**Neurostimulation for Long COVID** › Treating brain fog and fatigue
P.9

**Claude Shannon, Father of Machine Learning?** › He had the idea but didn't test it
P.23

**Finally, Full Color Comes to E-paper Displays** › How E Ink did it
P.30

FOR THE TECHNOLOGY INSIDER
FEBRUARY 2022

# IEEE Spectrum

## Electric Journey to a Metal World

IEEE

VOLUME 59 / ISSUE 2 FEBRUARY 2022

# Contents

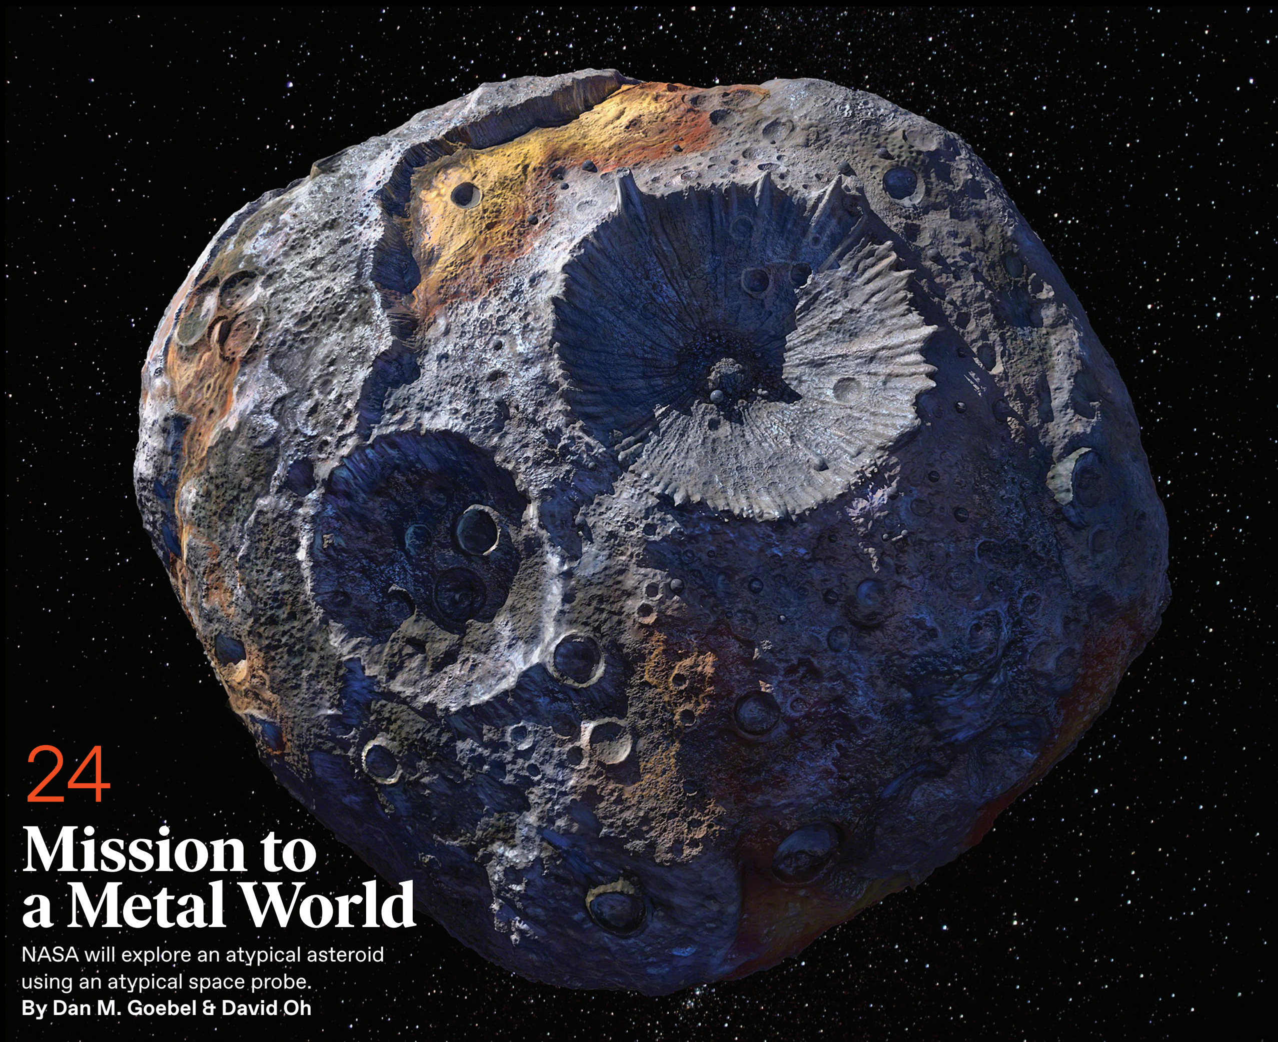






ON THE COVER:
Illustration by John MacNeill

JPL-CALTECH/ASU/NASA; BOTTOM: POCKETBOOK

CONTRIBUTORS

## JEFF FROLIK

Frolik is a professor at the University of Vermont who specializes in data communications. His coauthors in this issue, Paul Hines and Mads Almassalkhi, also of UVM, work on power-grid reliability. Their joint research merges the two fields, applying the Internet concepts of packetization and randomization to grid stabilization. They first applied this approach to electric-vehicle charging. Later, they realized it could be scaled up to power-grid management more broadly, which they describe on page 42.

## DAN M. GOEBEL

A research scientist at NASA's Jet Propulsion Laboratories, Goebel is spacecraft chief engineer for the upcoming Psyche mission. "Yeah, I'm Scotty," Goebel jokes. His coauthor and JPL colleague David Oh is project system engineer for the mission. Despite a fast-approaching August launch deadline, they found time to write "Mission to a Metal World" [p. 24]. "We have our heads down at this point," says Goebel. "If it launches and works great, we're going to be heroes—if it doesn't, we're going to be goats."

## EDZER HUITEMA

Huitema is chief technology officer (USA) for E Ink. He previously worked on display technologies at Apple and at startups Polyera and Polymer Vision. Ian French is chief technology officer (Taiwan) for E Ink. French was technical lead on the company's Kaleido project, described on page 30. Both he and Huitema began their display development careers at Philips Research.

## BEHROOZ REZVANI

Rezvani, a solid-state physicist, is the CEO of Neural Propulsion Systems, in Pleasanton, Calif., a driving-technology startup he cofounded in 2017. "In the mid-1980s I worked on an MRI system for simultaneously imaging two different elements, hydrogen and phosphorus, and that problem got me thinking about multiband, single-aperture signal processing," says Rezvani. He and his colleagues, Babak Hassibi, Fredrik Brännström, and Majid Manteghi, describe the result in "Letting Robocars See Around Corners," on page 36.


# IEEE Spectrum

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

**EDITORIAL DIRECTORS**
Harry Goldstein (Digital), h.goldstein@ieee.org
Glenn Zorpette (Development), g.zorpette@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

**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
**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

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

**CONTRIBUTING EDITORS** Robert N. Charette, Steven Cherry, Charles Q. Choi, Peter Fairley, Maria Gallucci, W. Wayt Gibbs, Mark Harris, Jeremy Hsu, Allison Marsh, Prachi Patel, Megan Scudellari, Lawrence Ulrich, Emily Waltz

**ADMINISTRATIVE ASSISTANT** Ramona L. Foster, r.foster@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
**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

**EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM***
Susan Hassler, *Chair*; David C. Brock, Robert N. Charette, Ronald F. DeMara, Shahin Farshchi, Jonathan Garibaldi, Lawrence O. Hall, Jason K. Hui, Leah Jamieson, Mary Lou Jepsen, Deepa Kundur, Michel Maharbiz, Somdeb Majumdar, Allison Marsh, Carmen Menoni, Sofia Olhede, Wen Tong, Maurizio Vecchione

**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

**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**

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

**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

**IEEE BOARD OF DIRECTORS**
**PRESIDENT & CEO** K.J. Ray Liu, president@ieee.org
**+1 732 562 3928 Fax: +1 732 981 9515**
**PRESIDENT-ELECT** Saifur Rahman
**TREASURER** Mary Ellen Randall
**SECRETARY** John W. Walz
**PAST PRESIDENT** Susan K. "Kathy" Land

**VICE PRESIDENTS**
Stephen M. Phillips, *Educational Activities*; Lawrence O. Hall, *Publication Services & Products*; David A. Koehler, *Member & Geographic Activities*; Bruno Meyer, *Technical Activities*; James E. Matthews, *President, Standards Association*; Deborah M. Cooper, *President, IEEE-USA*

**DIVISION DIRECTORS**
Franco Maloberti (I); Ruth A. Dyer (II); Khaled Ben Letaief (III); Manfred "Fred" J. Schindler (IV); Cecilia Metra (V); Paul M. Cunningham (VI); Claudio Cañizares (VII); Christina M. Schober (VIII); Ali H. Sayed (IX); Dalma Novak (X)

**REGION DIRECTORS**
Greg T. Gdowski (1); Barry C. Tilton (2); Theresa A. Brunasso (3); Johnson A. Asumadu (4); Bob G. Becnel (5); Timothy T. Lee (6); Robert L. Anderson (7); Antonio Luque (8); Enrique A. Tejera (9); Deepak Mathur (10)

**DIRECTOR EMERITUS** Theodore W. Hissey

**IEEE STAFF**
**EXECUTIVE DIRECTOR & COO** Stephen Welby
+1 732 562 5400, s.p.welby@ieee.org
**CHIEF INFORMATION OFFICER** Cherif Amirat
+1 732 562 6017, c.amirat@ieee.org
**CHIEF MARKETING OFFICER** Karen L. Hawkins
+1 732 562 3964, k.hawkins@ieee.org
**PUBLICATIONS** Steven Heffner
+1 212 705 8958, s.heffner@ieee.org
**CORPORATE ACTIVITIES** Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
**MEMBER & GEOGRAPHIC ACTIVITIES** Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
**STANDARDS ACTIVITIES** Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
**EDUCATIONAL ACTIVITIES** Jamie Moesch
+1 732 562 5514, j.moesch@ieee.org
**GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER**
Sophia A. Muirhead +1 212 705 8950, s.muirhead@ieee.org
**CHIEF FINANCIAL OFFICER** Thomas R. Siegert
+1 732 562 6843, t.siegert@ieee.org
**TECHNICAL ACTIVITIES** Mary Ward-Callan
+1 732 562 3850, m.ward-callan@ieee.org
**MANAGING DIRECTOR, IEEE-USA** Chris Brantley
+1 202 530 8349, c.brantley@ieee.org

**IEEE PUBLICATION SERVICES & PRODUCTS BOARD**
Lawrence O. Hall, *Chair*; Stefano Galli, Nazanin Bassiri Gharb, James Irvine, Clem Karl, Hulya Kirkici, Yong Lian, Fabrizio Lombardi, Peter Luh, Anna Scaglione, Gaurav Sharma, Isabel Trancoso, Peter Winzer, Bin Zhao, Weihua Zhuang

**IEEE OPERATIONS CENTER**
445 Hoes Lane, Box 1331
Piscataway, NJ 08854-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

**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. 2. 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/web/aboutus/whatis/policies/p9-26.html.

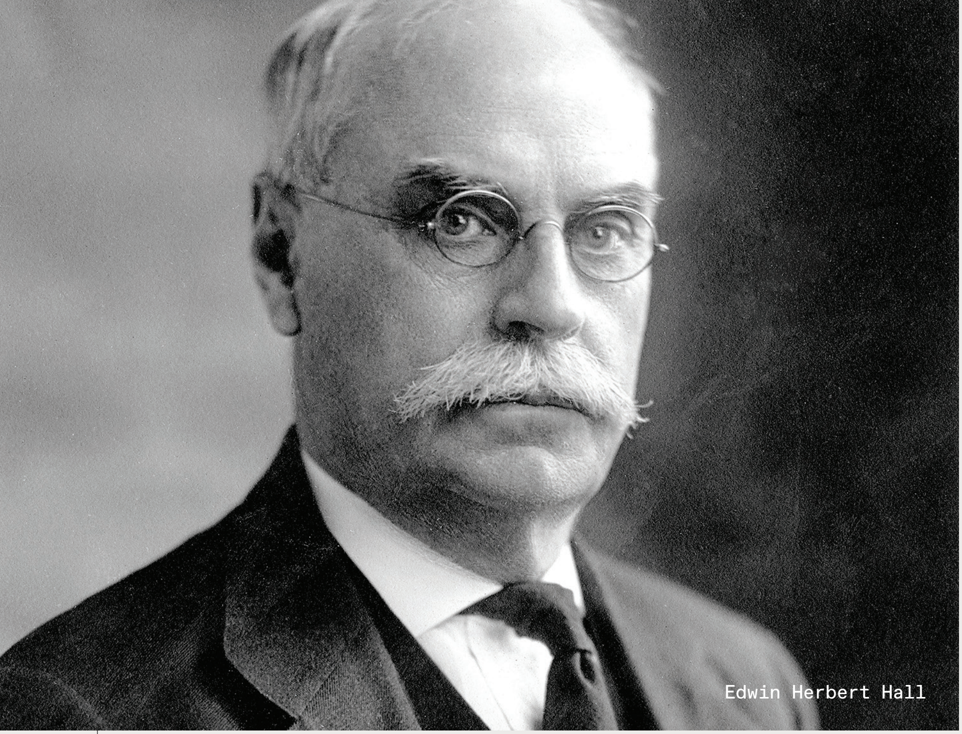



SPECTRAL LINES

# The Ripples From Edwin Hall's Effect › How a 143-year-old insight is shaping spacecraft propulsion

BY GLENN ZORPETTE

**Of the many** physicists who didn't win a Nobel Prize, Edwin Herbert Hall is among a select group for whom the failure seems a particular injustice. On 28 October 1879, the 23-year-old Hall made a discovery that still reverberates today (and which would eventually lead to Nobel Prizes for four *other* people). Working in a physics lab at the newly created Johns Hopkins University, Hall saw a needle in a galvanometer gauge shift, indicating an electric potential that had no obvious cause.

He had discovered a phenomenon that would soon be dubbed the Hall effect. Suppose you have current flowing through a flat ribbon of conductor. Now you place the ribbon in a magnetic field that is perpendicular to the plane of the conductor. Presto! The current flow is momentarily shifted, pushed by the magnetic field to one side of the ribbon. This causes a buildup of opposite charges on either side of the ribbon, and thus a difference in potential, which quickly balances the force created by the magnetic field. That potential manifests itself as a perpendicular voltage, which was what Hall measured.

Today, the Hall effect has major applications, both "real world" and out of this world. A standard way of detecting and measuring a magnetic field is with a device called a Hall effect sensor. Hundreds of millions of them are sold every year. But one of the most interesting applications of the Hall effect is in spacecraft propulsion. And now, as Dan M. Goebel and David Oh note in "Mission to a Metal World" (p. 24), for the first time ever a thruster based on the Hall effect is about to propel a spacecraft, called Psyche, on an interplanetary voyage.

In the early 1960s researchers in the United States and Soviet Union began publishing papers on experiments to use electrically generated plasmas to produce thrust. The United States soon abandoned its program for Hall thrusters in favor of a different kind, called ion thrusters. Work continued in the Soviet Union, though, resulting in a Hall thruster launched in 1971 aboard a Soviet weather satellite. It was the start of something big.

So far, more than 100 Soviet or Russian satellites have used Hall thrusters, mostly for orienting the spacecraft or maintaining its orbit. In the United States, interest was revived in 1992, when a team of researchers visited Russian labs as part of a U.S. program that produced thrusters based on Russian technology. Today, engineers are developing Hall thrusters in scores of labs all over the world.

Consider SpaceX's Starlink program, which is swarming satellites in low Earth orbit to provide broadband Internet services. Each one of the more than 1,800 satellites launched so far carries a Hall thruster, making SpaceX the largest operator by far of electric thrusters in the world.

A Hall thruster uses a power source to set up a potential difference between an anode and a cathode. A noble gas, such as xenon or krypton, is fed through holes in the anode into a channel, and electrons flowing from cathode to anode ionize the gas, forming a charged plasma. An applied magnetic field forces the electrons to circulate around the axis of the thruster near the exit, downstream of the anode. The transverse electron current rotating in the channel is called the Hall current. This current generates an electric field in the plasma reminiscent of that found by Hall, which accelerates the ions and shoots them out of the end of the channel, producing thrust.

Robert Goddard, Konstantin Tsiolkovsky, and Wernher von Braun all envisioned the use of some form of ion thrusters on long-duration space missions. Next August, as Goebel and Oh note in their article, four Hall thrusters loaded with xenon gas will begin pushing the Psyche spacecraft—all 2,600 kilograms of it—on a 2.4-billion-kilometer journey to the asteroid belt. Not bad for a rocket engine that will generate just 280 millinewtons of thrust—the weight of five U.S. quarters. And in 2024, a cluster of Hall thrusters will power the Gateway lunar space station, a key element in NASA's Artemis program to return astronauts to the moon.

You might call this a *moving* testament to Hall's genius. Maybe an even better one than a Nobel Prize. ■

JHU SHERIDAN LIBRARIES/GADO/GETTY IMAGES

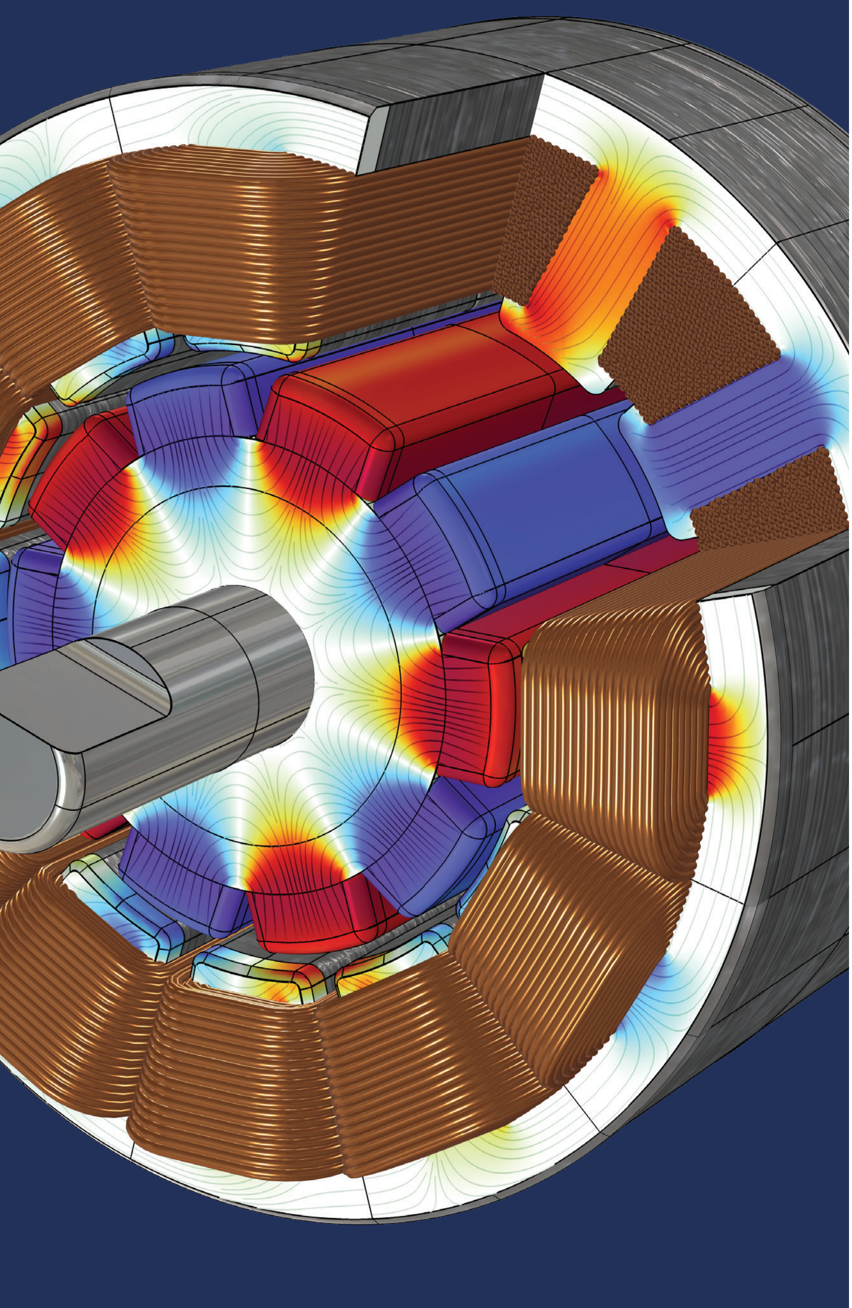

# News

A new kind of matter, a quantum time crystal, is represented in this artist's depiction.

GOOGLE QUANTUM AI

QUANTUM COMPUTING

## What's a Time Crystal? › And how do Google researchers use quantum computers to make them?

BY CHARLES Q. CHOI

**First conceived of** a decade or so ago, a quantum time crystal is a new kind of matter that bears an uncanny resemblance to a perpetual motion machine. Its parts can theoretically move in a repeating cycle without consuming energy for eternity, like a watch that runs forever without any batteries.

Scientists have strived to create this novel phase of matter for years. Now researchers at Google Quantum AI and their colleagues reveal they have created time crystals using Google's Sycamore quantum-computing hardware, findings they detailed online in November in the journal *Nature*.

