human labor (and error). But after 20 years and £17,500 (about US $3 million today), the British government still didn't have a working machine and finally pulled the plug. Babbage never completed either the Difference Engine or his later and more ambitious Analytical Engine. Only in 1991, in celebration of the bicentenary of Babbage's birth, did a team from London's Science Museum unveil a functional difference engine based on Babbage's drawings. ■

FOR MORE ON THE HISTORY OF THE DIFFERENCE ENGINE, SEE spectrum.ieee.org/pastforward-jun2022

SCIENCE MUSEUM GROUP