Whereas classical computers switch transistors either on or off to symbolize data as ones and zeros, quantum computers use quantum bits, or qubits, that because of the nature of quantum mechanics, can exist in a state of superposition where they are both 1 and 0 simultaneously. By linking qubits together via a quantum effect known as entanglement, a 300-qubit quantum computer theoretically could perform more calculations in an instant than there are atoms in the visible universe. In 2019, Google argued it used Sycamore to display "quantum primacy," finding answers to problems no classical computer could ever solve.

In the new study, the researchers used a 20-qubit system not for computation, but to realize time crystals. To find out more, we spoke with Google staff research scientist Kostyantyn Kechedzhi and Google senior research scientist Xiao Mi, who conducted much of the research on the theoretical and experimental sides, respectively. The conversation has been edited for length and clarity.

**What is a time crystal?**

**Kostyantyn Kechedzhi:** A crystal is a system of many atoms that, due to mutual interactions, organize themselves into a periodic pattern in space. A time crystal is a quantum system of many particles that organize themselves into a periodic pattern of motion—periodic in time rather than in space—that persists in perpetuity.

**What might you compare time crystals to in nature?**

**K.K.:** Persistent periodic motion is very familiar in nature. A system of two massive objects attracting each other due to gravity is the simplest example—the two objects move around the common center of mass following strictly periodic orbits. At first sight, this might seem like an example of a time crystal. However, the key novelty of a time crystal is the periodic motion of a system of many objects interacting with each other.

A system of many interacting objects demonstrates an entirely different behavior compared with just two massive objects orbiting each other—instead of repeated patterns, the motion continually changes. For instance, in the solar system, planets follow trajectories that are approximately periodic, but the true behavior of planets is chaotic, which means a small deviation of a planet from its path today will result in a completely reshaped trajectory over time, albeit billions of years.

The second law of thermodynamics postulates that systems of many interacting objects tend toward more disorder, which appears in contradiction with the strictly periodic motions of a time crystal. Nonetheless, a system of many interacting quantum objects could demonstrate periodic motion without violating the second law of thermodynamics due to a fundamentally quantum phenomenon called many-body localization.

**In your new work, you created a periodically driven many-body localized time crystal. This is a time crystal consisting of many parts whose activity is driven by an externally applied cyclic series of pulses. And by localization, you mean that physical laws act in the specific locale of this time crystal to help keep it stable and from dissipating energy?**

**K.K.:** Yes. A key property of a localized quantum system of many objects is that a sufficiently weak external pulse or force applied to any one of the objects will affect its neighbors but will not be felt across the entire system. In this sense, the system's response is localized. In contrast, in a chaotic system, a small disturbance is felt across the entire system. The phenomenon of localization prevents absorption of energy from the external drive.

**A time crystal is a quantum system of many particles that organize themselves into a periodic pattern of motion that persists in perpetuity.**

**How comparable are your time crystals to perpetual motion machines?**

**K.K.:** The time crystals observed in our experiment do not absorb any net energy from the pulses used to drive their behavior. This is perhaps why they are often compared to perpetual motion machines.

However, perpetual motion machines are expected to produce work without an energy source, which would violate the laws of thermodynamics. In contrast, the motion of a time crystal does not produce work without an energy source and therefore does not violate physical laws.

**Do your time crystals break down over time?**

**K.K.:** Our processor is not 100 percent isolated from the environment, and this weak coupling to the environment introduces a finite "extrinsic" lifetime of the time crystal. In other words, after a sufficiently long time, the order is lost and the periodic pattern no longer repeats itself.

**What applications might time crystals have?**

**K.K.:** A time crystal is, like ferromagnetism or superconductivity, an example of spontaneous symmetry breaking, or spontaneous order. For instance, a ferromagnet is essentially a system of much smaller magnets whose magnetic poles all point in a single direction, and in this sense are ordered. Symmetry is said to be "spontaneously" broken in such a state, since in normal matter the poles all point in random directions. Stable examples of spontaneous symmetry breaking, as seen in ferromagnetism or the vanishing electrical resistance of a superconductor, often have significant technological value.

Spontaneous symmetry breaking is associated with equilibrium. For instance, think of liquid water freezing into a crystal when brought to a stable low temperature. A remarkable property of the time crystal we observed is its spontaneous order despite it being driven out of equilibrium. This observation opens the door to identifying other out-of-equilibrium states of quantum matter with novel types of order.

**What about time crystals has proven difficult to research, and why?**

**K.K.:** The challenge is that the isolation of quantum matter from its environment is never perfect.

**Why use a quantum computer to help create time crystals?**

**Xiao Mi:** Quantum computers are the platform of choice to realize time crystals because they have precisely calibrated quantum logic gates.

**A quantum logic gate is the quantum computing version of the logic gates that conventional computers use to perform computations?**

**X.M.:** Yes. Quantum logic gates allow the many-body interactions that are necessary for time crystals to exist to be implemented with very high precision.

Previous studies on time crystals were all performed on so-called quantum simulators. These platforms lack the precision of quantum computers. As a result, many of these experiments were later found to be flawed due to unintended interactions.

**What did you show in your new study?**
**X.M.:** We engineered quantum circuits that have the types of interactions that are theoretically expected to lead to a time crystal. We then collected data from these quantum circuits and used a variety of techniques to verify that our data are consistent with time-crystalline behavior. This included three things:

1. Any decay or "melting" of time-crystalline order was induced only by external decoherence, not the internal dynamics of our system.
2. The signature of a time crystal was present regardless of the initial state of the system.
3. We could determine the boundary of the time-crystal phase—that is, where it "melted."

**What do you personally find most interesting about these results?**
**X.M.:** Understanding the behavior of interacting particles near the critical point of phase transitions—for example, the melting temperature of ice into water—is a longstanding problem in physics and still holds many unresolved puzzles for quantum systems. We were able to characterize the phase-transition point between the time crystal and quantum chaotic states. This is a very promising direction for early applications of quantum processors as a tool for scientific research, where modest-size systems of dozens or hundreds of qubits could already provide new experimental information about the nature of phase transitions.

**How might time crystals lead to better quantum computers?**
**X.M.:** Having an object like a time crystal that is stable against experimental interference may help design quantum states that are long-lived, a crucial task for the future improvement of quantum processors.

**Another time crystal was created using qubits by researchers at Delft University of Technology, in the Netherlands. How would you distinguish your work from theirs?**
**K.K.:** The Delft experiment implemented some of the protocols outlined in our earlier theoretical work, which distinguish many-body localized time crystals from so-called prethermal [discrete] time crystals observed in recent years. Whereas prethermal time crystals are characterized by finite intrinsic lifetimes, many-body localized time crystals are characterized by a diverging—that is, infinitely long—intrinsic lifetime.

The exceptional flexibility of our processor allowed us to demonstrate that time-crystal dynamics persist over a range of system parameters. One consequence of that is our observation of the phase transition between the time crystal and the chaotic behavior. The presence of the phase transition suggests that a time crystal is a distinct state of matter from the more ubiquitous chaotic many-body state, including prethermal time crystals.

Crucially, the protocol we describe in our new study is scalable—it can be easily applied to a larger quantum processor. This is a result of further theoretical analysis significantly improving our prior work on which the Delft experiment was based. In the future, I can see our experiment repeated on larger and larger systems.

**What specific directions do you think your research might go from here?**
**K.K.:** One of our goals is to develop our processor into a scientific tool for physics and chemistry. The key challenge is reducing the error in the device. This is key for future applications of quantum processors and realization of fault-tolerant quantum computation. It needs to be addressed through improvements in the hardware, algorithmic error-mitigation strategies, and fundamental understanding of the role of noise in many-body quantum dynamics. ■

JOURNAL WATCH

## Blockchains Talk Amongst Themselves

**There's been a Cambrian explosion of blockchains in recent years, which to date, have largely been developed in silos. It's also hard to share information between them. But now researchers are trying to build new channels to help these networks speak to each other.**

**"We've got these blockchains that are closed worlds that cannot interact with each other," says Stefan Schulte, a professor of data engineering at the Hamburg University of Technology.**

**Solving this problem is a booming area of research, though, and Schulte and his colleagues recently presented a potential workaround at the IEEE International Conference on Blockchain Computing and Applications. Their approach relies on blockchain relays, which are essentially smart contracts running on one blockchain that can verify events on another blockchain.**

**If a user wants to transfer an asset they first "burn" it on the source blockchain, typically by sending it to a nonexistent address. Also included in the transaction are details of the asset and which blockchain and user is the ultimate recipient. Third parties monitor the source blockchain for burn transactions and send them to the relay for a small reward. The relay then recreates the asset on the new blockchain.**

**Unfortunately, says Schulte, transaction fees have to date made this approach impractical. He says his group has trimmed the fees back by 92 percent. But when the "reduced" fees still amount to US $5 or more per transaction, the problem is far from solved.**
**—Edd Gent**

BIOMEDICAL

# Zapping the Brain and Nerves Could Treat Long COVID ›

## Pilot studies test electrical treatments for the still-mysterious malady

BY ELIZA STRICKLAND


Patients can self-administer some types of neurostimulation treatments, with minimal supervision via telemedicine.

SOTERIX MEDICAL

**New Year's Day** marked a dispiriting milestone for one New Jersey woman: 20 months of symptoms of COVID-19. The woman, who asked for anonymity to protect her medical privacy, suffers from a variety of neurological problems that are associated with long COVID, including brain fog, memory problems, difficulty reading, and extreme fatigue. In her search for help, she came across neurologists at New York University who were trying neurostimulation for long COVID patients. She signed up for experimental treatments five days per week that send gentle electric currents through her skull and into her cortex.

It might sound weird, she says, but the reality is quite mundane. "People ask me, 'You're putting electricity in your brain? Where do you go to do that?' And I say, 'I do it in my house. I just put on a headband and make a call.' "

The woman was part of a wave of people who started turning up at NYU's neurology clinic in the late spring of 2020, several months after the first wave of COVID-19 cases hit New York City. "They were saying, 'I can't function, I can't return to work,' " remembers Leigh Charvet, a professor of neurology at NYU Grossman School of Medicine. To make matters worse, doctors had little to offer these patients.

Even as the world grapples with new waves of acute illness, doctors are trying to understand and find treatments for long COVID, which can trouble patients for many months after their recovery from the initial infection. The syndrome, technically known as post-acute sequelae of SARS-CoV-2 infection (PASC), is associated with a long list of possible symptoms, including heart palpitations, breathing problems, and a wide variety of neurological issues. "We need to do so much work to understand what long COVID is," Charvet says. "But we also need to reach people now with something that we know is safe and deployable."

Neurostimulation refers to electrical stimulation of the brain or peripheral nerves with either implanted or external devices. When the pandemic hit, researchers who had been working on neurostimulation for other maladies looked for ways to help the response.

**If neurostimulation does help with the neurological symptoms of long COVID, it's not clear why.**

"This was a chance for neuromodulation to step up," says Marom Bikson, codirector of neural engineering at the City College of New York and cofounder of the neurotech company Soterix Medical, in Woodbridge, N.J., which supplied stimulation gear to several research groups.

While some researchers began investigating whether noninvasive neurostimulation could help with the acute phase of infection, others, including Charvet, took on long COVID. The trials so far have been very small, but the results have been promising enough to support larger studies to optimize the technology and test the efficacy of these treatments.

Charvet has tried a type of neurostimulation called transcranial direct current stimulation (tDCS) with a handful of people so far. A patient puts on an electrode-studded headband that's attached to a controller and calls the study coordinator, who provides a unique code to enable that day's 20-minute stimulation session. Charvet says the research so far has been "a testing ground—it's not scientific, it's not controlled." Patients have come to her for help with brain fog, fatigue, headaches, emotional dysregulation, and other problems, and she tweaks the treatment protocols based on each person's symptoms.

She's now planning a larger trial with NYU patients to optimize the technology for at-home treatments. The trial will debut a tDCS headband that also tracks heart-rate variability; she and her colleagues hope that this biomarker will serve as an indicator of the patient's response to treatment. They'll use a headset made by Soterix Medical that measures impedance in the electrodes and translates that signal into heart-rate data. "What drives us is that there's a tremendous unmet need," Charvet says. "And our patients are getting better."

At the Medical University of South Carolina, psychiatry professor Mark George tried a different neurostimulation approach in a pilot study of 20 patients that he began in late 2020; his study used an at-home device that stimulated the vagus nerve through the ear. Each patient got an iPad for telemedicine consultations, a stimulation device, and "a portable ICU" with wearables that measured heart rate, oxygen saturation, and blood pressure. George's patients did 1-hour sessions each morning and evening, six days per week.

"We showed you could do this kind of stimulation at home. The safety data was impeccable," George says. "And we saw reductions in brain fog, improvements in energy, some improvement in anxiety." He's now applying for funding for a larger study.

One of his patients, a woman in her 60s who asked to be identified only by her first name, Pam, says she suffered from brain fog, memory lapses, fatigue, and mood swings following her case of COVID-19, which sent her to the emergency room in April 2020. When she started the stimulation, she felt a lessening of the uncharacteristic depression and anger that had troubled her, she says. "When I started with the treatment, I felt a little brighter, more like myself," Pam says. "I think I was a little better mentally."

One of the challenges that researchers face as they investigate treatments for long COVID is the diversity of symptoms that patients report. George says his study deliberately took a "shotgun approach," enrolling patients with a variety of neurological symptoms and looking at which responded best to the treatment.

People may be classified as long COVID patients for a number of reasons, says Jennifer Frontera, a professor of neurology at NYU Grossman School of Medicine. She explains that some people experienced low levels of oxygen in their brains during their acute illness, while others may have immune systems that went into overdrive following COVID-19 infection. But others may be experiencing a worsening of underlying conditions such as mood disorders and dementia, and still others may be having symptoms that aren't related to their COVID-19 infections. "Many people are sitting around their houses," she says. "Some people are more in tune with their bodies and are noticing things they never noticed before."

To get a handle on the basics, Frontera and her colleagues studied the health impacts of the pandemic, surveying 1,000 people whose demographics roughly matched those of the United States in terms of age, gender, and ethnicity. They didn't ask participants if they'd been infected with COVID-19 until the end. They found that pandemic-related stress factors such as financial and relationship problems were equally predictive of anxiety, depression, and insomnia as a history of COVID-19 infection. However, a history of infection was more predictive of cognitive problems.

Frontera says the study's results don't undercut the severity of the long COVID epidemic. She notes that the study found that 25 percent of people with a history of COVID-19 had symptoms that persisted beyond a month. "If you translate that out to the population of the United States, that would be 6 million people," she says. She's most troubled by the cognitive symptoms, she says: "We don't have a medication for brain fog."

If neurostimulation does help with the neurological symptoms of long COVID, it's not clear why. Stimulation with tDCS has been shown to increase "plasticity" in the brain, or the ability of the brain to make new connections between neurons; neuroplasticity is associated with learning and rehabilitation after injury. Vagus-nerve stimulation has been shown to reduce inflammation in the body, which is a component of autoimmune disorders; if some long haulers are suffering from an overactive immune system, vagus-nerve stimulation could help.

The researchers are hoping that larger studies will shed light on how neurostimulation impacts the brains of people with long COVID. And if millions of people in the United States alone need treatment, this could be an unprecedented opportunity for research. Bikson of Soterix Medical notes that the industry of neurostimulation is just getting started. "We don't have Pfizers of neuromodulation," he says, "but you can only imagine what would happen if it shows an effect on long COVID." ■

HONDA MOTOR CO.

An e-rickshaw is parked alongside a model of the planned Honda Mobile Power Pack Exchanger.

TRANSPORTATION

# Battery Swapping Electrifies India's Rickshaws › A three-wheeled taxi fleet provides a foothold for the electrification of everything

BY JOHN BOYD

**Spurred on by** the triple threats of $CO_2$ emissions, air pollution, and fluctuating oil prices, governments around the world are encouraging auto manufacturers to replace vehicles using internal combustion engines with electrically driven power trains. India, with a surging economy and the world's dirtiest air, is no exception. The Indian government says it is undertaking "multiple initiatives to promote manufacturing and adoption of electric vehicles," such as providing tax incentives to manufacturers and subsidies to purchasers of e-vehicles, including electric three-wheelers, or e-rickshaws.

There are over 6 million three-wheelers used by the masses for cheap taxi rides on India's clogged roads, and the majority are powered by engines employing environmentally unfriendly compressed natural gas or lead-acid batteries. Honda Motor Co., hoping to take advantage of the central and local government incentives, is looking to tap into this shift to electric vehicles with a new mobile battery-sharing service. On 2 December, it established Honda Power Pack Energy India in Bengaluru, southern India, where the business will kick off in the first half of this year. The new subsidiary will also work with e-rickshaw OEM

A Honda researcher inserts a charged Honda Mobile Power Pack battery into an e-rickshaw.

HONDA MOTOR CO.

manufacturers and provide technical support to integrate its batteries—which will be made in India—into these vehicles.

"Transportation is still responsible for 24 percent of the world's direct $CO_2$ emissions from fuel combustion," says Professor Anita Ho-Baillie, a deputy director of the University of Sydney's Nano Institute. "So we need to seize every opportunity to electrify vehicles." And EVs have other advantages, she adds, "including immediate and smooth acceleration, while electric motors are quieter and have fewer moving parts compared to combustion engines."

Honda introduced its first mobile battery pack in 2017; in 2019, the automaker began testing it in motorbikes and scooters in Indonesia, Japan, and the Philippines. Testing in India began last February using 30 e-rickshaws driven for a total of 200,000 kilometers. "We learned a lot about the issues involved in turning it into a business," says Yoshihiro Nakajima, general manager of the mobile power pack project.

Not surprisingly, the three biggest issues were the lengthy time taken to charge the batteries, the short driving range they provide, and high cost. Honda's answer to the first two problems is the lithium-ion Honda Mobile Power Pack (MPP) battery, working in conjunction with the Honda Mobile Power Pack Exchanger. The exchanger consists of 12 charged MPP batteries in compartmentalized charging chambers that allow a depleted battery to be swapped for a fully energized replacement in a matter of minutes. A maximum of seven exchangers can be hooked together in a single installation.

In the Honda scheme, an e-rickshaw runs on four mobile power pack batteries. Each pack measures 29.8 by 17.7 by 15.6 centimeters, weighs just over 10 kilograms, and comes with handles. Nakajima says the replacement procedure lets the driver exchange two of the four batteries in one minute or less. Nakajima has even presented video evidence to back up this claim. The e-rickshaw driver touches the exchanger with a subscriber ID card, and the exchanger indicates which two charged batteries to remove. The driver plugs these into the nearby vehicle and inserts the depleted batteries into the exchanger for recharging. The procedure is repeated for the other two MPPs.

The concept benefits both users and suppliers, says Ho-Baillie. "It saves time and makes the drivers' job much easier because they don't have to stand around waiting for the batteries to be charged. As for the suppliers, portable battery packs mean they have more flexibility in terms of where they build their charging infrastructures."

Honda's latest iteration of the battery, MPP e, is rated at 50.3 volts, with a capacity of 1.3 kilowatt-hours, and a charging time of 5 hours. With Asian tropical climates in mind, the battery is designed to dissipate heat and prevent discharge deterioration in high temperatures. It is also resistant to water, vibration, and shock under normal operating conditions. A built-in control unit uses sensors to keep tabs on internal conditions and document usage patterns for the exchanger to assess.

"The cylindrical cells used in the battery pack are the same as those widely used in electronic products such as personal computers," says Nakajima. "Because their efficiency is improving year by year, we can get the benefits of cell evolution without changing the outer shape of the battery pack, and so maintain backward compatibility."

An e-rickshaw can travel just 60 km before needing to swap the batteries, which will presumably require numerous visits to charging stations throughout the day. To lessen the toll taken by what might otherwise be a draining chore, Honda will provide drivers with a smartphone app that displays a map indicating the locations of all the charging stations in the area, together with the number of charged batteries available at each station. And by initially limiting the service first to Bengaluru and then Mumbai, Nakajima says they can ensure there will be sufficient stations for the service to work smoothly from the start. He adds that during the earlier testing period, they found that having three stations within a 5-km radius met operational goals.

To tackle the issue of high cost, Honda points out that the specifications of the MPP conform to the safety standard UN R136 adopted by the United Nations. And to promote its common use and reduce costs via economies of scale, Honda has helped create two consortia: one covering the Japanese market with members Kawasaki, Suzuki, and Yamaha and the other for the European market, with KTM, Piaggio, and Yamaha on board.

The company is also exploring additional ways to expand the use of its battery pack, such as in small construction machines, in homes as backup power for electric appliances, as well as establishing battery-sharing systems for the civil engineering and construction industries.

Portable power packs also mean "you can 'transport' energy from and to different localities," says Ho-Baillie. "Should the suppliers want to charge the batteries using renewables, they can strategically build charging stations in places where it is windy and/or sunny. Charging the batteries using renewables instead of electricity from coal-fired-power stations would further reduce $CO_2$ and toxic emissions." ■

## Sun Roasted

The sun rising and an alarm blaring might be enough to rouse us from slumber. But for millions of us, the morning isn't off to a proper start until we've had a hot cup of coffee. Now, two Italian engineers have figured out how to use the sun to roast the coffee beans that end up as our morning brew. This setup, shown last October in Rome, can roast up to 50 kilograms of coffee beans in an hour. According to its inventors, if all the coffee harvested globally were roasted using the sun instead of heat generated by electricity or directly from the combustion of fossil fuels, it would keep roughly 4 billion kilograms of $CO_2$ from being released into the atmosphere annually.

ANDREW MEDICHINI/AP


TORREFAZIONE
SOLARE
PUROSOLE

# Hands On

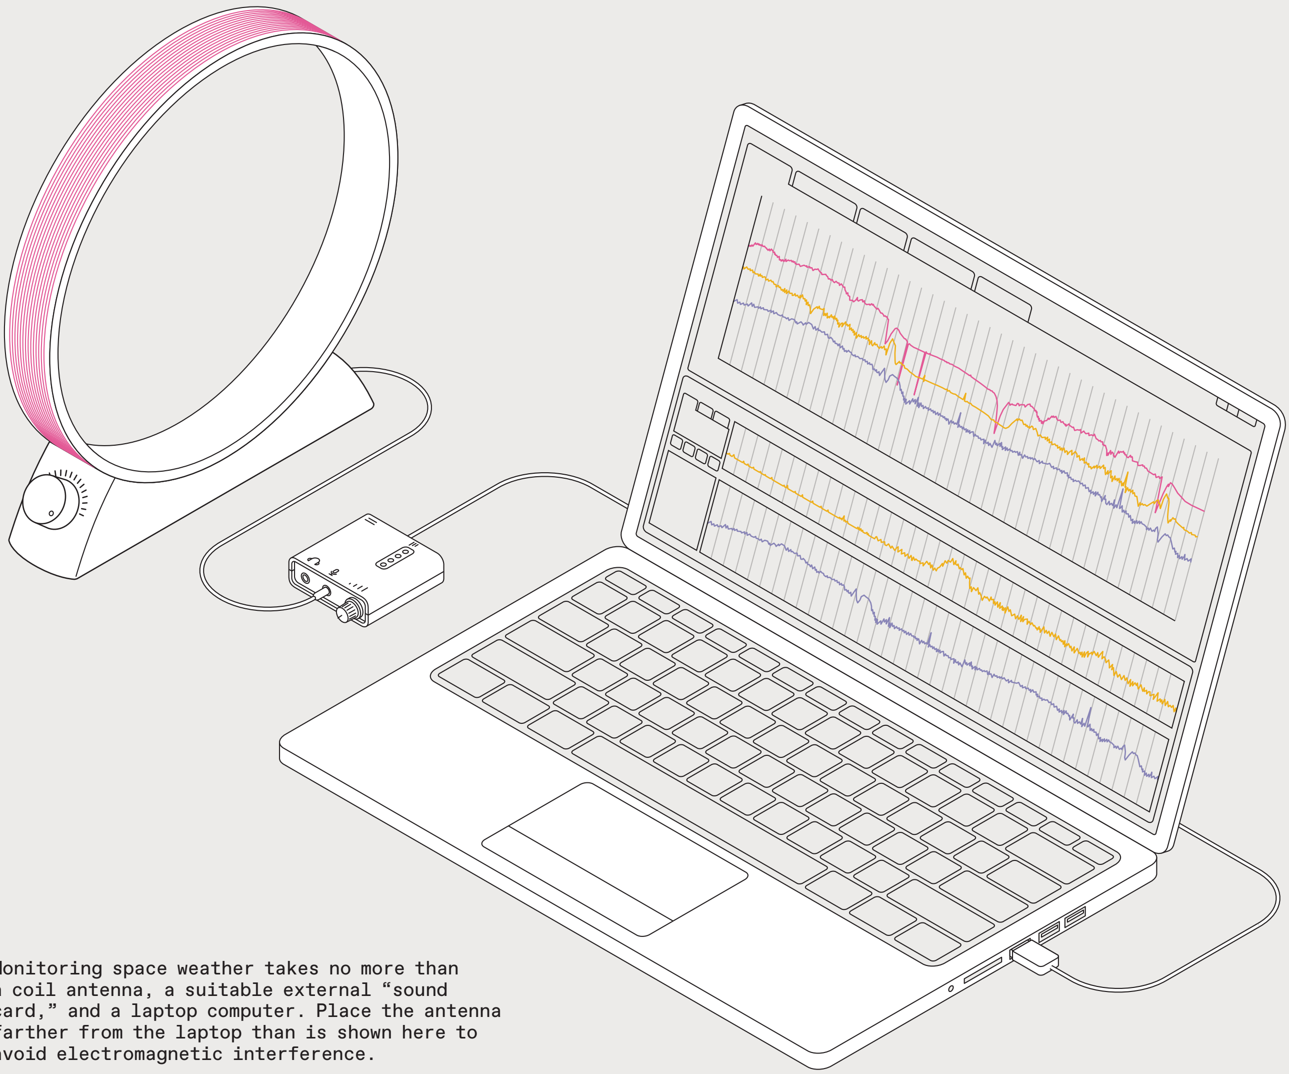

Monitoring space weather takes no more than a coil antenna, a suitable external "sound card," and a laptop computer. Place the antenna farther from the laptop than is shown here to avoid electromagnetic interference.

## A Barometer for Space Weather › Detect solar flares and gamma-ray bursts for less than $100

DAVID SCHNEIDER

**In the 1960s** and '70s, musicians would sometimes insert into their releases odd sounds that could be made intelligible only by rotating the vinyl record backward using your finger. If you suspect this is only an urban legend, load a digital version of Electric Light Orchestra's 1975 recording of "Fire on High" into an audio editor like Audacity and play it in reverse. You'll hear ELO drummer Bev Bevan very clearly say, "The music is reversible, but time…turn back, turn back, turn back."



*Illustrations by James Provost*

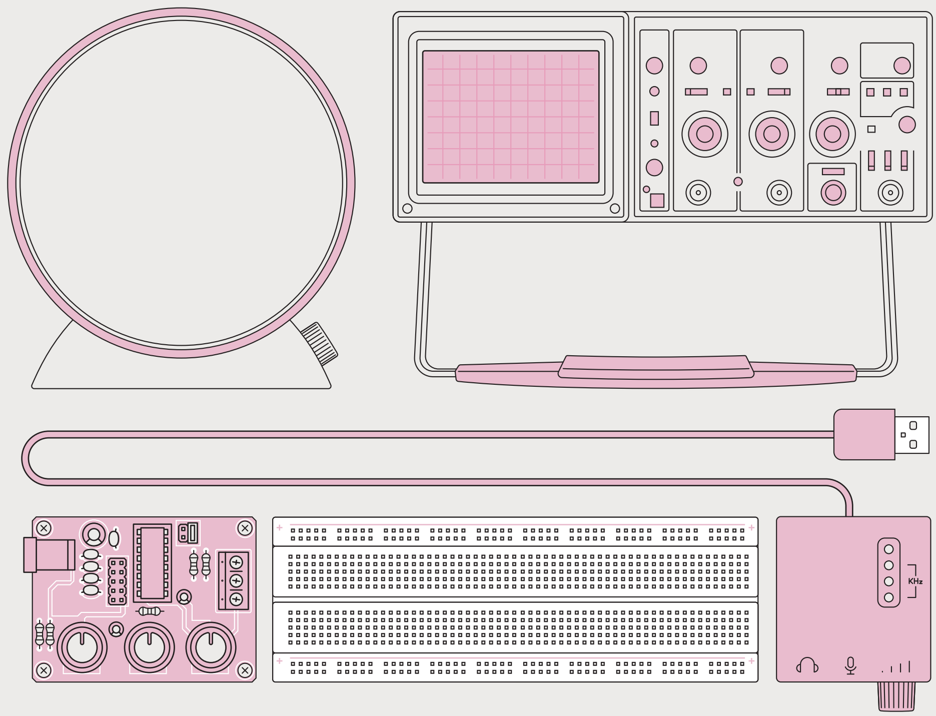

The solar-flare monitor consists of a coil antenna and external "sound card," which connects to a laptop computer. Tuning the coil antenna to an appropriate frequency also required a signal generator, a protoboard, and an oscilloscope.

In the 1980s, vinyl records gave way to compact discs, which weren't amenable to such "backmasking." But at least one CD of that era contains a hidden message: Virgin Records' 1983 release of the album *Tubular Bells*, recorded a decade earlier at Richard Branson's Manor Studio in Shipton-on-Cherwell, England.

You see, an hour's drive north from Shipton is a suburb of Rugby called Hillmorton, where at the time the British government operated a very-low-frequency (VLF) radio station to send messages to submarines. It seems the powerful emanations from this nearby station, broadcast at a radio frequency of just 16 kilohertz (within the audio range), were picked up by the electronic equipment at Branson's studio and recorded at a level too low for anyone to notice.

After learning of this, I purchased an old CD of *Tubular Bells*, ripped a WAV file of one track, and piped it into a software-defined-radio package. Tuning to 16 kHz and setting the SDR software to demodulate continuous-wave signals immediately revealed Morse code. I couldn't copy much of it, but I could make out many repetitions of VVV ("testing") and GBR (the station's call sign).

This inadvertent recording aptly demonstrates that VLF transmissions aren't at all hard to pick up. And these signals can reveal more than just the presence of a powerful radio transmitter nearby. The application I had in mind was to use changes in VLF-signal strength to monitor space weather.

That's possible because these VLF transmissions travel over large distances inside the globe-encircling waveguide that is formed by the Earth's surface and the ionosphere. Solar flares—and rare astronomical events called gamma-ray bursts—can alter the ionosphere enough to change how radio signals propagate in this waveguide. I hoped to use VLF broadcasts to track such goings-on.

There's a long history of amateur astronomers using VLF radio equipment to measure solar flares by the sudden ionospheric disturbances (SIDs) they spawn. Years ago, it was a challenge to build suitable gear for these observations, but it now takes just a few modest Amazon purchases and a laptop computer.

The first item needed is a simple coil antenna. The model I bought (US $35) actually contains two coils. One was connected to a variable capacitor, so it can be tuned to various AM-broadcast frequencies; the other coil, of just two

**There's a long history of amateur astronomers using VLF radio to measure solar flares by the ionospheric disturbances they spawn.**

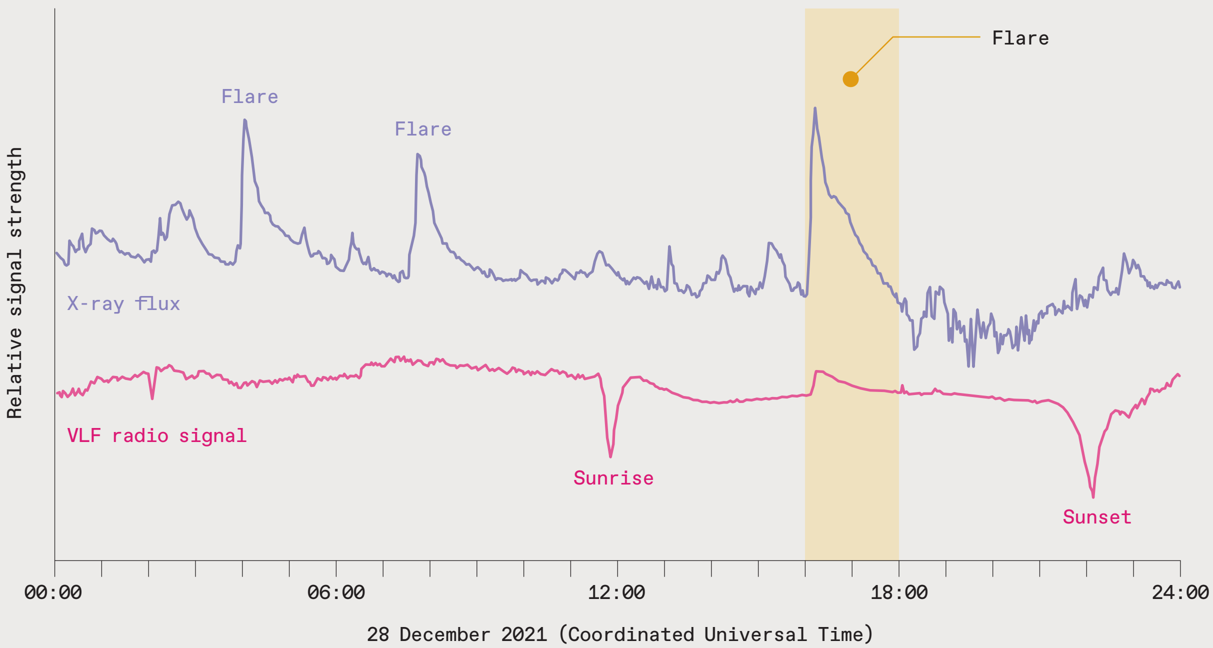


Within days of its construction, the monitor registered the signal from a solar flare. Two flares earlier that day, which were well documented in X-ray measurements taken by NASA's GOES-16 satellite in geosynchronous orbit [purple line], did not affect the VLF measurements [magenta line] because they occurred when the relevant part of the ionosphere was in darkness and shielded from the sun.

turns and inductively coupled to the first one, was wired to the output jack. I bypassed that two-turn coil and wired the jack directly to the wider coil, adding a couple of capacitors in parallel across it to lower its resonant frequency to 25 kHz.

To choose the right capacitors, I purchased a $9 signal generator, also on Amazon, temporarily connected that wider coil in series with a 1,000-ohm resistor, and applied a sinusoidal signal to this circuit. I used an oscilloscope to identify the frequency that caused the alternating voltage across the coil to peak. With some experimentation, I was able to find a couple of ceramic capacitors (nominally 0.11 microfarads in total) to place in parallel with the coil to set the resonant frequency near the broadcast frequency of some U.S. Navy VLF transmitters.

Using a scrounged 3.5-mm plug, I then plugged the modified coil antenna into the mic input of an external "sound card," having purchased one for $34 on Amazon that allowed a sampling rate of 96 kHz. This feature was key, because my plan was to tune in to a station that the U.S. Navy operates in Cutler, Maine, which goes by the call sign NAA and broadcasts at 24 kHz. Fans of Harry Nyquist will remember that you need to sample a signal at least two times per cycle to capture it properly. So a typical sound card that samples at 44 kHz wouldn't cut it.

The final thing I needed was suitable software. I experimented with two SDR packages (HDSDR and SDR Sharp), with my sound card taking the place of the usual radio dongle. While these packages displayed transmissions from NAA clearly enough, they didn't provide a good way to monitor variations in signal strength over time. But I soon discovered how to do that with Spectrum Lab, following an online tutorial explaining how to use this software to measure SIDs.

This combination of desktop AM antenna, external sound card, and Spectrum Lab software proved ideal. With it, I am not only able to monitor NAA, located about 1,400 kilometers from my home in North Carolina, I can also pick up the VLF station in LaMoure, N.D. (call sign NML), which transmits on 25.2 kHz. At times, I clearly receive the Jim Creek Naval station (NLK), near Oso, Wash., on 24.8 kHz and can even register the Navy's Aguada station in Puerto Rico, despite it transmitting at 40.75 kHz, far from my coil's resonant frequency.

The first few days of using this gear captured the expected pattern of daily variation in the signal from NAA, with sharp transitions when the sun rises and sets. Within a week, the sun became unusually active, producing three good-size flares in one day—as documented by NASA's Geostationary Operational Environmental Satellites, which measure X-ray flux in geosynchronous orbit. Two of those flares occurred when the East Coast was in darkness, so they had no effect on the relevant portion of the ionosphere or the signal strength I was monitoring. But the third, which took place at about 11 a.m. local time, showed up nicely.

It's rather amazing that with just $70 worth of simple electronics and a decade-old laptop, I can now monitor flares on the surface of the sun. One day I might see the effects of a gamma-ray burst taking place on a star in a distant galaxy, as a group at Stanford did in 2004. I'll probably have to wait years to detect one of those, though. In the meantime, I can entertain myself hunting for more radio signals inadvertently recorded at the Manor Studio in the '70s. Maybe I'll start those explorations, fittingly, with Van Morrison's 1978 album *Wavelength*. ■

# Geek Life

Also known as water bears, the largest tardigrades are just visible to the naked eye.

EYE OF SCIENCE/SCIENCE SOURCE

## Schrödinger's Tardigrade

### › Have researchers quantum-entangled hardy critters?

BY PHILIP E. ROSS

**Schrödinger's cat is** a thought problem that involves putting a cat in a box that would fill with poison if a radioactive atom splits. According to quantum mechanics, such a split can be said to have happened only if observed. And consequently, the cat must also be considered neither dead nor alive until observed.

Later researchers turned this thought problem into real experiments. One experiment used a resonator chilled nearly to absolute zero so that it became "entangled" across two quantum states, vibrating and not. Those two states were then shown to be superposed—the resonator was both vibrating and not vibrating at the same time.

But a resonator is not alive. To entangle a life-form, you have to cool it nearly to absolute zero without killing it. Bacteria have been so entangled. Now a group of scientists say they've entangled a multicellular organism, a microscopic tardigrade. The 11 researchers published their work last December in the non-peer-reviewed online preprint server arXiv.

The scientists study superconducting qubits for quantum computing. They wondered what would happen if they put a tardigrade on top of one of their qubits and brought the system to near absolute zero. A tardigrade is about as tough as an animalcule gets. Insult the thing and it goes dormant by curling up into a ball, called a tun. A tun is perhaps best characterized as a life that's been put on hold.

Not only did the tardigrade survive temperatures at which no metabolism can happen, the researchers believe they achieved entanglement between it and the qubit, a claim that has provoked controversy.

"We start with a superconducting qubit at energy state 0, comparable to an atom in the ground state; there's no oscillation—nothing is happening," says coauthor Rainer Dumke of the Center for Quantum Technologies, in Singapore. "We can use microwaves to supply exactly the right amount of energy for the right amount of time to raise this to level 1; this is like the second orbital in an atom. It is now oscillating.

"Or...we can add exactly that much energy but supply it for just half the time to raise the system to a quantum state of ½, which is the superposition state. In this state, it is at the same time oscillating and not oscillating."

The researchers found that the system consisting of the qubit and the tardigrade together occupied a lower energy state than either one alone would have occupied. They concluded that the two things had been entangled.

Criticism was swift. Ben Brubaker, a physicist turned journalist, argued on Twitter that other possibilities hadn't been ruled out—for example, quantum entanglement might have been achieved with only part of the tardigrade.

The authors admit that they could not perform the perfect experiment, which would involve measuring the tardigrade and the qubit independently, using two probes. Their tardigrade comes packaged with the qubit, forming a hybrid structure, and so two probes are hard to manage.

Kai Sheng Lee, another author of the tardigrade paper, says that the criticism is at least partially answered "when we introduce the second qubit." The presence of two superconducting qubits beside the tardigrade strengthens the case for entanglement, because here it seems the creature is in superposition with one qubit that's in the 0 state and also with the other qubit, which is in the 1 state.

For now, the question is likely to be unsettled until direct measurements of each element can be made, which Dumke suggests might be done by finding "a particular resonance frequency inside the tardigrade, then [using] this frequency to find what leads to a stronger entanglement." ■

# Crosstalk

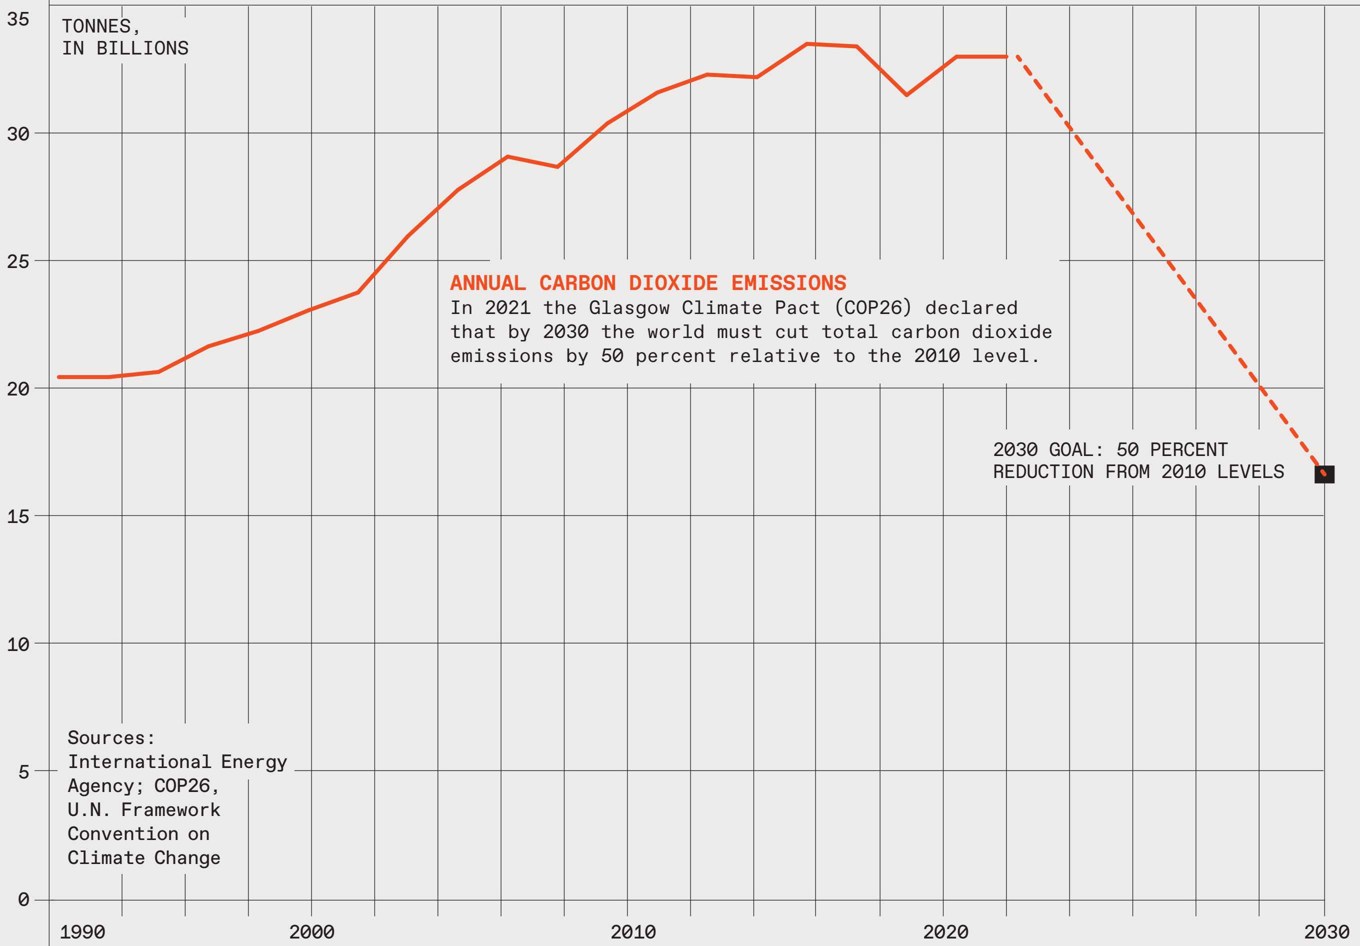


# Decarbonization Algebra

## The COP26 calls for impossibly steep cuts in carbon emissions

**Three months ago** the Glasgow Climate Pact (COP26) declared that by 2030 the world must cut total carbon dioxide emissions by 50 percent relative to the 2010 level, which was 30.4 billion tonnes. This would bring annual emissions to less than 20 billion tonnes, a level last seen more than 30 years ago.

What are the chances of that? Let's look at the arithmetic.

First, assume that all energy-consuming sectors share the cuts equally and that global energy demand stays constant (instead of increasing by 2 percent a year, as it did in the prepandemic decade). Today our best commercial batteries have energy densities of about 300 watt-hours per kilogram, less than 3 percent as much as kerosene; among some 25,000 planes in the global commercial fleet, there is not a single high-capacity electric or hydrogen-powered aircraft. A 50 percent cut in kerosene-fueled flying would mean that by 2030 we would have to build about 12,000 new airplanes with capacities of from 100 people (the Embraer 190) to 400 people (the Boeing 777-300ER), all powered by as-yet-nonexistent superbatteries or equally nonexistent hydrogen systems. That's what we'd need to fly

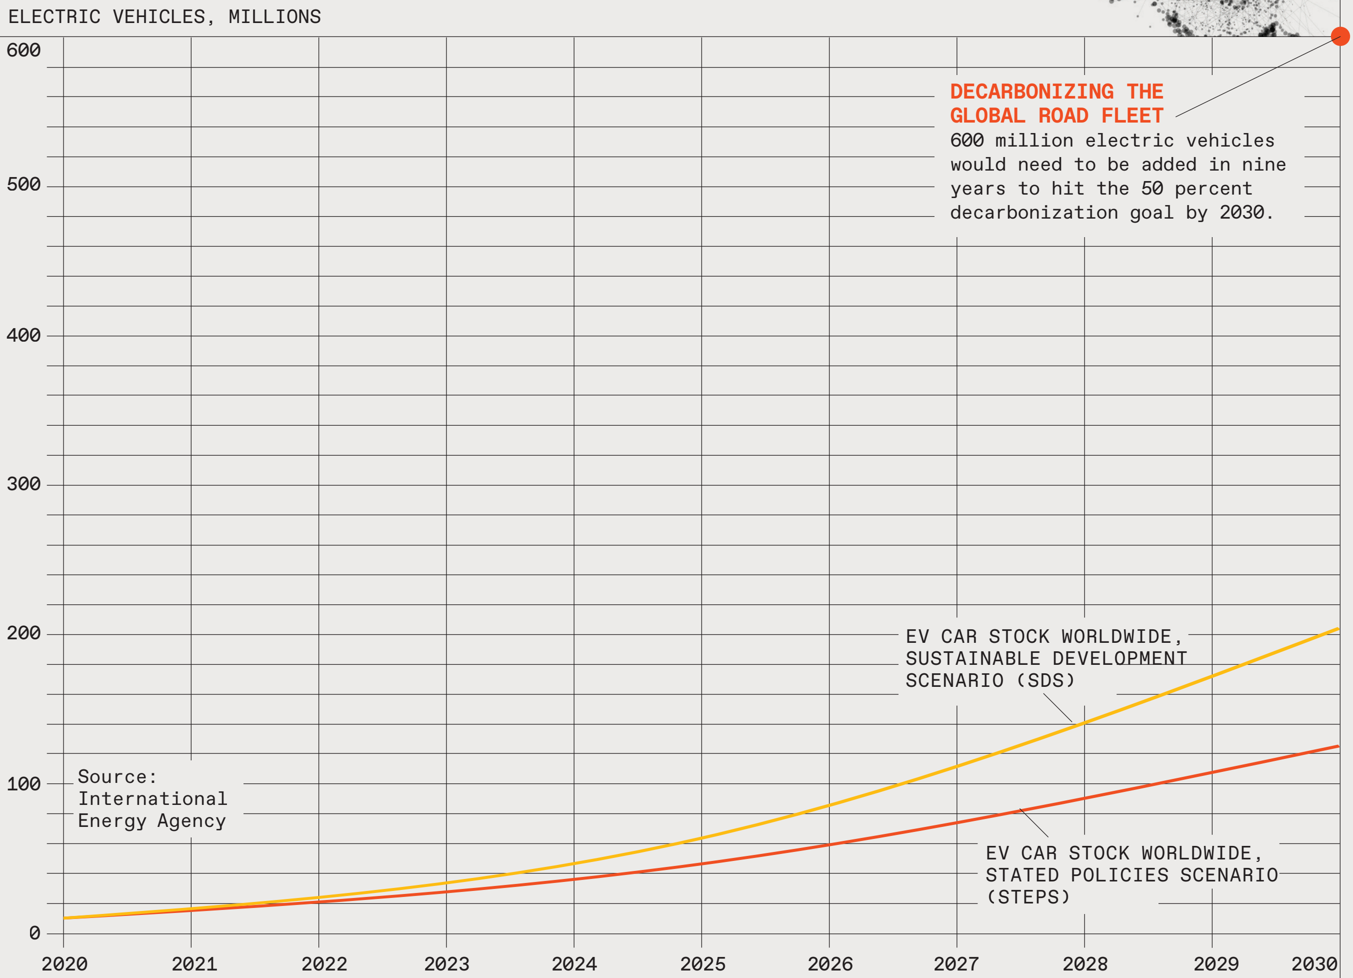


about 2.2 billion passengers a year, for a total of about 4.3 trillion carbon-free passenger-kilometers. What are the chances of that?

In 2019 the world produced 1.28 billion tonnes of pig (cast) iron in blast furnaces fueled with coke made from metallurgical coal. That pig iron was charged into basic oxygen furnaces to make about 72 percent of the world's steel (the rest comes mostly from electric arc furnaces melting scrap metal). Today there is not a single commercial steel-making plant that reduces iron ores by hydrogen. Moreover, nearly all hydrogen is now produced by the reforming of natural gas, and zero-carbon iron would require mass-scale electrolysis of water powered by renewable energies, something we still haven't got. A 50 percent cut of today's carbon dependence would mean that by 2030 we would have to smelt more than 640 million tonnes of iron–more than the annual output of all of the blast furnaces outside China–-by using green hydrogen instead of coke. What are the chances of that?

**Decarbonizing the global fleet of cars by 50 percent in nine years would require that we manufacture 66 million EVs a year, nearly as much as the total global production of all cars in 2019.**

In 2021 there were some 1.4 billion motor vehicles on the road, of which no more than 1 percent were electric. Even if the global road fleet were to stop growing, decarbonizing 50 percent of it by 2030 would require that we manufacture about 600 million new electric passenger vehicles in nine years—that's about 66 million a year, more than the total global production of all cars in 2019. In addition, the electricity to run those cars would have to come from zero-carbon sources. What are the chances of that?

To set goals that correspond to available technical capabilities while taking into account reasonable advances in the production and adoption of non-carbon energy sources, we must start with grade-school algebra. What are the chances of that? ■

PORTRAITS BY SEREGIO ALBIAC

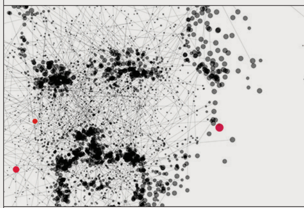

# Wi-Fi 7 Stomps on the Gas

Today's Wi-Fi 6 tops gigabit rates; the coming standard will be four times as fast


**Consumer technology is** often a story of revolutionary leaps followed by a descent into familiarity. The first computers advanced so quickly that new models went obsolete while they were still on store shelves. Today, any US $500 laptop will be relevant for a decade. A similar story can be told of smartphones, TVs, even cars.

Yet there is one technology that has escaped this trend: Wi-Fi.

Wi-Fi went mainstream with the 802.11g standard in 2003, which improved performance and reliability over earlier 802.11a/b standards. My first 802.11g adapter was a revelation when I installed it in my ThinkPad's PC Card slot. A nearby café jumped on the trend, making a midday coffee-and-classwork break possible. That wasn't a thing before 802.11g.

Still, 802.11g often tried your patience. Anything but an ideal connection left me staring at half-loaded Web pages. I soon learned which spots in the café had the best connection.

Wi-Fi 6, released in 2019, has maximum speeds of 600 megabits per second for the single band and 9,608 Mb/s on a single network. That's more than 175 times as fast as the 802.11g connection I used in 2003.

Those figures don't tell the whole story. Peak Wi-Fi speeds require support on each device for multiple "spatial streams"—that is, for multiplexed channels. Modern Wi-Fi can support up to eight spatial streams, but most consumer-grade Wi-Fi adapters support just one or two streams, to keep costs down. Fortunately, Wi-Fi 6 boosts the performance per stream enough to lift even entry-level Wi-Fi adapters above gigabit speeds.

That's key, as gigabit Internet remains the best available to most people across the globe. I'm lucky enough to have gigabit service, and I've tested quite a few Wi-Fi 6 devices that hurdle this performance bar. It renders gigabit Ethernet nearly obsolete, at least for most home use.

**Such extreme bandwidth is obviously overkill for Web browsing, but it's a necessity for streaming augmented- and virtual-reality content.**

Wi-Fi 6E, released in 2020, further improves the standard with a 6-gigahertz band that appears as a separate connection, just as 2.4- and 5-GHz bands have appeared separately on prior Wi-Fi networks. It's early days for Wi-Fi 6E, so device support is limited, but the routers I've tested were extremely consistent in hitting the peak potential of gigabit Internet.

Wi-Fi 6 already outperforms the Internet service available to most people. Yet the standard isn't letting off the gas. MediaTek plans the first demonstration of Wi-Fi 7 at CES 2022 (the standard is expected to be released in 2024). Wi-Fi 7 is expected to boost maximum bandwidth up to 40 gigabits per second, four times as fast as Wi-Fi 6. Such extreme bandwidth is obviously overkill for Web browsing, but it's a necessity for streaming augmented- and virtual-reality content.

This rapid improvement stands in contrast to the struggles in cellular networking. In theory, 5G can meet or beat the performance of Wi-Fi, but the reality often falls short.

The performance of 5G varies between markets. A report from OpenSignal found customers of Taiwan's FarEasTone can expect average download speeds of nearly 448 Mb/s. Verizon and AT&T customers in the United States average just 52.3 Mb/s.

The problem of inconsistent 5G is out of the consumer's hands. If you want faster Wi-Fi, you can make it happen by purchasing a new router and, possibly, an adapter for older devices. But if you want faster mobile bandwidth data, tough luck.

Perhaps cellular providers will get their act together and bring the best 5G speeds beyond dense urban centers. Until then, Wi-Fi is the way to go if you want maximum bandwidth without a cord. ■



*Illustration by Harry Campbell*

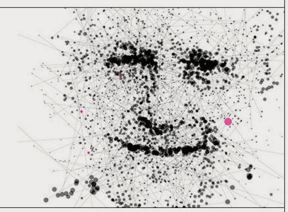

# Claude Shannon's Greatest Hits

His legacy goes far beyond information theory

**Among the great** engineers of the 20th century, who contributed the most to our 21st-century technologies? I say: Claude Shannon.

Shannon is best known for establishing the field of information theory. In a 1948 paper, one of the greatest in the history of engineering, he came up with a way of measuring the information content of a signal and calculating the maximum rate at which information could be reliably transmitted over any sort of communication channel. The article, titled "A Mathematical Theory of Communication," describes the basis for all modern communications, including the wireless Internet on your smartphone and the analog signal on a twisted-pair landline.

If information theory had been Shannon's only breakthrough, it would have been enough to secure his place in the pantheon. But he did a lot more.

**While working on his master's thesis at MIT, Shannon invented the logic gate.**

A decade before, while working on his master's thesis at MIT, he invented the logic gate. At the time, electromagnetic relays were used to build circuits that routed telephone calls or controlled complex machines. But there was no consistent theory on how to design or analyze such circuits. The way people thought about them was in terms of the relay coils being energized or not. Shannon showed that Boolean algebra could be used to move away from the relays and into a more abstract understanding of the function of a circuit. He used this algebra of logic to analyze, and then synthesize, switching circuits and to prove that the circuit would work as desired. In his thesis he invented the AND, OR, and NOT logic gates.

In 1950 Shannon published an article in *Scientific American* and also a research paper describing how to program a computer to play chess. He explained in detail how to design such a program. He discussed how data structures would be represented in memory, estimated how many bits of memory would be needed for the program, and broke the program down into things he called subprograms.

Shannon did all this at a time when there were *fewer than 10 computers* in the world. And they were all being used for numerical calculations. He began his research paper by speculating on all sorts of things that computers might be programmed to do beyond calculations, including designing relay and switching circuits, designing electronic filters for communications, translating between human languages, and making logical deductions. Computers do all these things today. Shannon noted that people believed that playing chess required "thinking." Therefore, he reasoned, it would be a great test case for whether computers could be made to think.

Shannon suggested it might be possible to improve his program by analyzing the games it had already played and adjusting the terms and coefficients in its evaluations of the strengths of board positions. There were no computers readily available to Shannon at the time, so he couldn't test his idea. But just five years later, in 1955, Arthur Samuel, an IBM engineer who had access to computers as they were being tested before being delivered to customers, was running a checkers-playing program that used Shannon's exact method to improve its play. And in 1959 Samuel published a paper about it with "machine learning" in the title—the first time that phrase appeared in print.

So, let's recap: information theory, logic gates, non-numerical computer programming, data structures, and, arguably, machine learning. Claude Shannon didn't bother predicting the future—he just went ahead and invented it, and even lived long enough to see the adoption of his ideas. Since his passing 20 years ago, we have not seen anyone like him. We probably never will again. ■

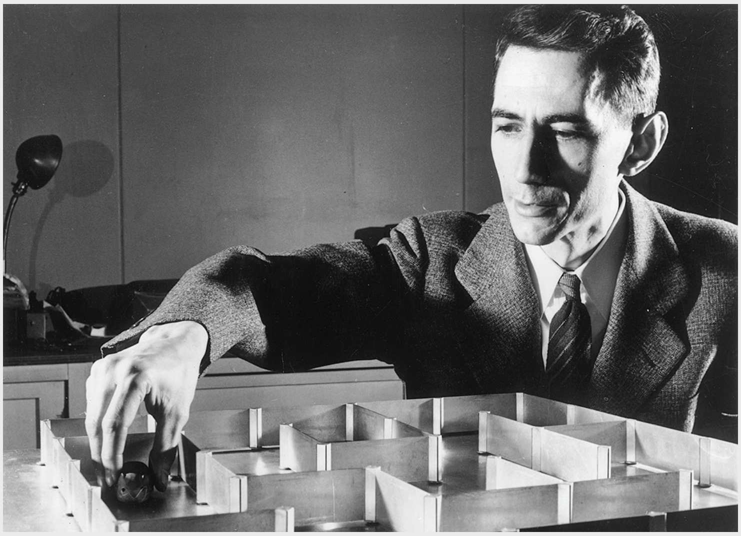

KEYSTONE/GETTY IMAGES

# MISSION TO A METAL WORLD

In August, NASA will launch a probe to study a strange metallic asteroid called Psyche

BY DAN M. GOEBEL & DAVID OH
Illustration by John MacNeill



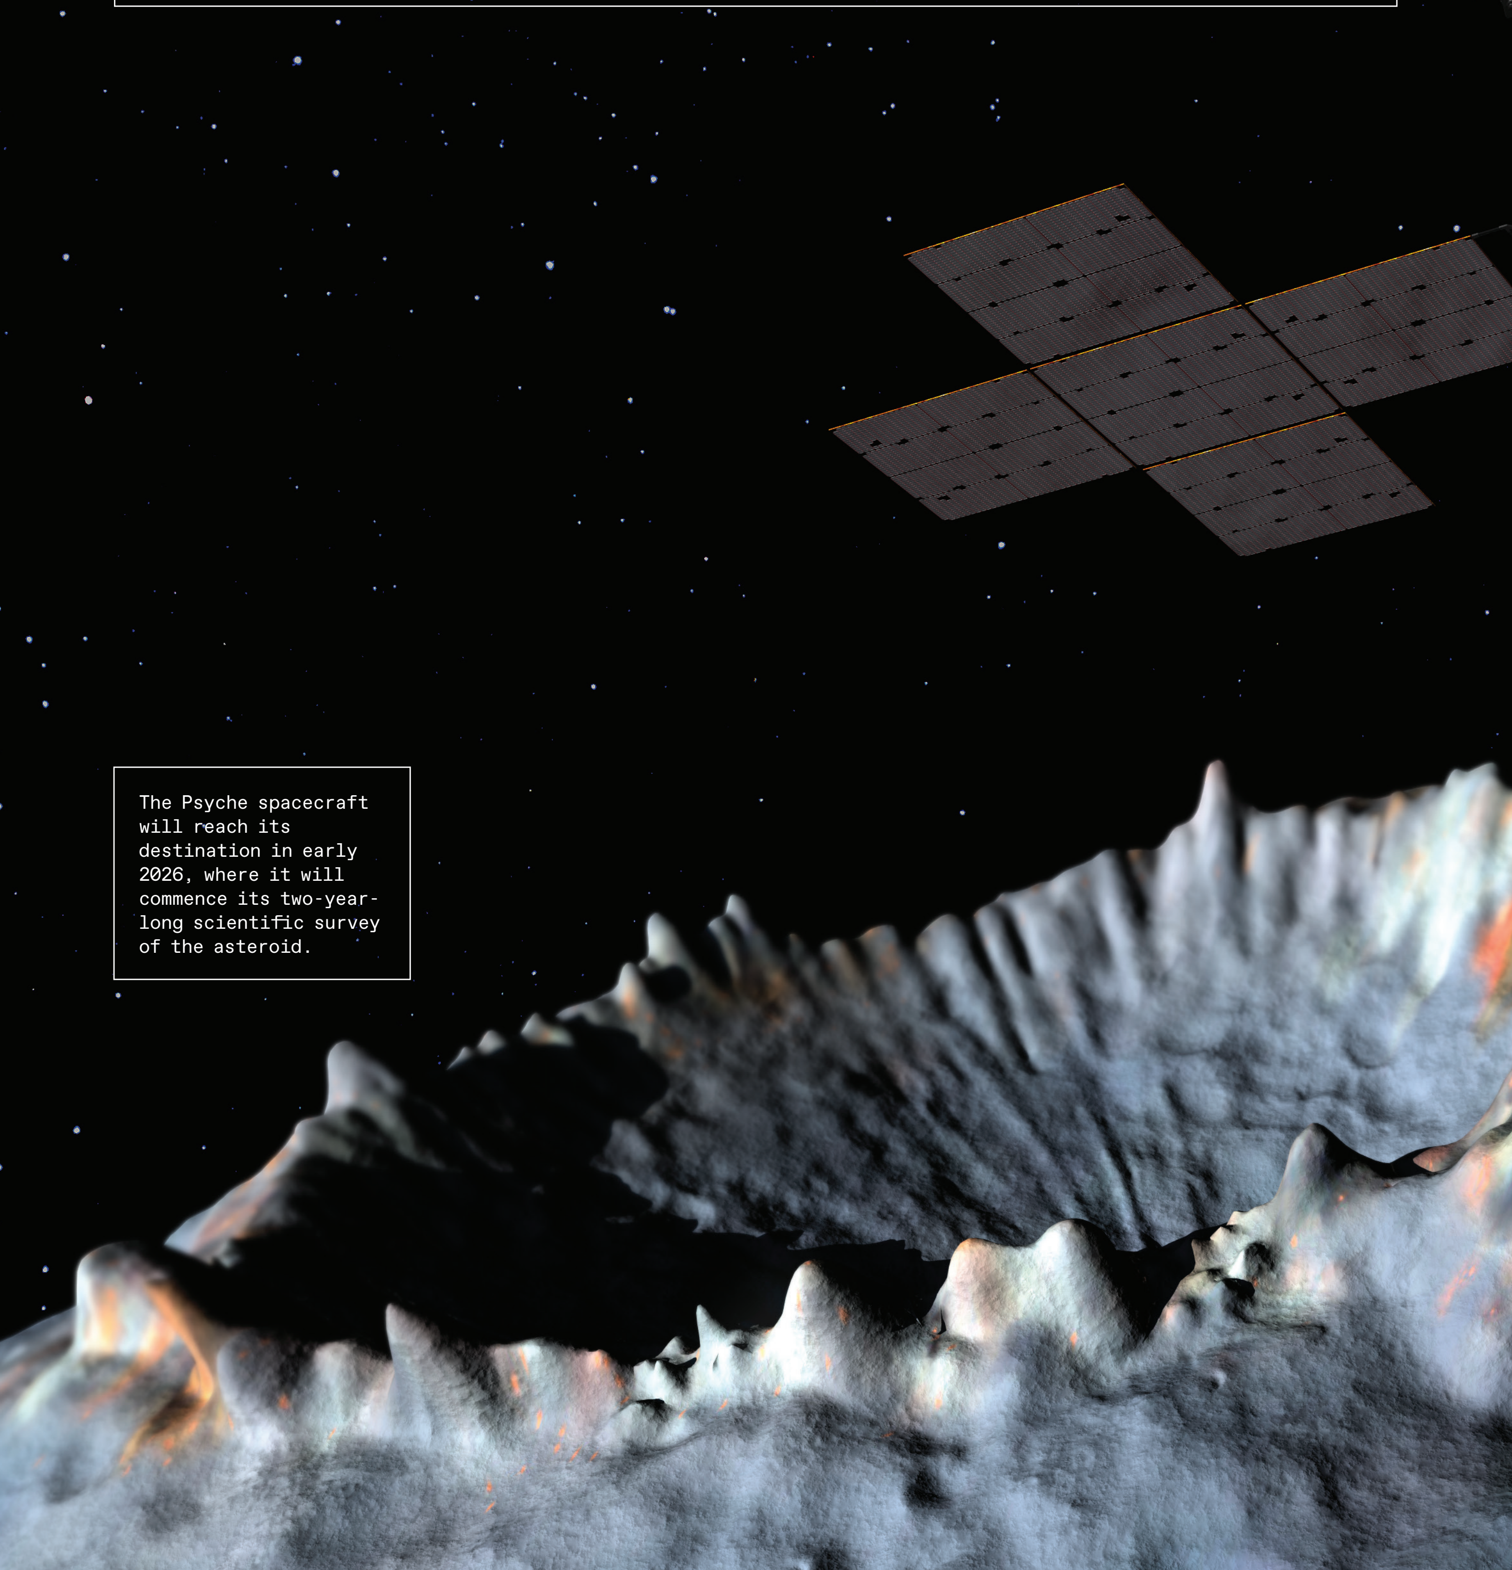

The Psyche spacecraft will reach its destination in early 2026, where it will commence its two-year-long scientific survey of the asteroid.

Left: A high-gain radio antenna shown here being tested at Maxar's facilities in Palo Alto, Calif., will provide the data communications throughout the mission. Center and right: Technicians at NASA's Jet Propulsion Laboratory work on the Psyche spacecraft.

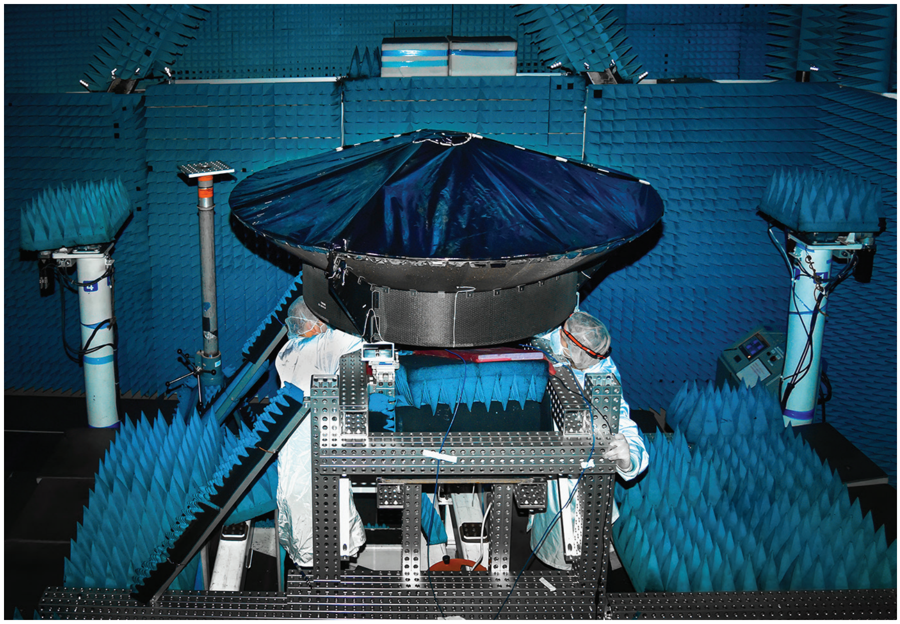

When our solar system was very young, there were no planets—only a diffuse disk of gas and dust circled the sun. But within a few million years, that churning cloud of primordial material collapsed under its own gravity to form hundreds, or maybe thousands, of infant planets. Some of those planetesimals, as astronomers call them, grew to be hundreds of kilometers across as they swept up more dust and gas within the swirling solar nebula.

Once they had attained such a size, heat from the decay of the radioactive elements within them became trapped, raising temperatures enough to melt their insides. The denser components of that melt—iron and other metals—settled to the center, leaving lighter silicates to float up toward the surface. These lighter materials eventually cooled to form mantles of silicate rock around heavy metallic cores. In this way, vast amounts of iron and nickel alloys were trapped deep inside these planetesimals, forever hidden from direct scrutiny.

Or were they?

At this time, the solar system was still relatively crowded despite its vast size. And over the next 20 million or so years, many planetesimals crossed paths and collided. Some merged and grew into even larger protoplanets, eventually forming what became the familiar planets we know today.

In each of those protoplanet collisions, the metallic cores were battered and remixed with silicate mantle material, later separating again after being melted by the heat of accretion. Some collisions had enough energy to completely obliterate a protoplanet, leaving behind debris that contributed to the asteroid belt that now exists between the orbits of Mars and Jupiter.

But a few protoplanets may have escaped either of these fates. Astronomers hypothesize that a series of "hit and run" impacts caused these bodies to lose most of their mantles, leaving behind only a small quantity of silicate rock and a large amount of metal. These materials combined to form a rare kind of world. If this theory is correct, the largest example would be an asteroid called 16 Psyche—named after the Greek goddess of the soul, Psyche, and because it was the 16th member of the asteroid belt to be discovered (in 1852).

16 Psyche is about as wide as Massachusetts and has metal-like density. This makes it large and dense enough to account for a full 1 percent of the total mass of the asteroid belt. Metal miners of the future may one day stake claims on it.

Psyche is also the name of a NASA mission to visit that asteroid. Led by Lindy Elkins-Tanton of Arizona State University and run by NASA's Jet Propulsion Laboratory, the Psyche mission will test astronomers' theories about planetary-core formation and composition while it explores a world with a landscape unlike any that space probes have visited so far.

The Psyche mission is scheduled to launch in August 2022, with the spacecraft reaching its destination more than three years later. What will it find there? Astronomers think we might see enormous surface faults from the contraction of freezing metal, glittering cliffs of green crystalline mantle minerals, frozen flows of sulfur lava, and vast fields of metal shards scattered over the surface from millennia of high-speed impacts. There will no doubt be plenty of surprises, too.

**The long journey this space probe** must make to reach its destination will be especially demanding. 16 Psyche resides in the outer part of the main asteroid belt, well beyond the orbit of Mars. The probe will begin circling the asteroid in January of 2026 and will study it for nearly two years.

Counterintuitively, arranging for a probe to orbit a small body like an asteroid is harder than orbiting a planet. Big planets have deep gravity wells, which allow spacecraft to enter orbit with a single low-altitude rocket burn. Small bodies have little gravity and provide essentially no gravitational leverage, so the spacecraft's propulsion system must do all the work.

Not long ago, NASA managed this maneuver successfully with its Dawn mission, which sent a probe to orbit the asteroids

FROM LEFT: MAXAR (2); JPL-CALTECH/NASA


Vesta and Ceres. The Dawn spacecraft used solar-electric propulsion—three highly efficient engines that convert electricity from solar arrays into thrust by ionizing a propellant gas and accelerating it though a high-voltage electric field.

When our team at JPL was designing the Psyche probe, we planned to do something similar. The main problem was figuring out how to do it without exceeding the mission's budget. Engineers at the Jet Propulsion Laboratory solved this problem by using what was for the most part existing technology, manufactured by Maxar, a company based in Westminster, Colo., that is one of the world's largest providers of commercial geosynchronous communication satellites, produced at a division located in Palo Alto, Calif.

The Psyche spacecraft is built on the "chassis" used for those satellites, which includes high-power solar arrays, electric-propulsion thrusters, and associated power and thermal control elements. In many ways, the Psyche spacecraft resembles a standard Maxar communications satellite. But it also hosts JPL's avionics, flight software, and the many fault-protection systems required for autonomous deep-space operation.

Making this concept work was difficult from the get-go. First, NASA management was rightfully wary of such cost-cutting measures, because the "faster, better, cheaper" model of missions mounted in the 1990s produced some spectacular failures. Second, using Earth-orbiting systems on the Dawn mission resulted in large cost overruns during the development phase. Finally, many people involved believed (erroneously) that the environment of deep space is very special and that the Psyche spacecraft would thus have to be very different from a communications satellite intended only to orbit Earth.

We and our many NASA colleagues addressed each of these issues by teaming with engineers at Maxar. We kept costs under control by using hardware from the company's standard product line and by minimizing changes to it. We could do that because the thermal environment in geosynchronous orbit wasn't in fact so different from what the Psyche probe would encounter.

Soon after launch, the Psyche spacecraft will experience the same relatively high solar flux that communications satellites are built for. It will also have to handle the cold of deep space, of course, but Maxar's satellites must endure similar conditions when they fly through Earth's shadow, which they do once a day during certain times of the year.

Because they serve as high-power telecommunications relays, Maxar's satellites must dissipate the many kilowatts of waste heat generated by their microwave power amplifiers. They do this by radiating that heat into space. Radiating lots of heat away would be a major problem for our space probe, though, because in the vicinity of 16 Psyche the flux of light and heat from the sun is one-tenth of that at Earth. So if nothing were done to prevent it, a spacecraft designed for orbiting Earth would soon become too cold to function this far out in the asteroid belt.

Maxar addressed this challenge by installing multilayer thermal blanketing all over the spacecraft, which will help to retain heat. The company also added custom louvers on top of the thermal radiators. These resemble Venetian blinds, closing automatically to trap heat inside when the spacecraft gets too cold. But plenty of other engineering challenges remained, especially with respect to propulsion.

To reduce the mass of propellant needed to reach the asteroid, the Psyche spacecraft will use solar-electric thrusters that accelerate ions to very high velocities—more than six times as high as what can be attained with chemical rockets. In particular, it will use a type of ion thruster known as a Hall thruster.

Soviet engineers pioneered the use of Hall thrusters in space during the 1970s. And we use four Russian-made Hall thrusters on the Psyche spacecraft for the simple reason that Maxar uses that number to maintain the orbits of their communications satellites.

Hall thrusters employ a clever strategy to accelerate positively charged ions [see sidebar, "How a Hall Thruster Works"]. This is different from what is done in the ion thrusters on the Dawn spacecraft, which used high-voltage grids. Hall thrusters, in contrast, use a combination of electric and magnetic fields to accelerate the ions. While Hall thrusters have a long history of use on satellites, this is the first time they will go on an interplanetary mission.

You might think that thrusting around Earth isn't any different from doing so in deep space. There are, in fact, some big differences. Remember, the power to run the thrusters comes from solar panels, and that power must be used as it is generated—there is no great big battery to store it. So the power available to run the thrusters will diminish markedly as the spacecraft moves away from the sun.

That's an issue because electric thrusters are usually designed to run best at their maximum power level. It turns out to be pretty easy to throttle them a little, maybe to about half their maximum output. For example, the Hall thrusters Maxar uses on its communications satellites can run at as much as 4.5 kilowatts when the satellite's orbit needs to be raised. For more routine station keeping, these thrusters run at 3 kW. We needed these thrusters to run at less than 1 kW when the spacecraft neared its destination.

The problem is that efficiency decreases when you do this kind of throttling. In that sense, a Hall thruster is like the engine in your car. But the situation is worse than in a car: The electrical discharge inside a thruster can become unstable if the power is decreased too much. The throttled thruster can even quit firing altogether—like a flameout in a jet engine.

But with some clever engineering, we were able to make modifications to how we run Maxar's thruster so that it could operate stably at power levels as low as 900 W. We then tested our reengineered thruster in facilities at NASA's Glenn Research Center and at JPL to prove to ourselves that it would indeed operate reliably for the full six-year Psyche mission.

**The Psyche probe will venture** more than three times as far from the sun as Earth ever does. Generating the 2 kW of power needed to operate the spacecraft and fire its thrusters when it reaches its destination requires an array of solar cells large enough to generate more than 20 kW near Earth. That's a lot of power as these things go.

Fortunately for NASA, the cost of solar power has dropped dramatically over the past decade. Today, the commercial satellites that beam television and Internet signals across the globe generate these power levels routinely. Their solar-power systems are effective, reliable, and relatively inexpensive. But they are designed to work while circling Earth, not at the outer edges of the asteroid belt.

When the Psyche mission was conceived in 2013, Maxar had successfully flown more then 20 spacecraft with power levels greater than 20 kW. But the company had never built an interplanetary probe. JPL, on the other hand, had years of experience operating equipment in deep space, but it had never built a power system of the size required for the Psyche mission. So JPL and Maxar combined forces.

The challenge here was more complicated than just dealing with the fact that sunlight at 16 Psyche is so dim. The solar cells on the Psyche spacecraft would also have to operate at temperatures much lower than normal. That's a serious issue because the voltage from such cells rises as they get colder.

When orbiting Earth, Maxar's solar arrays generate 100 volts. If these same arrays were used near 16 Psyche, they would produce problematically high voltages. While we could have added electronics to reduce the voltage coming out of the array, the new circuitry would be costly to design, build, and test for space. Worse, it would have reduced the efficiency of power generation when the spacecraft is far from the sun, where producing adequate amounts of power will be tough in any case.

Fortunately, Maxar already had a solution. When one of their communications satellites passes into Earth's shadow, it's powered by a bank of lithium-ion batteries about the size of what's found in electric cars. That's big enough to keep the satellite running while it is in darkness behind Earth, which is never for much longer than an hour. But the voltage from such batteries varies over time—perhaps from as low as 40 V on some satellites when the battery is deeply discharged all the way up to 100 V. To handle that variability, Maxar's satellites include "discharge converters," which boost voltage to provide power at a constant 100 V. These converters were flight proven and highly efficient—ideal to repurpose for Psyche.

The key was to rewire the solar array, lowering the voltage it produced in the vicinity of Earth to about 60 V. As the spacecraft moves away from the sun, the voltage will gradually rise as the arrays get colder until it reaches about 100 V at 16 Psyche. Maxar's discharge converters, normally attached to batteries, are connected to the solar array instead and used to provide the spacecraft with power at a constant 100 V over the entire mission.

This approach incurs some energy losses, but those are greatest when the spacecraft is close to Earth and power is abundantly available. The system will operate at its highest efficiency when the spacecraft nears 16 Psyche, where generating power will be a lot harder. It uses flight-proven

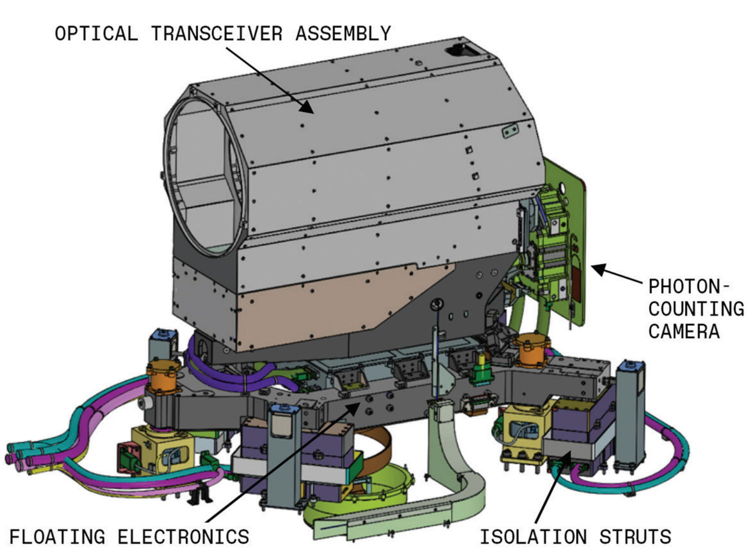


The Psyche mission will test equipment for sending and receiving data optically. This Deep Space Optical Communications (DSOC) system must be pointed with great precision and kept isolated from vibration.

JPL-CALTECH/ARIZONA STATE UNIVERSITY/NASA

## HOW A HALL THRUSTER WORKS

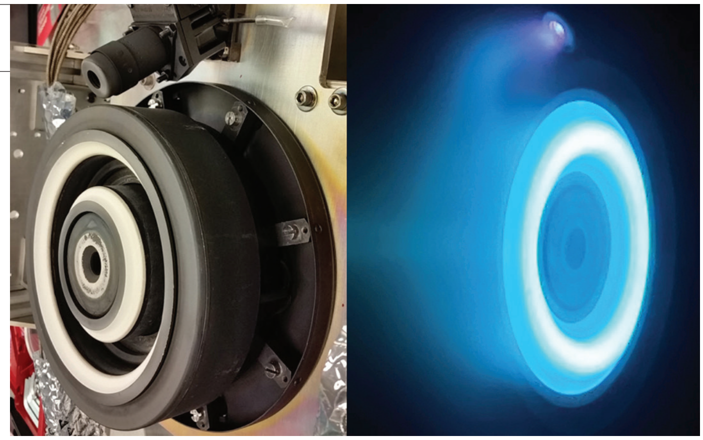

Hall thrusters use an electron discharge to create a plasma—a quasi-neutral collection of positive ions and electrons—not unlike what goes on in a fluorescent lamp.

These thrusters include a hollow cathode (negative electrode), placed outside the thruster body, and an anode (positive electrode) positioned inside a ring-shaped discharge chamber. If these electrodes were all there was, the power applied to the thruster would just go into making a current of electrons flowing from cathode to anode, emitting some blue glow along the way. Instead, a Hall thruster applies a radially directed magnetic field across its discharge channel.

The electrons emitted by the cathode are very light and fast. So this magnetic field impedes the flow of electrons to the anode, forcing them instead to go in circular orbits around the center line of the thruster. The positive xenon ions that are generated inside the discharge chamber accelerate toward the cloud of circling electrons, but these ions are too massive to be affected by the weak magnetic field. So they shoot straight out in a beam, sweeping up electrons along the way. The ejection of that material at high speed creates thrust. It's not much thrust—equal to about the weight of a few quarters—but applied steadily for months on end, it's enough to get the spacecraft zooming.

hardware and is far more economical than sophisticated systems designed to eke out peak power from a solar array throughout a deep-space mission.

**In addition to the set of scientific instruments** that will be used to study the asteroid, the Psyche spacecraft will also be carrying what NASA calls a "technology demonstration" payload. Like so many things at NASA, it goes by an acronym: DSOC, which stands for Deep Space Optical Communications.

DSOC is a laser-based communications system intended to outdo current radio technology by as much as a hundredfold. DSOC will demonstrate its capability by transmitting data at up to 2 megabits per second from beyond the orbit of Mars. One day similar technology may enable you to watch astronauts tromping around the Red Planet in high-definition video.

The DSOC instrument has a "ground segment" and a "flight segment," each of which includes both a laser transmitter and a receiver. The transmitter for the ground segment, a 7-kW laser, will be installed at JPL's Optical Communications Telescope Laboratory, located about 60 kilometers northeast of Los Angeles. A sensitive receiver, one capable of counting individual photons, will be attached to the 5.1-meter-wide Hale Telescope at Caltech's Palomar Observatory, located a similar distance northeast of San Diego.

The flight segment, the part on the spacecraft, contains the same type of equipment, but much scaled down: a laser with an average power of 4 watts and a 22-centimeter telescope. The flight segment sounds simple, like something you could cobble together yourself at home. In fact, it's anything but.

For one, it needs some rather elaborate gear to point it in the right direction. The Psyche spacecraft itself is able to keep the DSOC pointed toward Earth to within a couple of milliradians—about a tenth of a degree. Using built-in actuators, the DSOC then searches for the laser beacon sent from the ground. After detecting it, the actuators stabilize the pointing of the DSOC's own laser back at Earth with an accuracy measured in microradians.

The flight segment is able to point so steadily in the same direction because it's housed in a special enclosure that provides thermal and mechanical isolation from the rest of the spacecraft. DSOC also uses a long sun shield to eliminate stray light on its laser receiver, with a deployable aperture cover to ensure that the unit remains clean.

During DSOC operations in space, the spacecraft cannot use its thrusters or gimbal its solar arrays, which would introduce problematic movements. Instead, it will keep its attitude fixed solidly in one direction and will use its star-tracking system to determine what that direction is. The constraints on what the spacecraft can do at these times is not an impediment, though, because DSOC will be used only for tests during the first year of the mission, while traveling to just past the orbit of Mars. When the spacecraft reaches 16 Psyche, it will transmit data back to Earth over a microwave radio link.

**Having emerged from nearly a decade of planning,** and having traveled for more than three years, the Psyche spacecraft will finally reach its target in early 2026. There will no doubt be plenty of tension in the air when controllers at JPL maneuver the spacecraft into orbit, waiting the many minutes it will take signals to be returned to find out whether all went well in this distant corner of the asteroid belt.

If all goes according to plan, for the following two years this communications-satellite-turned-space-probe will provide scientists with a close-up look at this odd metallic world, having already demonstrated an advanced optical system for high-data-rate communications. These achievements will have been a long time coming for us—but we expect that what is learned will be well worth the many years we've put into trying to ensure that this mission is a success. ■

JPL-CALTECH/NASA


IEEE spectrum
Laser light diffraction
The Odyssey
by Homer
Rendered into English Prose by
Samuel Butler
POPULAR ELECTRONICS
APRIL 1955
25 CENTS
RADIO · TV · R/C · HI-FI · ELECTRONICS
USEFUL MINIATURE POWER SUPPLY
NOVICE 25 WATT TRANSMITTER
BUILDING LOW COST MICROPHONES
HOW TO SOLDER · TRANSFORMER SALVAGE
OHM'S LAW IN PHOTOS · VTVM OR VOM?
THE THEREMIN
MODERN MECHANIX
July
& INVENTIONS MAGAZINE
NOW 15¢
GIRO-CAR
SEE PAGE 37
Grappling With Death Under the Sea
Debunking Poison Gas War Scares
ALICE IN WONDERLAND
LEWIS CARROLL
EXPLORING the SUPERNATURAL!
FORBIDDEN WORLDS
10¢
TERROR ISLAND!
HELP! HELP!
THE HOUND OF THE BASKERVILLES
A.G.J.
CONAN DOYLE
The GREAT GATSBY
F·SCOTT·FITZGERALD
IN TEXAS WITH
DAVY CROCKETT
THE BUCKSKIN BOOKS
By JOHN T. Mc INTYRE
BOOX

By Edzer Huitema & Ian French

# E Ink's Technicolor Moment

## THE ROAD TO COLOR E-PAPER TOOK TWO DECADES

*Photo by Randi Klett*

**It was the end of 2008, October, right before the holiday shopping season. Talk-show host Oprah Winfrey released her highly anticipated Favorite Things list, with the Amazon Kindle topping the gadget category. • This is the moment that the concept of electronic paper, or e-paper, went mainstream. • But this black-and-white, reflective display that always appeared to be on was invented well before the Amazon Kindle made it famous. Its story began a decade earlier, in 1997, at the MIT Media Lab, when it was created by two students, J.D. Albert and Barrett Comiskey, who were inspired by their professor Joseph Jacobson.**

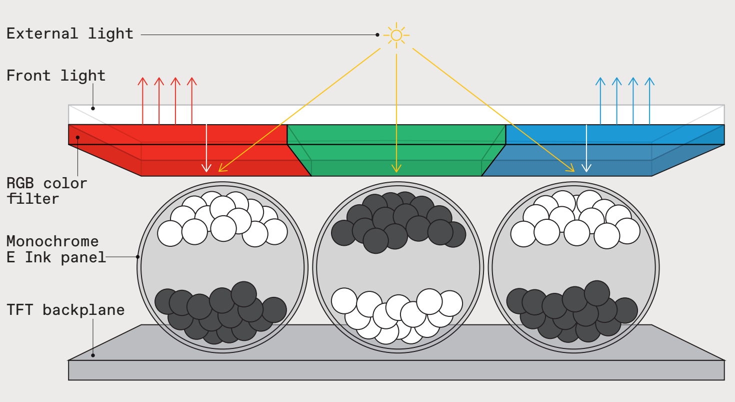


In E Ink's Triton and Kaleido displays, color filters turn light reflected from white particles into red, green, and blue subpixels. This approach, however, reduces resolution and brightness, limiting the popularity of the first generation of the technology.

From the very beginning, e-paper seemed magical. It was easy on the eyes, even outdoors and in bright sunlight, where other portable displays became unreadable. It could go weeks between charges while mobiles equipped with other displays barely made it through a day (some of them still barely make it through a day). Yet its limitation was obvious—images could appear only in black and white. In a world that hadn't seen a monochrome display in a very long time—TVs made the switch in the 1960s, computer monitors in the late '80s—a monochrome display was definitely quaintly old school.

So, since the initial development of electronic ink, as the basic technology behind e-paper is known, and even more with the release of the Kindle, a big question hung over e-paper: When would we see this magical display in brilliant, blazing color?

It's not that people hadn't been trying. Electronic ink researchers had been pursuing color e-paper for years, as had other researchers around the world, in universities, corporate research labs, and startups. They came up with some early products that targeted shelf labels for brick-and-mortar retail stores and also for signage. But these added just one color to a black-and-white screen—red or yellow—and that wasn't anybody's idea of a full-color display. Indeed, more than a decade after that first Kindle, and more than two decades after e-paper was invented, full-color e-paper had still not reached the consumer market.

Why did it take so long for e-paper to make that Wizard-of-Oz transition from black and white to color? Over the years, researchers tried several approaches, some taking technologies from more traditional displays, others evolving from the original e-paper's unique design. Qualcomm, for example, spent billions pursuing an approach inspired by butterfly wings. Overall, the path to successful color e-paper is a classic, if tortuous, tale of tech triumph. Read on to find out why this seemingly straightforward challenge was only realized just two years ago at E Ink, where we are chief technical officers.

Today, E Ink's full-color ePaper is in consumer hands, in products including e-readers and smartphones and note-taking devices, and from roughly a dozen manufacturers. These include the Guoyue Smartbook V5 Color, the HiSense A5C Color Smartphone, the Onyx Boox Poke 2 Color, and the PocketBook Color. Only one other full-color electronic paper product has been announced—DES (Display Electronic Slurry) from China's Dalian Good Display. At this writing, no devices using DES have shipped to consumers, though a handful of journalists have received samples and two Kickstarter campaigns feature products designed to use the display.



*Illustration by James Provost*

**The challenge stemmed** from the nature of the technology. Black-and-white electronic ink is a straightforward fusion of chemistry, physics, and electronics, and it does pretty much what traditional ink and paper does. E Ink's version is made of microcapsules of negatively charged black particles and positively charged white particles—the same pigments used in the printing industry today—floating in clear liquid. Each microcapsule is about the width of a human hair.

To manufacture our ePaper display, we start by making batches of this electronic ink, then use it to coat a plastic substrate some 25 to 100 micrometers thick, depending on which product it's intended for. We then cut the rolls of coated film into the desired display size and add thin-film transistors to create electrodes above and below the ink layer, which is sandwiched between protective sheets, and, possibly, touch panels or front lights.

To produce an image, an ePaper device applies different voltages to the top and bottom electrodes to create an electric field. At the top, the voltage is close to zero, and at the bottom it alternates among –15, 0, or 15. Every time the image on the screen needs to change, a specific sequence of voltages applied to the bottom electrode moves the particles from their previous position to the position needed to show the correct color for the new image. This update time typically takes less than half a second.

Bringing white particles to the top of the display creates the appearance of "paper"; black ones create "ink." But the particles don't have to sit at the very top or very bottom; when we stop generating that electric field, the particles stop in their tracks. This means we can create a mixture of black-and-white particles near the top of the display—appearing as shades of gray.

The software that determines the timing and the voltages applied to each electrode is complex. The choices depend on what was previously displayed at that pixel. If a black pixel in one image will be black again in the next image, for example, no voltage needs to be applied at that spot. We also have to be careful with the transitions; we don't want a previous image to linger, yet we don't want an abrupt change to cause the screen to flash. These are but a few of the factors we took into consideration when designing the algorithms, called waveforms, that we use to set the sequence of voltages. Designing them is as much art as science.

**To bring color into the equation greatly compli**cates the waveforms. Black and white is a simple dichotomy, given that an electric field can create either a positive or a negative charge. That approach can't accommodate full-color digital paper. We needed something entirely new.

The PocketBook Inkpad 3 Pro, introduced in 2021, uses the second-generation E Ink Kaleido Plus color display.

POCKETBOOK

We started exploring options in the early 2000s. One of our first commercially launched color products, in 2010, used a color filter—an array of squares printed onto a layer of glass placed on top of the standard black-and-white ink layer. When we applied a charge to move the white particles to the surface at a selected spot, the light would bounce back to the viewer through the red, green, or blue filter above it. It was an obvious approach: All of the colors visible to humans can be created with combinations of red, green, and blue light, which is why most of today's most common display technologies, like LCDs and OLEDs, use RGB emitters or color filters.

We called our product E Ink Triton. While an electronic textbook did launch with the technology, the main thing this effort taught us was what would *not* work for the consumer market. Its resolution was simply too low and the colors not bright enough for people who were used to the high resolution of tablet computers or print magazines.

The brightness problem stemmed from the fact that unlike LCDs and OLEDs, which, respectively, use a backlight or emit light directly, E Ink's displays are fully reflective. That is, light from an outside source goes through the transparent cover, hits the ink layer, and bounces back to the viewer's eyes. This arrangement is great for outdoor use, because reflective displays are enhanced rather than washed out by bright sunlight. And the displays are good for eye comfort, because they don't shine light directly at a user. But with a reflective system, every layer between the ink and eye absorbs or scatters some of the light. Adding that color filter layer, it turned out, caused significant dimming.

In addition, using a color filter to split monochrome pixels into three colored pixels reduced the overall resolution. A display originally having a resolution of 300 pixels per inch, with an addition of a three-color filter, now has a resolution of 100 pixels per inch. This was not as much of an issue for a 32-inch display used as a sign—pixel sizes could be larger, and big letters don't require high resolution. But it was a real problem for small fonts and line drawings on handheld devices.

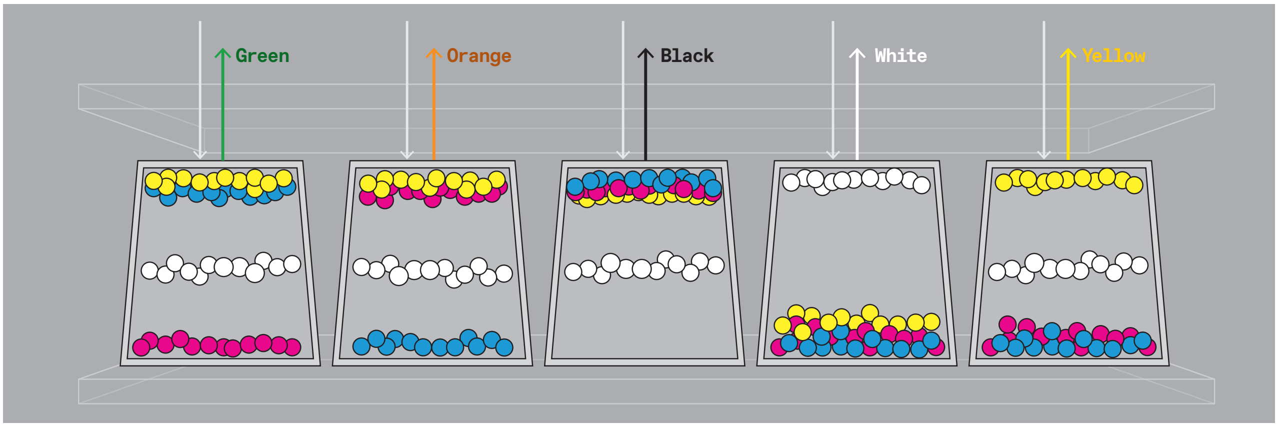


E Ink's Advanced Color ePaper (ACeP) uses four different types of pigment particles, varying in size and charge. The system applies varying electric fields to push and pull them to different positions in each trapezoidal Microcup to create the desired colors.

While our researchers were coming up with this filtered display, others in our labs focused on a different approach, called multipigment, that didn't rely on color filters. However, that approach requires far more complicated chemistry and mechanics.

**Multipigment e-paper** also shares fundamentals with its monochrome predecessors. However, instead of only two types of particles, there are now three or four, depending on the colors chosen for a particular application.

We needed to get these particles to respond uniquely to electric fields, not simply be attracted or repelled. We did a few things to our ink particles to allow them to be better sorted. We made the particles different sizes—larger particles will generally move more slowly in liquid than smaller ones. We varied the charges of the particles, taking advantage of the fact that charge is more analog than digital. That is, it can be very positive, a little positive, very negative, or a little negative. And a lot of gradations in between.

Once we had our particles differentiated, we had to adapt our waveforms; instead of just sending one set of particles to the top as another goes to the bottom, we both push and pull them to create an image. For example, we can push particles of one color to the top, then pull them back a little so they mix with other particles to create a specific shade. Cyan and yellow together, for example, produce green, with white particles providing a reflective background. The closer a particle is to the surface, the greater the intensity of that color in the mix.

We also changed the shape of our container, from a sphere to a trapezoid, which gave us better control over the vertical position of the particles. We call these containers Microcups.

For the three-particle system, now on the market as E Ink Spectra and used primarily in electronic shelf labels (ESLs), we put black, white, and red or black, white, and yellow pigments into each Microcup. In 2021, we added a fourth particle to this system; our new generation uses black, white, red, and yellow particles. These are great for generating deeply saturated colors with high contrast, but these four colors cannot be combined to create full-color images. This technology was first launched in 2013 for retail ESLs. Companies have built E Ink screens into millions of these tags, shipping them throughout the world to retailers such as Best Buy, Macy's, and Walmart. Similar electrophoretic shelf labels that use displays from China's DKE Co. have since come on the market.

For our true, full-color system, which we call Advanced Color ePaper (ACeP), we also use four particles, but we have dropped the black and rely on white—our paper—along with cyan, magenta, and yellow, the colors used in inkjet printers. By stopping the particles at different levels, we can use these particles to create up to 50,000 colors. The resulting display renders colors like those in newspapers or even watercolor art.

E Ink launched ACeP as E Ink Gallery in 2016. Again, it wasn't appropriate for consumer devices, because of slow refresh rates. Also, as it's a reflective display without a backlight, the colors were too muted for consumers accustomed to bright smartphone and tablet displays. For now, it has been geared predominantly toward use in retail signs in Asia.

**Realizing we still** weren't hitting the consumer-market sweet spot with our color displays, our R&D team went back to take another look at Triton, the system that used RGB color filters. What worked and what didn't? Were there modifications we could make to finally produce a color e-reader that consumers would want?

We knew the filters were sapping brightness. We were pretty sure we could significantly reduce this loss by getting the filters closer to the electronic ink.

We also wanted to increase the resolution of the displays, which meant a much finer color-filter array. To get a resolution more in line with what consum-



*Illustrations by James Provost*

ers are accustomed to, we had to shoot for at least 200 pixels per square inch. That's about twice the density we were able to achieve with our first round of Triton displays.

Compared with the complexity of formulating inks with a variety of charges, as we had done in developing ACeP, you might think this would have been easy. But it ended up requiring a new technology to print the color filters on the glass substrate.

We had created our earlier filters by printing semi-transparent red, green, and blue ink on glass. But this glass was an added layer. So we decided to print directly onto the plastic film that holds the top electrode, adding this step after our display modules were nearing the end of the assembly process. This arrangement would get the filters as close to the electronic ink as possible. It would also allow us to increase resolution, because aligning the filters with the display pixels could be done more precisely than was possible when using a separate surface.

We found the type of printer we needed at the German company Plastic Logic, a partner of E Ink since the early days of the company. But this printer was intended for use in an R&D lab, not for high-volume production. The processes it used had to be converted to operate in a different, production-ready machine.

We also needed to figure out new printing patterns for the color filter. These are the actual shapes and arrangements of the red, blue, and green filters. We had found through working on Triton that printing the filters as a simple square grid was not the best option, as the pattern could be visible during certain image transitions. And so the hunt for the perfect pattern was on. We went through many iterations, considering the angle at which light hit the display, as this angle could easily shift the color seen by the user. We evaluated a grid, straight printed lines, long lines, and a host of other designs, and settled on a pattern of short lines.

Because this is a reflective display, the more light hitting the display, the brighter it is. The research team decided to add a front light to the display, something that was not part of Triton, working hard to ensure that the light rays hit the ink layer at an angle that maximizes reflectivity. Using a front light increases energy consumption, of course, but it's worth it in this case.

As a result, E Ink's new color technology, E Ink Kaleido, has significantly more saturated colors and a better contrast ratio than E Ink Triton. And finally, a full-color electronic-ink display was ready for use in consumer products.

The first official batch of Kaleido displays rolled off the manufacturing line in late 2019. We began shipping to customers soon after, and you can now see the technology in products like the Hisense A5C, the iFlytek Book C1, and the PocketBook Color, all of which were launched in 2020. A second generation of Kaleido, called Kaleido Plus, began shipping in early 2021, with products released by Onyx and PocketBook and more launching soon. This update improved color saturation thanks to adjustments made in the printing pattern and the light guides for the front light.

**We have a few things** to work on. Light efficiency, the fraction of incoming light that makes its way back out to the user's eyes, is good but it could be better. We are continuing to work on our film layers to further cut this loss.

By continuing to refine our printing pattern, we are also working to improve resolution by using denser circuitry in the electronics that sit below the ink layer and turn voltages on and off to move the charged particles.

We are also continuing to work on our filterless, multipigment electronic-ink technology. We expect to release a new generation for use in signage soon, and it will include brighter colors and faster page updates. Someday we might even be able to move this into consumer devices.

When E Ink's researchers set out exploring color electronic ink in the early 2000s, they thought it would be a matter of a few years to fruition, given our expertise with the technology. After all, black-and-white e-paper took only 10 years from concept to commercialization. The road to full color turned out to be much longer. But, just like Dorothy in *The Wizard of Oz,* we finally made it over the rainbow. ■

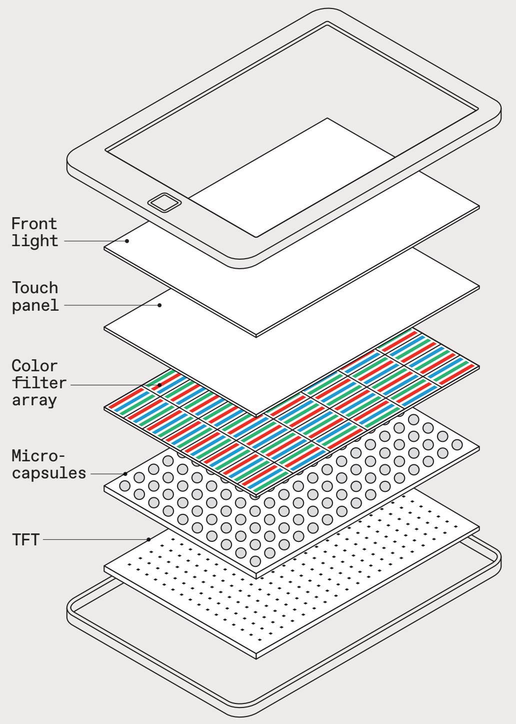


For its Kaleido color display, E Ink included a front light and patterned the color filters as a series of short lines to improve brightness, color saturation, and contrast.

Seeing around the corner is simulated by modeling an autonomous vehicle approaching an urban intersection with four high-rise concrete buildings at the corners. A second vehicle is approaching the center via a crossing road, out of the AV's line of sight, but it can be detected nonetheless through the processing of signals that return either by reflecting along multiple paths or by passing directly through the buildings.

Using several bands of radar at once can give cars a kind of second sight

# LETTING ROBOCARS SEE AROUND CORNERS

An autonomous car needs to do many things to make the grade, but without a doubt, sensing and understanding its environment are the most critical. A self-driving vehicle must track and identify many objects and targets, whether they're in clear view or hidden, whether the weather is fair or foul.

Today's radar alone is nowhere near good enough to handle the entire job—cameras and lidars are also needed. But if we could make the most of radar's particular strengths, we might dispense with at least some of those supplementary sensors.

BY BEHROOZ REZVANI, BABAK HASSIBI, FREDRIK BRÄNNSTRÖM & MAJID MANTEGHI

Illustration by Chris Philpot

Multipath reflections and through-building transmission allow the autonomous vehicle [red circle, at right in each diagram] to begin detecting the second vehicle [red rectangle, bottom of each diagram] around the 0.45-second mark, at which time the second vehicle remains firmly occluded by the bottom-left building. Because both frequency bands produce “ghost targets” [blue circles] due to reflections and multiple paths, the system employs a Bayesian algorithm to determine the true targets and remove the ghosts. The algorithm uses a combination of ray tracing and fusion of the results over time across both the UHF and C bands.

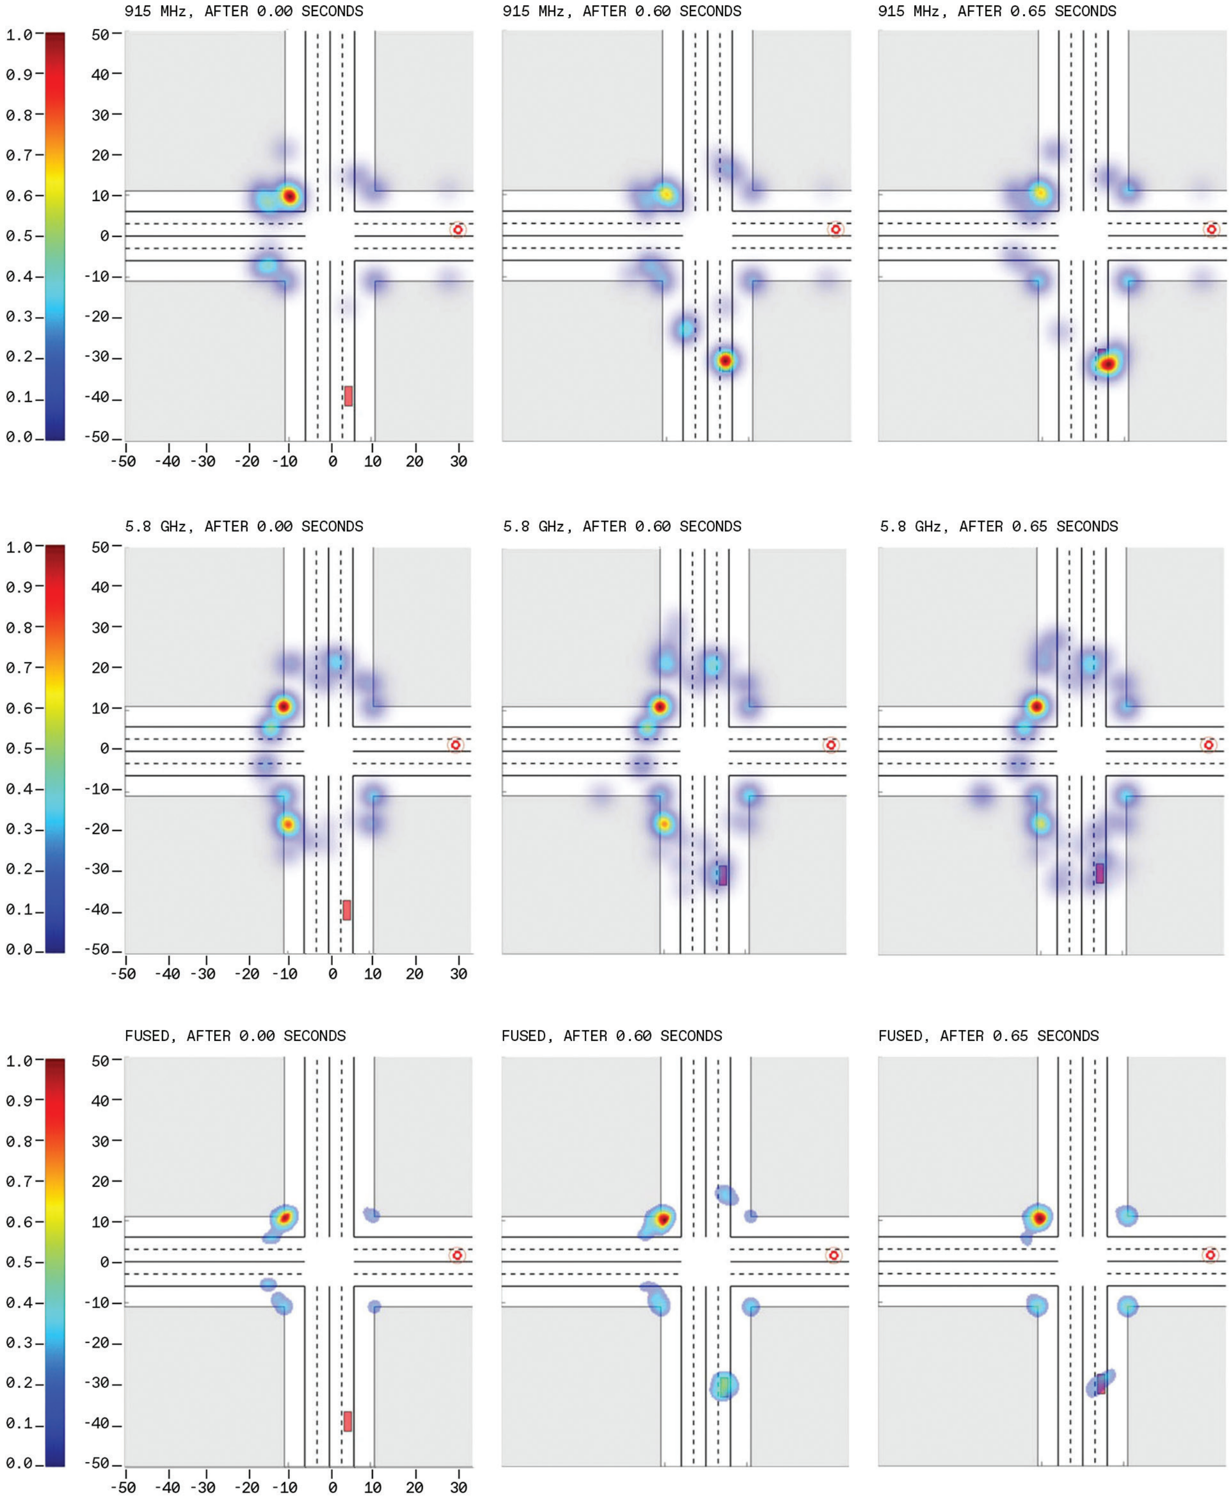


NEURAL PROPULSION SYSTEMS

Conventional cameras in stereo mode can indeed detect objects, gauge their distance, and estimate their speeds, but they don't have the accuracy required for fully autonomous driving. In addition, cameras do not work well at night, in fog, or in direct sunlight, and systems that use them are prone to being fooled by optical illusions. Laser scanning systems, or lidars, do supply their own illumination and thus are often superior to cameras in bad weather. Nonetheless, they can see only straight ahead, along a clear line of sight, and will therefore not be able to detect a car approaching an intersection while hidden from view by buildings or other obstacles.

Radar is worse than lidar in range accuracy and angular resolution—the smallest angle of arrival between two distinct targets that's needed to resolve one from another. But we have devised a novel radar architecture that overcomes these deficiencies, making it much more effective in augmenting lidars and cameras.

Our proposed architecture employs what's called a sparse, wide-aperture multiband radar. The basic idea is to use a variety of frequencies, exploiting the particular properties of each one, to free the system from the vicissitudes of the weather and to see through and around corners. That system, in turn, employs advanced signal processing and sensor-fusion algorithms to produce an integrated representation of the environment.

We have experimentally verified the theoretical performance limits of our radar system—its range, angular resolution, and accuracy. Right now, we're building hardware for various automakers to evaluate, and recent road tests have been successful. We plan to conduct more-elaborate tests to demonstrate around-the-corner sensing in early 2022.

RAIN ATTENUATION AGAINST RAIN RATE FOR RADAR AND LIDAR

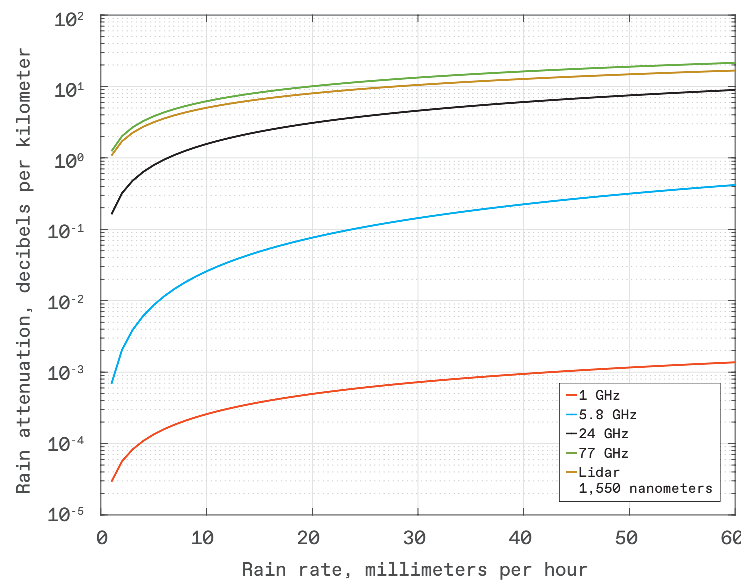


Seeing in the rain is generally much easier for radar than for light-based sensors, notably lidar. At relatively low frequencies, a radar signal's loss of strength is orders of magnitude lower.

**Each frequency band** has its strengths and weaknesses. The band at 77 gigahertz and below can pass through 1,000 meters of dense fog without losing more than a fraction of a decibel of signal strength. Contrast that with lidars and cameras, which lose 10 to 15 decibels in just 50 meters of such fog.

Rain, however, is another story. Even light showers will attenuate 77-GHz radar as much as they would lidar. No problem, you might think—just go to lower frequencies. Rain is, after all, transparent to radar at, say, 1 GHz or below.

This works, but you want the high bands as well, because the low bands provide poorer range and angular resolution. Although you can't necessarily equate high frequency with a narrow beam, you can use an antenna array, or highly directive antenna, to project the millimeter-long waves in the higher bands in a narrow beam, like a laser. This means that this radar can compete with lidar systems, although it would still suffer from the same inability to see outside a line of sight.

For an antenna of given size—that is, of a given array aperture—the angular resolution of the beam is inversely proportional to the frequency of operation. Similarly, to achieve a given angular resolution, the required frequency is inversely proportional to the antenna size. So to achieve some desired angular resolution from a radar system at relatively low UHF frequencies (0.3 to 1 GHz), for example, you'd need an antenna array tens of times as large as the one you'd need for a radar operating in the K (18- to 27-GHz) or W (75- to 110-GHz) bands.

Even though lower frequencies don't help much with resolution, they bring other advantages. Electromagnetic waves tend to diffract at sharp edges; when they encounter curved surfaces, they can diffract right around them as "creeping" waves. These effects are too weak to be effective at the higher frequencies of the K band and, especially, the W band, but they can be substantial in the UHF and C (4- to 8-GHz) bands. This diffraction behavior, together with lower penetration loss, allows such radars to detect objects around a corner.

One weakness of radar is that it follows many paths, bouncing off innumerable objects, on its way to and from the object being tracked. These radar returns are further complicated by the presence of many other automotive radars on the road. But the tangle also brings a strength: The widely ranging ricochets can provide a computer with information about what's going on in places that a beam projected along the line of sight can't reach—for instance, revealing cross traffic that is obscured from direct detection.

To see far and in detail—to see sideways and even directly through obstacles—is a promise that radar has not yet fully realized. No one radar band can do it all, but a system that can operate simultaneously at multiple frequency bands can come

NEURAL PROPULSION SYSTEMS

The truck and the car are fitted with wide-aperture multi-band radar from Neural Propulsion Systems, the authors' company. Note the very wide antenna above the windshield of the truck.

NEURAL PROPULSION SYSTEMS

very close. For instance, high-frequency bands, such as K and W, can provide high resolution and can accurately estimate the location and speed of targets. But they can't penetrate the walls of buildings or see around corners; what's more, they are vulnerable to heavy rain, fog, and dust.

Lower frequency bands, such as UHF and C, are much less vulnerable to these problems, but they require larger antenna elements and have less available bandwidth, which reduces range resolution—the ability to distinguish two objects of similar bearing but different ranges. These lower bands also require a large aperture for a given angular resolution. By putting together these disparate bands, we can balance the vulnerabilities of one band with the strengths of the others.

Different targets pose different challenges for our multiband solution. The front of a car presents a smaller radar cross section—or effective reflectivity—to the UHF band than to the C and K bands. This means that an approaching car will be easier to detect using the C and K bands. Further, a pedestrian's cross section exhibits much less variation with respect to changes in his or her orientation and gait in the UHF band than it does in the C and K bands. This means that people will be easier to detect with UHF radar.

Furthermore, the radar cross section of an object decreases when there is water on the scatterer's surface. This diminishes the radar reflections measured in the C and K bands, although this phenomenon does not notably affect UHF radars.

Another important difference arises from the fact that a signal of a lower frequency can penetrate walls and pass through buildings, whereas higher frequencies cannot. Consider, for example, a 30-centimeter-thick concrete wall. The ability of a radar wave to pass through the wall, rather than reflect off of it, is a function of the wavelength, the polarization of the incident field, and the angle of incidence. For the UHF band, the transmission coefficient is around –6.5 dB over a large range of incident angles. For the C and K bands, that value falls to –35 dB and –150 dB, respectively, meaning that very little energy can make it through.

**The tangled return paths of radar are also a strength because they can provide a computer with information about what's going on sideways—for instance, in cross traffic that is obscured from direct inspection.**

**A radar's angular** resolution, as we noted earlier, is proportional to the wavelength used; but it is also inversely proportional to the width of the aperture—or, for a linear array of antennas, to the physical length of the array. This is one reason why millimeter waves, such as the W and K bands, may work well for autonomous driving. A commercial radar unit based on two 77-GHz transceivers, with an aperture of 6 cm, gives you about 2.5 degrees of angular resolution, more than an order of magnitude worse than a typical lidar system, and too little for autonomous driving. Achieving lidar-standard resolution at 77 GHz requires a much wider aperture—1.2 meters, say, about the width of a car.

Besides range and angular resolution, a car's radar system must also keep track of a lot of targets, sometimes hundreds of them at once. It can be difficult to distinguish targets by range if their range to the car varies by just a few meters. And for any given range, a uniform linear array—one whose transmitting and receiving elements are spaced equidistantly—can distinguish only as many targets as the number of antennas it has. In cluttered environments where there may be a multitude of targets, this might seem to indicate the need for hundreds of such transmitters and receivers, a problem made worse by the need for a very large aperture. That much hardware would be costly.

One way to circumvent the problem is to use an array in which the elements are placed at only a few of the positions they normally occupy. If we design such a "sparse" array carefully, so that each mutual geometrical distance is unique, we can make it behave as well as the nonsparse, full-size array. For instance, if we begin with a 1.2-meter-aperture radar operating at the K band and put in an appropriately designed sparse array having just 12 transmitting and 16 receiving elements, it would behave like a standard array having 192 elements. The reason is that a carefully designed sparse array can have up to 12 × 16, or 192, pairwise distances between each transmitter and receiver. Using 12 different signal transmissions, the 16 receive


antennas will receive 192 signals. Because of the unique pairwise distance between each transmit/receive pair, the resulting 192 received signals can be made to behave as if they were received by a 192-element, nonsparse array. Thus, a sparse array allows one to trade off time for space—that is, signal transmissions with antenna elements.

**In principle,** separate radar units placed along an imaginary array on a car should operate as a single phased-array unit of larger aperture. However, this scheme would require the joint transmission of every transmit antenna of the separate subarrays, as well as the joint processing of the data collected by every antenna element of the combined subarrays, which in turn would require that the phases of all subarray units be perfectly synchronized.

None of this is easy. But even if it could be implemented, the performance of such a perfectly synchronized distributed radar would still fall well short of that of a carefully designed, fully integrated, wide-aperture sparse array.

Consider two radar systems at 77 GHz, each with an aperture length of 1.2 meters and with 12 transmit and 16 receive elements. The first is a carefully designed sparse array; the second places two 14-element standard arrays on the extreme sides of the aperture. Both systems have the same aperture and the same number of antenna elements. But while the integrated sparse design performs equally well no matter where it scans, the divided version has trouble looking straight ahead, from the front of the array. That's because the two clumps of antennas are widely separated, producing a blind spot in the center.

In the widely separated scenario, we assume two cases. In the first, the two standard radar arrays at either end of a divided system are somehow perfectly synchronized. This arrangement fails to detect objects 45 percent of the time. In the second case, we assume that each array operates independently and that the objects they've each independently detected are then fused. This arrangement fails almost 60 percent of the time. In contrast, the carefully designed sparse array has only a negligible chance of failure.

**Seeing around the corner** can be depicted easily in simulations. We considered an autonomous vehicle, equipped with our system, approaching an urban intersection with four high-rise concrete buildings, one at each corner. At the beginning of the simulation the vehicle is 35 meters from the center of the intersection and a second vehicle is approaching the center via a crossing road. The approaching vehicle is not within the autonomous vehicle's line of sight and so cannot be detected without a means of seeing around the corner.

At each of the three frequency bands, the radar system can estimate the range and bearing of the targets that are within the line of sight. In that case, the range of the target is equal to the speed of light multiplied by half the time it takes the transmitted electromagnetic wave to return to the radar. The bearing of a target is determined from the incident angle of the wavefronts received at the radar. But when the targets are not within the line of sight and the signals return along multiple routes, these methods cannot directly measure either the range or the position of the target.

We can, however, infer the range and position of targets. First we need to distinguish between line-of-sight, multipath, and through-the-building returns. For a given range, multipath returns are typically weaker (due to multiple reflections) and have different polarization. Through-the-building returns are also weaker. If we know the basic environment—the position of buildings and other stationary objects—we can construct a framework to find the possible positions of the true target. We then use that framework to estimate how likely it is that the target is at this or that position.

As the autonomous vehicle and the various targets move and as more data is collected by the radar, each new piece of evidence is used to update the probabilities. This is Bayesian logic, familiar from its use in medical diagnosis. Does the patient have a fever? If so, is there a rash? Here, each time the car's system updates the estimate, it narrows the range of possibilities until at last the true target positions are revealed and the "ghost targets" vanish. The performance of the system can be significantly enhanced by fusing information obtained from multiple bands.

We have used experiments and numerical simulations to evaluate the theoretical performance limits of our radar system under various operating conditions. Road tests confirm that the radar can detect signals coming through occlusions. In the coming months we plan to demonstrate round-the-corner sensing. We expect that these features will enable a form of driving safer than we have ever known. ■

Illustration by Dan Page

The rules of the Internet can also balance electricity supply and demand

# Packetizing the Power Grid

By Mads Almassalkhi, Jeff Frolik & Paul Hines

**Bad things happen** when demand outstrips supply. We learned that lesson too well at the start of the pandemic, when demand for toilet paper, disinfecting wipes, masks, and ventilators outstripped the available supply. Today, chip shortages continue to disrupt the consumer electronics, automobile, and other sectors. Clearly, balancing the supply and demand of goods is critical for a stable, normal, functional society.

That need for balance is true of electric power grids, too. We got a heartrending reminder of this fact in February 2021, when Texas experienced an unprecedented and deadly winter freeze. Spiking demand for electric heat collided with supply problems created by frozen natural-gas equipment and below-average wind-power production. The resulting imbalance left more than 2 million households without power for days, caused at least 210 deaths, and led to economic losses of up to US $130 billion.

Similar mismatches in supply and demand contributed to massive cascading blackouts in August 2003 in the northeastern United States and Canada, in July 2012 in India, and in March 2019 in Venezuela.

The situation is unlikely to get better anytime soon, for three reasons. First, as countries everywhere move to decarbonize, the electrification of transportation, heating, and other sectors will cause electricity demand to soar. Second, conventional coal and nuclear plants are being retired for economic and policy reasons, removing stable sources from the grid. And third, while wind and solar-photovoltaic systems are great for the climate and the fastest-growing sources of electric generation, the variability of their output begets new challenges for balancing the grid.

So how can grid operators keep supply and demand balanced, even as they shut down old, dirty power plants, ramp up variable generation, and add new electric loads? There are a few possibilities. One is

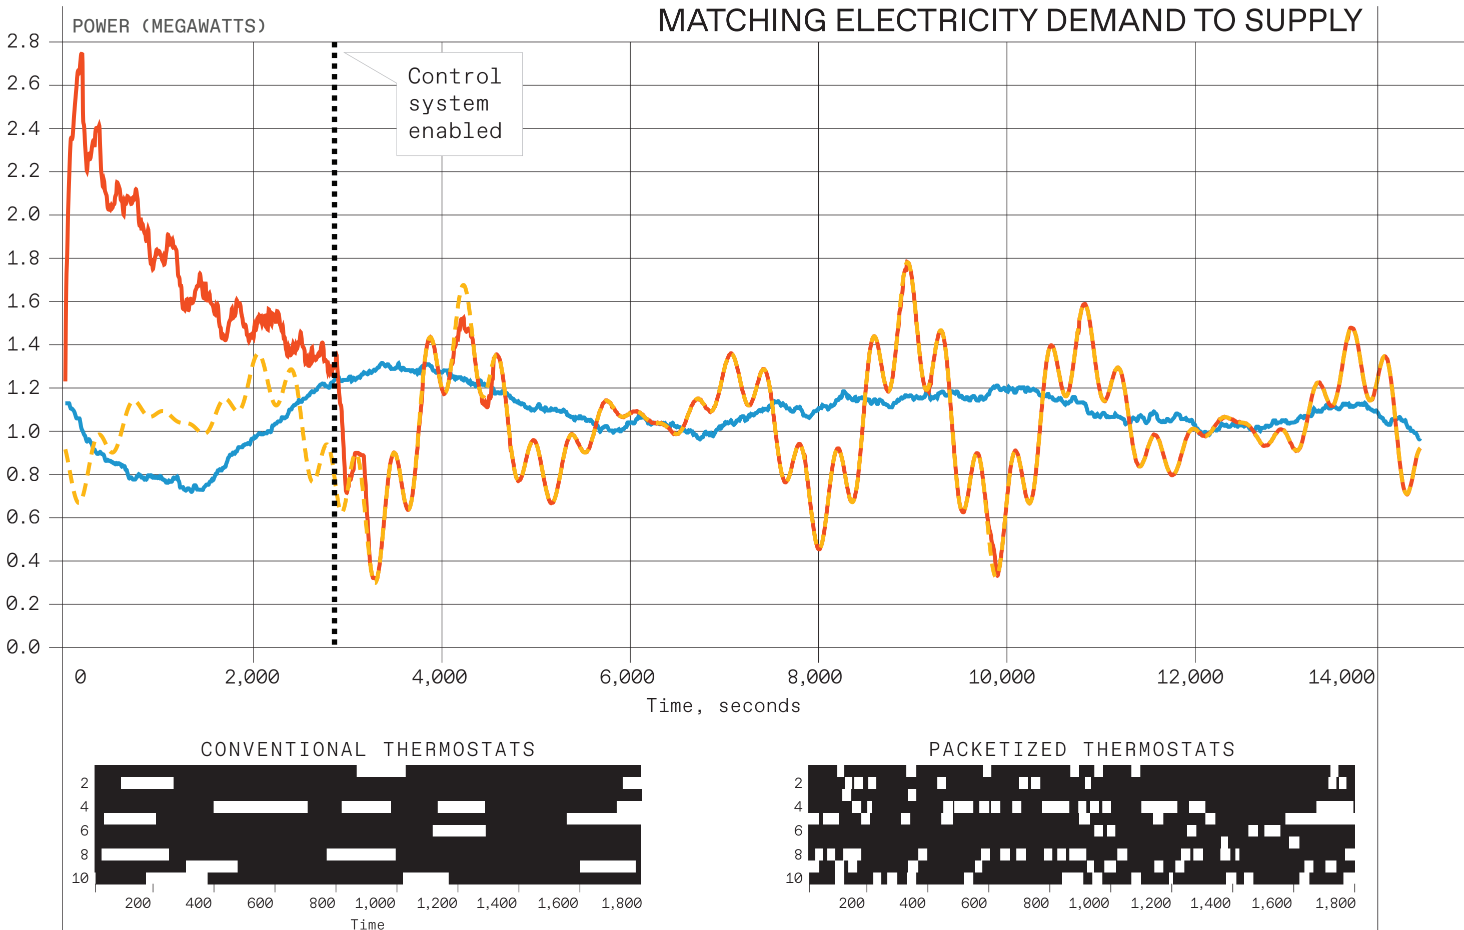


Packetized energy management (PEM) allows the power grid to flexibly handle a varying supply of renewable energy. In a simulation, the aggregated load from 1,000 electric water heaters [solid orange line] almost exactly matches renewable energy supply [dashed gold line] after packetized control is switched on [vertical dotted line]. The blue line shows the uncoordinated load from 1,000 electric water heaters for reference. The black-and-white images show patterns of off and on for a set of 10 water heaters controlled by conventional thermostats and by PEM-enabled thermostats.

SOURCE: MADS ALMASSALKHI, JEFF FROLIK & PAUL HINES

to do a modernized version of what we have done in the past: Build giant, centralized infrastructure. That would mean installing vast amounts of energy storage, such as grid-scale batteries and pumped-hydro facilities, to hold the excess renewable power being generated, and interconnecting that storage with high-voltage transmission lines, so that supply can meet demand across the grid. China is a leader in this approach, but it's incredibly expensive and requires an enormous amount of political will.

We think there's a better way. Instead of drastically scaling up power-grid infrastructure, our work at the University of Vermont has focused on how to coordinate demand in real time to match the increasingly variable supply. Our technology takes two ideas that make the Internet fundamentally scalable—packetization and randomization—and uses them to create a system that can coordinate distributed energy. Those two data-communication concepts allow millions of users and billions of devices to connect to the Internet without any centralized scheduling or control. The same basic ideas could work on the electrical grid, too. Using low-bandwidth connectivity and small controllers running simple algorithms, millions of electrical devices could be used to balance the flow of electricity in the local grid. Here's how.

**Electricity demand on** the grid comes from billions of electrical loads. These can be grouped into two broad categories: commercial and industrial loads, and residential loads. Of the two, residential loads are far more dispersed. In the United States alone, there are over 120 million households, which collectively account for about 40 percent of annual electricity consumption. But residential customers generally don't think about optimizing their own electricity loads as they go about their day. For simplicity's sake, let's call these residential loads "devices," which can range from lights and televisions to water heaters and air conditioners.

The latter devices, along with electric-vehicle chargers and pool pumps, are not only large electric loads (that is, greater than a 1-kilowatt rating), but they're also flexible. Unlike lighting or a TV, which you want to go on the instant you throw the switch, a flexible device can defer consumption and operate whenever—as long as there's hot water for your shower, your pool is clean, your EV has enough charge, and the indoor temperature is comfortable.

Collectively, there is a lot of flexibility in residential electricity loads that could be used to help balance variable supply. For example, if every household in California and New York had just one device that could consume power flexibly, at any time, the power grid would have the equivalent of around 15 gigawatts of additional capacity, which is more than 10 times the amount currently available from utility-scale battery storage in these states.

Here's what flexibility means when it comes to operating, say, a residential electric water heater. While heating water, a typical unit draws about 4.5 kilowatts. Over the course of a normal day, the appliance is on about a tenth of the time, using about 10.8 kilowatt-hours. To the homeowner, the daily cost of operating the water heater is less than US $2 (assuming a rate of about 15¢ per kWh). But to the utility, the cost of electricity is highly variable, from a nominal 4¢ per kWh to over $100 per kWh during annual peak periods. Sometimes, the cost is even negative: When there is too much power available from wind or solar plants, grid operators effectively pay utilities to consume the excess.

To reduce demand during peak periods, utilities have long offered demand-response programs that allow them to turn off customers' water heaters, air conditioners, and other loads on a fixed schedule—say, 4 p.m. to 9 p.m. during the summer, when usage is historically high. If all we want to do is reduce load at such times, that approach works reasonably well.

However, if our objective is to balance the grid in real time, as renewable generation ebbs and flows unpredictably with the wind and sun, then operating devices according to a fixed schedule that's based on past behavior won't suffice. We need a more responsive approach, one that goes beyond just reducing peak demand and provides additional benefits that improve grid reliability, such as price responsiveness, renewable smoothing, and frequency regulation.

**How can grid** operators coordinate many distributed, flexible kilowatt-scale devices, each with its own specific needs and requirements, to deliver an aggregate gigawatt-scale grid resource that is responsive to a highly variable supply? In pondering this question, we found inspiration in another domain: digital communication systems.

Digital systems represent your voice, an email, or a video clip as a sequence of bits. When this data is sent across a channel, it's broken into packets. Then each packet is independently routed through the network to the intended destination. Once all of the packets have arrived, the data is reconstructed into its original form.

How is this analogous to our problem? Millions of people and billions of devices use the Internet every day. Users have their individual devices, needs, and usage patterns—which we can think of as demand—while the network itself has dynamics associated with its bandwidth—its supply, in other words. Yet, demand and supply on the Internet are matched in real time without any centralized scheduler. Likewise, billions of electrical devices, each with its own dynamics, are connecting to the power grid, whose supply is becoming, as we noted, increasingly variable.

Recognizing this similarity, we developed a technology called packetized energy management (PEM) to coordinate the energy usage of flexible devices. Coauthor Hines has a longstanding interest in power-system reliability and had been researching how transmission-line failures can lead to cascading outages and systemic blackouts. Meanwhile, Frolik, whose background is in communication systems, had been working on algorithms to dynamically coordinate data communications from wireless sensors in a way that used very little energy. Through a chance discussion, we realized our intersecting interests and began working to see how these algorithms might be applied to the problem of EV charging.

Shortly thereafter, Almassalkhi joined our department and recognized that what we were working on had greater potential. In 2015, he wrote a winning proposal to ARPA-E's NODES program—that's the U.S. Department of Energy's Advanced Research Projects Agency–Energy's Network Optimized Distributed Energy Systems program. The funding allowed us to further develop the PEM approach.

Let's return to the electric water heater. Under conventional operation, the water heater is controlled by its thermostat. The unit turns on when the water temperature hits a lower limit and operates continuously (at 4.5 kW) for 20 to 30 minutes, until the water temperature reaches an upper limit. The pair of black-and-white graphs at the bottom of "Matching Electricity Demand to Supply" shows the on and off patterns of 10 heaters—black for off and white for on.

Under PEM, each load operates independently and according to simple rules. Instead of heating only when the water temperature reaches its lower limit, a water heater will periodically request to consume a "packet" of energy, where a packet is defined as consuming power for just a short period of time—say, 5 minutes. The coordinator (in our case, a cloud-based platform) approves or denies such packet requests based on a target signal that reflects grid conditions, such

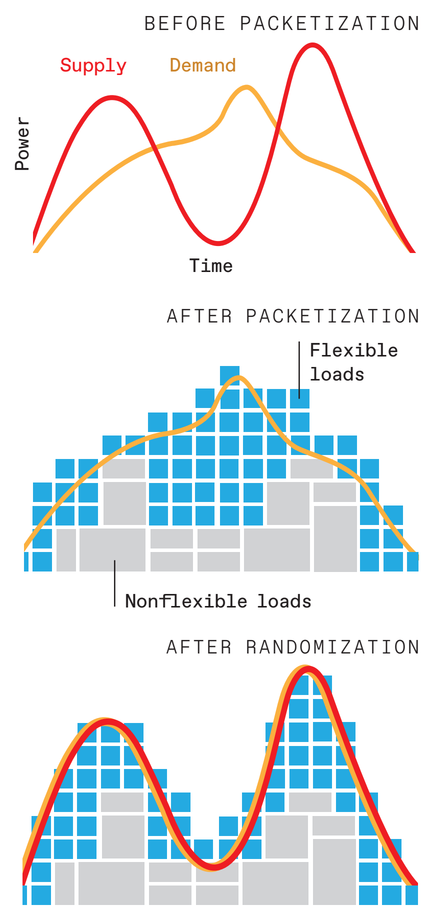


Electricity supply and demand can sometimes diverge in dramatic ways [top]. The middle and bottom graphs show how packetization and randomization of flexible electricity loads allows demand to match the available supply.

UNIVERSITY OF VERMONT

as the availability of renewable energy, the price of electricity, and so on. The top graph in "Matching Electricity Demand to Supply" shows how PEM consumption closely follows a target signal based on the supply of renewable energy.

To ensure that devices with a greater need for energy are more likely to have their requests approved, each device adjusts the rate of its requests based on its needs. When the water is less hot, a water heater requests more often. When the water is hotter, it requests less often. The system thus dynamically prioritizes devices in a fully decentralized way, as the probabilities of making packet requests are proportional to the devices' need for energy. The PEM coordinator can then focus on managing incoming packet requests to actively shape the total load from many packetized devices, without the need to centrally optimize the behavior of each device. From the customer's perspective, nothing about the water heater has changed, as these requests occur entirely in the background.

These same concepts can be applied to a wide range of energy-hungry devices. For example, an EV charger or a residential battery system can compare the battery's current state of charge to its desired value—equivalent to its need for energy—translate this into a request probability, and then send a request to the PEM coordinator, which either accepts or denies the request based on real-time grid or market conditions. Depending on those conditions, it might take somewhat longer for a battery to fully charge, but the customer shouldn't be inconvenienced.

In this way, flexible energy devices communicate using the common, simple language of energy-packet requests. As a result, the coordinator is agnostic to the type of device making the request. This device-agnostic coordination is similar to net neutrality in data communications. In general, the Internet doesn't care if your packet carries voice, video, or text data. Similarly, PEM doesn't care if the device requesting a packet is a water heater, a pool pump, or an EV charger, so it can readily coordinate a heterogeneous mix of kilowatt-scale devices.

**Right now, bottom-up,** device-driven technologies like PEM are not widely deployed. Instead, most of today's demand-response technologies take a top-down approach, in which the coordinator broadcasts a control signal to all devices, telling them what to do. But if every device is told to do the same thing at the same time, things can go wrong very quickly, as the power consumption of the devices becomes synchronized. Imagine the effect of millions of air conditioners, water heaters, and EV chargers turning on (or off) at once. That would represent gigawatt spikes—as if a large nuclear power plant were turning on or off with the flip of a switch. A spike that large could cause the grid to become unstable, which could trigger a cascading blackout. That's why most utilities today split devices into groups to limit spikes to the order of tens of megawatts. However, actively managing these different groups beyond a few annual peak events is a challenge for top-down approaches.

But if each device works to meet its own unique need for energy, then packet requests (and resulting power use) are inherently randomized, and as a result, synchronization becomes much less of a concern.

The top-down approach also makes it difficult to take into account customer preferences for hot water, charged cars, and cool homes on hot days. If we are going to coordinate energy devices to make the grid work better, we need to make sure that we do it in a way that is essentially unnoticeable and automatic for the consumer.

This controller connects to a residential electric water heater and uses simple algorithms to request "packets" of energy from a cloud-based coordinator to maintain a suitable temperature.

Now, consider how PEM accounts for an individual customer's preferences in the case of the water heater. If the water temperature drops below its lower limit and the heater isn't already consuming a packet of energy, it can temporarily "opt out" of the PEM scheme and turn on until the temperature recovers. The water heater will inform the PEM coordinator of this change in its operating mode, and the coordinator will simply update its accounting of the aggregate demand. The impact of this single load on the total is small, but for the customer, having the guarantee of hot water when needed builds trust and ensures ongoing participation.

PEM's device-driven approach also makes things easier for the coordinator because it doesn't need to centrally monitor or model each device to develop an optimized schedule. The coordinator only needs to monitor grid and market conditions, reply to the live stream of incoming packet requests, and keep a record of the "opted out" devices—the coordinator manages just three set of numbers, in other words.

**To increase the** impact of our work, we decided to commercialize PEM in parallel with our research and founded Packetized Energy in 2016. The company has deployed its cloud-based energy coordination platform in several utility-sponsored pilot projects in the United States and Canada. These projects each started by retrofitting existing electric water heaters with a smart thermostat that we designed, developed, and had UL-certified. We have also demonstrated PEM with EV chargers, residential batteries, and thermostats. Our first customer was our hometown Vermont utility, Burlington Electric Department. In 2018, BED began the nation's first 100 percent renewable-powered water heater program, which has now expanded to include EV chargers.

Our projects have yielded some promising results. "A Real-Time Demo of Load Coordination" shows how PEM coordinated the load from 208 residential water heaters in Vermont and South Carolina over a typical 2-hour period.

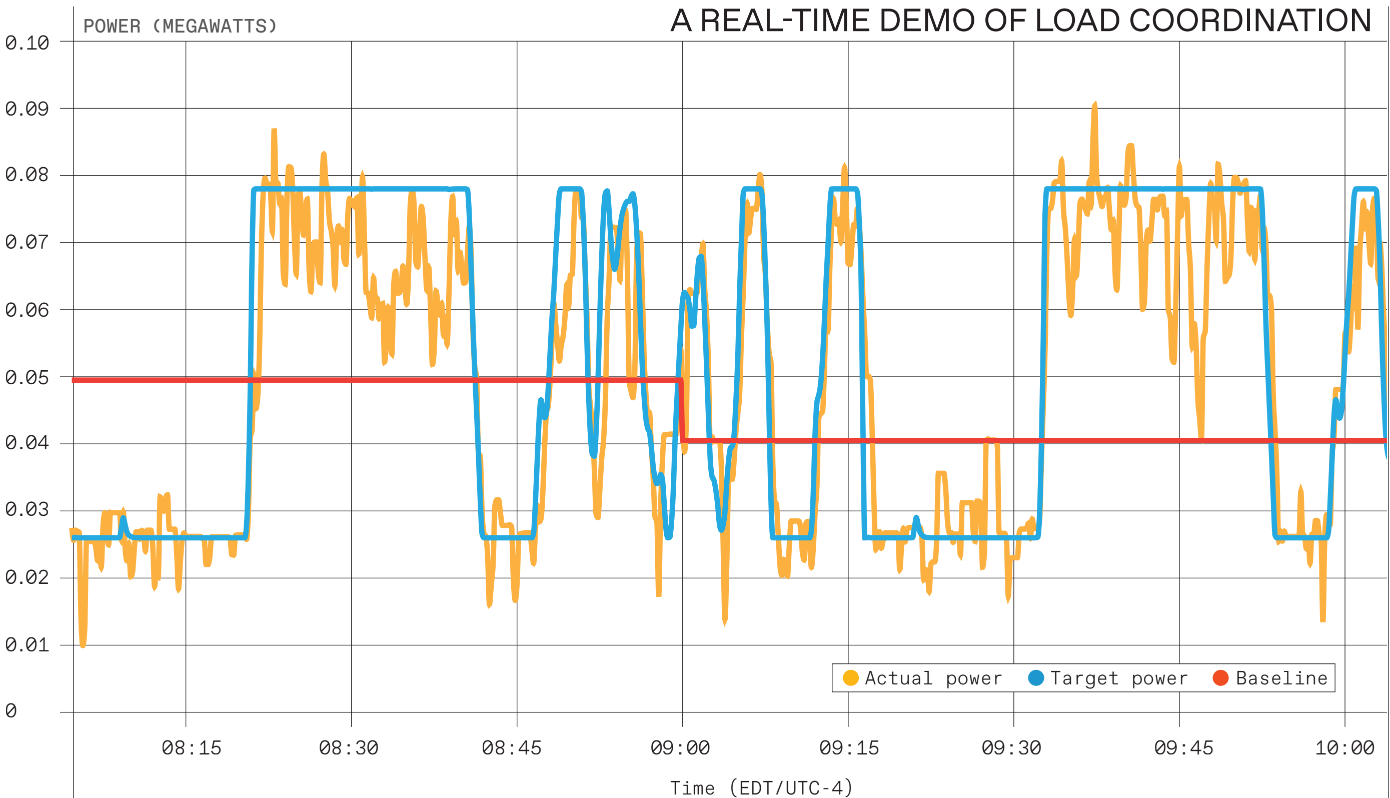


On 21 August 2021 over the course of more than 2 hours, a load of 208 residential water heaters was rapidly coordinated from 25 kilowatts to 80 kW—from about half the baseline load to about twice that load.

The heaters [orange line] followed a rapidly changing target [blue line] that ranged from about half the nominal load to about twice that load.

As systems scale to thousands of packetized devices, the asynchronous packet requests will appear as a continuous signal. Our simulations show that at this scale, any gaps between the target and the actual will disappear. The aggregate load is at least as responsive as the reaction times of a modern natural-gas power plant—and you don't have the expense of building, operating, and maintaining the physical plant.

**Falling costs for** sensors and microcontrollers are leading to the rapid growth of the Internet of Things. Combined with smart-home technology, IoT makes it possible to imagine a world in which all energy devices—loads, energy storage, and generators—are actively coordinated to keep the grid stable and take full advantage of renewable energy. But challenges do lie ahead.

First, there are few standards today to guide manufacturers interested in device-level coordination and no real incentives for them to adopt any particular approach. This has resulted in a proliferation of proprietary technologies that address the same fundamental problem. Here, again, we can draw inspiration from the Internet: Proprietary solutions are unlikely to scale up to the point of addressing the energy problems at hand. New initiatives driven by industry such as EcoPort (formerly CTA 2045) and Matter (formerly Connected Home over IP) hold promise for secure, low-latency communications with devices made by different manufacturers. IEEE technical committees, working groups, and task forces are also playing supporting roles, such as the IEEE Power and Energy Society's Smart Buildings, Loads, and Customer Systems technical committee. We hope that in the future these efforts will seamlessly support the device-driven "packetization" concepts described here, and not just serve traditional top-down communication and control architectures.

What's also needed are incentives for electricity customers to shift their energy usage. Right now, the daily cost of electricity for a residential water heater is about the same, regardless of when the heater turns on. There's no financial benefit to the homeowner to run the water heater when renewable energy supply is high or the wholesale electricity price is low. Regulators, utilities, and others will need to rethink and redesign incentives and flexible-demand programs to ensure that the contributions and rewards are fair and equitable across all customers. They will also need to educate consumers about how the program works.

There is plenty of precedent for solving such technical and policy challenges. A public system that is fair, responsive, accessible, reliable, resilient, and scalable sounds a lot like the Internet. Packetized energy management, with its core design modeled on the Internet's data communications, would deliver those same important benefits. As we transition to a new kind of grid, based on distributed and renewable generation, we'll need new technology and new paradigms. Fortunately, we have a time-tested model that is showing us the way. ■

UNIVERSITY OF VERMONT AND PACKETIZED ENERGY TECHNOLOGIES

BY ALLISON MARSH

# Past Forward

This prototype cochlear implant had a design flaw that influenced the subsequent successful commercialization of the technology.

## Electric Hearing

**Rod Saunders lost** his hearing in a car accident in 1976. Nearly two years later, Graeme Clark, an Australian ear, nose, and throat surgeon, and Brian Pyman, a surgeon assisting him, inserted into Saunders's auditory canal a prototype of a multichannel cochlear implant. The implant was designed to wirelessly connect to an external microphone and speech-processing unit. Several weeks later, Saunders came back to the hospital to determine if the implant's electrical stimulation of the auditory nerve could replicate sound that he could hear. It didn't work. A few days later, it still didn't work. Third time was the charm: After engineers fixed a loose connection in the speech-processing unit, Saunders heard a rhythm and sang along to a recording of "Waltzing Matilda." By December he was able to distinguish speech using electrical signals alone. Today hundreds of thousands of people rely on cochlear implants to hear. Even so, the technology remains controversial within the Deaf community. ■

**FOR MORE ON THE HISTORY OF COCHLEAR IMPLANTS, SEE** spectrum.ieee.org/pastforward-feb2022

RYAN HERNANDEZ/MUSEUM OF APPLIED ARTS AND SCIENCES